



*Searchable
Abstracts
Document*

SIAM Conference on
**Computational Science
and Engineering**
(CSE23)

February 7–March 3, 2023

This document was current as of February 7, 2023. Abstracts appear as submitted.



3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.

Telephone: 800-447-7426 (U.S. & Canada) +1-215-382-9800 (Worldwide)
meetings@siam.org

IP1**Model Reduction for Nonlinear Control Systems Based on (Differential) Balancing and Data**

We present the standard balancing theory for nonlinear systems, which is based on an analysis around equilibrium points. Its extension to the contraction framework offers computational advantages, and is presented as well. We provide definitions for controllability and observability functions and their differential versions which can be used for simultaneous diagonalization procedures, providing a measure for importance of the states, as can be shown by the relation to the Hankel operator. In addition, we propose a data-based model reduction method based on differential balancing for nonlinear systems whose input vector fields are constants by utilizing its variational system. The difference between controllability and reachability for the variational system is exploited for computational reasons. For a fixed state trajectory, it is possible to compute the values of the differential Gramians by using impulse and initial state responses of the variational system. Therefore, differential balanced truncation is doable along state trajectories without solving nonlinear partial differential equations.

Jacqueline M. Scherpen
Rijksuniversiteit Groningen
j.m.a.scherpen@rug.nl

IP2**Epidemic Control using Group Testing, Compressed Sensing and Machine Learning**

Testing populations using for instance Quantitative Polymerase Chain Reaction (qPCR) testing is a key component in controlling an epidemic. The testing can be significantly fastened by pooling samples together. This pool testing has its theoretical foundations in Group testing (GT). There are limitations in sampling rates and performance of algorithms in GT that the series of works stated below, attempt to improve through compressed sensing (CS) and machine learning (ML) approaches. Firstly, we propose a CS-based testing approach with a practical measurement design and a tuning-free and noise-robust algorithm for detecting infected persons. Due to nonnegativity of virus loads and an appropriate noise model, the compressed sensing problem can be solved with the non-negative least absolute deviation regression (NNLAD) algorithm. NNLAD requires the same number of tests as current state of the art GT methods and provably detects a small number of infected persons among a possibly large number of people. Moreover, CS approaches also allow to recover the viral load, but the underlying noise models are not common in CS and not well understood. Therefore, in our second improvement work, we study hybrid approaches that combine methods from CS and GT on various noise models. We compare the performance of such approaches with classical decoders from CS and GT. Our results show that combined strategies can improve the error rates and provide viral load estimation. These combined strategies exposed the strengths and limitations of each of these methods. In our final improvement attempt, we sought to exploit the strengths and overcome the limitations of these approaches by developing a decoding pipeline which terminates with a ML decoder using various noise models and their combinations. This framework allows for the ML decoder to learn a data-driven nonlinear decision boundaries that improves upon the GT and CS predictions. COVID-19 pandemic was used as a case study, but these results will apply generally to all epidemics

with similar characteristics to COVID.

Bubacarr Bah
Medical Research Council Unit The Gambia
bubacarr.bah1@lshtm.ac.uk

IP3**Cultivating a Culture of Inclusion**

Science, technology, engineering, and mathematics (STEM), being the backbone to progress and growth throughout the world, require a variety of skillsets and mindsets from diverse individuals. It is essential for organizations to emphasize inclusivity, diversity, equity, and accountability in its workforce to maximize the potential of the organization. The impact of an inclusive workforce includes improved organizational learning and development, competitive advantage, and better problem-solving capabilities, among other advantages. This presentation explores one organizations inclusion journey, strategies used, and benefits of a diverse and inclusive workforce.

Tony Baylis
Lawrence Livermore National Laboratory
baylis3@llnl.gov

IP4**Splitting Methods Revisited**

Operator-splitting techniques are useful in scientific computing, image processing, and optimization. To solve challenging problems from science and engineering, the goal is to exploit separable structures by solving in parallel individual sub-problems which are then coordinated by performing a simple algebraic step (typically, a projection onto a linear subspace). The popular ADMM and the Progressive Hedging algorithms are instances of such an approach, when applied to large-scale optimization. While parallelism is the strength of splitting methods, their weak points are the adjustment of certain proximal parameter and the lack of a fully implementable stopping test. These difficulties, which date back to Douglas-Rachford's method, stem from the very design of these approaches, which were created to find zeroes of operators. We discuss how to endow some splitting methods with an optimization perspective that introduces knowledge about the objective function throughout the iterative process. The resulting new family of methods la bundle incorporates a projective step to coordinate parallel information while allowing the variation of the proximal parameters and the construction of a reliable stopping test. A Bundle Progressive Hedging algorithm, derived from the general theory, illustrates the interest of proposal.

Claudia Sagastizabal
IMECC-Unicamp and CEMEAI
sagastiz@unicamp.br

IP5**How Open-Source is Changing Scientific Computing in Industry and Research: Lessons Learned from 25 years of Deal.II**

Open source software has had a profound impact on many areas of life and it is likely to continue to shape the way we work, live, and innovate in the future. In the Computational Science and Engineering community we have witnessed giant steps forward in the last decades, thanks

to wider and wider adoption of open-source tools and softwares. This talk will revolve around one of the most successful examples of open-source community-driven development in CSE: the successor of the Differential Equations Analysis Library – deal.II – a modular and flexible adaptive finite element library that provides tools for massively parallel multi-physics simulations. Deal.II is a 25-year-old worldwide open-source collaboration, driven by a large group of developers from all over the places, with thousands of downloads each month, and it is the basis for thousands publications. It provides researchers and engineers with a high-quality tool for solving PDEs, it facilitates collaboration, and it promotes transparency and reproducibility in research, enabling small businesses and entrepreneurs to compete with larger companies, and empowering individuals to learn, create, and innovate in ways that were previously not possible. I will introduce the library, its history, and its design principles, together with some examples of applications in research and industry, and I will discuss how the deal.II library has had a significant impact on the field of CSE.

Luca Heltai

SISSA: International School for Advanced Studies
Trieste, Italy
luca.heltai@sissa.it

IP6

AI for Science and the Implication for Mathematics

For many years, progress in scientific modeling has been hindered by the curse of dimensionality (CoD). Deep learning has provided a new effective tool for overcoming the CoD problem. This has given birth to a new paradigm, namely the AI assisted paradigm, for scientific research. In this talk, we will discuss the main issues and review some of the major progresses that have been made in AI for Science. We will also discuss some of the new opportunities that arise for applied mathematics.

Weinan E

Peking University
weinan@math.princeton.edu

IP7

The Data Behind Human Appearance and Motion: What are the MetaHumans?

Humans as individuals present an ongoing data generation systems. When we talk, when we express emotions, when we engage in expressions that influence our blood flow, when we move our eyebrows that fold the skin and form the wrinkles on our forehead, These are only illustrative activities continuously generating a comprehensive dataset that defines our appearance and motion. This dataset is ever-growing and expanding on a daily basis as we age, leading to change of our skin properties and maturing of our facial muscle structure. In this talk we will dig into which data is being captured and how it is analyzed in order to build digital models that can represent the data rich appearance and motion of individuals. We will go through principles from the domain of 3D and computer graphics, all in order to explain how vastly complex human facial simulators can be stored and driven on computers, cellphones, gaming consoles and similar memory and computationally limited hardware. We will introduce MetaHumans with their DNA file format, as a new data standard applied for representation of an individual's face. MetaHuman applications span across the entertainment industry, diverse artificial intelli-

gence use cases, all the way to the currently relevant and often confusing metaverse concepts.

Uros Sikimic

3Lateral, Serbia
uros.sikimic@3lateral.com

IP8

Fast Solution Methods for Wave Propagation Problems: From Classical Domain Decomposition Solvers to Learning

Wave propagation and scattering problems are of huge importance in many applications in science and engineering - e.g., in seismic and medical imaging and more generally in acoustics and electromagnetics. Large scale simulations of those applications are one of the hard problems from computational point of view since requires an interplay between the parsimonious but sufficiently accurate discretisation methods and more sophisticated solution methods. Our aim is to show on one side, how classical domain decomposition methods developed in the latest years coupled with carefully chosen discretisations can help in this endeavour. On the other side, we would like to propose some openings towards the very attractive package of approximation-solution-optimisation offered by the new methods in scientific machine learning.

Victorita Dolean Maini

University of Strathclyde, UK and Universite Côte dAzur
victorita.dolean@strath.ac.uk

SP1

James H. Wilkinson Prize for Numerical Software: The BLAS-Like Library Instantiation Software

The BLAS-Like Library Instantiation Software (BLIS) is a framework for instantiating high-performance dense linear algebra (DLA) operations on new CPU architectures quickly and effectively. BLIS has been ported to dozens of architectures, often as the first high-performance, open-source implementation. However, BLIS is also much more than just a BLAS implementation. Numerous intellectual and technical innovations within the BLIS framework make it possible to instantiate a far wider range of operations beyond traditional DLA, and to (re-)combine algorithmic pieces in myriad ways without a combinatorial explosion of complexity or effort. In this talk, I discuss the history and development of BLIS and the scientific achievements which led to its success, as well as some of the nuts and bolts of how BLIS does and will continue to enable such a diverse repertoire of functionality.

Devin Matthews

Southern Methodist University
damatthews@smu.edu

Field G. Van Zee

Department of Computer Sciences
The University of Texas at Austin
field@cs.utexas.edu

SP2

SIAM Activity Group on Computational Science and Engineering Best Paper Prize

To follow.

Andrew Horning

MIT
phros23@gmail.com

SP3

SIAM Activity Group on Computational Science and Engineering Early Career Prize

To follow.

Kaibo Hu
University of Oxford
kaibo.hu@maths.ox.ac.uk

SP4

2023 SIAM/ACM Prize in Computational Science and Engineering

To follow.

Carol S. Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

MS0

Mathematical Challenges at the Heart of One of Earths Lungs

The Southern Ocean is known as one of Earths lungs as it is both a massive carbon sink and an oxygen factory. If that wasnt enough, it also soaks up much of the excess heat produced by human activities. But the Southern Ocean is changing rapidly in response to warming temperatures and a key scientific question is how long it can continue to slow climate change. Numerical models are the sole means to predict future scenarios for the complex Southern Ocean environment, which is governed by rich multi-scale dynamics spanning microscale turbulent mixing to the evolution of floating ice shelves as large as countries. Breakthroughs in understanding Southern Ocean dynamics are needed to generate higher confidence in models, and hence better resilience for global communities against climate change. I will present a selection of the mathematical challenges at the heart of this grand scientific challenge for our generation.

Luke Bennetts
University of Adelaide
luke.bennetts@adelaide.edu.au

MS0

Math: The Secret Key to Unlock Solutions to Climate Change

There are many angles to how our planet experiences climate change, and the grand challenge to understand, adapt to, and mitigate its effects will require every tool in our toolbox. Thankfully, applied mathematics is being put to use to prepare Earth and its neighborhoods for challenges of a warming world. This talk will focus on how math enables us to study the flow of water around the globe, tracking individual cloud droplets, continent-sized weather events, ocean gyres, and ice sheet flow with precision. Changes in water patterns lead to deluges and droughts, so we use math to recreate and analyze ocean and atmosphere behavior using the worlds largest supercomputers.

Katherine J. Evans

Oak Ridge National Laboratory
evanskj@ornl.gov

MS0

In Control of Life

All organisms, including humans, are continuously under attack. The attacks can e.g. be intrusions by foreign pathogens, or they can result from some internal part of the living system that goes rogue and threatens the stability of the entire system. Still many of us enjoy stable and healthy existences over extended time-periods. The goal of biomedicine is to support and prolong those healthy states. I will describe how, over the last decades, detailed and massive measurements on living systems, in combination with systems biology models, have revolutionized our understanding of some fundamental life processes. Furthermore we will look at several examples of how, in biomedicine, this new knowledge has led to precise and successful interventions to support and control human health. Finally, I will discuss what the future might hold in store for us in terms of controlling and prolonging Life.

Magnus Fontes
Roche
magnus.fontes@roche.com

MS0

Mathematics: From Smoothies to the Stars

Ever wonder how blenders manage to chop fruit and smash ice into tiny particles to make the perfect smoothie? Or perhaps youve looked up at the night sky and pondered how scientists claim to know what a planet is made of, even though it is hundreds of light-years away? Both of these problems, as well as countless others, can be studied and explained using applied mathematics. Novel mathematical models allow scientists to simulate complex, physical processes that describe how our world works. In this talk, we will discuss some of the many ways that maths can be applied to further our understanding of the universe, from everyday tasks such as making a smoothie, to unlocking the secrets of the cosmos by searching for signs of life.

Caoimhe Rooney
NASA, US and Mathematigals,
rooney@baeri.org

MS0

Digital Twins: The Personalized Future of Computing for Complex Systems

Could you imagine one day having a dynamically evolving virtual replica of yourself that your doctor could use to drive personalized decisions to optimize your health and well-being? In engineering, these personalized dynamic computational models are known as digital twins, and are already being used to drive predictive maintenance decisions for aircraft and aircraft engines. This talk will discuss the computational science that goes into creating a digital twin, and will discuss exciting new directions for digital twins in engineering, geosciences and medicine.

Karen E. Willcox
UT Austin

kwillcox@oden.utexas.edu

MS0

Designing the Smart Grid of the Future

The power grid has been one of the most significant engineering achievements in the history of mankind. It has kept our lights on for over a century, but is currently facing immense transformations. To combat climate change, countries increasingly rely on renewable energy generated by solar and wind, but these sources of energy are not always available. How do we keep the lights on in an affordable way, and how do we manage the increasing risk of blackouts? To facilitate this transition, the electricity network needs a radical transformation to deal with increased levels of usage, variability, and decentralization. Indeed, the future smart grid will look dramatically different from the current electricity network, just like today's internet looks nothing like the 20th-century telephone network. In addition to the many technical, economic, regulatory, and societal challenges, the design and analysis of the future smart grid requires the development of novel mathematical foundations. This talk will discuss some of the challenges that arise across the entire spectrum of mathematics, including geometry, optimization, dynamical systems, and probability.

Bert Zwart

Centrum Wiskunde & Informatica, Amsterdam
and Eindhoven University of Technology, The Netherlands
a.p.zwart@tue.nl

MS1

Edge Computing as a Missing Link in the Post Moore Era

In the first part of this talk we discuss the concept of extreme performance HPC where the applications often include latencies below 100 ms or even below 10 ms. To facilitate low latency computation has to be placed in the vicinity of the end users by utilizing the concept of Edge Computing. We present the novel failure resilience mechanisms applied to Edge systems considering timeliness, hyper-heterogeneity and resource scarcity. We discuss our machine learning based mechanism that evaluates the failure resilience of a service deployed redundantly on the Edge infrastructures. Our approach learns the spatiotemporal dependencies between Edge server failures and combines them with the topological information to incorporate link failures by utilizing the concept of the Dynamic Bayesian Networks (DBNs). In the second part we discuss the potential impact of Edge Computing in the current rise of highly specialized architectures ranging from neuromorphic to quantum computing. As we experience the paradigm shift from generalized architectures of the Von Neumann era to highly specialized architectures in the Post-Moores law era we expect the coexistence of multiple types of architectures specialized for different types of computation. We define the two of the architectures that attracted the most interest in the research community, where we witness not only theoretical developments but also first implementations and practical use cases.

Ivona Brandic

Vienna University of Technology

ivona.brandic@tuwien.ac.at

MS1

Robust Science Roadmap: Challenges in Software Systems and High Throughput Computing Applications

While trust in and reproducibility of scientific artifacts are generally application-dependent and require correctness, numerical, and sensibility evaluations, the performance reproducibility is more architecture- and runtime-dependent (i.e., heterogeneous architectures and systems causing deadlocks, non-deterministic MPI runtime, unpredictable behaviors due to the increasingly complex software stacks). Moreover, systems have been increasing in complexity over the past two decades. Accelerators, SoCs, and other configurations are exponentially increasing the diversity of hardware resources. In this talk we will discuss how all these system architecture features should be leveraged to facilitate robust science by moving in new research directions. Specifically, we will outline how system architecture is not an isolated area but must be considered in conjunction with system software; hardware and system software configurations can be voluminous and very detailed in terms of organization; and documenting hardware and system software configurations can also exhibit inconsistencies. We will also present (a) case studies in which, at the execution phase, computation and frequent nondeterminism are not easy to measure, verify, or reproduce, and (b) tentative solutions to these case studies that require engaging the system architecture and system software properties (e.g., by establishing consistent metrics and methods).

Ewa Deelman

Information Science Institute
University of Southern California
deelman@isi.edu

Michela Taufer

The University of Tennessee
taufer@utk.edu

MS1

Environmental Energy-Aware Computing applied High Throughput Applications

Environmental research comes in many different flavors and on a variety of topics. Many of these requires highest performance from today's supercomputers, which are demanding corresponding high levels of energy and thus affect our environment. As a consequence, we need to balance the needs and costs for environmental computing, such that more computing is available at lower energy consumptions. This leads to innovative approaches in energy-efficient computing with heterogeneous architectures. The talk provides use cases on environmental research and showcases some of the ideas pursued at Leibniz Supercomputing Centre (LRZ) for energy reduction, such as hot-water cooling and waste heat usage, applied to high throughput applications.

Dieter Kranzmueller

Ludwig-Maximilians-Universität München
kranzmueller@ifi.lmu.de

MS1

High-throughput Applications in Computational

Medicine

Computational medicine combines recent advances in computing, to build predictive models for diagnosis and treatment of diseases, and combines computer simulations with artificial intelligence. While these complementary approaches typically require different computing technologies (traditional HPC vs GPU computing for machine learning), both of them often need an additional level of high-throughput computing (HTC) to achieve meaningful results. In the case of modeling and simulation approaches, an important step in development of models which can be translated into clinical practice, is the verification, validation and uncertainty quantification (VVUQ). This process requires running the model with a large number of parameters, which calls for high-throughput computing. In the case of AI-based models, the typical HTC workload results from hyperparameter optimization in machine learning, architecture search in deep neural networks, or large-scale studies in federated learning. Finally, the models in computational medicine can be used to run large studies called in-silico trials, which aim to experiment with various models and treatment strategies on a set of real or virtual patients. Examples of the tools which support such studies on HPC and cloud environments are the Model Execution Environment or CloudVVUQ library.

Maciej Malawski
AGH University of Science and Technology
m.malawski@sanoscience.org

MS3

Blessing of Nonconvexity in Factorized Models

Factorized models, from low-rank matrix recovery to deep neural networks, play a central role in many modern machine learning problems. Despite their widespread applications, problems based on factorized models are deemed difficult to solve in their worst case due to their inherent nonconvexity. Our talk is inspired by the recent observations in the optimization and machine learning communities that many realistic and practical instances of factorized models are far from their worst case scenarios. We study a natural nonconvex and nonsmooth formulation of two prototypical factorized models, namely low-rank matrix factorization and deep linear regression, where the goal is to recover a low-dimensional model from a limited number of measurements, a subset of which may be grossly corrupted with noise. On the negative side, we show that this problem does not have a benign landscape: with high probability, there always exists a true solution that is not a global minimum of the loss function. However, on the positive side, we show that a simple subgradient method with small initialization is oblivious to such problematic solutions; instead, it converges to a balanced solution that is not only close to the ground truth but also enjoys a flat local landscape. Lastly, we empirically verify that the desirable optimization landscape of factorized models extends to other robust learning tasks, including deep matrix recovery and deep ReLU networks with L1-loss.

Salar Fattahi
University of Michigan
fattahi@umich.edu

MS4

How Dark Is the Unilluminable Room?

The 1958 Christmas issue of *The New Scientist* contained

two pages with puzzles posed by Sir Roger Penrose (and his father S. L. Penrose). One of these puzzles asks the reader to design a smooth closed reflecting surface (mirror) which contains two regions with the property that a source of light placed in one cannot be seen from the other. This "room" has become known as the unilluminable room and there are now numerous fascinating solutions to the problem. The original puzzle assumes that the light is described by rays (a so-called billiards problem) so that the light cannot bend around corners. Here we model light by solving the Helmholtz equation with a point source in one of the regions of the room and study (among other things) how dark the other region of the room actually is as we change the frequency of the light-source.

Daniel Appelo
Michigan State University
appeloda@msu.edu

MS4

On the Stable Discretization of the Vectorial Equations of Stellar Oscillations

We consider the time-harmonic equation of stellar oscillations, which is a Galbrun-like equation with additional rotational and gravitational terms. The main difficulty of this equation is that it is not only indefinite, but it contains differential operators with opposite signs, and recently its well-posedness was analyzed in [M. Halla, T. Hohage, On the well-posedness of the damped time-harmonic Galbrun equation and the equations of stellar oscillations, *SIAM J. Math. Anal.* 53(4) (2021), pp. 40684095]. In this talk we discuss how to construct and analyze convergent discretizations. To this end we report a new weak T-compatibility framework and apply it to certain H^1 -conforming finite element methods. The main assumption on the finite element spaces is a uniformly stable invertibility of the divergence operator, which is ensured under certain conditions on the meshes and the polynomial degree. Thus we obtain a family of reliable finite element methods, for which we present computational results. Details can be found in the report [M. Halla, C. Lehrenfeld, P. Stocker, A new T-compatibility condition and its application to the discretization of the damped time-harmonic Galbrun's equation, arXiv:2209.01878].

Martin Halla
Max Planck Institute for Solar System Research
University of Goettingen
m.halla@math.uni-goettingen.de

MS4

Stabilisation Using Control Theory of Spurious High Frequencies for High-Order Discretisations of the Wave Equation

The objective of this work is to propose and analyze numerical schemes to solve transient wave propagation problems that are exponentially stable (i.e. the solution decays to zero exponentially fast). Applications are in data assimilation strategies or the discretisation of absorbing boundary conditions. More precisely the aim of our work is to propose a discretization process that enables to preserve the exponential stability at the discrete level as well as a high order consistency when using a high-order finite element approximation. The main idea is to add to the wave equation a stabilizing term which damps the high-frequency oscillating components of the solutions such as spurious waves. This term is built from a discrete multiplier analy-

sis that proves the exponential stability of the semi-discrete problem at any order without affecting the order of convergence.

Sébastien P. Imperiale
Inria
sebastien.imperiale@inria.fr

MS4

Dispersion of Torsional Surface Waves in a Three-fold Concentric Compounded Cylinder with Imperfect Interface

The present paper investigates the torsional wave propagation in a threefold concentric pre-stressed compounded cylinder with imperfect contact conditions. The three-dimensional linearized theory of elastic waves and the piecewise homogeneous body model has been employed to formulate the problem. The mathematical modelling has been carried out in two independent cases. In the first case a solid cylinder encased in a hollow cylinder embedded in an infinite elastic medium has been considered. Whereas second case comprises of hollow cylinder of finite thickness in place of solid cylinder. By means of Murnaghan potential the mechanical characteristics of the three materials have been used. Further, the dispersion relations for both the cases have been obtained in terms of the Bessel and modified Bessel functions. In order to validate the present findings, two particular cases have derived which matches with the previous works. The first case is obtained by removing the outermost cylinder while the second case has been derived by removing the imperfection in addition to that. To summarize the computations, a complete numerical simulation has been carried out, and graphical illustrations have been shown to aid the mathematical analyses.

Sumit K. Vishwakarma
Indian School of Mines, Dhanbad, Jharkhand,
India-826004
sumo.ism@gmail.com

MS5

Community Software Usage on the UK National Supercomputing Service, ARCHER2 from Analysis of Slurm Data

In this presentation we analyse the use of software by different research communities on the UK national supercomputing service, ARCHER2. We use data from the Slurm scheduler along with a tool that matches usage to software packages that has been developed over the past two UK national supercomputing services. We see clear differences between how different research communities approach developing, supporting and using software on large scale facilities from the analysis for example, some communities have a small number of general-purpose software packages used by many researchers, other communities have a much broader range of software packages used by smaller numbers of researchers. We provide some initial thoughts as to why these differences exist and their potential implications from a researcher and service productivity perspective. The analysis tool has recently been extended to provide insight into differences in energy use by different communities and software packages and we present initial findings from this data and how it can potentially support the transition of large-scale HPC facilities to net zero emissions.

Andrew Turner

EPCC, UK
a.turner@epcc.ed.ac.uk

Michael Bareford
University of Edinburgh, UK
m.bareford@epcc.ed.ac.uk

MS5

Productivity of Large Community Codes on the Supercomputer Fugaku

Fugaku is a supercomputer jointly developed by RIKEN and Fujitsu. It features Arm-based A64FX CPU, TOFU-D high-performance network interconnect, etc. to It has been made available for shared use since March, 2021, and many achievements have already been made with Fugaku in various fields. In this talk, I will show you some efforts to make the supercomputer more usable and productive in the viewpoint of the software environment. First, we have been working to support important open-source software (OSS) packages including libraries, tools, and applications, some of which are optimized for Fugaku. These OSS packages are provided via Spack. Second, we are working to port and support commercial software on Fugaku in cooperation with Fujitsu and software vendors. Gaussian is the first example and available since Nov. 2022. Finally, we contribute to develop the Arm software ecosystem for HPC. For example, we started a work of improving the LLVM compiler by implementing more effective optimization feature for A64FX. With these efforts, users are and will be able to use Fugaku to achieve better productivity and performance.

Hitoshi Murai
RIKEN
h-murai@riken.jp

MS5

Simulations of Human-Scale Mars Lander Descent Trajectories

To achieve human exploration on Mars, new approaches to entry, descent, and landing are necessary in order to support delivery of substantially larger payloads to the surface. In this presentation, large-scale computational fluid dynamics (CFD) simulations of a human-scale Mars lander configuration descending through the Martian atmosphere with the use of retropropulsion are described. The CFD solver is coupled with a trajectory analysis package which is used to provide the necessary control inputs to maintain the desired flight path. Simulations are performed using thousands of GPUs on the Summit system located at the Oak Ridge Leadership Computing facility. An overview of the GPU-based implementation is presented, and aerodynamic performance and characteristics of the highly dynamic flowfield are compared with those observed at static freestream conditions as studied in prior computational campaigns.

Eric Nielsen
NASA Langley
eric.j.nielsen@nasa.gov

Brad Robertson, Hayden Dean, Zach Ernst
Georgia Tech
bradford.robertson@asdl.gatech.edu, hdean@gatech.edu, zernst@gatech.edu

Kevin Jacobson, Gabriel Nastac, Aaron Walden

NASA Langley
 kevin.e.jacobson@nasa.gov, gabriel.c.nastac@nasa.gov,
 aaron.c.walden@nasa.gov

Patrick Moran
 NASA Ames
 patrick.moran@nasa.gov

Alexandra Hickey
 Georgia Tech
 ahickey7@gatech.edu

MS5

Adapting Experimental Science Tools and Techniques to Computational Sciences

As computational mathematics, science, and engineering problems become larger, more ambitious, and more complex, it is increasingly important to develop and use tools and techniques that ensure that computational science research is based on a strong foundation of general, low-level scientific best practices. In this talk, I will relate my experience of transitioning from working in the worlds of experimental and observational sciences to the world of computational sciences as well as my experience adapting experimental tools and techniques to computational research. In particular, I will discuss the role of lab notebooks in experimental sciences as well as the challenges associated with adapting lab notebooks to computational sciences research. Rather than develop this adaptation in isolation, lab notebooks will be addressed within the larger scope of computational laboratory environments designed and implemented for executing reproducible computational science studies with solid scientific rigor.

Jared O'Neal
 Argonne National Laboratory
 joneal@anl.gov

MS5

Nwchemex: Challenges Faced in Designing An Electronic Structure Program for the Exascale

For over 20 years, the original NWChem package has been synonymous with massively parallel computational chemistry. Unfortunately the exascale-era version of the high-performance computing (HPC) landscape is dramatically different from the HPC landscape which was present when NWChem was originally designed. Of specific note: HPC hardware is increasingly heterogeneous, C/C++ has replaced Fortran as the dominant HPC language, and object-oriented programming is now the prevailing design paradigm. Ultimately, in order to adopt NWChem to the exascale era, the decision was made to write an entirely new program from scratch. The result is NWChemEx. NWChemEx is written in modern object-oriented C++ (presently C++ 17, with the intent to adopt newer C++ standards as compiler support allows). All user-facing C++ application programming interfaces (APIs) are mirrored in Python 3. From a software architecture standpoint, NWChemEx takes separation of concerns to an extreme. The entire package is written as a series of plugins, where each plugin consists of self contained code modules. To the extent possible, the entire code is written using a single instruction, multiple data design paradigm, in turn allowing parallelism to be treated separately from the implementation of new functionality. This talk will explain the architecture and design of NWChemEx in more detail including the rationale underpinning our decisions and the

challenges faced in implementing NWChemEx.

Ryan Richard
 Iowa State University
 ryanmrichard1@gmail.com

MS6

High Order Approximation of Non Linear Hyperbolic Systems by Mean of a Virtual Finite Element

There exist many ways of approximating non linear hyperbolic problems: finite difference (FD), continuous finite elements (CFEM), discontinuous Galerkin method (dG), finite volume (FV), Residual distribution methods (RD), etc. All rely, in a manner or another, on an explicit representation of the numerical solution, either on a stencil (FD), or within an element (all the others), and then from this one compute the elements of the scheme by differentiating basis functions, or reconstructing the variables by means of polynomials. If this strategy is very successful, this paper intends to go beyond, by following ideas arising from the Virtual Finite Elements methods (VEM) introduced by Bezzi and his collaborators a few years ago. Here also, even for polygonal meshes, the solution is thought as globally continuous. No explicit representation of the solution is introduced, only its polynomials approximation on the polygonal elements that compose the mesh, and some moments of the solution. We will describe several methods to achieve this goal, for the compressible Euler equations. Our method has several common points with the Active Flux method developed recently by Roe et al. It has also several differences, in particular a wider flexibility in the choices of variables, as well as the shape of the elements. Acknowledgement: We acknowledge several important discussions with W. Barsukow (CNRS, Talence, France)

Remi Abgrall
 University of Zurich
 remi.abgrall@math.uzh.ch

Mario Ricchiuto
 INRIA
 mario.ricchiuto@inria.fr

MS6

Subcell Limiting Strategies for Split-Form DG Schemes

In this talk, we present ideas from the recently published paper [Rueda-Ramrez, A. M., Pazner, W., Gassner, G.J. Subcell Limiting Strategies for Discontinuous Galerkin Spectral Element Methods. Computers and Fluids (2022). <https://arxiv.org/abs/2202.00576>.] and explain how to construct a compatible subcell finite volume scheme that can be seamlessly blended with a high order spectral element discontinuous Galerkin method. We focus on the discretization of (non-linear) conservation laws, such as the KPP equation and the compressible Euler equation. Starting with an entropy-dissipative split-form discontinuous Galerkin scheme on collocated Legendre-Gauss-Lobatto (LGL) nodes, we show how to carefully design a finite volume type discretization on the LGL subcell grid such that it is (i) still provably entropy-dissipative/stable/conservative and (ii) compatible to the high order DG scheme in the sense that we can blend between the low order finite volume and the high order DG discretization on a per element level, or a per node level. This new hybrid scheme enables us to do shock capturing and preserve positivity of the solution. We present sev-

eral strategies including heuristic troubled cell indicators and positivity preserving mechanisms, as well as sophisticated positivity techniques from Flux Corrected Transport schemes including Invariant Domain Preservation a posteriori stabilization and discuss their advantages and disadvantages.

Gregor Gassner, Andres Rueda-Ramirez
University of Cologne
ggassner@uni-koeln.de, aruedara@uni-koeln.de

Will Pazner
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
pazner@pdx.edu

MS6

Kinetic Energy Preservation and Entropy Stable Modal Flux Reconstruction Formulations

The flux reconstruction method has gained popularity in the research community as it recovers promising high-order methods through modally filtered correction fields, such as the Discontinuous Galerkin (DG) method, on unstructured grids over complex geometries. Under a class of energy stable flux reconstruction (ESFR) schemes also known as Vincent-Castonguay-Jameson-Huynh (VCJH) schemes, the flux reconstruction method allows for larger time-steps than DG while ensuring linear stability on linear elements. Alternatively, for nonlinear problems, split forms emerged as the popular approach proving nonlinear stability for unsteady problems on coarse unstructured grids; albeit only having been proved for the strong form DG scheme and numerically shown for the g_2 -lumped Lobatto strong form ESFR scheme for the Euler equations. Incorporating split forms with VCJH schemes, alike Ranocha and Abe, generally lead to unstable discretizations. Over the past year we have developed a new Nonlinearly Stable Flux Reconstruction approach based on modal uncollocated ESFR schemes in split forms that ensure energy and entropy stability. This new approach simplifies to VCJH type schemes for linear problems, but for nonlinear problems, it ensures nonlinear stability and the correct orders of convergence. In this work, we show analytically that with a suitable interface flux the scheme guarantees kinetic energy preservation for the stated Nonlinearly Stable Flux Reconstruction framework.

Siva Nadarajah, Alexander Cicchino
McGill University
siva.nadarajah@mcgill.ca, alexander.cicchino@mail.mcgill

MS6

Accurate Numerical Solutions of the Three-Dimensional Bratu's Problem by Simple Discretization

In this manuscript, a simple discretization algorithm for computing the different solutions to the three-dimensional Bratu problem is developed. The proposed numerical scheme is applied to compute/approximate the different critical values of the transition parameter. The aim of the study is to provide highly-accurate reliable results, given the scarcity of proven accurate results in the literature. Numerical simulation shows that the proposed algorithm has the capability of accurately computing the solutions of the three-dimensional Bratu problem, as well as the critical values of the transition parameter. Furthermore, the

bifurcated behavior of the solution is presented by investigating the infinity norm of the solution for different values of the transition parameter.

Helmi Temimi
Gulf University for Science and Technology, Kuwait
temimi.h@gust.edu.kw

MS6

Quantification of Numerical Uncertainty via Non-linear Dynamics Approach

For the last three decades I and collaborators have worked extensively on nonlinear behavior of spatial & temporal discretizations related to the predictability and reliability of nonlinear simulations via the dynamical system approach. The so-called stiff source term problem of LeVeque & Yee (1990) is one well-known example. This problem manifests itself as discontinuities that propagate with the wrong speed when the numerical solution is under resolved and/or too diffusive in both space and time. The construction of numerical methods for (a) stable and accurate simulation of turbulence with strong shocks, and for (b) obtaining the correct propagation speed of discontinuities in the presence of stiff reacting terms share one important ingredient - minimization of numerical dissipation while maintaining numerical stability. The dual requirement to achieve both numerical stability and minimal numerical dissipation are often conflicting for existing shock capturing schemes for flows with turbulence and discontinuities or stiff source terms. During the last decade, Wang et al., Yee et al., Sjogreen & Yee have developed highly accurate structure preserving numerical methods with very small numerical dissipation that maintain nonlinear stability for compressible turbulence with strong shocks. This talk gives a brief overview of the quantification of numerical uncertainty via a nonlinear dynamics approach to complement existing uncertainty quantification in numerical simulations.

H.C. Yee
NASA Ames Research Center
Helen.M.Yee@nasa.gov

MS7

Sparse Matrix Kernels for Parallel Graph Algorithms

Sparse matrix algebra has emerged as an expressive and powerful framework for designing parallel algorithms for graph problems including all-pairs shortest paths, betweenness centrality, and connectivity. We employ this framework to design scalable routines for two key sparse kernels arising in parallel minimum spanning forest (MSF) and single-source shortest path (SSSP) algorithms. In particular, we target the hooking step of the Awerbuch-Shiloach MSF algorithm and the nearest neighbor search leveraged by the low hop emulator SSSP construction (Andoni, Stein, Zhong, STOC 20). We introduce a multilinear kernel that operates on an adjacency matrix and two vectors to eliminate a write bottleneck common to both steps. This kernel updates vertices by simultaneously using information from an edge and its two adjacent vertices. We implement these routines with Cyclops, a distributed-memory library for generalized sparse tensor algebra. We analyze the parallel scalability of our implementations on the Stampede2 supercomputer.

Tim Baer
University of Illinois at Urbana-Champaign

tbaer2@illinois.edu

MS7

Randomized Linear Algebra can Accelerate Spectral Graph Partitioning

Spectral partitioning uses eigenpairs of the graph Laplacian to compute a partitioning. The eigencomputation is the most expensive part, but low accuracy is often sufficient for partitioning. We show that a randomized eigensolver can speed up the partitioning time by a factor 5-50X, without much difference in quality for certain types of graphs. Our approach can easily be adapted to spectral clustering.

Erik G. Boman

Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Jennifer A. Loe

Sandia National Laboratories
jloe@sandia.gov

Heliezer J. Espinoza

California Polytechnic State University, Pomona
dhespinoza@cpp.edu

MS7

Parallel Graph Algorithms for Fixed-Parameter Tractable Problems

Although many useful graph problems, such as subgraph isomorphism, are NP-hard in general, they are solvable in polynomial time for certain families of sparse graphs. As these algorithms might still require significant resources, providing parallel algorithms can further raise the scale of problems solvable within a given time limit. For example, subgraph isomorphism becomes tractable on planar graphs and other apex-minor-free families of graphs. Moreover, searching for maximum cliques becomes tractable in planar and other bounded-degeneracy graphs. This talk discusses our recent advances in parallel algorithms for these problems. Our algorithms must avoid large sequential graph traversals and instead quickly break apart the graphs into independent parts. We provide algorithms whose work grows only linearly with the size of the target graph and whose depth grows poly-logarithmically with the size of the target graph. These results imply more work-efficient algorithms for other problems, such as vertex connectivity in planar graphs.

Lukas Gianinazzi

ETH Zurich
lukas.gianinazzi@inf.ethz.ch

MS7

Observability of Consensus of Multi-Agent Networks Under Signed Laplacian Dynamics

Signed multi-agent networks have been an area of recent interest in networks and control community as they involve both cooperative and competitive interactions among its agents. This work establishes the observability conditions of consensus problem of such networks when their agents agrees on signed Laplacian dynamics on a continuous time-scale. The interconnection topology is time-invariant. Under leader-follower framework, where, the agents of the network are classified into leaders and followers, the nec-

essary and sufficient observability conditions are obtained. Graphical interpretations are provided to gain more insight on the interactions among the agents of the network and network topology. In addition, the developed results are generic, in the sense that they are not only applicable to signed networks but also to unsigned networks, since unsigned networks are the particular case of signed networks that only contain cooperative agents. Furthermore, we prove that, if the signed multi-agent network is structurally balanced, then the obtained observability conditions will become equivalent to that of the corresponding unsigned multi-agent network, under some special leaders selection. Numerical examples are given to illustrate the obtained theoretical results.

Muhammed Rafeek K V

Research Scholar, Department of Mathematics,
Central University of Karnataka, India
rafeekkv997@gmail.com

Vijayakumar S Muni

Assistant Professor. Department of Mathematics
Sri H.D . Devegowda Government First Grade College.
vijayakumarmuni0@gmail.com

G. Janardhana Reddy

Assistant Professor
Department of Mathematics, Central University of
Karnataka.
gjr@cuk.ac.in

MS7

Theoretically and Practically Efficient Parallel Nucleus Decomposition

This paper studies the nucleus decomposition problem, which has been shown to be useful in finding dense substructures in graphs. We present a novel parallel algorithm that is efficient both in theory and in practice. Our algorithm achieves a work complexity matching the best sequential algorithm while also having low depth (parallel running time), which significantly improves upon the only existing parallel nucleus decomposition algorithm (Sariyuce et al., PVLDB 2018). The key to the theoretical efficiency of our algorithm is the use of a theoretically-efficient parallel algorithms for clique listing and bucketing. We introduce several new practical optimizations, including a new multi-level hash table structure to store information on cliques space-efficiently and a technique for traversing this structure cache-efficiently. On a 30-core machine with two-way hyper-threading on real-world graphs, we achieve up to a 55x speedup over the state-of-the-art parallel nucleus decomposition algorithm by Sariyuce et al., and up to a 40x self-relative parallel speedup. We are able to efficiently compute larger nucleus decompositions than prior work on several million-scale graphs for the first time.

Jessica Shi

Massachusetts Institute of Technology
jeshi@mit.edu

MS8

European Research Council (ERC) Panel: How and Why Applying to ERC Funding?

The ERC's mission is to support excellent frontier research through competitive funding to creative researchers of any nationality and age, to run projects based across Europe

or associated countries. The purpose of this special panel, co-organized by the Committee for Applications and Interdisciplinary relations of EMS, is to encourage mathematicians working in the areas connected with applications to prepare very competitive proposals and submit them to ERC for evaluation. During the first part of this session, Maria Gonzalez, the ERC Mathematics Panel coordinator, will give an overview of the funding opportunities offered by the ERC as well as some general information on funded projects and the application and evaluation process. This will be complemented during the second part by the testimonies of Annalisa Buffa and Daniel Peterseim, AdG and CoG grantees respectively.

Maria Gonzalez
European Research Council Executive Agency
ERCEA.B.4 Physical Sciences and Engineering
maria-teresa.gonzalez-cerveron@ec.europa.eu

Daniel Peterseim
Universität Augsburg
daniel.peterseim@uni-a.de

Annalisa Buffa
Institute of Mathematics
École polytechnique fédérale de Lausanne, Switzerland.
annalisa.buffa@epfl.ch

MS9

A Systematic Approach to Constructing Preconditioners for the Hp-Version Mass Matrix on Unstructured and Hybrid Finite Element Meshes

We present a systematic approach to building uniform preconditioners for the mass matrix on unstructured meshes composed of higher dimensional elements using only preconditioners for lower dimensional simplices. In particular, we show that the resulting preconditioners are automatically uniform with respect to both the mesh size, h , and polynomial degree, p , and that the preconditioners can be implemented efficiently by exploiting the structure of the preconditioners on the lower dimensional elements. We illustrate the approach by developing preconditioners for prismatic elements and for the challenging case of hybrid meshes of hexahedra, prisms, and tetrahedra.

Mark Ainsworth
Brown University
mark.ainsworth@brown.edu

MS9

NekRS, Next Generation Spectral Element Navier-Stokes Solver

This talk is about our next generation CFD code nekRS targeting CPUs and accelerators like GPUs. An overview of the project will be provided including on-going developments and various performance studies.

Stefan Kerkemeier
Argonne National Laboratory
stgeke@gmail.com

MS9

Efficient Exascale Discretizations: High-Order Finite Element Methods

Efficient exploitation of exascale architectures requires re-

thinking of the numerical algorithms used in many large-scale applications. These architectures favor algorithms that expose ultra fine-grain parallelism and maximize the ratio of floating point operations to energy intensive data movement. One of the few viable approaches to achieve high efficiency in the area of PDE discretizations on unstructured grids is to use matrix-free / partially-assembled high-order finite element methods, since these methods can increase the accuracy and/or lower the computational time due to reduced data motion. In this paper we provide an overview of the research and development activities in the Center for Efficient Exascale Discretizations (CEED), a co-design center in the Exascale Computing Project that is focused on the development of next-generation discretization software and algorithms to enable a wide range of finite element applications to run efficiently on future hardware. CEED is a research partnership involving more than 30 computational scientists from two US national labs and five universities, including members of the Nek5000, MFEM, MAGMA and PETSc projects. We discuss the CEED co-design activities based on targeted benchmarks, miniapps and discretization libraries and our work on performance optimizations for large-scale GPU architectures.

Tzanio V. Kolev
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
tzanio@llnl.gov

MS9

Dft-Fe: Massively Parallel Hybrid CPU-GPU Based Computational Methodologies for First Principles Material Modelling Using Finite-Element Discretization of Density Functional Theory

Kohn-Sham density functional theory (DFT) calculations have been instrumental in providing many crucial insights into materials behaviour and occupy a sizable fraction of the world's computational resources today. However, the stringent accuracy requirements required to compute meaningful material properties, in conjunction with the asymptotic cubic-scaling computational complexity of the underlying eigenvalue problem, demand enormous computational resources. Thus, these calculations are routinely limited to material systems with at most a few thousand electrons. In this talk, recent advances in enabling fast and accurate large-scale real-space Kohn-Sham DFT calculations -via- the development of DFT-FE, a massively parallel open-source finite-element (FE) based DFT code on hybrid CPU-GPU architectures will be discussed. DFT-FE employs adaptive FE discretisation that handles pseudopotential and all-electron calculations while accommodating periodic, non-periodic and semi-periodic boundary conditions. The talk will highlight novel HPC-centric numerical strategies that significantly reduce the data movement costs and increase arithmetic intensity on evolving hybrid architectures, considerably delaying the onset of cubic scaling computational complexity to system sizes comprising 30,000 electrons. These advances have wide-ranging implications for tackling various scientific problems involving large-scale material systems in the areas of energy storage, catalysis, alloy design etc.

Phani S. Motamarri
Indian Institute of Science Bangalore
phanim@iisc.ac.in

Sambit Das
Department of Mechanical Engineering

University of Michigan Ann Arbor
dsambit@umich.edu

Vikram Gavini
Mechanical Engineering, University of Michigan
vikramg@umich.edu

MS9

Accelerating Very High-Order Finite Element Operations with GPU Matrix-Cores

We will discuss the challenges in achieving high performance for the DOE ECP CEED bake-off problems on the latest NVIDIA and AMD server grade GPUs. In particular we will consider matrix-free algorithms for very high-order finite element and spectral element discretizations of elliptic problems. We will introduce specialized asymmetric streaming benchmarks that establish attainable throughput goals for these memory bound finite element operations. We will finally illustrate the importance of exploiting intrinsic matrix instructions to reach empirically calibrated device memory bandwidth rooflines at very high order.

Tim Warburton
Virginia Tech
tcew@vt.edu

Noel Chalmers, Damon Mcdougall
AMD Research
noel.chalmers@amd.com, damon.mcdougall@amd.com

Jason R. Wilson
Virginia Tech
jasonwil@math.vt.edu

MS10

Extremum Seeking Control Systems as a Mean to Study Real Time Dynamic Optimization Phenomena in Biology: An Example of Soaring Birds Optimized Flight

Extremum seeking control (ESC) systems have been around for decades. They are very powerful special type of systems as they automatically steer a given dynamical system to the extremum (minimum/maximum) of an objective function that we do not have its mathematical closed form a priori. In fact, ESC systems only require access to measurements of said unknown objective function. We hypothesize that ESC systems are able to capture, describe and replicate many biological optimization phenomena. Unlike many dynamic optimization methods and optimal control solvers used in modeling/studying biological optimization phenomena, ESC systems are real-time, stable and operable through access to measurements (sensing); this matches the nature of biological optimization phenomena. As a proof of concept, we provide an example of work we done on capturing and mimicking soaring birds optimized flight maneuvers. Our results show the effectiveness of ESC systems as a mean to study real time dynamic optimization phenomena in biology.

Sameh Eisa
University of California, Irvine
eisash@ucmail.uc.edu

MS10

Quantifying the Shape and Growth of Colliding Tissues via Continuum Models and Bayesian In-

ference

Although tissues are usually studied in isolation, this situation rarely occurs in biology, as cells, tissues, and organs, coexist and interact across scales to determine both shape and function. Here, we take a quantitative approach combining data from recent experiments, mathematical modelling, and Bayesian parameter inference, to describe the self-assembly of multiple epithelial sheets by growth and collision. We use two simple and well-studied continuum models, where cells move either randomly or following population pressure gradients. After suitable calibration, both models can reproduce the main features of single tissue expansions. However, our findings reveal that whenever tissues are not isolated from others, and hence interactions become relevant, the random motion assumption can lead to unrealistic behaviour. Under this setting, a model accounting for population pressure from different cell populations is more appropriate and can help to compare with experimental measures. Finally, we discuss how tissue shape and pressure affect multi-tissue complex collisions.

Carles Falco
Mathematical Institute
University of Oxford
falcoigandia@maths.ox.ac.uk

MS10

Bacterial Communication in Space - An Interplay of Nonlinear Dynamics and Spatial Inhomogeneity

Many bacterial species use the so-called Quorum sensing, a communication system, to coordinate their behaviour, e.g. for pathogenicity or other substantial changes of their life-style. Gene regulation systems are the basis for such systems. They often contain positive and negative feedback loops. This may allow under certain conditions for bistability. As the communication involves the diffusion or transport of molecules, we find more interesting phenomena due to the combination of diffusion and non-linear reactions. We aim in understanding this from both, the mathematical and the biological perspective, including a short view on treatments like classical antibiotics but also Quorum quenching - a kind of inhibition of Quorum sensing.

Christina Kuttler
The Technical University of Munich
kuttler@ma.tum.de

MS10

Modeling Physical Limits in Cell Migration

Migrating cells choose their preferential direction of motion in response to different signals and stimuli sensed by spanning their external environment. However, the presence of dense fibrous regions, lack of proper substrate, and cell overcrowding may hamper cells from moving in certain directions or even from sensing beyond regions that practically act like physical barriers. I will first describe a non-local kinetic model in which the sensing radius is not constant, but depends on position, sensing direction and time as the behaviour of the cell might be determined on the basis of information collected before reaching physically limiting configurations. I will then describe a continuum model to deal with the transmigration through basal membranes.

Luigi Preziosi
Politecnico di Torino

Corso Duca degli Abruzzi 24 10129 Torino (Italy)
luigi.preziosi@polito.it

Nadia Loy
Politecnico di Torino
nadia.loy@polito.it

Chiara Giverso
politecnico di torino
chiara.giverso@polito.it

MS10

Multilevel Markov Chain Monte Carlo Sampling Methods for Bayesian Inverse Problems in Cardiac Electrophysiology

In modeling of cardiac arrhythmias, an important problem is to use medical data tune the model parameters to represent an individual patient. Since the amount of measurement data is limited, not only a optimal parameter fit is needed, but also an estimate on the uncertainty of this optimal fit. In this talk, we study the use of multilevel Monte Carlo methods to sample the posterior parameter distribution of a Mitchell-Schaeffer model for cardiac arrhythmia. To validate the computational procedure, we will work synthetic data, simulated via the same model, to which we will add measurement noise with different sizes and covariance structures. The goal is to recover the (spatially homogeneous) model parameters that were used to generate the data, and to assess the importance of the size and the correct quantification of measurement noise. To measure similarity between measurement data and model output, we use the LAT as a summary statistic, effectively resorting to an Approximate Bayesian Computation. Since PDE models for biological systems typically contain lots of modeling error, we also study the effect of model error on the posterior distribution.

Giovanni Samaey
Department of Computer Science, K. U. Leuven
giovanni.samaey@cs.kuleuven.be

MS11

Algorithms for Extension Along Normals

Function extension is a useful technique for efficiently evaluating a particular solution of a constant coefficient PDE on a smooth domain, by evaluating a convolutional integral in "free space". This strategy avoids specialized quadrature in the volume and utilizes high throughput fast algorithms that are available in simple geometries. Recently, methods for function extension have been proposed which leverage the relative simplicity of one dimensional function extrapolation: extrapolation is performed only in the direction normal to the domain boundary and interpolation is performed tangentially. We propose a high order and efficient framework for extrapolation in this style which is suitable for both uniform and adaptive discretizations. Examples will be presented in 2 and 3 dimensions.

Travis Askham
New Jersey Institute of Technology
askham@njit.edu

MS11

An FMM Accelerated Poisson Solver for Complicated Geometries in the Plane Using Function Ex-

tension

We present a potential theory-based adaptive solver for the Poisson equation in complicated geometries in the plane, accelerated with a fast multipole method (FMM). The solution is given explicitly as the sum of a volume potential and a double layer potential. To simplify the evaluation of the volume potential over the complicated domain, we extend the source data to a geometrically simpler domain, which contains the original geometry. It remains only to solve the homogeneous Laplace equation, with modified boundary data, to obtain the full solution. This is done with existing fast and accurate boundary integral methods. The novelty of our solver is the scheme used for creating the extension. The source data is represented on boxes in an adaptive quad-tree. For leaf boxes intersected by the boundary we extend the source data locally to a subset of its colleagues. The extension is created by sampling data from neighboring boxes and using a box-independent precomputed interpolation matrix. The resulting potential induced by the extended source data can then be efficiently computed with existing volume-integral FMMs. We demonstrate speed, robustness and high-order convergence through several examples, including piecewise smooth domains.

Fredrik Fryklund
New York University, U.S.
nf2235@nyu.edu

Leslie Greengard
New York University
Flatiron Institute
greengar@cims.nyu.edu

MS11

Monolithic AMG Preconditioners for Stokes Equations

Advanced discretizations and complex meshes are becoming increasingly commonplace in a range of coupled physics and engineering applications, from weather prediction to fluid flow to graph problems. These discretizations also place increased demands on the underlying solvers, focusing on key properties of the PDE models - e.g., underlying conservation laws. While algebraic multigrid methods effectively tackle some PDEs and some discretizations on unstructured meshes, the presence of higher-order basis functions and multiple coupled unknowns makes the algebraic coarsening a challenging endeavor. This talk will highlight several new monolithic algebraic multigrid preconditioning approaches targeting high-order and non-standard discretizations of Stokes flow problems on structured and unstructured meshes. We will demonstrate how the robust monolithic AMG preconditioners can be constructed to precondition higher-order Taylor-Hood and Scott-Vogelius discretizations.

Alexey Voronin
Department of Computer Science
University of Illinois Urbana-Champaign
voronin2@illinois.edu

MS12

Neural Ideograms: Unsupervised Learning of Geometric Symbols via Kendall Shape Space VAEs

A picture speaks a thousand words, which is poetically beautiful, but highly problematic from a practical point of view. That is why in our daily lives we work with symbols,

ideograms in general, to convey information as efficiently as possible. In this work, we show that the idea of compressing information into geometric symbols is not only computationally feasible with artificial neural networks, but also that it allows for learning effective and interpretable neural representations. Kendall shape VAEs can learn shapes that are translation, rotation, and scale invariant in an unsupervised manner. We qualitatively inspect the learned symbols and observe that they consistently represent classes of images that are equivalent by group actions and semantic meaning. In a quantitative comparison to other types of VAEs, we show that Kendall shape VAEs are more efficient in learning compressed representations and that the learned representations convey more information for downstream tasks such as classification.

Erik Bekkers

Amsterdam Machine Learning Lab (AMLab)
Informatics Institute, University of Amsterdam
e.j.bekkers@uva.nl

MS12

Topological Structures in Neural Activity

To be updated.

Benjamin Dunn

Norwegian University of Science and Technology
benjamin.dunn@ntnu.no

MS12

Graph-Based Nearest Neighbour Search in Hyperbolic Spaces

The nearest neighbour search (NNS) problem is widely studied in Euclidean space, and graph-based algorithms are known to outperform other approaches for this task. However, hyperbolic geometry often allows for better data representation in various domains, including graphs, words, and images. In this paper, we show that graph-based approaches are also well suited for hyperbolic geometry. From a theoretical perspective, we rigorously analyze the time and space complexity of graph-based NNS, assuming that an n -element dataset is uniformly distributed within a d -dimensional ball of radius R in the hyperbolic space of curvature -1 . Under some conditions on R and d , we derive the time and space complexity of graph-based NNS and compare the obtained results with known guarantees for the Euclidean case. Interestingly, in the dense setting ($d \ll \log n$) and under some assumptions on the radius R , graph-based NNS has lower time complexity in the hyperbolic space. This agrees with our experiments: we consider datasets embedded in hyperbolic and Euclidean spaces and show that graph-based NNS can be more efficient in the hyperbolic space. We also demonstrate that graph-based methods outperform other existing baselines for hyperbolic NNS. Overall, our theoretical and empirical analysis suggests that graph-based NNS can be considered a default approach for similarity search in hyperbolic spaces.

Alexander Kolpakov

University of Neuchâtel
kolpakov.alexander@gmail.com

Liudmila Prokhorenkova

Yandex Research
ostroumova-la@yandex-team.ru

Dmitry Baranchuk

Moscow State University

dmitry.baranchuk@graphics.cs.msu.ru

Nikolay Bogachev

Institute for Information Transmission Problems
nvbogach@mail.ru

Yury Demidovich

Moscow Institute of Physics and Technology
yuradem9595@mail.ru

MS12

From Neural Manifolds to Neural Lie Groups: Uncovering Group Structure in Neural Data with Geometric Machine Learning

An emerging set of findings in sensory and motor neuroscience is beginning to illuminate a new perspective on neural coding. Across sensory and motor regions of the brain, neural circuits appear to mirror the geometric and topological structure of the systems they represent—either in their synaptic structure, or in the implicit manifold generated by their activity. This suggests a general computational strategy that is employed throughout the brain to preserve the geometric structure of data throughout stages of information processing. In parallel, there has been a growing recognition of the importance of respecting the geometry of data in deep neural network architectures for learning useful representations for downstream tasks. This has given rise to the nascent sub-field of Geometric Deep Learning. There is high potential for synergy between these two emerging fields. In this talk, I present results applying novel geometric machine learning methods to the analysis of neural data, revealing latent group structure in neural population codes and providing explicit parameterizations of Riemannian neural manifolds. These findings suggest fruitful territory for the application of differential geometry and group theory to the understanding of the neural representations that subserve perception and action.

Sophia Sanborn

University of California, Santa Barbara
sanborn@ucsb.edu

MS12

Lie Groups and Hierarchical Composition in Neural Representations

Mounting evidence in both neuroscience and deep learning indicates that an understanding of neural representation is an essential component of robust intelligence. In the natural world, sense data is highly structured, arising from the symmetry and geometry of the space in which objects lie. In vision for example, objects undergo transformations such as translation, rotation, scaling, and deformation. Further, visual scenes themselves are composed of many objects, objects are composed of parts, and so on down to low-level image features. Thus, understanding how to represent transformations and rich part-whole relationships in neural architectures has been a longstanding challenge in both deep learning and visual neuroscience. However, this aim has been complicated by the fact that sensory data alone is not sufficient to disambiguate between distinct world states. In this work, I propose a model for learning compositional representations of visual scenes which leverage Lie group structure, hierarchical composition, and bayesian inference. I demonstrate the many useful properties of this model on a variety of tasks such as object classification, scene segmentation, visual analogy,

and video compression. Further, I connect this work to many well-studied phenomena in primate visual cortex.

Christian Shewmake

University of California, Berkeley
cshewmake2@berkeley.edu

MS13

A Fast and Accurate Domain-Decomposition Non-linear Reduced Order Model Using Shallow Masked Autoencoders

Training reduced order models (ROMs) from data typically requires access to high-dimensional full order model (FOM) simulation data. However, for so-called "extreme-scale" problems, the storage of such high-dimensional FOM simulation data renders ROM training infeasible. Domain-decomposition (DD) alleviates this issue by solving the FOM on smaller subdomains, thereby generating training data of more manageable sizes. Model reduction can then be applied to each subdomain, and the separate ROMs can be reassembled to compute a global ROM for the DD FOM. A promising model reduction approach for the DD problem is the so-called nonlinear-manifold ROM (NM-ROM). NM-ROM has provided improved accuracy over linear-subspace ROMs (LS-ROMs), particularly for advection-dominated problems. NM-ROM approximates the FOM state in a nonlinear-manifold, which is learned from training a shallow, sparse-masked autoencoder using the FOM simulation data. The shallow, sparse architecture of the autoencoder allows for hyper-reduction to be applied, yielding computational speedup. In this talk, the DD formulation of the FOM and the application of NM-ROM to each subdomain are discussed. The results of the DD NM-ROM approach with hyper-reduction are numerically compared to DD LS-ROM with hyper-reduction for the 2D steady-state Burgers' Equation.

Alejandro N. Diaz
Rice University
and5@rice.edu

MS13

Intelligent Fire Evacuation Navigation Algorithm-Based on Internet of Things

With the huge complex electrical systems in the hospitality sector, problems such as imbalanced electrical loads, overloaded circuits, faulty fuses, etc., are common and may cause fire and eventually risk human life especially in high-rise buildings. The interior environments in such infrastructure can be complicated and diverse, which makes it important to safeguard people in the event of a fire and get them to the nearest safe exit quickly. Hence, in this work, we propose a Q-learning based approach to find a deadline-aware adaptive emergency navigation strategy in buildings with WSNs, which informs each customer about a hazard-avoided evacuation path to successfully reach safe exit gates within the specified deadline under all circumstances. For this, the proposed approach analyzes the intensity of fire to predict the specific limited evacuation time and then, an online navigation strategy is proposed based on a real-time adaptive routing algorithm to maximize and ensure the user's safety. The proposed approach is evaluated by conducting extensive simulations and prototype experiments.

Akila K.
B

Anna University Regional Campus Coimbatore
akila.karthikeyan25@gmail.com

MS13

Derivative Informed Neural Operators for PDE-Constrained Optimization under Uncertainty

We present a novel framework for solving PDE-constrained optimization problems under uncertainty through the use of derivative informed neural operators. When optimizing the performance of systems governed by PDEs, uncertainty in the model parameters lead to optimization under uncertainty (OUU) problems, where the optimization cost is defined in terms of risk measures of a quantity of interest. OUU problems are often orders of magnitude more expensive to solve compared to their deterministic counterparts due to the need to evaluate the risk measure by stochastic integration. This can require many evaluations of the governing PDE at every optimization iteration. To this end, we propose to use a reduced-basis network to approximate the PDE mapping from the input spaces of the uncertain parameter and the optimization variable to the output state, where derivative information of the PDE is used to determine the reduced-bases. We also incorporate the PDE derivatives into a Sobolev-type training loss to ensure that the neural operator has accurate derivatives with respect to the optimization variable, such that they are amenable to derivative based optimization algorithms when solving the OUU problem. In this talk, we will discuss the construction of such neural operators and its use as surrogates for OUU. We also demonstrate its performance and computational speed-ups over a suite of numerical examples.

Dingcheng Luo

Oden Institute for Computational Engineering and Sciences,
The University of Texas at Austin, Austin, TX, USA
dc.luo@utexas.edu

Thomas O'Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Peng Chen, Omar Ghattas
The University of Texas at Austin
peng@cc.gatech.edu, omar@oden.utexas.edu

MS13

Context-Aware Learning of Stabilizing Controllers in the Scarce Data Regime

Stabilizing dynamical systems in science and engineering is a challenging task, especially in edge cases and limit states where typically little data are available. In this work, we propose a data-driven approach that guarantees finding stabilizing controllers from as few data samples as the dimension of the unstable dynamics, which typically is orders of magnitude lower than the state dimension of the system. The key is learning stabilizing controllers directly from data without learning models of the systems, which would require larger numbers of data points. Numerical experiments with chemical reactors and fluid dynamics behind obstacles demonstrate that the proposed approach stabilizes systems after observing fewer than five data samples even though the dimension of states is orders

of magnitude higher.

Steffen W. R. Werner, Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
steffen.werner@nyu.edu, pehersto@cims.nyu.edu

MS13

Learning Stabilizing and Optimal Controllers from Data: An Informativity Approach

Learning accurate models of complex dynamical systems is a daunting task that requires large amounts of data. At the same time, robust control theory teaches us that near-perfect modelling is not required in all situations. Often, control policies can be obtained from data as long as the uncertainty is "not too large". Therefore, taking into account the eventual purpose of the data is of paramount importance when learning dynamics. In this talk, I will introduce the concept of data informativity for control. This concept captures the richness of data that is required for various data-based control problems. We will then discuss necessary and sufficient conditions under which a given dataset is informative for stabilization and for optimal control. These conditions also lead to new methods to design (optimal) controllers directly from given data via semidefinite programming.

Henk van Waarde
University of Groningen
h.j.van.waarde@rug.nl

MS14

Efficient Eigensolvers for Ab Initio Nuclear Physics Calculations

Microscopic calculations of the structure of atomic nuclei using the Configuration Interaction (CI) approach require computing the lowest eigenvalues and eigenvectors of a very large but extremely sparse symmetric matrix. The size of these matrices can be in the (tens of) billions, which require efficient algorithms and implementations on current HPC platforms. Many-Fermion Dynamicsnuclear (MFDn) is a CI code that obtains the lowest eigenpairs in large-scale ab-initio nuclear physics calculations, using either a Lanczos iterative eigensolver or Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) solver. Rapid convergence of the iterative LOBPCG algorithm is achieved by a suitable initialization in combination with an efficient (both in term of convergence and in terms of computational cost) preconditioner. MFDn is a hybrid MPI/OpenMP Fortran 95 code, with OpenACC for GPU offloading. I discuss some of the challenges and solutions for an efficient GPU implementation, and its performance on Perlmutter at NERSC. For large-scale calculations, data movement and in particular inter-node communication becomes a bottleneck, both with the Lanczos algorithm and with LOBPCG. On GPU systems, we have therefore also implemented a matrix-free version of the solver; furthermore, we are exploring the feasibility of compressing (and decompressing) the vectors in order to reduce the MPI overhead.

Pieter Maris
Iowa State University

pmaris@iastate.edu

MS14

Finding Quasiparticles in Dynamically Correlated Systems

Capturing the dynamics of electronic excitations in realistic systems containing more than a few electrons is one of the outstanding theoretical challenges. Dynamical quantum correlations mediate interactions and couplings between multiple excited states in materials and represent an important driver of their optoelectronic characteristics. Practical simulations resort to a quasiparticle picture in which the large-scale (and practically intractable) many-body problem is downfolded onto an effective single or few quasiparticle systems. This low-dimensional representation requires space-time non-local operators (capturing the renormalization of one- and two-body interactions) and leads to non-linear eigenvalue problems. I will discuss the first-principles techniques for studying individual excited states within the quasiparticle picture in molecules and large-scale condensed systems and outline the challenges associated with these approaches. I will exemplify these methods by exploring the correlated phenomena for localized moire states in twisted bilayer graphene and quantum defects.

Vojtech Vlcek
University of California Santa Barbara
vlcek@ucsb.edu

MS14

Perspective on High-Performance Algorithms for Eigenvalue Problems in Physical Simulations

Eigenvalue problems (EVPs) are ubiquitous in physical simulations and often represent a formidable computational bottleneck when simulating systems of growing size. Many forms of EVPs are commonly encountered in these simulations, ranging from large scale linear EVPs stemming from fundamental laws of physics, to non-linear EVPs which arise from dimensionality reduction and require application of robust numerical optimization algorithms to find physically-realizable solutions. In this work, we will examine several applications of large scaling EVPs in physical simulations which will be discussed in this minisymposium, the physically-motivated algorithmic motifs which make them computationally tractable, and the high-performance algorithms and software used to solve them on large computing resources.

David B. Williams-Young
Lawrence Berkeley National Laboratory
dbwy@lbl.gov

MS14

Nucleon Structure from a Light-Front Hamiltonian Approach

Nucleon structure encodes key information on the strong interaction, one of the four fundamental interactions in nature. Understanding the structure of the nucleon in terms of the quantum many-body problem of its constituents is at the frontier of particle and nuclear physics. Light-front quantum field theory provides a viable framework for accessing nucleon structure through large scale eigenvalue problems of the Hamiltonian of Quantum Chromodynamics (QCD), the fundamental theory of the strong

interaction. In this talk, I will introduce our parallel application named BLFQ (Basis Light-front Quantization) through which the nucleon structure problem is formulated as a large sparse symmetric eigenvalue problem of the QCD Hamiltonian. I will discuss the computational challenges in constructing and solving the eigenvalue problem of the QCD Hamiltonian as well as the features from the QCD Hamiltonian which can be exploited to boost the numerical efficiency. Finally I will demonstrate the numerical results as well as the scaling of BLFQ on high-performance computing platforms.

Xingbo Zhao
Chinese Academy of Sciences
xbzhao@impcas.ac.cn

Siqi Xu, Chandan Mondal
Institute of Modern Physics, Chinese Academy of Sciences
xsq234@impcas.ac.cn, mondal@impcas.ac.cn

Yang Li
University of Science and Technology of China
leeyoung1987@ustc.edu.cn

James Vary
Iowa State University
jvary@iastate.edu

MS15

Advances in Volume Penalization Methods for Multiphase Flows and FSI

The volume penalization method (VP), also known as the Brinkman penalization method, reduces the numerical complexity of solving PDEs defined over irregular domains by embedding them inside larger, regular domains. Most commonly, the VP method is used to solve FSI problems, although it has also been used to solve scalar advection-diffusion PDEs. The original VP technique was inspired by Brinkman's work and treats solids embedded in a fluid as porous media. The velocity boundary condition on the fluid-solid interface, which is of Dirichlet type, is imposed through a volumetric feedback force that is inversely proportional to the penalty parameter κ . Since both fluids and solids are represented on a single Eulerian grid, the VP method can be implemented in parallel codes much more easily than the dual grid (Lagrangian-Eulerian) IB method. VP method research has largely focused on understanding the accuracy of the penalized Dirichlet solution. There is a little research on (1) developing efficient solvers for the volume penalized fluid solvers, especially when the system becomes stiff as $\kappa \rightarrow 0$; and (2) imposing spatially varying Neumann and Robin boundary conditions on geometrically complex interfaces. In this talk we discuss recent advances in the VP method that address both of these issues.

Anneet P.S. Bhalla, Ramakrishnan Thirumalaisamy
San Diego State University
asbhalla@sdsu.edu, rthirumalaisam1857@sdsu.edu

Kaustubh Khedkar
Department of Mechanical Engineering
San Diego State University
kkhedkar9879@sdsu.edu

MS15

Lubricated Immersed Boundary Method with Ap-

plication to Fiber Bundles

Fluid-mediated near contact of elastic structures is a recurring theme in biofluids. The thin fluid layers that arise in applications such as the flow of red blood through blood vessels are difficult to resolve by standard computational fluid dynamics methods based on uniform fluid grids. A key assumption of the lubricated immersed boundary method, which incorporates a subgrid model to resolve thin fluid layers between immersed boundaries, is that the average velocity of nearby boundaries can be accurately computed from under-resolved simulations to bridge between different spatial scales. Here, we present a one-dimensional numerical analysis to assess this assumption and quantify the performance of the average velocity as a multiscale quantity. We explain how this analysis leads to more accurate formulations of the method and present examples from two-dimensional simulations, including applications to filament bundles.

Thomas Fai
Brandeis University
tfai@brandeis.edu

Yoichiro Mori
University of Pennsylvania
y1mori@math.upenn.edu

MS15

Multirate Sharp-Interface Methods for Fluid-Structure Interaction

This talk will describe ongoing work to create sharp-interface immersed boundary (IB) type methods for simulating the dynamics of fluids and immersed elastic structures. In such fluid-structure systems, it is common for the intrinsic timescales of the fluid and structure to differ, sometimes by orders of magnitude. However, it can be challenging to decouple these time scales in conventional IB formulations of fluid-structure interaction (FSI). This talk will outline a new sharp-interface approach to FSI that leverages a coupling scheme based on the immersed interface method. We use a penalty method to couple the fluid and structure that makes it straightforward to advance fluid and structure variables at different time step sizes, and we investigate the impact of multi-rate schemes on the accuracy and efficiency of the methodology in a range of benchmark cases and large-scale applications.

Boyce E. Griffith
University of North Carolina - Chapel Hill
boyceg@unc.edu

Ebrahim Kolahdouz
Simons Foundation - Flatiron Institute, U.S.
New York, NY
ekolahdouz@flatironinstitute.org

Qi Sun
University of North Carolina - Chapel Hill
qisun@unc.edu

MS15

Immersed Boundary Double Layer Method for Flows with Rigid Bodies

The Immersed Boundary (IB) method is useful for problems that involve fluid-structure interactions or complex geometries. By using a regular Cartesian grid that is inde-

pendent of the geometry, the IB framework yields a robust scheme that can efficiently handle immersed deformable structures. The IB method has also been adapted to problems with prescribed motion. IB methods for these problems traditionally involve penalty forces or they are formulated as constraint problems. In the latter approach, one must find the unknown forces by solving an equation that corresponds to a poorly conditioned first-kind integral equation. This operation can require a large number of iterations of a Krylov method, and since a time-dependent problem requires this solve at each time step, this method can be prohibitively inefficient without preconditioning. In this talk, we introduce a new, well-conditioned IB formulation for flows with rigid bodies, which we call the Immersed Boundary Double Layer (IBDL) method. In this formulation, the equation for the unknown boundary distribution corresponds to a well-conditioned second-kind integral equation that can be solved efficiently with a small number of iterations of a Krylov method without preconditioning. Furthermore, the iteration count is independent of both the mesh size and boundary point spacing. Additionally, while the original constraint method applies only to Dirichlet problems, the IBDL formulation can also be used for Neumann problems.

Brittany J. Leathers
University of California, Davis
brittany.leathers@gmail.com

Robert D. Guy
Mathematics Department
University of California Davis
guy@math.ucdavis.edu

MS15

A Discrete Leray Projection for Solving a Semi-Implicit Immersed Boundary Method in a Staggered Grid

In many real-world applications involved with interaction problems between a fluid and an immersed interface, the effects of high inertia and elasticity are very important. In this talk, we present an efficient and stable immersed boundary (IB) method for solving the motion of elastic interface, in particular, when the fluid has high inertia and the interface has strong elasticity. Our contributions are three folds. First, an iteration-free semi-Lagrangian approach is used in Navier-Stokes equations so that the high inertial effect of fluid is more stably simulated. Second, the elastic interfacial force is treated semi-implicitly allowing to handle strong elasticity and also to construct a resulting linear system. Last but not least, in solving the original linear system, we transform the 3-by-3 block matrix to a reduced 2-by-2 block matrix after eliminating the fluid pressure term. Such rank reduction is possible by applying the discrete Leray projection operator in a staggered MAC grid. By virtue of this feature, we call our method as a reduced immersed boundary method (rIBM). The equivalence between the rIBM and the original problem is proved in a theorem. Then an efficient algorithm for finding two unknowns is suggested using the Schur complement. We provide numerical results which show that the proposed rIBM improves both the numerical stability and the computational speed when solving more realistic problems.

Sangbeom Park, Soyeon Bak
Department of Mathematics, Kyungpook National University
Daegu 41566, Republic of Korea
piaoxf76@gmail.com, jiya525@knu.ac.kr

Philsu Kim
Kyungpook Nat'l Univ
kimps@knu.ac.kr

Yunchang Seol
Sungkyunkwan University
ycseol@skku.edu

MS17

Assessing the Performance of Data-Based and Physics-Based Model Order Reduction Techniques for nonlinear problems

Model order reduction aims to reduce the complexity of solving high-fidelity problems by projecting the entire system of equations onto a lower-dimensional subspace. In this contribution, we consider two ways of generating these projection functions data-based and physics-based approaches. In a data-based method called Proper Orthogonal Decomposition (POD), Singular value decomposition (SVD) is applied to training data to obtain the projection function. In contrast, in the physics-based approach, named Linear manifold(LM), the dynamic eigenmodes of the system are extended using modal derivatives that can capture the effect of nonlinear kinematics. In this contribution, we intend to model quasi-statics of the same high-fidelity problem to observe the difference between these methods. To this extent, we propose a residual parameter in the reduced space for both these methods and an additional mode selection algorithm for the physics-based method(LM). As a start, we assess the performance of both these methods on problems involving geometric non-linearity. The results showed that the displacement error for these methods for problems involving simple loading scenarios falls below 1 %, and an approx. time gain of 30 % with the original FE calculation. The difference in these methods has been visible in complex loading scenarios, where LM takes less number of modes in comparison to POD to reach an error below 1 %, but the time gain remained the same.

Phaniram Babbepalli, Joris Remmers
Eindhoven University of Technology
v.v.s.p.r.babbepalli@tue.nl, j.j.c.remmers@tue.nl

Olaf Van Der Sluis
Eindhoven University of Technology
Philips Research
o.v.d.sluis@tue.nl

MS17

Towards Efficient Dynamical Low-Rank Approximation of Collisional Kinetic Equations

Kinetic equations are of great importance to application areas such as plasma physics and rarefied gas dynamics. In regimes where collisions are very strong, efficient model reductions to hydrodynamic systems can be derived from asymptotic arguments such as the Chapman-Enskog expansion, reducing the dimensionality from 6 to 3. At the same time, dynamical low-rank approximation (DLRA) algorithms are particularly promising for kinetic equations because of their high dimensionality and the ‘tensor-friendly’ nature of their terms. However, it is not obvious how to combine the hydrodynamic model reduction approach with dynamical low-rank approximation algorithms, because the zeroth-order solution—the Maxwellian—is not a low-rank function of its arguments. We present a DLRA strategy based on an alternative low-

rank decomposition, focusing on a model plasma kinetic equation, which can capture the Maxwellian limit with a very small numerical rank. Our approach yields multiple orders of magnitude speedups on a 2D2V problem, while preserving the system's dynamics across a range of regimes.

Jack Coughlin
University of Washington
johnbc@uw.edu

MS17

Tensor Rank Reduction via Coordinate Flows

We present a new tensor rank reduction method that leverages coordinate flows and can greatly increase the efficiency of high-dimensional tensor approximation algorithms. The idea is very simple: given a multivariate function, determine a coordinate transformation so that the function in the new coordinate system has smaller tensor rank. We restrict our analysis to linear coordinate transformations, which give rise to a new class of functions that we refer to as tensor ridge functions. By leveraging coordinate flows and tensor ridge functions, we develop an optimization method based on Riemannian gradient descent that yields a quasi-optimal linear coordinate transformation for tensor rank reduction. The theoretical results we present for rank reduction via linear coordinate transformations can be generalized to larger classes of nonlinear transformations. We demonstrate the effectiveness of the proposed new tensor rank reduction method on prototype function approximation problems, and in computing the numerical solution of the Liouville equation in dimension three and five

Alec Dektor
University of California, Santa Cruz
adektor@ucsc.edu

Daniel Venturi
UC Santa Barbara
venturi@ucsc.edu

MS17

Low-Rank Approximations for Dynamic Particle Transport Problems

Dynamical low rank methods applied to particle transport problems have several open questions that I will address in this talk. Firstly, there is the question of whether or not typical problems actually have low rank structure and how this varies as a system evolves. We will answer this with examples from common test problems. One interesting result is that if the particle transport problem is decomposed into different octants based on the direction of particle travel there are different ranks in each octant. We show how this can be exploited to have more efficient calculations. Furthermore, we demonstrate that the adaptive low rank method recently developed can be used to find approximately the correct rank during a calculation.

Ryan McClarren
University of Notre Dame
USA
rmcclarr@nd.edu

MS17

A Hybrid AMR Low-Rank Tensor Approach for

Solving the Boltzmann Equation

The Boltzmann equation describes the time evolution of a particle distribution function in six-dimensional position-velocity phase space. The exponential growth in computational complexity often challenges a grid-based approach to modeling the Boltzmann equation as the dimensionality grows. To mitigate such issues, scalable low-rank tensor decomposition techniques have recently been developed with applications to high-dimensional PDEs. Despite the remarkable progress made in the community, low-rank structures in the phase-space are not evident in realistic engineering systems with complex geometries (e.g., electric propulsion systems and fusion reactors), where discontinuities, shocks, complex boundary conditions, and material-dependent physics (e.g., collisions, fusion reactions, ionization/excitation, charge-exchange processes) pose formidable challenges. In this talk, we propose a novel hybrid algorithm where quad-tree adaptive mesh refinement (AMR) is applied in real space while a low-rank approximation is applied in the velocity space. The AMR algorithm efficiently handles challenges pertaining to complex structures in real space, while the low-rank formulation targets dimensionality challenges in the velocity space. We present preliminary results on the new algorithm applied to challenging multi-dimensional gas kinetics problems.

William T. Taitano
Los Alamos National Laboratory
taitano@lanl.gov

Samuel J. Araki
Jacobs Inc.
Air Force Research Laboratory
samuel.araki.ctr@us.af.mil

MS18

A Nonconforming Substructuring Approach for Thermo-Mechanical Analysis of Micro-Electronics

Thermo-mechanical analysis of micro-electronics requires the coupled evaluation of multiple chips and solders to a single PBC. This problem is expensive to evaluate through full order finite-element (FE) models, as each chip can easily contain more than one million degrees-of-freedom. Even classical model order reduction methods struggle in this setting as reducing the full system model directly might be overly expensive. A solution to this would be to apply substructuring methods, where each chip and solder connection is reduced separately and then coupled together. However, in this case the large interfaces between the chips and solder, which leads to relatively large reduced order models. We propose the use of an inexact substructuring approach where the interfaces are described through a penalty formulation. This inexact interface enables additional flexibility on two fronts: on the one hand we are not bound by a consistent FE discretization on matching surfaces, and on the other hand we do not need to have conforming deformation modes on the interfaces. We exploit this framework to analyze transient thermal loads on a structural response of a PCB with multiple chips attached. We demonstrate the low computational cost of this approach obtained by the independence of the ROM size on the interface discretization, and the flexibility with which design changes can be evaluated due to the possibility of inconsistent discretizations on the interfaces.

Babak Bozorgmehri
DMMS Lab, Flanders Make at KU Leuven

Department of Mechanical Engineering, KU Leuven
babak.bozorgmehri@kuleuven.be

Frank Naets

Department of mechanical engineering, KU Leuven
frank.naets@kuleuven.be

MS19

Convergence Analysis of Multi-Step One-Shot Methods for Linear Inverse Problems

We consider general linear inverse problems for parameter identification where the corresponding forward and adjoint problems are solved iteratively. More precisely, we are interested in one-shot methods, which iterate at the same time on the forward/adjoint problem solution and on the inverse problem unknown. We analyze two variants of the so-called multi-step one-shot methods and establish sufficient conditions on the descent step for their convergence, by studying the eigenvalues of the block matrix of the coupled iterations. We illustrate numerically the convergence of these methods for a Helmholtz inverse problem and compare them with the classical usual and shifted gradient descent methods. In particular, we observe that very few inner iterations on the forward problem are enough to guarantee good convergence of the inversion algorithm.

Marcella Bonazzoli

Inria
Institut Polytechnique de Paris
marcella.bonazzoli@inria.fr

Housseem Haddar
CMAP, Ecole Polytechnique
housseem.haddar@inria.fr

Tuan Anh Vu
Inria
Institut Polytechnique de Paris
vtanh@math.ac.vn

MS19

A Matrix-Free Parallel Two-Level-Deflation Preconditioner for the Two-Dimensional Helmholtz Problems

We propose a matrix-free parallel two-level-deflation preconditioner combined with the Complex Shifted Laplacian preconditioner (CSLP) for the two-dimensional Helmholtz problems. The Helmholtz problem, widely studied in seismic exploration, is hard to solve both in terms of accuracy and convergence, due to the scalability issues of the numerical solvers. For large-scale applications, high-performance parallel scalable methods are also indispensable. In our method, we use the preconditioned Krylov subspace methods to solve the linear system obtained from finite-difference discretization. The CSLP preconditioner is approximately inverted by one parallel geometric multigrid V-cycle. Motivated by the observation that the eigenvalues of the CSLP-preconditioned system shift towards zero for large wavenumbers, deflation with multigrid vectors and further high-order vectors were incorporated to obtain wave-number-independent convergence. We also compare Galerkin coarsening method and high-order rediscrretization on the coarse grid. The matrix-vector products and the inter-grid operations are implemented based on the finite-difference grids without constructing the coefficient matrix. These adjustments lead to direct improve-

ments in terms of memory consumption. Numerical experiments show that wavenumber independence has been obtained for medium wavenumbers. The matrix-free parallel framework shows satisfactory parallel performance and weak scalability.

Jinjiang Chen

Technische Universiteit Delft
Technische Universiteit Delft
j.chen-11@tudelft.nl

Vandana Dwarka
Delft University of Technology
v.n.s.r.dwarka@tudelft.nl

Cornelis Vuik
Technical University of Delft
Delft Institute of Applied Mathematics
c.vuik@tudelft.nl

MS19

A Hybridizable Discontinuous Galerkin Method with Characteristic Variables for Helmholtz Problems

We will present a new hybridizable discontinuous Galerkin method, called CHDG, for solving time-harmonic scalar wave propagation problems. This method relies on a standard discontinuous Galerkin scheme with upwind numerical fluxes. Auxiliary unknowns corresponding to characteristic variables are defined at the interface between the elements, and the physical fields are eliminated to obtain a reduced system. The reduced system can be written as a fixed-point problem that can be solved with stationary iterative schemes. Numerical results with 2D benchmarks will be presented to study the performance of the approach. Compared to the standard HDG approach, the properties of the reduced system are improved with CHDG, which is more suited for iterative solution procedures with e.g. GMRES or CGN. Reference: <https://arxiv.org/abs/2212.11529>

Axel Modave

CNRS - ENSTA-ParisTech
axel.modave@ensta-paris.fr

Théophile Chaumont-Frelet
Inria Sophia-Antipolis
theophile.chaumont@inria.fr

MS19

Scalable Machine Learning to Analyze Rocket Combustion Data

Scalable machine learning algorithms were used to analyze hybrid rocket combustion experiments. Recently, combustion tests with different paraffin-based fuels have been performed at the German Aerospace Center (DLR) and the whole process has been recorded with a high-speed video camera. This has led to a larger number of images that needs to be automatically analyzed [Ruetters et al., Clustering of paraffin-based hybrid rocket fuels combustion data. *Exp. Fluids*, 61:4 (2020)]. Since data analysis of millions of images is very demanding with respect to computing time and memory requirement, HeAT, the Helmholtz Analytics Toolkit, is used on a parallel cluster at DLR to reduce the computing time. HeAT is an open source Python library (<https://github.com/helmholtz-analytics/heat>) for parallel machine learning. It is jointly developed by

Juelich Research Center, Karlsruhe Institute of Technology and DLR [Goetz et al., HeAT-a Distributed and GPU-accelerated Tensor Framework for Data Analytics. 2020 IEEE International Conference on Big Data (2020), pp. 276–287]. In this talk, machine learning algorithms are used to identify different flow phases and to detect anomalies [Ruettggers and Petrarolo, Local Anomaly Detection in Hybrid Rocket Combustion Tests. *Exp. Fluids*, 62:136 (2021)]. Furthermore, we present results on the parallel scaling behaviour of Heat compared to other existing parallel machine learning libraries.

Alexander Ruettggers

German Aerospace Center (DLR)
Institute for Simulation and Software Technology
alexander.ruettggers@dlr.de

Anna Petrarolo

German Aerospace Center (DLR)
Institute of Space Propulsion
anna.petrarolo@dlr.de

Philipp Knechtges

German Aerospace Center (DLR)
philipp.knechtges@dlr.de

MS19

Analysis of a Domain Decomposition Method for Convected Helmholtz Like Equation

In this work, we are interested in solving a convected Helmholtz like equation. This type of PDE occurs in several contexts (acoustic flow, Schrödinger equation, Wave-ray method...), and it raises similar difficulties as the classical Helmholtz equation. Indeed, in the case of constant parameter, one can show that with an appropriate change of variable the PDE can be reformulated as the classical Helmholtz equation. This reformulation allows to get several byproducts, such as the construction of absorbing boundary conditions, a PML formulation, and in particular to apply the Fourier analysis to study the convergence properties of an iterative Schwarz algorithm. In this talk, we will detail the Fourier analysis and emphasise the impact of the boundary conditions to end the computational domain.

Antoine Tonnoir

National Institute of Applied Sciences of Rouen
atonnoir@insa-rouen.fr

Martin J. Gander

University of Geneva
martin.gander@unige.ch

MS20

On the Surface Effect and the Imposition of the Bc in Peridynamics

Peridynamics is a non-local continuum theory: two inter-related problems affect it at the boundary of the domain: the surface effect and a difficulty in the imposition of the boundary conditions. Since the points near the boundary are characterised by an incomplete neighbourhood we observe an unrealistic variation of the stiffness properties in the most external layers of the domain, the surface effect. Two main methods are used to correct it: the definition of modified bonds in the external layers of the domain or the introduction of a fictitious layer of nodes around its

boundary. The imposition of the boundary conditions is often achieved by using a fictitious layer, in the case of displacement bc, or by distributing the external tractions as body forces applied to a certain number of nodes on or close to the boundary. Imposing the boundary conditions as one would do in a local model, generates large fluctuations of the solution. There is no general agreement on the way to distribute the boundary loads or displacements over some finite layer. We will present 1-2-3D examples of the application of a recently proposed version of the Taylor-based extrapolation method that makes use of the nearest-node strategy. The method uses a fictitious layer of points around the body where the displacements are determined as functions of the displacements of their closest true nodes by means of multiple Taylor series expansions. The examples will include as well dynamic cases.

Ugo Galvanetto

University of Padova
ugo.galvanetto@unipd.it

MS20

Algebraic Multigrid Solver for Nonlocal Equations

The naive discretization of nonlocal operators leads to matrices with significant density, as compared to classical PDE equations. This makes the efficient solution of nonlocal models a challenging task. In this presentation, we will discuss on-going research into assembly using hierarchical matrices and algebraic multigrid solution techniques that are suitable for nonlocal models.

Christian Glusa

Sandia National Laboratories
caglusa@sandia.gov

MS20

A Hybrid Meshfree Discretization to Improve the Numerical Performance Of Linearized Peridynamic Elasticity And Corrosion Models

The numerical computation of spatial integrals is crucial in determining the accuracy and efficiency of peridynamic models. In this work, we propose a meshfree discretization scheme based on a collocation approach, which takes a significantly smaller number of neighborhood nodes as compared to the classical particle-based scheme, in order to achieve a certain accuracy while maintaining the same length scale. Our proposed discretization procedure is employed within the regions of the computational domain, where the field variable can be approximated by smooth basis functions inside the neighborhood. In the hybrid approach, our method is then combined with the classical scheme, the latter being applied within the regions subject to discontinuities and material damage. We illustrate the increase in efficiency of our proposed discretization scheme due to the reduced number of neighborhood nodes and evaluate the convergence behavior. As numerical examples of the hybrid approach, we adopt the brittle fracture problem by means of bond-based peridynamics as well as the peridynamic corrosion model in two and three dimensions, respectively.

Alexander Hermann

Institutes of the Helmholtz-Zentrum Hereon

alexander.hermann@hereon.de

MS20

Peridynamics Computations at the Exascale

Peridynamics is a nonlocal reformulation of classical continuum mechanics suitable for material failure and damage simulation, which has been successfully demonstrated as an effective tool for the simulation of complex fracture phenomena in many applications. However, the nonlocal nature of peridynamics makes it highly computationally expensive, compared to classical continuum mechanics, which often hinders large-scale fracture simulations. In this talk, we will present a GPU-enabled, performance portable, and exascale-capable peridynamics code designed to run on United States Department of Energy's supercomputers, in particular Frontier (currently ranked #1 in the TOP500 list) and Summit (currently ranked #4 in the TOP500 list), the two fastest supercomputers in the United States, both at Oak Ridge National Laboratory.

Pablo Seleson

Oak Ridge National Laboratory
selesonpd@ornl.gov

MS21

Asymptotic Analysis Applied to Small Volume Inverse Shape Problems

We consider two inverse shape problems coming from diffuse optical tomography and inverse scattering. For both problems, we assume that there are small volume subregions that we wish to recover using the measured Cauchy data. We will derive an asymptotic expansion involving their respective fields. Using the asymptotic expansion, we derive a MUSIC-type algorithm for the Reciprocity Gap Functional, which we prove can recover the subregion(s) with a finite amount of Cauchy data. Numerical examples will be presented for both problems in two dimensions in the unit circle.

Giovanni Granados

Purdue University
ggranad@purdue.edu

MS21

Regularization of the Factorization Method with Applications

In this talk, we discuss a new regularized version of the Factorization Method. The Factorization Method uses Picard's Criteria to define an indicator function to image an unknown region. In most applications, the data operator is compact which gives that the singular values can tend to zero rapidly which can cause numerical instabilities. The regularization of the Factorization Method presented here seeks to avoid the numerical instabilities in applying Picard's Criteria. This method allows one to image the interior structure of an object with little a priori information in a computationally simple and analytically rigorous way. Here we will focus on an application of this method to diffuse optical tomography which will prove that this method can be used to recover an unknown subregion from the Dirichlet-to-Neumann mapping.

Isaac Harris

Purdue University

harri814@purdue.edu

MS21

Single Mode Multi-Frequency Factorization Method for the Inverse Source Problem in Acoustic Waveguides

This talk is to address the inverse source problem with a single propagating mode at multiple frequencies in an acoustic waveguide. The goal is to provide both theoretical justifications and efficient algorithms for imaging extended sources using the sampling methods. In contrast to the existing far/near field operator based on the integral over the space variable in the sampling methods, a multi-frequency far-field operator is introduced based on the integral over the frequency variable. This far-field operator is defined in a way to incorporate the possibly non-linear dispersion relation, a unique feature in waveguides. The factorization method is deployed to establish a rigorous characterization of the range support which is the support of source in the direction of wave propagation. A related factorization-based sampling method is also discussed. These sampling methods are shown to be capable of imaging the range support of the source. Numerical examples are provided to illustrate the performance of the sampling methods, including an example to image a complete sound-soft block.

Shixu Meng

Academy of Mathematics and Systems Science
Chinese Academy
shixumeng@amss.ac.cn

MS21

Using New Sets of Eigenvalues to Image Obstacle Density

The Linear Sampling Method (LSM) aims to image the shape of defects in a known medium from measurements of scattered fields and generally provides satisfactory solutions if multi-static (far field) data is available. However, this method becomes less effective for cluttered defects (for instance a network of cracks). In this talk, we shall present a recently introduced imaging method based on the LSM formalism but that exploits the spectrum of an eigenvalue problem. This eigenvalue problem is specifically designed by the introduction of a modified background so that its analysis in terms of defects size is possible. We prove that the eigenvalues of this problem can be identified from multi-static data. We then exploit this fact to construct an indicator function of the defects density based on the sensitivity of the eigenvalues to the presence and size of these defects. We shall illustrate the obtained inverse algorithm and its performances on some synthetic numerical experiments.

Fabien Pourre

ENSTA Paris
fabien.pourre@inria.fr

Houssem Haddar

CMAP, Ecole Polytechnique
houssem.haddar@inria.fr

Lorenzo Audibert

EDF R&D PRISME

lorenzo.audibert@edf.fr

MS22

Machine Learning-Enhanced Refinement and Agglomeration Strategies for Polygonal and Polyhedral Methods

In this talk we discuss how to enhance the accuracy and performance of Polyhedral Finite Element Methods based on designing suitable Machine Learning-aided numerical algorithms to handle the process of grid refinement and agglomeration. More specifically, we propose new strategies to handle polytopal grid refinement, to be employed within an adaptive framework. Specifically, Convolutional Neural Networks are employed to classify the shape of an element so as to apply ad-hoc refinement criteria or to enhance existing refinement strategies at a low online computational cost. We test the proposed algorithms considering two families of finite element methods that support arbitrarily shaped polytopal elements, namely the Virtual Element method and the Polytopal Discontinuous Galerkin method. In the second part of the talk ML-aided grid agglomeration techniques are presented. Mesh agglomeration strategies are important both within adaptive refinement algorithms and to construct multilevel algebraic solvers. We propose to use Graph Neural Networks (GNNs) to automatically perform grid agglomeration. GNNs have the advantage to process naturally and simultaneously both the graph structure of mesh and the geometrical information. We assess the performance of the proposed agglomeration algorithm and demonstrate its effectiveness when employed within multigrid solvers in a Polytopal Discontinuous Galerkin framework.

Paola F. Antonietti

MOX-Dipartimento di Matematica
Politecnico di Milano
paola.antonietti@polimi.it

MS22

Divergence-Conforming Pressure Robust Methods for Fluid Flow

This talk presents new matrix finite elements for approximating viscous stresses in fluids. These stress finite elements have shear continuity. Such stress spaces pair naturally with velocity approximations in $H(\text{div})$, the space of vector fields with square integrable components and divergence. This work can be viewed as natural continuation of a series of developments by multiple authors in the treatment of the incompressibility constraint using the Sobolev space $H(\text{div})$. The stress finite element space can be seen as arising from a nonstandard Sobolev space $H(\text{curl div})$. We will show that structure-preservation properties like mass conservation and pressure robustness are immediate in the newly introduced mixed finite element methods.

Jay Gopalakrishnan

Portland State University
gjay@pdx.edu

MS22

Fem Approximations of the Operator Lyapunov Equation Have Low Rank

We present a low-rank adaptive hp-finite element algorithm for computing a low rank approximation to the solution of the Lyapunov operator equation. Our adaptive finite element

algorithm is based on the auxiliary subspace error estimation technique. The use of hp-adapted finite element methods in this context is justified by showing that the eigenfunctions of the solution operator are A-analytic functions. Here A denotes the coefficient operator for the Lyapunov equation. A spectral theoretic setting of Babuska and Osborn is used both to define measures of the approximation error as well as to prove the reliability of the error estimator. On the example of the Laplace operator on the dumbbell domain we achieve eight figures of accuracy for computing the trace of the solution of the Lyapunov equation using a finite element space of dimension of only $1e4$ degrees of freedom. Even more surprising is the observation that hp-refinement has an effect of reducing the rank of the numerical approximation of the solution.

Luka Grubisic

University of Zagreb
luka@math.hr

Harri H. Hakula

Aalto University
Harri.Hakula@aalto.fi

MS23

Close Evaluation of Layer Potentials

I will discuss propulsion by flexible appendages, as in swimming or flying. A small-amplitude flow model combined with an efficient Chebyshev PDE solver and conformal mapping techniques enables efficient numerical solutions. These solutions elucidate the performance benefits garnered by optimizing flexibility characteristics of appendages. The method also generalizes to multiple flapping bodies through the so-called prime function. New results shed light on flow-mediated collective behavior, such as schooling and flocking.

Shilpa Khatri

Applied Mathematics
University of California, Merced
skhatri3@ucmerced.edu

MS23

Steklov-Poincaré Analysis of the Basic Three-Domain Stent Problem

The Steklov-Poincaré problem was previously considered in the artery lumen and wall setting with a single interface. Here the analysis is expanded to incorporate solute behavior in the presence of a fixed-volume, solid, simple stent. In this geometry, a third domain is added to the two-domain structure of wall and artery. Through this intersecting domain volume setting there are three interfaces: lumen-wall, stent-lumen, and wall-stent. Steady-state incompressible Navier-Stokes equations are used to explain the behavior of blood through the lumen, while advection-diffusion dynamics are considered for the solute mechanics across the lumen, wall, and stent. Having a fixed blood velocity value, Steklov-Poincaré decomposition of the advection-diffusion equations is applied locally to each of the interfaces. To unify these instances on a global scale, their overall intersection is explored in a smaller manifold, reducing the problem to one previously solved by Quarteroni, Veneziani, and Zunino. Through finite element analysis (FEM), the solution is discretized and found to be convergent. Finally, computational simulations of different types of stent of one, three, and five stent rings of different thickness, placed between straight and curved volumes of inner and outer

cylindrical meshes, were performed using NGSolve, confirming the convergence of the solution and its relation to the coarseness of the mesh.

Irving Martinez
Emory University
irving.martinez@emory.edu

Alessandro Veneziani
MathCS, Emory University, Atlanta, GA
ale@mathcs.emory.edu

MS23

Propulsion by Flexible Appendages and Hydrodynamic Schooling

I will discuss propulsion by flexible appendages, as in swimming or flying. A small-amplitude flow model combined with an efficient Chebyshev PDE solver and conformal mapping techniques enables efficient numerical solutions. These solutions elucidate the performance benefits garnered by optimizing flexibility characteristics of appendages. The method also generalizes to multiple flapping bodies through the so-called prime function. New results shed light on flow-mediated collective behavior, such as schooling and flocking.

Nick Moore
Colgate University
nickmoore83@gmail.com

MS23

Reduced-Order Model of Hydrodynamically Interacting Flapping Swimmers

Fish schools are examples of collective motion in which hydrodynamic interactions play a role in dynamical behavior and self-organization principles. However, the long-time evolution of hydrodynamically interacting collectives is challenging to investigate due to the persistent influence of long-lived vortical structures, and the high-resolution requirements of direct numerical simulation at large Reynolds numbers. Reduced-order models have therefore played an important role in theoretical investigations. We introduce a new reduced-order model of swimmers that self-propel by flapping, i.e., executing a prescribed periodic rigid body motion. The model is an extension of a discrete-time dynamical system developed by Oza, Ristroph and Shelley in which flapping swimmers interact through periodically shed vortices. Our work allows a variable separation distance between swimmers and also describes more faithfully the wing-wing interactions of our dynamical system. The model is used to investigate conditions under which hydrodynamic interactions lead to stable swimming configurations. Preliminary numerical results are presented.

Jose L. Pabon
New Jersey Institute of Technology.
jlp43@njit.edu

MS24

High Dynamic Range Tomography via Modulo Radon Transform

The topic of high dynamic range tomography is starting to gather attention due to recent advances in hardware technology. The issue is that registering high-intensity projec-

tions that exceed the dynamic range of the detector cause sensor saturation. While existing methods typically rely on the fusion of multiple exposures, we propose a one-shot computational imaging solution based on the novel Modulo Radon Transform (MRT). The MRT generalizes the conventional Radon Transform and is obtained via computing modulo of the line integral of a two-dimensional function at a given angle. In this way, the MRT encodes folded Radon projections and avoids information loss arising from saturation or clipping effects. In this talk, we rigorously introduce the Modulo Radon Transform and propose a sequential reconstruction algorithm for finitely many Modulo Radon projections, which is backed by mathematical guarantees. Our recovery strategy makes use of a property we call compact λ -exceedance, which is motivated by practice; in applications the object to be recovered is of finite extent and the measured quantity has approximately compact support. Our theoretical results are illustrated by numerical simulations, where we report the reconstruction of target functions with Radon projections 1000-times larger in amplitude than the assumed modulo threshold. This talk is based on joint work with Ayush Bhandari (Imperial College London) and Felix Kraemer (Technical University Munich).

Matthias Beckmann
University of Hamburg
Department of Mathematics
matthias.beckmann@uni-bremen.de

MS24

Instabilities of Relative Pose Estimation and the Role of RANSAC

Structure-from-motion in computer vision leads to problems in 3D geometry. These are cast as multivariate polynomial systems, which need to be solved many times for one scene reconstruction. In this talk, I will present a framework for quantifying the numerical stability of such systems. It relates the 3D world scene to the 2D image data. Configurations with infinite condition number are characterized both in the world and in the image. In particular, this gives explicit formulas for certain discriminant hypersurfaces. We connect the theory to standard estimation procedures in computer vision, and argue that there exists a real opportunity for improvements. Based on joint work with Hongyi Fan (Cognex) and Ben Kimia (Brown).

Joe Kileel
University of Texas at Austin, U.S.
jkileel@math.utexas.edu

MS24

Robust Rotation Averaging via Quadratic Programming

We propose a novel quadratic programming formulation for estimating the corruption levels of the relative camera orientations, and use these estimates to solve rotation averaging problem. Our objective function exploits the cycle consistency of the relative rotations and we thus refer to our method as detection and estimation of structural consistency (DESC). This general framework can be extended to other algebraic and geometric structures. Our formulation has the following advantages: it can tolerate corruption as high as the information-theoretic bound, it does not require a good initialization for the estimates of absolute orientations, it has a simple interpretation, and under some mild conditions the global minimum of our objective

function exactly recovers the corruption levels. We demonstrate the competitive accuracy of our approach on both synthetic and real data experiments of rotation averaging.

Yunpeng Shi
Princeton University
yunpengs@princeton.edu

MS24

Counting Objects by Diffused Index: Geometry-Free and Training-Free Approach

Counting objects is a fundamental but challenging problem. In this paper, we propose diffusion-based, geometry-free, and learning-free methodologies to count the number of objects in images. The main idea is to represent each object by a unique index value regardless of its intensity or size, and to simply count the number of index values. First, we place different vectors, referred to as seed vectors, uniformly throughout the mask image. The mask image has boundary information of the objects to be counted. Secondly, the seeds are diffused using an edge-weighted harmonic variational optimization model within each object. We propose an efficient algorithm based on an operator splitting approach and alternating direction minimization method, and theoretical analysis of this algorithm is given. An optimal solution of the model is obtained when the distributed seeds are completely diffused such that there is a unique intensity within each object, which we refer to as an index. We refer to this approach as Counting Objects by Diffused Index (CODI). We explore scalar and multi-dimensional seed vectors. For Scalar seeds, we use Gaussian fitting in histogram to count, while for vector seeds, we exploit a high-dimensional clustering method for the final step of counting via clustering. We present counting results in various applications such as biological cells, agriculture, concert crowd, and transportation. Some comparisons with existing methods are presented.

Maryam Yashtini
Department of Mathematics and Statistics
Georgetown University
my496@georgetown.edu

MS25

Extreme-Scale CSE Challenges: Progress via Community Solutions

Extreme-scale CSE projects involve many people and skill sets, and complex hardware and software environments. Technical excellence is essential in these environments, but success also hinges on collaboration, complementarity, and coordination of efforts within and across teams and communities. In this talk, we discuss efforts in the Exascale Computing Project to produce community-driven ecosystems and team structures that support our abilities to holistically pursue the goals of extreme-scale computational science. We highlight ongoing efforts to create hierarchical software organizations and stacks, and emerging efforts to improve scientific software development and use via the application of social and cognitive sciences to software teams and communities.

Michael A. Heroux
Sandia National Laboratories
St. John's University

maherou@sandia.gov

MS25

Combining HPC, AI and RDM: Challenges and Approaches

The availability of large amount of digital data, computational power combined with new approaches of data analysis and so called artificial intelligence created the potential of new scientific discoveries in almost all scientific fields. However, to realize the potential a number of challenges need to be addressed. The amount, quality and accessibility of data has to be suitable. In short, the FAIR principles need to be implemented in a proper manner and the meta data has to be suitable and sufficient for the selected machine learning method. Many of the challenges can only be addressed if the IT infrastructure offers suitable and comprehensive support in the fields of HPC, AI, and RDM. This talk will present the approach taken by RWTH Aachen University and its partners to address the challenges. The approach includes 1. Engagement in several projects of the German Initiative for Research Data Management (NFDI) with a focus on engineering. 2. Creation of a data integration platform Coscine implementing the FAIR principles and FAIR Digital Object interfaces. 3. Process Mining technologies to discover similarities between data sets and reveal otherwise undetected workflows. 4. Data stewards in large, long time collaborative projects (CRCs) 5. Integration of HPC and RDM platforms to exploit data analytics capabilities of HPC. Besides the scientific content the talk will cover the technical, financial, and human resources of the activities.

Matthias Müller, Christian Terboven, Marcel Nellesen
RWTH Aachen University
mueller@itc.rwth-aachen.de,
terboven@itc.rwth-aachen.de, nellesen@itc.rwth-aachen.de

M. Amin Yazdi
RWTH Aachen University
yazdi@itc.rwth-aachen.de

Marius Politze, Ilona Lang, Benedikt Heinrichs, Florian Claus, Thomas Eifert
RWTH Aachen University
politze@itc.rwth-aachen.de, lang@itc.rwth-aachen.de,
heinrichs@itc.rwth-aachen.de, claus@itc.rwth-aachen.de,
eifert@itc.rwth-aachen.de

MS25

h3-Open-BDEC: Innovative Software Infrastructure for Scientific Computing in the Exascale Era by Integrations of (Simulation + Data + Learning)

We propose an innovative method for computational science for sustainable promotion of scientific discovery by supercomputers in the Exascale Era by combining (Simulation + Data + Learning (S+D+L)). In May 2021, we started operation of the Wisteria/BDEC-01 system with 33+PF at University of Tokyo. It is a Hierarchical, Hybrid, Heterogeneous (h3) system, which consists of computing nodes for CSE with A64FX and those for Data Analytics/AI with NVIDIA A100 GPUs. We develop a software platform h3-Open-BDEC for integration of (S+D+L) and evaluate the effects of integration of (S+D+L) on the Wisteria system. The h3-Open-BDEC is designed for extracting the maximum performance of the supercomputers with minimum energy consumption focusing on (1) innovative method for numerical analysis with

high-performance/high-reliability/power-saving based on the new principle of computing by adaptive precision, accuracy verification and automatic tuning, (2) Hierarchical Data Driven Approach (hDDA) based on machine learning, and (3) Software for heterogeneous systems, such as Wisteria/BDEC-01. Integration of (S+D+L) by h3-OpenBDEC enables significant reduction of computations and power consumption, compared to those by conventional simulations.

Kengo Nakajima

Information Technology Center, The University of Tokyo
RIKEN Center for Computational Science (R-CCS)
nakajima@cc.u-tokyo.ac.jp

MS25

On the Road to Brain-Scale Datasets and Applications for Post-Exascale Supercomputers

Several machine learning methods, such as Graph Convolutional Networks, would manipulate graphs with dozens of billions of nodes. Such brain-scale datasets are represented by very large sparse adjacency matrices, often associated with Laplacian or transition matrices. Moreover, rectangular skinny matrices may store features of nodes. Analyses of those data, such as rankings, homophily evaluations, and clusters detections, for example, are often required by those applications and methods. Moreover, the highly hierarchical architectures of the supercomputers allowing such computation, from multicore chips having network on chip, to distributed nodes and data storage units, are not always well-adapted for such computation manipulating non-structured data. Understanding the behaviors of such methods is important and lead to large number of experiments, asking to have fast accesses to many different datasets. In this talk, we first present two dataset generators, directly computed in parallel, allowing to experiment very large problems without any I/O. We introduce an open source software which generates very large datasets inspired from an easy brain structure, experimented on distributed platform and supercomputers. Then, as an example of applications on such graphs, we present results of PageRank methods experimented on several supercomputers. We conclude with analyses of obtained performances with respect of the dataset structures and of the supercomputer architectures

Serge G. Petiton

University of Lille, and. CNRS
serge.petiton@univ-lille.fr

Xinzhe Wu

Jülich Supercomputing Centre
xin.wu@fz-juelich.de

Nicolas Hochart

CRSITAL, University of Lille
nicolas.hochart@polytech-lille.net

MS26

Exploring Mesh Partitioning with the Coupe Partitioning Platform

Coupe is a dedicated mesh partitioner, written in Rust. It implements, in parallel using shared memory, different partitioning models for load balancing mesh based scientific simulations. We will present how coupling a number of partitioning algorithms along with geometric partitioning and topological refinement can become an alternative to

classical multi-level graph partitioners.

Cedric Chevalier

CEA, France
cedric.chevalier@cea.fr

Sébastien Morais

CEA, France
sebastien.morais@cea.fr

Franck Ledoux

CEA
franck.ledoux@cea.fr

Hubert Hirtz

CEA, France
hubert.hirtz@cea.fr

MS26

Design and Implementation of Multi-Threaded and Hybrid Parallel Graph Partitioning Algorithms in Scotch v7

Graph partitioning is a ubiquitous problem which has many applications in scientific computing. Due to the ever increasing size of the problems to solve, many parallel implementations of graph partitioning algorithms have been proposed in the literature, whether for shared-memory multiprocessors or distributed-memory multicomputers. This paper describes the design and implementation of multi-threaded algorithms in version 7 of the Scotch partitioning package. These algorithms concern both the formerly sequential version, Scotch, and the parallel, distributed-memory version, PT-Scotch. Notably, a hybrid parallel (MPI+threads) graph coarsening algorithm is proposed, which is used to accelerate the classic multi-level partitioning framework. In order to provide scientific reproducibility, deterministic algorithms have been implemented whenever possible, and can be selected at the user's choice. Concurrent execution is made possible by the encapsulation of partitioning tasks within execution contexts. While a global linear speedup is out of reach due to the many remaining non-parallel sections, the multi-threaded algorithms evidence high scalability themselves, and provide for a significant improvement in run time, without any loss in partition quality.

François Pellegrini

l'université de Bordeaux and INRIA
francois.pellegrini@u-bordeaux.fr

MS26

Recent Advances in Streaming (Hyper)Graph Partitioning

Partitioning a (hyper)graph into balanced blocks such that few edges run between blocks is a key problem for large-scale distributed processing. Currently there is a gap in the space of available algorithms. On the one hand, streaming algorithms have been adopted to partition massive graph data on small machines. In the streaming model, vertices arrive one at a time and then are directly assigned to a block. These algorithms partition huge graphs quickly with little memory, but produce partitions with low solution quality. On the other hand, there are offline (shared-memory) multilevel algorithms that produce partitions with high quality but also need a machine with enough memory to partition huge networks. In this talk, we present recent advances in streaming algorithms for the

problem. First, we present a buffered streaming approach: this model allows to read more than one node and its neighborhood at the time. This enables our algorithm to leverage multilevel techniques, and thus significantly improve solution quality while surprisingly also enhancing the overall complexity of the algorithm. On the other hand, we present a shared-memory streaming multi-recursive partitioning scheme that performs recursive multi-sections on the fly without knowing the overall input graph to compute hierarchical partitionings. If the topology of a distributed system is known, it is possible to further optimize the communication costs by mapping partitions onto processing elements.

Marcelo Fonseca Faraj
Heidelberg University
marcelofaraj@informatik.uni-heidelberg.de

Christian Schulz
Heidelberg University, Germany
christian.schulz@informatik.uni-heidelberg.de

MS26

Pebbling Game and Alternative Basis for High Performance Matrix Multiplication

Matrix multiplication is one of the most extensively used kernels in scientific computing. Although sub-cubic algorithms exist, most high performance implementations are based on the classical $\Theta(n^3)$ matrix multiplication. Designing an algorithm that obtains even modest improvements in performance over existing implementations, requires carefully addressing challenges such as reducing computation costs, communication costs, and memory footprint. We provide the first high performance general matrix-matrix multiplication that utilizes the alternative basis method on Strassen's algorithm. We reduce the basis transformations overheads and decrease the memory footprint of the bilinear phase by using the pebbling game optimization scheme, consequentially improving both arithmetic and communication costs. Our algorithm outperforms DGEMM on feasible matrix dimensions starting at $n = 1024$. It obtains an increasing speedup up to nearly $\times 2$ speedup for larger matrix dimensions. Joint work with Oded Schwartz.

Noa Vaknin
The Hebrew University
noa.vaknin1@mail.huji.ac.il

Oded Schwartz
Hebrew University of Jerusalem
odedsc@cs.huji.ac.il

MS27

Solving the Integro-Differential Equation of Supercooled Liquid Dynamics Using Machine Learning

We introduce a machine-learning approach to predict the complex non-Markovian dynamics of supercooled liquids from their static averaged quantities. Our method is based on a theoretical framework that uses as input and output system-averaged quantities and describe the dynamics as an integro-differential equation. Compared to a particle resolved approach, our method is easier to apply in an experimental context where particle specific information is not available. First, using a deep neural network we predict the self intermediate scattering function of the binary mixtures we investigate. While its performances

are excellent when training data across all the temperature range are available, we also show that the model retains some transferability being able to make decent predictions at temperatures lower than the one it was trained for, or when we use it for similar systems. Then, we develop an evolutionary strategy that is able to predict the elusive memory function underlying the integro-differential equation describing the observed dynamics. This method, which is much easier than any Laplace inversion of the memory equation, lets us conclude that the memory function of supercooled liquids can be effectively parameterized as the sum of two stretched exponentials, which physically corresponds to two dominant relaxation modes.

Simone Ciarella
École normale supérieure, Paris
simoneciarella@gmail.com

MS27

Adaptive Marking for Adaptive Mesh Refinement via Reinforcement Learning

Adaptive mesh refinement in the context of adaptive finite element methods (AFEM) requires an often overlooked and time-consuming offline effort: tuning parameters that control which elements are marked for refinement. To automate this effort and improve efficiency of adaptive methods, we recast adaptive mesh refinement as a partially-observed Markov decision process that can be optimized using methods from reinforcement learning. This recasting delivers a tractable optimization framework that can be tuned once by automated training and then deployed successfully in contexts outside the training regime. We use the Poisson equation to demonstrate various applications of our framework including representative h- and hp-refinement AFEM simulations on non-convex polyhedra. Our experiments indicate that superior marking policies remain undiscovered for many canonical AFEM applications. An unexpected observation is that marking policies trained on one family of PDEs can be robust enough to perform well on problems far outside the training family. For instance, we show that a simple hp-refinement policy optimized on 2D problems can be used for 3D problems without significant performance loss. Extensions of these ideas to transient PDEs and multi-objective optimization frameworks will also be discussed.

Andrew Gillette
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
gillette7@llnl.gov

Brendan Keith
Brown University
brendan_keith@brown.edu

Socratis Petrides
Lawrence Livermore National Laboratory
petrides1@llnl.gov

MS27

Mutual Information Based Uncertainty Quantification in Scientific Machine Learning

Neural networks (NNs) are currently changing the computational paradigm on how to combine data with mathematical laws in physics and engineering in a profound way, tackling challenging inverse and ill-posed problems not solvable with traditional methods. However, quan-

tifying errors and uncertainties in NN-based inference is more complicated than in traditional methods. We have presented a comprehensive framework that includes uncertainty modeling, new and existing solution methods, as well as evaluation metrics and post-hoc improvement approaches in the review work *Uncertainty Quantification in Scientific Machine Learning: Methods, Metrics, and Comparisons*. In this talk, we will present a new approach for UQ in scientific machine learning based on the mutual information theory. Some numerical examples are tested to demonstrate the applicability and reliability of the new method.

Ling Guo
Shanghai Normal University
lguo@shnu.edu.cn

MS27

Bayesian Learning of Reduced-Order Dynamics

Two probabilistic methods for the learning of reduced-order dynamics will be discussed. The first method is the Bayesian reduced-order operator inference, a non-intrusive, glass-box approach that inherits the formulation structure of projection-based, reduced-state governing equations yet without requiring access to the full-order solvers. The reduced-order operators are learned using Bayesian inference with Gaussian priors and recovered as posterior Gaussian distributions conditioning on projected state data, which provides a quantification of modeling uncertainties and a naturally embedded Tikhonov regularization. The second method employs deep kernel learning, a manifold Gaussian process with a deep neural network embedded inside, for the data-driven reduced-order modeling from high-dimensional measurements given by noise-corrupted images. Such a probabilistic deep learning model is utilized for both dimensionality reduction and the representation of reduced-order dynamics. Numerical results show the effectiveness of deep kernel learning in the denoising and uncertainty quantification of reduced models. The first method is a joint work with S. A. McQuarrie and K. E. Willcox (UT Austin), while the second with N. Botteghi and C. Brune (UTwente).

Mengwu Guo
University of Twente
m.guo@utwente.nl

MS27

Posterior Sampling Methods with Stochastic Gradients

I will discuss the usage of stochastic gradients in sampling methods for Bayesian statistical models. When sampling from Bayesian posterior distributions, widely used schemes such as Langevin Dynamics or Hamiltonian Monte Carlo quickly become unfeasible as the data sets (and models) grow. This is because their dynamics is governed by the gradient of the log-posterior, whose computational complexity is linear in the size of the data set. I will present some work on using stochastic gradients in the samplers, obtained by data set subsampling, in order to decrease computational cost (an approach that is well-known from the neighbouring field of loss function optimization in neural network training). In particular, I will focus on the use of Metropolis correction of the bias introduced by gradient noise and highlight efficiency related issues.

Rene Lohmann

University of Edinburgh
r.lohmann@ed.ac.uk

MS28

Deep Kernel Learning of Dynamical Models from High-Dimensional Noisy Data

This work proposes a Stochastic Variational Deep Kernel Learning method for the data-driven discovery of low-dimensional dynamical models from high-dimensional noisy data. The framework is composed of an encoder that compresses high-dimensional measurements into low-dimensional state variables, and a latent dynamical model for the state variables that predicts the system evolution over time. The training of the proposed model is carried out in an unsupervised manner, i.e., not relying on labeled data. Our learning method is evaluated on the motion of a pendulum – a well studied baseline for nonlinear model identification and control with continuous states and control inputs – measured via high-dimensional noisy RGB images. Results show that the method can effectively denoise measurements, learn compact state representations and latent dynamical models, as well as identify and quantify modeling uncertainties.

Nicol Botteghi, Mengwu Guo
University of Twente
n.botteghi@utwente.nl, m.guo@utwente.nl

Christoph Brune
Department of Applied Mathematics
University of Twente
c.brune@utwente.nl

MS28

Deep Learning Based Surrogate Models for Flow over Non-Parametric Perforated Structures

In many engineering applications design optimization is a time-consuming task which, combined with the limited time in the design phase, often leads to sub-optimal solutions. The use of deep learning based surrogate models can accelerate the design phase with close to real-time simulations. In this talk we will focus on deep learning surrogate models for its use in the design of offshore wind turbine support structures with perforations. Here, the development of surrogate model faces three main challenges: first, we have incompressible turbulent flows with transient inflow conditions; second, the geometries resulting from arbitrary number/shape of perforations are complex and mesh generation might be time-consuming; and third, we need surrogate models that are able to deal with non-parametric geometries, e.g. geometries that cannot be parametrized by the same parametric variables used in the training phase. Here we propose surrogate models based on Convolution Neural Networks (CNN) for flow over objects with arbitrary geometry together with the use of unfitted Finite Element method, the Shifted Boundary Method, that addresses the three challenges.

Oriol Colomés
Delft University of Technology
j.o.colomesene@tudelft.nl

Alexander Heinlein
Delft University of Technology
Delft Institute of Applied Mathematics
a.heinlein@tudelft.nl

Quinten Star
Delft University of Technology
j.quintenstar@gmail.com

MS28

Latent Space Data Assimilation with Uncertainty Quantification Using Deep Learning

In recent years, neural networks have shown great performance in reduced order modeling for PDEs. With these computational speed-ups, real-time solutions to problems where conventional methods often struggle, such as data assimilation, is now possible. In this talk, we will present how a neural network-based forward solver is embedded into a Bayesian data assimilation framework. Specifically, we will show how a particle filter approach can be sped-up by several orders magnitude without sacrificing accuracy by using neural network surrogates. The general idea is to reduce the dimension of the high-fidelity state using an autoencoder and then perform time-stepping in the low-dimensional latent space using a transformer neural network. The data assimilation then takes place in the latent space instead of the high-fidelity space. To do so efficiently, it is helpful to incorporate certain properties into the latent space, such as distributional and smoothness to enable sampling. The result of the data assimilation is a latent space posterior distribution that characterizes the assimilated latent state. This information is then transformed back to the high-fidelity state to get the final estimates and uncertainties. The methodology is showcased on a leak detection problem in pipelines. This problem is modeled by a set of nonlinear hyperbolic PDEs with a discontinuous source term, which makes it difficult to solve for conventional methods in a satisfying timeframe.

Nikolaj T. Mücke
Centrum Wiskunde & Informatica (CWI) and Utrecht University
nikolaj.mucke@cwi.nl

Cornelis W. Oosterlee
CWI, Centrum Wiskunde & Informatica, Amsterdam
c.w.oosterlee@cwi.nl

Sander Bohte
CWI, Centrum Wiskunde & Informatica
s.m.bohte@cwi.nl

Benjamin Sande
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cwi.nl

MS28

A Convolutional Graph Neural Network Approach to Model Order Reduction for Nonlinear Parametrized PDEs

The development of efficient reduced order models (ROMs) from a deep learning perspective enables users to overcome the limitations of traditional approaches. Convolutional autoencoder structures gained popularity in the ROM community as an extension to linear compression procedures [K. Lee and K. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. JCP, 2020] and [S. Fresca, L. Dede, and A. Manzoni. A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-

dependent parametrized PDEs. JSC, 2021]. One drawback of these approaches is the lack of geometrical information when dealing with complex domains defined on unstructured meshes. The present work proposes a framework for nonlinear model order reduction based on Graph Convolutional Autoencoders (GCA) to exploit emergent patterns in different physical problems, including those that show a bifurcating behavior [F. Pichi, B. Moya, and J. S. Hesthaven. A convolutional graph neural network approach to model order reduction for non-linear parametrized PDEs. In preparation, 2022]. The proposed structure exploits the autoencoder paradigm to extract the evolution of the latent space while handling consistently the underlying geometrical domain and alleviating the learning process through pooling and unpooling operations. We present the capability of this novel architecture w.r.t. classical benchmarks and bifurcating phenomena.

Federico Pichi
École Polytechnique Fédérale de Lausanne
federico.pichi@epfl.ch

Beatriz Moya
Universidad de Zaragoza
beam@unizar.es

Jan S. Hesthaven
EPFL
jan.hesthaven@epfl.ch

MS29

Physics-Based and Data-Driven Models for the Response to Momentum-Injection Actuation on a Stalled Airfoil

Active flow control techniques can be useful for improving the performance of aerodynamic structures, such as by providing transient lift enhancement, or mitigating the effects of flow separation or airfoil stall. This work employs data-driven and physics-based modeling techniques to capture and predict the response to burst-type momentum-injection actuation near the leading edge of a stalled NACA0009 airfoil, utilizing data from direct numerical simulations. The response to this type of actuation can be decomposed into two components: a short-time response that is characterized by an initial decrease followed by an increase in the lift, and a long-time response that can be sensitive to the instantaneous wake state at the onset of actuation. Theoretical models are developed using assumptions from classical unsteady aerodynamic theory, which provide insight into the form that the data-driven models should take. The data-driven models, which are identified using variants of dynamic mode decomposition, are capable of capturing both the short- and long-time response of the system to actuation. We additionally show how physics-based approaches can extract spectral content from systems with limited data available, with applications for vortex interactions in wake flows.

Katherine Asztalos
Illinois Institute of Technology
kasztalo@hawk.iit.edu

MS29

Time-Localized Resolvent Analysis of Fluid Flows: Numerical Methods and Analytic Approximations

This talk will discuss extensions of resolvent analysis that allow for the identification of time-localized forcing and

response structures corresponding to high energy amplification. We discuss two methods for enabling such analysis. In the first, the resolvent analysis optimization problem is formulated in a space-time setting, and is then modified to include an L1-norm term to promote temporal sparsity. The resulting optimization problem can be formulated as a nonlinear eigenproblem, which in turn can be solved numerically using a generalized inverse power method. In the second approach, forcing or response modes are assumed to take the form of Gabor wavelets in time, allowing for resolvent gains to be computed as a function of wavelet parameters. We relate the two approaches by demonstrating that temporal wavepackets with this structure (corresponding to Fourier modes modulated by Gaussian envelopes) naturally emerge when applying the sparsity-promoting methodology. We additionally demonstrate that this assumed temporal structure can be analyzed analytically in certain situations. Results from both approaches will be compared for statistically-stationary turbulent channel flow, and for a turbulent Stokes boundary layer, where the mean velocity profile varies with time.

Scott Dawson, Barbara Lopez-Doriga
Illinois Institute of Technology
sdawson5@iit.edu, blopezdorigacostales@hawk.iit.edu

Eric Ballouz, H Jane Bae
California Institute of Technology
eballouz@caltech.edu, jbae@caltech.edu

MS29

A Co-Kurtosis Based Dimensionality Reduction with Neural Network Based Reconstruction of Chemical Kinetics in Reacting Flows

Dimensionality reduction aims to shrink the feature space of high-dimensional data while effectively retaining the information and dynamics of the original system. The widely used principal component analysis (PCA) achieves this for combustion data by transforming the original thermochemical state space into a low-dimensional manifold with eigenvectors of the data covariance. However, this may not effectively capture the stiff chemical dynamics when the reaction zones are spatiotemporally localized. Alternatively, a co-kurtosis PCA (CoK-PCA), wherein the principal components are obtained from singular value decomposition (SVD) of the matricized co-kurtosis tensor, demonstrated greater accuracy in capturing the stiff dynamics. However, both methods incur significant errors due to a linear reconstruction of data onto the original manifold. Nonlinear methods such as artificial neural networks (ANNs) can improve reconstruction accuracy over linear methods, thereby allowing further dimensionality reduction of the original manifold. We use homogeneous reactor data of premixed ethylene-air to investigate (1) the efficacy of CoK-PCA-ANN relative to PCA-ANN and (2) ANN reconstruction relative to linear reconstruction by comparing reconstruction errors of the thermo-chemical state, species production, and heat release rates. Our results show that, while ANN outperforms linear reconstruction in general, the proposed CoK-PCA-ANN captures the stiff dynamics better than PCA-ANN.

Dibyajyoti Nayak
Indian Institute of Science, Bengaluru
dibyanayak@iisc.ac.in

Anirudh Jonnalagadda
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru

anirudhj@iisc.ac.in

Uma Balakrishnan, Hemanth Kolla
Sandia National Laboratories
ubalacr@sandia.gov, hnkolla@sandia.gov

Konduri Aditya
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
konduriadi@iisc.ac.in

MS29

Data-Driven Reduction of Chemical Input-Output Libraries in Simulations of Hypersonic Flows in Non-Equilibrium

Hypersonic flows are of great interest in a wide range of aerospace applications and are a critical component of many technological advances. Accurate simulations of these flows in thermodynamic (non)-equilibrium (accounting for high temperature effects) rely on detailed thermochemical gas models. While accurately capturing the underlying aerothermochemistry, these models dramatically increase the cost of such calculations. In this paper, we present a novel model-agnostic machine-learning technique to extract a reduced thermochemical model of a gas mixture from a library. A first simulation gathers all relevant thermodynamic states and the corresponding gas properties via a given model. The states are embedded in a low-dimensional space and clustered to identify regions with different levels of thermochemical (non)-equilibrium. Then, a surrogate surface from the reduced cluster-space to the output space is generated using radial-basis-function networks. The method is validated and benchmarked on a simulation of a hypersonic flat-plate boundary layer with finite-rate chemistry. The gas properties of the reactive air mixture are initially modeled using an open-source library. Substituting the library with the light-weight, machine-learned alternative improves the performance of the solver by 50% while maintaining overall accuracy.

Clement Scherding
Sorbonne Université
clement.scherding@sorbonne-universite.fr

Georgios Rigas
Imperial College London
g.rigas@imperial.ac.uk

Peter Schmid
KAUST
peter.schmid@kaust.edu.sa

Denis Sipp
ONERA, France
denis.sipp@onera.fr

Taraneh Sayadi
Sorbonne Université
taraneh.sayadi@sorbonne-universite.fr

MS29

Unravelling the Linear and Nonlinear Mechanisms in Flow Around a Stalled Airfoil

The Navier-Stokes equations linearised around the mean (time-averaged) flow can successfully identify the energetic frequencies in variety of flows. One particular example is

the flow around a stalled airfoil where linear analysis identifies two important frequencies: the vortex shedding in the wake and the Kelvin-Helmholtz instability in the shear layer. Although the frequencies are a good match with experimental data obtained from particle image velocimetry (PIV), the structures computed from linear analysis do not agree with those computed directly from the PIV data. The discrepancy can be explained by the limitations of the linear analysis which only accounts for energy exchange mechanisms between the mean and fluctuations. Inter-scale energy exchanges are neglected resulting in structures which extend too far into the wake. In this talk, the nonlinear energy exchanges are unravelled by isolating the key triadic interactions underpinning the flow. In addition, the regions of the flow where these nonlinear energy exchanges take place are identified in order to develop suitable eddy viscosity models that mimic nonlinear energy transfer.

Sean Symon
University of Southampton
sean.symon@soton.ac.uk

MS30

Application of Co-Kurtosis Based Reduced Order Models for Combustion Datasets

Principal Component Analysis (PCA) is a frequently employed dimensionality reduction technique in numerical simulations of combustion processes. However, it could be argued that the co-variance eigenvectors of the thermochemical state space used for PCA may be unable to fully capture information about important, extreme-valued localized chemical dynamics like ignition-kernel-formation. In such situations, a useful approach may be to consider a dimensionality reduction procedure where the required principal vectors are computed from high-order joint statistical moments and may enable improved identification of the state space directions corresponding to stiff dynamics. This talk presents one such procedure, namely the co-kurtosis PCA (CoK-PCA), that is based on the fourth-order joint statistical moment i.e. the co-kurtosis tensor. A priori analysis of combustion datasets corresponding to spontaneous ignition of premixed ethylene-air in a simple homogeneous reactor and ethanol-fueled homogeneous charged compression ignition (HCCI) engine demonstrates that the co-kurtosis based reduced manifold is able to represent the original thermo-chemical state in the regions where chemical reactions are important with significant accuracy.

Anirudh Jonnalagadda
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
anirudhj@iisc.ac.in

Shubham Kulkarni, Akash Rodhiya
Indian Institute of Science
shubham.a.kulkarni@sony.com, akashrodhiya@iisc.ac.in

Hemanth Kolla
Sandia National Laboratories
hnkolla@sandia.gov

Konduri Aditya
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru

konduriadi@iisc.ac.in

MS30

Time-Dependent Subspaces for Reduced-Order Modeling (ROM) of Compressible Reacting Flows with Non-Trivial Boundary Conditions

The performance of the dynamically bi-orthogonal (DBO) decomposition for the reduced order modeling of compressible reacting flows is assessed. DBO is an on-the-fly low-rank approximation technique, in which the instantaneous composition matrix of the reactive flow field is decomposed into a set of orthonormal spatial modes, a set of orthonormal vectors in the composition space, and a factorization of the low-rank correlation matrix. This approach bypasses the need to solve the full-dimensional set of species transport equations to generate high-fidelity data as is commonly performed in data-driven dimension reduction techniques such as the principal component analysis (PCA). Because of these features, DBO can adapt on-the-fly to intrinsic and externally excited transient changes in state of the transport variables. A demonstration case of DBO-based ROM of reacting transport equations exhibiting strongly transient combustion phenomena including flame propagation, end-gas auto-ignition and developing detonation is presented using a 39-species gasoline surrogate mechanism. A-posteriori comparisons against the data generated via full-rank direct numerical simulation (DNS) as well as the PCA reduction of the DNS data is carried out to highlight the effectiveness and accuracy of the DBO based ROM of compressible reacting flows.

Swapnil Desai
Sandia National Laboratories
ssdesai@sandia.gov

Seshu Yamajala
Stanford University
seshu@slac.stanford.edu

Hessam Babae
University of Pittsburgh
h.babae@pitt.edu

Jacqueline Chen
Sandia National Laboratories
jhchen@sandia.gov

MS30

Adaptive, Physics-Based Digital Twins of Practical Combustion Systems

The combination of simulation and experimental data is key for developing accurate and reliable Reduced Order Models that can serve as digital twins (DT) of practical combustion systems to forecast combustion evolution in real-time and act as soft sensors. Dimensionality reduction in combination with non-linear regression has proven powerful in building reduced-order models from Computational Fluid Dynamics simulations. In this work, we demonstrate the development of a physics-based digital twin for a furnace operating in flameless combustion conditions, in a wide range of operating conditions deriving from the modification of the equivalence ratio, the air injection velocity (controlling the internal recirculation degree) and the fuel composition (from pure methane to pure hydrogen). Integrating heterogeneous data streams in the DT development process and continuous improvement is also

discussed, using a combination of methodologies based on sparse sensing and data assimilation.

Laura Donato
 Université Libre de Bruxelles
 Università di Pisa
 laura.donato@ulb.be

Alberto Procacci, Axel Coussement
 Université Libre de Bruxelles
 alberto.procacci@ulb.be, axel.coussement@ulb.be

Chiara Galletti
 Università di Pisa
 chiara.galletti@unipi.it

Alessandro Parente
 Université Libre de Bruxelles
 alessandro.parente@ulb.be

MS30

Topological Characteristics of Low-Dimensional Manifolds in Reduced-Order Modeling of Turbulent Combustion

Dimensionality reduction techniques are used in turbulent combustion to find low-dimensional manifolds (LDMs) in high-dimensional reacting systems to build reduced-order models (ROMs), where only the low-dimensional parameters are transported. We achieve a substantial reduction in the number of parameters needed to visualize, describe and predict complex systems, but some topological properties of LDMs can hinder their practical application. For instance, ROMs often require training a nonlinear regression model to predict physical quantities of interest from the reduced representation, and accurate reconstruction of minor species as well as projected source terms (both of which can exhibit large gradients) is challenging. In this talk, we explore the challenges that remain in ROM of turbulent combustion. We focus on the topological issues that the LDMs present. We discuss our recent advances in characterizing manifold quality, generating improved manifold topologies and improving nonlinear regression performance. We demonstrate novel quantitative tools for characterizing the quality of LDMs from the perspective of non-uniqueness and steep gradients in the physical quantities of interest. We show applications of the manifold assessment tools in optimization algorithms that yield improved manifolds. We discuss novel local kernel regression models that achieve better predictive performance than the current state-of-the-art models.

James C. Sutherland
 Department of Chemical Engineering
 The University of Utah
 james.sutherland@chemeng.utah.edu

Kamila Zdybal
 Université Libre de Bruxelles
 kamilazdybal@gmail.com

MS31

Contact Problems in Porous Media

In this talk, we consider a Biot contact problem, i.e. Biot equations with Signorini contact conditions. The corresponding two-field variational inequality problem is analyzed and a finite element discretization is pre-

sented. Robust numerical discretization, and corresponding parameter-robust a posteriori analysis is of paramount importance in applications involving porous media. An a posteriori analysis in the spirit of is therefore presented for our Biot contact problem. Numerical results underline the theory and demonstrates that our adaptive strategy leads to optimal convergence rates.

Fleurianne Bertrand
 University of Twente
 f.bertrand@utwente.nl

Lothar Banz
 University of Salzburg
 lothar.banz@sbg.ac.at

MS31

A Numerical Method for Subglacial Cavity Formation

Water-filled cavities form in the interface between the ice and the bedrock as an ice sheet or glacier slides along an obstacle. Over sufficiently long timescales, ice is assumed to flow as a viscous flow with a nonlinear rheology. Subglacial cavity formation is an important mechanism in several glaciological applications: 1) the determination of a relationship between sliding speed and basal shear stress (the sliding law), 2) the evolution of subglacial hydrology systems, which act as porous media under an ice sheet through which meltwater flows, 3) the filling and drainage of subglacial lakes, which are essentially large cavities under ice sheets coupled to the subglacial hydrology system. In this talk, we present a numerical model for subglacial cavitation which couples the Stokes equations to a free boundary equation. The opening and closing of a cavity is a time-dependent problem in which the ice detaches and reattaches from the bedrock. Detachment is achieved in the model by solving the Stokes equations with contact boundary conditions; as a result, a Stokes variational inequality must be solved at each time step.

Gonzalo Gonzalez de Diego
 University of Oxford
 gonzalezdedi@maths.ox.ac.uk

MS31

Adaptive Meshing Strategies for Sea-Ice Simulation

A very influential simplification of the general governing equation of sea-ice behaviour is the so-called dynamical sea-ice model by Hibler. This model has been intensively discussed in the last half-century and is used in many climate models. Many models use finite-difference methods to approximate the model equations and therefore are restricted to quadrilateral meshes. Therefore in recent years, there was some development to use finite-element methods for this kind of simulation. This approach allows the use of more general meshes, in particular triangular meshes. We want to extend this development to the next step. We present an a posteriori error estimator, which allows us to derive an adaptive mesh refinement algorithm. This is particularly useful for saving computing power and guarantees good approximation in cases of low regularity solutions. This is especially useful for higher order methods which accelerate the convergence.

Henrik Schneider
 University of Duisburg-Essen

henrik.schneider@uni-due.de

Fleurianne Bertrand
University of Twente
f.bertrand@utwente.nl

MS31

Stress and Flux Equilibration for Biot's Consolidation Problem in Porous Media

The Biot consolidation model of flow and deformation in porous media involves the fluid flux and solid stress as important variables which one may focus on in the numerical approximation. While there are variational formulations which focus on flux and stress as independent variables, one may also reconstruct these quantities from the pressure-displacement formulations which are more commonly used. Such reconstructions can be obtained by an equilibration technique involving Raviart-Thomas elements on a vertex patch decomposition. This talk will focus on the H^1 -conforming Taylor-Hood finite element combination, consisting of polynomial degrees $k+1$ for the displacements and the fluid pressure and k for the total pressure from [J.J. Lee, K.A. Mardal, R. Winther, Parameter-robust discretization and preconditioning of Biot's consolidation model. *SIAM J. Sci. Comput.* 39, A1–A24 (2017)]. It is concerned with the equilibration technique imposing weak symmetry on the stresses studied in [F. Bertrand, G. Starke, A posteriori error estimates by weakly symmetric stress reconstruction for the Biot problem. *Comp. & Maths. with Appl.* 91, 3–14 (2021)]. Numerical results for the resulting a posteriori error estimator will be presented to illustrate the practical viability of the method.

Gerhard Starke
University of Duisburg-Essen
Faculty of Mathematics
gerhard.starke@uni-due.de

Fleurianne Bertrand
University of Twente
f.bertrand@utwente.nl

MS32

Tackling Application Portability in Julia

In research software engineering we face a plethora of hardware (new and old) that are being used to solve scientific problems at many scales. From the researches new laptop, over institutional resources, to national supercomputers. Scientific applications need to be portable across these architectures, but also exploit specific capabilities of the hardware to reach maximum performance. The Julia programming language is a high-level programming language used for technical programming (science, engineering, ...) and many scales. It provides capabilities to program a range of accelerators from different vendors, as well as supporting hardware based reduced precision and vectorization. In this talk we will explore the interplay between writing general programs that are specialized by the compiler, using examples from computational fluid dynamics. The two focus topics are (1) the challenge of supporting reduced precision and (2) targeting multiple GPU vendors from a single-source program.

Valentin Churavy
Massachusetts Institute of Technology

vchuravy@mit.edu

MS32

Towards Everything-Flexible Climate Models in Julia: Productivity Meets Performance

Weather and climate models range in complexity from simple and conceptual to large high-performance computing applications. The middle ground of intermediate-complexity models has received inadequate attention: Such models are simple enough for rapid prototyping to test new concepts, but simultaneously complex enough for research results to be meaningful for operational weather and climate models. We present `SpeedyWeather.jl`, a spectral atmospheric model which is developed as a computational playground for research. Using Julia, its development itself is an investigation into how productivity, flexibility and performance can be combined for next-generation weather and climate models. With minimal code redundancies `SpeedyWeather.jl` supports different number formats and numerical precisions; different architectures; different physical equations; different spatial grids across a wide range of horizontal resolutions. Among those, a 16-bit dynamical core that runs efficiently on GPUs and automatic differentiation are in development. Julia's ability to abstraction leverages code composability. With multiple dispatch, types like the number format or the spatial grid are not hard-coded and become essentially a parameter to the model. Such high modularity in Julia is key for research models as new ideas can be quickly combined and tested, without penalising performance.

Milan Kloewer
University of Oxford
milan.kloewer@physics.ox.ac.uk

Maximilian Gelbrecht
Potsdam Institute for Climate Impact Research
Humboldt-Universität zu Berlin
gelbrecht@pik-potsdam.de

Alistair White
Potsdam Institute for Climate Impact Research
Technical University of Munich
alistair@pik-potsdam.de

Sam Hatfield
European Centre for Medium-Range Weather Forecasts
samuel.hatfield@ecmwf.int

MS32

High Performance Particle Simulations and Simulation Analysis Tools with `CELLLISTMAP.JL`

N-body simulations and trajectory analysis rely on the calculation of attributes that depend on pairwise particle distances within a cutoff. Interparticle potential energies, forces, distribution functions, neighbor lists, and distance-dependent distributions, for example, must be calculated. Cell lists are widely used to avoid computing distances outside the cutoff. However, efficient cell list implementations are difficult to customize. Here, we provide a fast and parallel implementation of cell lists in Julia that allows the mapping of custom functions dependent on particle positions in 2 or 3 dimensions. Arbitrary periodic boundary conditions are supported. Automatic differentiation and unit propagation can be used. The implementation provides a framework for the development of

new analysis tools and simulations with custom potentials. The performance of resulting computations is comparable to state-of-the-art implementations of neighbor list algorithms and cell lists, available in specialized software. Examples are provided for the computation of potential energies, forces, distribution of pairwise velocities, neighbor lists and other typical calculations in molecular and astrophysical simulations. The Julia package is freely available at <http://m3g.github.io/CellListMap.jl>. Interfacing with Python and R with minimal overhead is possible.

Leandro Martínez

Department of Physical Chemistry
University of Campinas - UNICAMP
lmartine@unicamp.br

MS32

High-Performance Computing with TRIXI.JL

Clean and understandable code, an extensible design, and a fast implementation are common goals for scientific software development. They are, however, difficult to achieve simultaneously, since they often require contradictory implementation strategies. In this talk, we will discuss how the Julia programming language can help to overcome some of the typical challenges that arise when trying to develop a high-performance research code that can be easily used and extended by other scientists. We will demonstrate how Trixi.jl, an adaptive high-order numerical simulation package for conservation laws, was designed as a software library with modular components, but with the necessary infrastructure code to make it usable as a standalone framework. The code extensibility afforded by the multiple dispatch mechanism of the Julia language is leveraged to attain high-performance implementations without having to sacrifice accessibility for new users. Finally, we will show how the seamless interoperability with C libraries allows us to scale our Julia code to thousands of cores.

Michael Schlottke-Lakemper

RWTH Aachen University
m.schlottke-lakemper@acom.rwth-aachen.de

MS33

Uncertainty Quantification for Random Field Quantities Using Multi-Fidelity Karhunen-Love Expansions with Active Learning

Sampling-based uncertainty quantification is often computationally prohibitive when using high-fidelity (HF) models. One approach to alleviate this cost is by building multi-fidelity (MF) surrogate models that combine a small number of HF simulations with a larger number of less expensive lower-fidelity (LF) simulations. We focus on using Karhunen-Love expansions (KLEs) as surrogate models, which are based on low-rank spectral approximations that can effectively characterize random field (e.g., spatio-temporal) quantities. We build KLEs in a multi-fidelity manner by retaining the connections to the physical sources of uncertainty through the use of polynomial chaos expansions inside an LF KLE and an additive correction KLE. Predictions can then be generated rapidly for any input parameter value, and greatly accelerate the UQ computations. When additional high-fidelity simulations can be performed to further improve the surrogate model, we propose an active-learning framework to strategically select new simulation points in the input parameter space. We achieve this by using a Gaussian Process (GP) to represent the current MF surrogate error, and use Expected Improve-

ment to find the point with the highest expected surrogate error for performing the next HF simulation. We show the advantages of the overall method on benchmark problems and field quantities from a turbulent round jet application.

Aniket Jivani

University of Michigan
ajivani@umich.edu

Cosmin Safta

Sandia National Laboratories
csafta@sandia.gov

Beckett Zhou

University of Bristol
beckett.zhou@bristol.ac.uk

Xun Huan

University of Michigan
xhuan@umich.edu

MS33

Gaussian Process Regression with Kernels Learned from Data

The design performance given by the manufacturers may usually differ from the operational performance, due to the variability of some operational parameters. Usually, the design development is divided into two different phases. The first phase is to determine the pre-optimal design. Through the use of numerical software, the best design is chosen considering some ideal performances to be achieved. The second phase is to certify through full-scale experiments that the design determined previously is valid. In this way, the post-optimal design is specified. This second phase is the most costly one. In that respect, manufacturers seek to lower the utilization of full-scale experiments. Optimal Uncertainty Quantification (OUQ) is a powerful mathematical tool which can be used to rigorously bound the probability of exceeding a given performance threshold for uncertain operational conditions or system characteristics. This mathematical framework can be very computationally demanding. The use of a metamodel is highly desirable. Moreover, the robustness of the bounding obtained in this framework will strongly depend on the quality of this metamodel. Therefore, one needs to build a metamodel which is quick to evaluate and as accurate as possible. An algorithm, named Spectral Kernel Ridge Regression, is introduced to design kernels from available data in Gaussian process regression surrogate modeling techniques. This algorithm is illustrated on several numerical examples.

Eric M. Savin

École Centrale de Paris
eric.savin@centralesupelec.fr

MS33

Optimization-Based Confidence Interval Construction for Ill-Posed Inverse Problems in the Physical Sciences

The difficulty of uncertainty quantification (UQ) for inverse problems in the physical sciences is compounded by various statistical challenges. These problems are often ill-posed, and even with linear forward models, the rank-deficient forward matrices may only be implicitly available through computationally demanding simulators. Typically, some

of these challenges can be addressed via regularization, but this biases inference and thus affects UQ guarantees. Thus, standard statistical solutions to these types of problems are often either intractable, unreliable or fail to use relevant information. We propose that some of these challenges can be handled by simply incorporating physical constraints, but this approach leads to challenges of its own. In this talk, I present optimization-based confidence intervals related to the strict bounds interval methodology, addressing some of the aforementioned challenges by incorporating physical constraints. In particular, I present the intervals through the lens of the log-likelihood ratio, allowing connections to constrained inference. Finally, I discuss a few physical science applications.

Michael Stanley
Carnegie Mellon University
mcstanle@andrew.cmu.edu

MS34

Lasso Monte Carlo, a Novel Method for High Dimensional Uncertainty Quantification

High-dimensional uncertainty quantification (UQ) is an active area of research, applicable to virtually all domains of science and engineering. In particular in the field of nuclear engineering, simulation codes use nuclear data libraries with tens of thousands of uncertain values, that need to be propagated to the output of the simulations. The most common methods used for UQ are Monte Carlo and surrogate-modelling. The former method is dimensionality independent but has a slow convergence, while the latter method has been shown to yield large computational speedups with respect to Monte Carlo. However, surrogate models are biased and suffer from the curse of dimensionality, becoming costly to train for high-dimensional problems. We present a new technique, Lasso Monte Carlo (LMC), which combines surrogate models and the multilevel Monte Carlo technique, to perform UQ in high-dimensional problems. We prove the unbiasedness of LMC, and that it can converge faster than simple Monte Carlo. The method is numerically tested on a variety of problems: high-dimensional linear functions, the Sobol function where LMC is compared to PCE, the FPUT lattice problem, and on a nuclear physics problem of calculations for spent nuclear fuel involving over ten thousand dimensions. The LMC method has faster convergence than simple Monte Carlo in all tests, and reduces the computational cost of UQ by up to a factor of 5.

Arnau Albà
Paul Scherrer Institut (PSI)
ETH Zurich
arnau.albajacas@psi.ch

Andreas Adelman
Paul Scherrer Institut
andreas.adelman@psi.ch

Dimitri Rochman
Paul Scherrer Institut (PSI)
dimitri-alexandre.rochman@psi.ch

MS34

Low-Dimensional Models with Nonlinear Parametrizations for Systems of Transport-Dominated Evolution Equations

Latent dynamics of transport-dominated evolution phe-

nomena such as described by hyperbolic conservation laws typically exhibit nonlinear structures that make traditional linear dimensionality reduction methods in low-dimensional subspaces inefficient. In this presentation, we build on Dirac-Frenkel variational methods to propagate forward in time nonlinear parametrizations such as deep neural networks that efficiently describe the latent dynamics of transport-dominated problems. The key contribution is an extension of previous work to handle systems of equations. Each variable, e.g., pressure and velocity, is described by a separate nonlinear parameterization, which are subsequently coupled together. Numerical results demonstrate that the proposed approach requires few degrees of freedom to accurately describe and predict transport-dominated dynamics.

Jules Berman
NYU
jmb1174@nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS34

Information-Theoretic Approaches for Model Identifiability in Inverse Problems

An information-theoretic estimator is proposed to assess the global identifiability of statistical models in a practical setting. No assumptions are made on the structure of the statistical model or the prior distribution while constructing the estimator. The estimator has the following notable advantages: first, no controlled experiment or data is required to conduct the practical identifiability analysis; second, different forms of uncertainties, such as model-form, parameter, or measurement can be taken into account; third, the identifiability analysis is global, rather than being restricted to a local region of the parameter space. If an individual parameter is unidentifiable, it can belong to an identifiable subset such that parameters within the subset have a functional relationship with one another and thus have a combined effect on the statistical model. The practical identifiability framework is extended to highlight the dependencies between parameter pairs that emerge a posteriori, with the goal of finding identifiable parameter subsets. The applicability of the proposed approach is demonstrated using a linear Gaussian model and a non-linear methane-air reduced kinetics model.

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

Sahil Bhola
University of Michigan, Ann Arbor
sbhola@umich.edu

MS34

Automating Hyperparameter Tuning for Deep Neural Network Training

Deep neural networks (DNN) have become central to scientific machine learning (SciML) due to their ability to approximate a wide range of functions and thereby act as surrogate models. Despite their success and prevalence, DNNs are notoriously difficult to train. In this talk, we will address the DNN training challenge of choosing hy-

perparameters, such as the learning rate and regularization parameter, which often are chosen by trial-and-error. Our approach will automate the choice of hyperparameters during training using iterative sampling techniques. We will demonstrate the efficacy of our approach on various tasks vital to SciML, including surrogate modeling and dimensionality reduction.

Elizabeth Newman
Emory University
elizabeth.newman@emory.edu

Julianne Chung
Department of Mathematics
Emory University
julianne.mei-lynn.chung@emory.edu

Matthias Chung
Emory University, U.S.
matthias.chung@emory.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

MS35

Performance-Portable Matrix-Free Finite Element Solvers with SYCL

In this talk, we discuss and demonstrate how performance-portable matrix-free finite element solvers can be developed in SYCL and how the algorithms can be adapted to match the performance characteristics of the target hardware. For that, we present a performance model that can guide the choice of algorithm and optimization strategies to apply. The performance model takes into consideration both implementation aspects and hardware performance metrics such as arithmetic peak performance. Finally, we will present performance results for different CPU and GPU architectures, including AMD Milan and Intel Icelake CPU nodes, and NVIDIA A100 and AMD Instinct MI100 accelerators. The numerical experiments together with the performance model show that our SYCL implementation is able to achieve a reasonable fraction of peak performance on a wide variety of platforms, however it may require not only different kernels implementations for different devices but also different algorithms.

Igor Baratta
Department of Engineering
University of Cambridge
ia397@cam.ac.uk

Chris Richardson
University of Cambridge
cni12@cam.ac.uk

Garth Wells
Department of Engineering
University of Cambridge
gnw20@cam.ac.uk

MS35

Vision and Scope of the SYCL Minisymposium

SYCL is a heterogeneous parallel programming model, stewarded by the Khronos Group. The open standard of-

fers portable heterogeneous parallel programming in modern C++. For scientific and engineering simulation and analysis codes, SYCL therefore offers the potential for applications to be performance portable, i.e., for the codes to attain high levels of performance across a range of CPU and GPU devices from different vendors, in a productive way. This talk will briefly introduce the key concepts in the SYCL programming model itself, so this minisymposium will be accessible to all no matter their prior experience of SYCL. Heterogeneous programming models need to be able to describe the underlying hardware, and enable programmers to write programs which control the location of their data, and the location of computation on that data. I will discuss how these abstractions in the SYCL programming model enable portability and create the potential for performance, and hence strive for performance portability. I will finish this talk with a survey of how the current SYCL ecosystem stands with respect to this goal. This will set the stage for this minisymposium, where SYCL will be explored in a wide range of different application areas.

Tom Deakin
Bristol University
United Kingdom
tom.deakin@bristol.ac.uk

MS35

Performance-Portable Earthquake Simulation with SeisSol and SYCL

SeisSol is a high-performance simulation package for large-scale simulation of earthquakes and seismic phenomena. It solves the seismic wave equations in various material models using high-order discontinuous Galerkin with ADER time stepping. To tackle the demands of managing model variations and achieving high performance on various CPU and GPU architectures, SeisSol uses a custom-built DSL and code generator, YATeTo (Yet Another Tensor Toolkit), as an abstraction layer for implementation. We will report on recent work and results obtained with a SYCL backend for YATeTo and we will discuss benefits of SYCL over OpenMP for the implementation of dynamic rupture simulation in SeisSol.

Michael Bader, Ravil Dorozhinskii
Technical University of Munich
bader@in.tum.de, ravil.dorozhinskii@tum.de

MS35

Making HiCMA Hardware-Agnostic with SYCL

This talk reports on user experiences and performance analysis of porting the Hierarchical Computations on Manycore Architectures (HiCMA) software to SYCL. We target systems including x86 and hardware accelerators from various vendors. We assess performance portability using scientific applications powered by HiCMA, while achieving high user-productivity.

Hatem Ltaief
KAUST, Saudia Arabia
hatem.ltaief@kaust.edu.sa

Sameh Abdulah
Extreme Computing Research Center (ECRC)
King Abdullah University of Science and Technology
sameh.abdulah@kaust.edu.sa

Rabab Alomairy

KAUST
rabab.omairy@kaust.edu.sa

Yuxi Hong
King Abdullah University of Science and Technology
yuxi.hong@kaust.edu.sa

Huda Ibeid
Intel
huda.ibeid@intel.com

David E. Keyes
King Abdullah University of Science and Technology
(KAUST)
david.keyes@kaust.edu.sa

MS36

The C^* -Dynamical System of Amenable Quasi-Lattice Ordered Groups and Positive Elements

Let (G, P) be a quasi-lattice ordered group. In previous work, the author constructed a universal covariant representation (A, U) for (G, P) in a way that avoids some of the intricacies of the other approaches. Then showed if (G, P) is amenable, true representations of (G, P) generate C^* -algebras which are canonically isomorphic to the C^* -algebra generated by the universal covariant representation. In this paper, we show that amenability of (G, P) can be established by showing that the universal covariant representation of the C^* -dynamical system is faithful on the positive elements of the algebra B_V .

Mamoon Ahmed
Princess Sumaya University for Technology
m.ahmed@psut.edu.jo

MS36

Adaptive Continuation Methods for Nonlinear PDEs in Deforming Domains

Many problems in mechanics require the solution of partial differential equations within deforming domains; for example, fluid flows containing drops and bubbles, or flows within elastic tubes and channels. The deformation of the domain typically leads to nonlinearities within the system, even when the bulk equations are linear, which can give rise to complex bifurcation structures that connect multiple solution branches. Another aspect of such systems is that fine-scale features of the solutions are often localised in space and time, meaning that mesh adaptivity is advantageous for efficient numerical solution. Traditional continuation methods designed to follow solution branches are generally formulated from discrete, rather than continuous, systems, which presents problems if the discretisation is not fixed over the parameter range of interest. In this talk, I will explain the generic implementation of spatially adaptive continuation methods within the software library oomph-lib (<https://www.oomph-lib.org>) and how these methods can be used to explore the complex solution structures that arise when solving nonlinear PDEs in deforming domains.

Andrew Hazel
The University of Manchester
andrew.hazel@manchester.ac.uk

MS36

Torsional Flows of Convection in Rotating Fluid

Spheres: Computation of the Critical Region, the Periodic Solutions, and Their Stability

It is known that the onset of convection in rotating fluid spheres and spherical shells usually gives rise to rotating waves, which can travel in the prograde or retrograde direction relative to the frame of reference rotating with the bulk of the fluid. It was discovered recently that axisymmetric periodic regimes can also be preferred at low Prandtl, Pr, and Ekman, Ek, numbers. These flows are known as torsional. In order to determine the parameter-space region where the axisymmetric flows are the first bifurcated solutions from the conduction state, the curves of double Hopf points corresponding to simultaneous transitions to azimuthal wave numbers $(m_1, m_2) = (0, 1), (1, 1), (0, 2)$, etc. were computed. These curves form the skeleton of the bifurcation diagram, separating the regions of different preferred azimuthal wave numbers. Their intersections are triple Hopf points, several of which were found. It turned out that the region of interest was limited by the curves $(m_1, m_2) = (0, 1)$ and $(0, 2)$. The torsional solutions emerging from the conduction state were computed for several pairs (Pr, Ek) inside the above mentioned region, and their stability to azimuthal dependence was studied using Floquet multipliers.

Juan Sanchez Umbría
Universitat Politècnica de Catalunya, Spain
juan.j.sanchez@upc.edu

Marta Net
Departament Física Aplicada
Universitat Politècnica de Catalunya
marta.net@upc.edu

MS36

Recent Developments in Numerical Continuation and Bifurcation for Nonlinear PDEs

Numerical continuation and bifurcation methods have become important tools to study PDEs in 2D and higher space dimension, with some general purpose toolboxes available, for instance `BifurcationKit.jl` [Veltz, <https://hal.archives-ouvertes.fr/hal-02902346>, 2020], `coco`-interfaces for PDEs [Dankowicz and Schilder, Recipes for continuation, SIAM 2013], and `pde2path` [Uecker, Numerical continuation and bifurcation for nonlinear PDEs, SIAM 2021]. The main applications are in pattern formation for classical semilinear PDEs, e.g., reaction diffusion systems. Here we explain the numerical challenges for geometric PDEs such as constant mean curvature surfaces, for instance liquid bridges and triply periodic surfaces, and some higher order problems such as shape equations for vesicles, and present some recent solutions, and bifurcation results for such problems obtained with `pde2path`.

Hannes Uecker, Alexander Meiners
Department of Mathematics
University of Oldenburg
hannes.uecker@uni-oldenburg.de, alexander.meiners@uni-oldenburg.de

MS36

Pattern Formation in a Nonlocal Model for Microbial Motility

We study a model of aggregation of organisms in two spatial dimensions. The uninhibited motion of the organisms

is represented by a transport equation for the organism density, while the sensing of the environment, for instance by auditory clues or chemotaxis, is represented by spatial convolutions. The sensing mechanism includes a cut-off at low signal intensity and a saturation at high intensity. We investigate the patterns that arise at critical values of the parameters that regulate the sensing, as well as their spectra and the secondary bifurcations they produce.

Lennaert van Veen

UOIT, Canada

Lennaert.vanVeen@ontariotechu.ca

Pietro-Luciano Buono

University of Ontario Institute of Technology

Oshawa, Ontario Canada

Pietro-Luciano.Buono@uoit.ca

MS37

Non-Intrusive Implementation of Multiscale Finite Element Methods

The multiscale finite element method (MsFEM) is a finite element method (FEM) that allows to solve partial differential equations (PDEs) with highly oscillatory coefficients on a coarse mesh, i.e. a mesh with elements of size much larger than the characteristic scale of the heterogeneities. To do so, MsFEMs use pre-computed basis functions, adapted to the differential operator, thereby taking into account the small scales of the problem. Standard FEM software is based on problem-independent (e.g. polynomial) basis functions, so that the resolution of the FEM can be automated for a large variety of problems. Such software cannot easily be adapted to deal with specific basis functions for each new problem; the MsFEM approach is intrusive. In this talk, we shall propose adaptations of the MsFEM that allow for a non-intrusive approximation strategy: our MsFEM leads to an effective PDE with slowly varying coefficients, hence can be solved by means of standard FEM software. We hope this will facilitate the implementation of the MsFEM outside academic environments. The support of DIM Math INNOV and Inria is gratefully acknowledged.

Rutger Biezemans

MATERIALS project-team

Inria Paris & Ecole des Ponts ParisTech, France

rutger.biezemans@inria.fr

MS37

Optimal Computational Boundary Conditions for Linear Elliptic Equations in Random Media

We consider a model problem of conductivity in a random heterogeneous medium taking the form of a uniformly elliptic equation with random coefficient. We will present in this talk two recent works proposing high accuracy methods to compute the solution based on optimal computational boundary conditions. First, we consider the case of highly oscillatory coefficient fields, putting the problem in the homogenization regime. We will describe a method for computing the homogenised matrix based on representative volume elements with periodic boundary conditions. We will compare two different approaches: periodising the realisations versus periodising the law of the coefficients, and explain how the second approach leads to a higher accuracy. This is a joint work with Marc Josien, Felix Otto and Qiang Xu. Second, we will consider the problem of computing the electrical field generated by a charge distribution

localized on scale ℓ in a situation where the medium is only known in a box of diameter $L > \ell$ around the support of the charge. We study the artificial Dirichlet boundary condition proposed in [Wang, Otto, Ju, Optimal artificial boundary conditions based on second-order correctors for three dimensional random elliptic media] in the case where the medium is a sample from a stationary ensemble with slowly decaying correlations. This is a joint work with Lihan Wang and Quinn Winters.

Nicolas Clozeau

Institute of Science and Technology

Vienna, Austria

nicolas.clozeau@ist.ac.at

Felix Otto

Max Planck Institute for Mathematics in the Sciences

Leipzig, Germany

Felix.Otto@mis.mpg.de

Lihan Wang

MPI Leipzig

lihan.wang@mis.mpg.de

Marc Josien

CEA Cadarache

marc.josien@cea.fr

Qiang Xu

Lanzhou University

xuq@lzu.edu.cn

Quinn Winters

Tum university

quinn.winters@tum.de

MS37

Multiscale Discretization of Nonlinear Schroedinger Equations with Application to Bose-Einstein Condensates

The Gross-Pitaevskii equation (GPE) is a nonlinear Schroedinger equation which is used in quantum physics to model the dynamics of Bose-Einstein condensates (BECs). It is well known that this equation has important time invariants such as the total energy of the system. Preserving the energy under numerical discretization can be of great significance in many practical situations. In this talk we consider numerical approximations of the GPE based on multiscale approaches. More precisely, we choose a generalized finite element space which is based on the localised orthogonal decomposition method and which allows to capture the energy with a 6th order accuracy. Paired with energy-preserving time integrators we demonstrate how such an approach can lead to an efficient solver for the GPE and so for simulating the dynamics of BECs on larger time scales.

Christian Döding

Faculty of Mathematics

Ruhr Universität Bochum

christian.doeding@ruhr-uni-bochum.de

MS37

Fully Discrete Heterogeneous Multiscale Method for Parabolic Problems with Multiple Temporal

and Spatial Scales

Problems with multiple spatial and temporal scales occur in a variety of different considered phenomena and materials like saltwater intrusion or storage of radioactive waste products. In the parabolic case, where the microscopic scales are additionally periodic, the heterogeneous multi-scale method provides a framework for efficiently solving such problems. We consider the critical case where the cell problem is also a parabolic one and we replace the periodic boundary and time conditions by Dirichlet boundary and initial values. Further, we give a detailed a priori error analysis of the fully discretized, i.e., in space and time for both the macroscopic and the cell problem, method.

Daniel Eckhardt

Karlsruhe Institute of Technology
daniel.eckhardt@kit.edu

MS37

Convergence of High-Order Numerical Schemes for the Euler Equations via Dissipative Weak Solutions

It is well known that for the Euler system, a large class of solutions remain smooth only on a certain maximal time interval. Dissipative weak solutions (DWS) have demonstrated their capability analytically and numerically to be the right solution concept to the Euler equations of gas dynamics when strong solutions are not attainable. In this talk, I will talk about the notion of DWS in the context of Euler equations. Furthermore, in this talk, I will present a convergence analysis of the MUSCL scheme and a high-order finite method, the spectral difference method. Namely, we show that structure preserving properties, such as positivity preservation and entropy inequality hold. We demonstrate how to ensure them and prove the convergence of our multidimensional high-order scheme via dissipative weak solutions. Lastly, we demonstrate our theoretical findings with numerical simulations of turbulent flows.

Maria Han Veiga

University of Michigan
mhanveig@umich.edu

MS38

Regularly Hyperbolic Action Principles and Their Discontinuous Galerkin Discretization

Discontinuous Galerkin methods provide a flexible and general approach to approximating initial-boundary value problems defined by first order Friedrichs systems. However, physical systems are often described by action principles associated with Euler-Lagrange equations in second order form. Although in many cases it is possible to associate these equations with first order systems, the change in form requires additional initial and boundary conditions, so that their solutions will agree only if certain (possibly nonlinear) constraints are satisfied. In this talk we will present a construction of discontinuous Galerkin discretizations which can be directly applied to the second order Euler-Lagrange equations. We consider as our model systems the so-called regularly hyperbolic systems introduced by Christodoulou (*The Action Principle and Partial Differential Equations*, 2000). These possess leading order terms satisfying an energy equality which we can use to define numerical fluxes and establish stability of the schemes. Some examples from gravitational wave theory are given to

illustrate the construction.

Thomas M. Hagstrom

Southern Methodist University
Department of Mathematics
thagstrom@smu.edu

MS38

Local Time-Stepping Schemes for Linear Wave Equations

For the spatially discretized linear acoustic wave equation, stability of explicit time integration schemes such as the leapfrog scheme can only be guaranteed under a CFL condition of the form $\tau \lesssim h_{\min}$, where τ denotes the time-step size and h_{\min} the diameter of the smallest element in the underlying mesh. In the case of locally refined meshes, where only a few mesh elements are small compared to the remaining coarse ones, this condition is the main bottleneck for the efficiency of explicit schemes. To overcome this issue, we introduce local time-stepping (LTS) schemes which rely on the leapfrog scheme on the coarse part of the mesh. On the fine part we employ two variants: a stabilized leapfrog-Chebyshev scheme leading to an explicit LTS scheme and θ -schemes resulting in (locally) implicit LTS schemes. For the space discretization we focus on a symmetric interior penalty discontinuous Galerkin discretization. After the construction of these schemes we show that their CFL condition only depends on the coarse part of the mesh and that they converge with second order in time. Moreover, we present some numerical examples confirming our theoretical results.

Marlis Hochbruck, Constantin Carle

Institute for Applied and Numerical Mathematics
Karlsruhe Institute of Technology
marlis.hochbruck@kit.edu, constantin.carle@kit.edu

MS38

Elastodynamic Response of Sh Wave Dispersion in a Multi-Layered Concentric Cylinders Composed of Reinforced and Piezo-Materials

The present study fundamentally focuses on analyzing the limitations and transference of horizontally polarized Shear waves in a four-layered compounded cylinder. It comprises of concentric cylinders of infinite length composed of self-reinforced, fibre-reinforced, piezo-magnetic, and piezo-electric materials. The entire structure is assumed to be pre-stressed along the azimuthal direction. In order to make the structure sensitive to the application pertaining to sensors and actuators, the PM and PE cylinders have been categorically placed in the outer part of the geometry. Whereas in order to provide stiffness and stability to the structure, the inner part consists of self-reinforced and fibre-reinforced media. The common boundary between each of the cylinders has been essentially considered as imperfectly bounded. At the interface of PE and PM media, mechanical, electrical, magnetic, and inter-coupled types of imperfections have been exhibited. The closed-form of dispersion relation has been deduced for two contrast cases i.e. EOMS and ESMO circuit conditions. Dispersion curves have been plotted to illustrate the salient features of parameters like normalized imperfect interface parameters, initial stresses, and radii of the concentric cylinders. Graphs have been plotted to compare the effects of these parameters using two consecutive modes. This theoretical study may be implemented to improvise the performance

of surface acoustic wave (SAW) sensors and actuators.

Sunita Kumawat

Birla Institute of Technology & Science Pilani, Hyderabad
sunita14kumawat@gmail.com

Sumit Viswakarma

Bits-Pilani, Hyderabad

sumitkumar@hyderabad.bits-pilani.ac.in

MS39

Fast and Scalable Computation of Reduced-Order Nonlinear Solutions for PDEs

We develop a method for fast and scalable computation of reduced-order nonlinear solutions (RONS). RONS is a framework to build reduced-order models for time-dependent partial differential equations (PDEs), where the reduced-order solution depends nonlinearly on time-varying parameters. With RONS we obtain an explicit set of ordinary differential equations (ODEs) to evolve the parameters. These ODEs minimize the instantaneous error between dynamics of the governing PDE and dynamics of the reduced-order solution. Additionally, conserved quantities of the PDE are easily enforced in the reduced solution using the RONS framework. For a reduced-order model with n parameters, naive calculation of ODEs produced by RONS requires evaluating $\mathcal{O}(n^2)$ integrals. By exploiting the structure of the RONS equations and using symbolic computation, we reduce the computational cost to $\mathcal{O}(K^2)$ where $K \ll n$. With this approach we apply RONS to problems which require many parameters in the reduced-solution, including simulation of vortex dynamics in turbulent fluid flow and the Fokker-Planck equation in high dimensions.

William Anderson

Department of Mathematics
North Carolina State University,
wmander3@ncsu.edu

Mohammad Farazmand

Department of Mathematics

North Carolina State University

farazmand@ncsu.edu

MS39

Spectral/hp Element Methods for Industry: Some Challenges on the Way

We present the evolution of spectral/hp element methods over the past several years, and their recent successful deployment to industrial problems. To this end, we present some of the milestones that we developed to achieve this result, from high-order mesh generation to numerical stability, robustness, and computational efficiency. We then present the success story of the high-fidelity Large-Eddy Simulation (LES) of a real automotive car, namely the Elemental Rp1 model. The simulation presents the common challenges of an industry-relevant simulation, namely high Reynolds number and complex geometry.

Gianmarco Mengaldo

National University of Singapore
College of Design and Engineering

mpegim@nus.edu.sg

MS39

On the Robustness of Time Integration Schemes Coupled with Adaptive High-Order Spatial Discretizations with the Sbp Property

We report the performance of explicitly Runge–Kutta schemes that are equipped with time step adaptivity and possibly coupled with the relaxation procedure to achieve entropy stability for computational fluid dynamics. The goal is to verify the robustness of the overall adaptivity procedure and the computational costs for practical and complex flow problems discretized with an *hp*-adaptive collocated entropy stable discontinuous Galerkin solver.

Rasha Al Jahdali

King Abdullah University of Science and Technology

Thuwal

rasha.aljahdali@kaust.edu.sa

Lisandro Dalcin

King Abdullah University of Science and Technology

Extreme Computing Research Center

dalcinl@gmail.com

Matteo Parsani

King Abdullah University of Science and Technology (KAUST)

matteo.parsani@kaust.edu.sa

MS39

A Convergent Finite-Volume Scheme for an Alternative Navier-Stokes System

The ultimate target for any numerical scheme aspiring to solve a PDE is convergence. That is, by refining the grid, ever more accurate approximate solutions are obtained. This property is the sole guarantee that a numerical solution can be trusted to track the true solution. For consistent approximations of linear PDEs, the celebrated Lax-Richtmyer Theorem, states the equivalence of stability, e.g. a single L^2 bound, and convergence. One may nurture the hope that a non-linear bound, such as a bounded entropy, has the same effect for non-linear conservation laws. Unfortunately, this is generally not the case. In this talk, I will discuss why proving convergence (to any reasonably defined solution) for the compressible Navier-Stokes equations is such a challenging problem and show that it is a consequence of the derivation of the equations. Drawing from these insights, I will present an alternative Navier-Stokes model with the same modelling capabilities as the standard Navier-Stokes system (and point to references where it has been validated). Furthermore, and without any a priori stability assumptions, I will outline a proof showing that the numerical solutions to a finite-volume scheme converges to a weak solution. That is, I will have demonstrated the existence of weak solutions to the alternative Navier-Stokes system. Finally, I will show some numerical results produced by the scheme.

Magnus Svard

University of Bergen

Norway

magnus.svard@uib.no

MS39

Entropy-Split Multidimensional Summation-By-

Parts Discretization of the Euler and Navier-Stokes Equations

High-order Hadamard-form entropy conservative or stable spatial discretizations of the Euler and Navier-Stokes equations are considerably more expensive than the divergence-form discretization. In search of a more efficient entropy conservative scheme, we investigate and analyze properties of a discretization based on the idea of entropy-splitting. The main ingredients of the scheme are Harten's entropy functions, diagonal-E summation-by-parts operator with diagonal norm matrix, and entropy conservative or stable simultaneous approximation terms (SATs). We show that the scheme is high-order accurate and entropy conservative on general curvilinear unstructured grids for the Euler equations. In the special cases of constant density or zero heat conductivity, the scheme remains entropy stable for the Navier-Stokes equations if appropriate viscous SATs are used. Due to the splitting of the inviscid fluxes, the method is not conservative in the sense of Lax-Wendroff. To mitigate the effects of this loss of conservation, we propose a strategy to enforce local conservation in the vicinity of discontinuities without loss of entropy stability. To support the theoretical findings, several numerical investigations involving both smooth and discontinuous solutions are presented. Preliminary efficiency comparison studies in terms of the time required to compute the spatial residual terms are also included.

Zelalem Arega Worku
University of Toronto Institute for Aerospace Studies
zelalem.worku@mail.utoronto.ca

David Zingg
University of Toronto Institute for Aerospace Studies
dwz@utias.utoronto.ca

MS40

Approximation As a Paradigm for Designing Parallel Graph Algorithms

We describe a paradigm for designing parallel algorithms by approximation techniques. Instead of solving a problem exactly, for which parallel algorithms may not exist, we seek a solution with provable approximation guarantees. Furthermore, we design these algorithms to be concurrent. We discuss several matching and edge cover problems for which such algorithms have been designed, and describe their use in solving problems in sparse matrix computations and load balancing problems in quantum chemistry.

Alex Pothén
Purdue University
Department of Computer Science
apothén@purdue.edu

MS40

Intelligent Rendezvous Point Selection for Improving Network Performance in SDIoT

Software-defined networking is a critical enabler of the Internet of Things. The wireless sensor network is a network of specialized embedded devices that have limited power and storage which provide sensing services to the Software Defined Internet of Things (SDIoT). The radio unit in a sensor node consumes the most energy, hence, an efficient routing strategy is essential to enhance the network lifetime. Most IoT-based WSN environments suffer from traditional clustering and routing methods, which cause data

loss, energy hole issues, and sensor premature death. If it did not use adaptive learning methodologies and continued to use the existing protocol in IoT environment, the network's performance would degrade. In order to overcome these issues, a clustering algorithm based on double Q-learning algorithm is proposed to improve network lifetime and decrease delay in the IoT with uncertain communication links. The simulation findings demonstrate that it outperforms existing clustering and routing algorithms, particularly in terms of performance parameters such as network lifetime, data collection delay and learning rate.

Vishnuvarthan Rajagopal
Research Scholar
Anna University Regional Campus, Coimbatore
rvvarthan@gmail.com

Bhanumathi Velusamy
Assistant Professor
Anna University Regional Campus, Coimbatore
vbhanu_02@yahoo.com

MS40

How Good Is Approximately Solving Graph Problems As Width-Constrained Positive Linear Programs in Parallel?

Graph problems, like those of s-t path and densest subgraph, are frequently employed to solve a wide range of applied problems. A number of these graph problems can be expressed either approximately or exactly as a width-constrained positive linear program (LP), and this formulation may also admit further approximation by using low-hop emulators or hopsets. Furthermore, many of these problems can be solved with a $(1 + \epsilon)$ approximation factor in a highly parallel way. In recent years, novel algorithms using these techniques have been developed and analyzed theoretically, but their practical implications are not well understood in the literature. In this talk, we will explore some of the practical implications of these algorithms for large-scale distributed graph problems. In particular, we have implemented and analyzed the parallel $(1 + \epsilon)$ -approximate with $\tilde{O}(n)$ parallel depth s-t path algorithm from Andoni, Stein, and Zhong (STOC 2020) and the Multiplicative Weights Update (MWU) method from Mahoney et. al (ICALP 2016) that finds a $(1 + \epsilon)$ -approximate solution to a positive LP in $\tilde{O}(\epsilon^{-3})$ time (this is a joint work with Caleb Ju and Serif Yesil.) We show that, for certain types of problems, the novel methods trades vastly superior performance for a small decrease in accuracy, while for others the gains are insubstantial. We also present some implementation-specific enhancements.

Mina Sun
University of Illinois at Urbana-Champaign
ms65@illinois.edu

MS40

GraphZeppelin: Linear Sketching for Dynamic Graph Connectivity

Existing graph stream processing systems must store the graph explicitly in RAM which limits the scale of graphs they can process. The graph semi-streaming literature offers algorithms which avoid this limitation via linear sketching data structures that use small (sublinear) space, but these algorithms have not seen use in practice to date. In this talk I will present GraphZeppelin, a streaming graph system for computing the connected components of

a dynamically changing graph. I will show how GraphZepelin uses linear sketching to efficiently process graphs too large to store explicitly in RAM, and why existing linear sketching algorithms for this and other graph problems fall short in practice. I will show how these sketching techniques scale incredibly well in the distributed setting, with provably near-optimally low network communication overhead. Finally, I will discuss what these results mean for the semi-streaming model and what is required for linear sketching algorithms to be practical.

David Tench
Rutgers
dtench@pm.me

MS40

Algebraic and Humble Programming

We first recall the concept of a "humble programmer" and contrast it to that of a "hero programmer". Classically, the former focuses on achieving high productivity, while the latter focuses on extracting peak performance on a given system. Given that both the complexity of programming novel architectures increases, and that the trend of producing heterogeneous systems that combine multiple such complex architectures is becoming normal, humble programming that achieves scalable performance close to peak on a variety of architectures is becoming a necessity—yet remains challenging. This talk introduces the free and open-source Algebraic Programming (ALP) paradigm as such a candidate solution, and ALP/GraphBLAS specifically. Here, programmers must annotate their code with algebraic information, which the compiler then exploits in optimization, the detection of programmer errors, and unscalable behaviour. Beyond sparse linear algebra, we additionally extend ALP to cover dense linear algebra as well as vertex-centric programming. Finally, we demonstrate the ALP platform's efficiency and scalability, thus arriving at a humble programming model for both numerical linear algebra, graph algorithms, and beyond.

Albert Jan Yzelman
Huawei Zürich Research Center
albertjan.yzelman@huawei.com

MS41

Nektar++: Re-Design of a Spectral/HP Element Solver for Streaming and Vectorized Architectures

The software design ethos of Nektar++ mimics the mathematical construction of the spectral/hp element method. This involves a hierarchy of libraries that: i) represent polynomial expansions in a standardised element, ii) a map from a standard to a local element and iii) a library which contained geometric properties for integration and differentiation in the local element. To construct a piecewise expansion over a domain of multiple local elements we then built a library of multiple local expansions which contains connectivity information for a CG or DG scalar field. We also chose to hold the local expansion coefficients and operators such as differentiation and integration over the whole computational domain in this last library. We are currently revisiting this library structure since it is not appropriate on modern streaming architectures such as GPUs. A necessary step of our re-design is to remove the underlying operators and scalar field storage from the top level library level. Our new design is essentially inverting the current top level concept so that the operators hold the necessary expansion information about connectivity and is

passed a storage field that contains multiple scalar fields in different layouts depending on the target architecture. In this presentation we discuss our strategy to take our legacy code, maintain backwards compatibility for our current users whilst evolving this new structure for efficient use of GPUs and vectorised CPUs.

Chris Cantwell
Imperial College London
c.cantwell@imperial.ac.uk

David Moxey
Department of Engineering
King's College London
david.moxey@kcl.ac.uk

Giacomo Castiglioni
Imperial College London, United Kingdom
g.castiglioni@imperial.ac.uk

Spencer Sherwin
Imperial College London
s.sherwin@imperial.ac.uk

MS41

Performance Optimizations of High-Order Matrix-Free Algorithms for CPU Architectures

My talk will present efficient algorithms for high-order finite element and discontinuous Galerkin discretizations on CPU architectures, which are implemented in the deal.II finite element library. The core algorithm is the matrix-free evaluation of the discretized partial discretization, computing the cell and face integrals of high order continuous and discontinuous finite elements on general curved meshes on the fly with sum-factorization techniques. I will contrast the algorithm selection for achieving a high throughput on GPUs and CPUs, with a focus on the latter in terms of SIMD vectorization and cache behavior. As modern implementations involve an arithmetic intensity of one to five Flop/byte, the most pressing limit is the memory bandwidth. I will present available options to reduce the memory access and increase data locality, such as computation of metric terms on the fly and the fusion of vector operations into the operator evaluation. The performance will be assessed in the context of multi-grid algorithms with point-Jacobi and fast-diagonalization smoothers as well as iterative solvers with simple preconditions. These node-level optimizations will then be evaluated in terms of their strong-scaling efficiency on super-computer scale, with a particular focus on implicit time stepping with stage-parallel Runge-Kutta methods.

Martin Kronbichler
University of Augsburg
martin.kronbichler@uni-a.de

MS41

3D Full Waveform Inversion with Tetrahedral Adapted Meshes and High-Order Hexahedron Spectral Elements Based on Firedrake

This presentation focuses on the implementation of open-source software for full waveform inversion (FWI). FWI is used mainly for imaging the earth's subdomains, but also has applications in noninvasive neuroimaging and high-resolution ultrasound breast imaging. The 3D FWI is an inverse approach that requires several iterations of costly acoustic or elastic wave propagations and gradi-

ent calculations. Here we describe an open-source FWI software implementation, based on the Firedrake Framework, that runs on unstructured triangular and tetrahedral meshes and employs higher-order mass lumped elements, and we compare its performance with high-order spectral element methods on quadrilateral and hexahedron structured meshes. For the unstructured mesh, we reduce degrees of freedom based on an automatic mesh refinement that adapts to subsurface material heterogeneity. We illustrate our results of FWI on different realistic geophysical examples while discussing computational performance on different computer architectures. This research was carried out in association with the ongoing R&D project registered as ANP 20714-2, Software technologies for modeling and inversion, with applications in seismic imaging (University of Sao Paulo / Shell Brasil / ANP).

Alexandre Olander, Thiago D. dos Santos
University of São Paulo
olender@usp.br, santos.td@gmail.com

Daiane I. Dolci
Universisty of São Paulo
dolci@usp.br

Bruno S. Carmo
University of São Paulo
bruno.carmo@usp.br

MS41

Exascale Workflows

With leadership class computing facilities commissioning powerful computers that are getting closer to exascale, users of these facilities are trying to solve much harder and larger problems as fast as possible. There are numerous challenges faced by users in adapting their code, algorithms and workflow in preparation for future exascale systems. Looking at state of the art pre-exascale systems in the U.S., we can see that these systems increasingly use accelerators or General Purpose Graphic Processing Units (GPGPUs). Developing scalable algorithms and implementing them efficiently on these systems is a challenging task, partly due to complicated memory hierarchies of GPGPUs and the complexity of programming models involved. With exascale computers reaching large number of processors, users also have to pay extra attention to how to decompose their problem and communicate efficiently among different processors in these distributed memory parallel computers. In this talk, we go through the challenges faced when porting spectral element codes to efficiently utilize leadership class computing systems in preparation for exascale machine and how we manage to solve them for exascale and beyond.

Thilina Rathnayake
University of Illinois
thilinarmtb@gmail.com

MS42

Statistical Analysis of Neonatal Mortality in Ghana: A Case Study in Ashanti Region

The study's goal was to identify the variables influencing newborn mortality in Ghana. A 2017 survey was used for the analysis, with the major events occurring between 2012 and 2017. Using R software, an analysis using logistic regression approaches was performed to assess the impact of socioeconomic and proximate factors on infant mortality.

The dependent variable of the study was whether the infant was alive or dead, while the independent variables were divided into socioeconomic and proximate factors. According to the findings, some proximate factors have a significant connection with whether the infant is dead or living. The child's weight had a strong relationship with whether the child was alive or dead. The sex of the newborn was the final proximate factor that showed a significant connection with the child being dead or alive. In the Ashanti Region, most of the neonatal deaths were due to low birth weight. Also, 62.7% of the babies that died were males. The study revealed that if the significant elements under the proximate factors are considered, newborn mortality will be reduced.

Elizabeth D. Amankwah
Kwame Nkrumah University Of Science and Technology
lizzy51199@gmail.com

Juliet Richardson, Isaac Adjei
Kwame Nkrumah University of Science and Technology
julietrichardson761@gmail.com, : isaac.adjei@gmail.com

MS42

Stochastic Collocation for Neural Field Equations with Random Data

Neural fields equations model the activity of the brain cortex subject to an external stimulus. These spatially-extended dynamical systems are specified by modelling brain characteristics such as the synaptic kernel and firing rate of neurons, which are generally only partially available in experiments. Tools from uncertainty quantification are key to understanding neuronal behaviour and making predictions of how the brain responds to external stimuli, crucial problem in healthcare and biological applications. Motivated by these considerations, we build a numerical scheme for neural fields with random data (input, synaptic kernel, firing rate, or initial conditions) producing a pointwise approximation of the expectation and variance of neural activity at any space-time coordinate. Our approach consists in casting the system into a set of ODEs on a Banach space, and using a two steps scheme to approximate the solutions mean and variance. The implemented algorithm performs a spatial projection followed by a stochastic projection. In this talk, we explain and motivate possible choices for the projectors (interpolator and orthogonal) and the corresponding quadrature rules to approximate the integrals. In particular, we demonstrate optimal convergence of stochastic collocation analytically and present numerical results.

Francesca Cavallini
Vrije Universiteit Amsterdam
f.cavallini@vu.nl

MS42

Uncertainty Quantification in Data-Driven Model Learning in Biology

Equation learning aims to infer differential equation models from data. While a number of studies have shown that differential equation models can be successfully identified when the data are sufficiently detailed and corrupted with relatively small amounts of noise, the relationship between observation noise and uncertainty in the learned differential equation models remains unexplored. We demonstrate that for noisy data sets there exists great variation in both the structure of the learned differential equation models as

well as the parameter values. We explore how to combine data sets to quantify uncertainty in the learned models, and at the same time draw mechanistic conclusions about the target differential equations. We generate noisy data using a stochastic agent-based model and combine equation learning methods with approximate Bayesian computation (ABC) to show that the correct differential equation model can be successfully learned from data, while a quantification of uncertainty is given by a posterior distribution in parameter space.

Simon Martina Perez
University of Oxford
martinaperez@maths.ox.ac.uk

MS42

Modeling Cell Behavior with Hybrid Cellular Potts and Continuum Models

To form the patterns and behaviors that we see in multicellular development, cells must carefully coordinate their behavior through biophysical and biochemical cues. Numerical modeling and theory are essential tools for analyzing the mechanism of such coordinated, collective cell behavior. To do so, single-cell models must be sufficiently detailed so they correctly capture essential aspects of individual cells and do not oversimplify. At the same time, single-cell models must be sufficiently simple and computationally efficient so they can be upscaled to multicellular systems. Here I will present a series of our recent hybrid cellular Potts models for modeling individual cell behavior, and show how these can be used to study the coordinated cell behavior that is seen in biological development. I will discuss single cell models used to analyze observations such as anomalous cell migration patterns of immune cells, mechanical cell and extracellular matrix interactions, and models of anisotropic force generation. I will end by presenting strategies for using these hybrid CPMs to study collective cell behavior.

Roeland Merks
Universiteit Leiden
merksrmh@math.leidenuniv.nl

MS42

An Adaptive, Two-Scale Scheme for Precipitation-Dissolution Models in Porous Media with Evolving Pore-Scale Geometry

We consider a two-scale model for precipitation and dissolution processes in a porous medium. Such processes determine a micro-structural evolution, leading to free boundaries separating the fluid from the solid inside each pore. This evolution impacts the flow, which, on its turn, affects the solute transport, the precipitation/dissolution processes and, ultimately, the micro-scale geometry. To avoid difficulties related to the evolving geometry inside pores, we propose a phase-field pore-scale model. Employing formal homogenization, a two-scale model is derived, in which the pore-scale and the Darcy-scale model components are coupled through the cell problems, providing the effective parameters. For the resulting model, we discuss an adaptive two-scale scheme. It involves iterations between the scales, linear iterations for solving the nonlinear cell problems, an adaptive selection of the elements wherein the effective parameters are computed, and adaptive mesh refinement.

Manuela Bastidas

Diabatix, Belgium
manumnlb@gmail.com

Carina Bringedal
University of Stuttgart
SimTech, IWS, Pfaffenwaldring 61, 70569 Stuttgart
carina.bringedal@hvl.no

Iuliu Sorin Pop
Hasselt University
Faculty of Sciences
sorin.pop@uhasselt.be

MS43

An Accelerated, High-Order Direct Solver for the Lippmann-Schwinger Equation

Scattering problems in variable media are often best-handled by first reformulating the problem as an integral equation, known as the Lippmann-Schwinger equation. This has numerous advantages, such as better conditioning, automatic enforcement of radiation boundary conditions, and mitigation of dispersion errors. However, the resulting coefficient matrices are dense which poses a challenge in the development of efficient solvers, especially in applications with backscattering where iterative solvers can take prohibitively large numbers of iterations. In this talk, we present a new direct solver for the Lippmann-Schwinger equation in the plane. The solver requires $O(N^{3/2})$ floating point operations, where N is the total number of degrees of freedom in the problem and is capable of achieving high accuracy. We demonstrate that our solver is capable of solving problems hundreds of wavelengths in size to high accuracy in just a few hours on a workstation.

Abinand Gopal
Yale University
abinand.gopal@yale.edu

Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

MS43

A Fast Multilevel Fourier Method for the Potential Evaluation

A fast multilevel Fourier method is constructed for the potential evaluation with discrete and continuous sources. The method is applicable to non-oscillatory radially symmetric kernels. The method is dimension independent, though it suffers from the curse of dimensionality naturally. The performance of method is illustrated with several numerical examples.

Shidong Jiang
Department of Mathematical Sciences
New Jersey Institute of Technology
sjiang@flatironinstitute.org

Leslie Greengard
Simons Foundation and Courant Institute
New York University

lgreengard@flatironinstitute.org

MS43

Nonlinear Waves and Entropy Stable Iteration

Long time simulations of nonlinear wave equations have applications in the study of light propagation in optical fibers, surface gravitation and deep water waves. A common feature of such equations is the presence of high order spatial derivatives, rendering implicit methods necessary for their simulation. The equations considered here are accompanied by an entropy conservation principle. It is known that simulations based on entropy conservative discretizations display favorable error growth (typically linear in time) compared to non-conservative alternatives (typically quadratic). This result assumes that the nonlinear systems of equations arising from the implicit scheme are solved exactly, or nearly so. However, in practice the discrete solution is approximated using iterative methods that are terminated after a finite number of iterations. Since iterative methods have not been derived with entropy conservation in mind, entropy conservation is generally violated, thus compromising the favorable error growth. We discuss techniques based on relaxation and modifications of existing methods for reclaiming entropy conservation with iterative solvers that otherwise violate this principle. Certain iterative methods, as well as relaxation techniques, impact on the flow of time. This can be exploited to improve the quality of long time simulations. Numerical experiments indicate that accurate solutions can be obtained with comparably large tolerances on the iterations.

Viktor Linders, Philipp Birken
Lund University
viktor.linders@math.lu.se, philipp.birken@math.lu.se

MS43

Differentiable Poisson Solver for Cosmological Simulations

Rapid advances in deep learning have brought not only a myriad of neural network models, but also breakthroughs in automatic differentiation (AD) tools and computational accelerators like GPUs that benefit scientific researches. For example, analyses of future cosmological large-scale structure observational data sets requires performant and differentiable simulations as forward models. We develop fast algorithms for the Poisson solver in cosmological simulations, and implement them in JAX and CUDA. Using the adjoint method and the reversed time integration, we are able to greatly reduce the memory cost of the reverse mode AD. Our work enables larger and more accurate forward modeling, and thus will improve gradient based optimization and inference.

Libin Lu, Yin Li
Flatiron Institute
llu@flatironinstitute.org, yinli@flatironinstitute.org

MS44

Deep Dive - Exascale System Software and Programming Models

HPC applications are rising in complexity, to address problems such as climate, medicine, and digital twins, and to combine numerical simulation with HPDA and AI. Increased system performance and energy efficiency leads to complex system architectures, with accelerators, mul-

iple kinds of memory device, and processing in memory and/or storage. These developments put a high demand on the system software and programming environment. This talk will discuss the approaches of three related research projects in Europe: DEEP-SEA, IO-SEA and RED-SEA. The projects address Software for Exascale Architectures, driven by application use cases, from the perspectives of heterogeneity and the modular supercomputer architecture, I/O and storage, and the interconnect. This talk will outline the main software research directions. Performance portability and productivity are key challenges, requiring compiler and runtime intelligence, and a separation of the source code from its optimizations. Heterogeneous memory systems are targeted through APIs and automated data placement. All layers, including resource management middleware, are being updated to support malleability, i.e. to allow jobs to dynamically adjust their resource usage. Data management and hierarchical storage are being updated to span NVMe and NVRAM down to the least active data on tape. This talk will highlight these and other research directions, the results achieved so far and the impact on the future of European HPC.

Paul Carpenter
Barcelona Supercomputing Centre
paul.carpenter@bsc.es

MS44

Deep Dive - Exascale I/O and Storage Systems

The future arrival of Exaflop supercomputers raises several major challenges for the HPC storage systems. Besides the scale of such systems (hundreds of thousands of nodes), the volume of data created by applications running on them, and the need for highest energy efficiency, root causes include new hardware and new paradigms involved in supercomputers, evolution of software methods, use cases that were previously impossible to handle and only now become feasible. This talk will identify these new constraints and challenges and propose tracks and solution in order to address them and build future, highly efficient Exascale storage systems. In particular, solutions involving non-Posix interfaces like Object Stores, I/O resource scheduling and storage hierarchies will be analysed and investigated.

Philippe Deniel
CEA
philippe.deniel@cea.fr

MS44

Deep Dive - Next Generation Interconnects

Next generation of interconnect networks defined will have to scale to hundreds of thousands of nodes to enable Exascale computing and target advanced HPC, HPDA and AI applications. The RED-SEA RD project will leverage the BXI technology from ATOS, defined as the key production-proven European Interconnect, which features smart Network Interface Cards (NICs) capable of achieving high bandwidth communications. Other KPI (Key Performance Indicators) such as scalability, reliability or congestion management are also central to cope with issues raised in an Exascale environment. Future environments are likely to feature multiple NICs thus increasing substantially bandwidth capabilities that are yet to be taken advantage of for BXI networks. To abstract these network complexities to the developer, MPI (Message Passing Interface) has become the de facto standard in the scientific community. In this presentation, we propose to in-

roduce our implementation of the multi-rail feature in the MPC SW framework, embedding CEA's implementation of the MPI standard. In particular, we will explain the software architecture and design choices that were made to efficiently support the protocols needed to communicate on multiple network cards while taking advantage of the hardware capabilities of the BXI NICs. Performance results on BXI networks with multiple NICs will be shown demonstrating the benefits of multiplexing communications for very large message sizes.

Gilles Moreau, Hugo Taboada
CEA
gilles.moreau@cea.fr, hugo.taboada@cea.fr

Marc Perache
CEA, DAM, DIF, F-91297 Arpajon, France
marc.perache@cea.fr

MS44

Exascale - Introduction and Principal Challenges

The enormous computational performance of Exascale machines offers new opportunities for industrial and scientific users, but also presents them with difficult challenges. Many of these challenges stem from the heterogeneity of the hardware. On the processing side, the main source of computer performance is accelerators (e.g., general-purpose graphics cards) that are combined with CPUs (which in turn already contain a variety of arithmetic/logic units and vector processing engines). On the memory side, not only are cache hierarchies deepening, but the diversity of technologies and devices in the overall memory and storage subsystems is growing: from fast SRAM to tape, almost every type of memory and storage technology in existence is in use. HPC users should be aware of the heterogeneity of modern system architectures, but should not be left alone in the face of it. Ideally, most of the hardware complexity should be handled by the middleware and software stack, while interfaces to applications become hardware agnostic and standardised. This approach facilitates the long-term sustainability of applications, as well as their portability across platforms. However, today we are far from this situation. To improve it, the SEA projects are developing a software stack with the functionality and tools needed to cope with heterogeneous resources and share them more dynamically between users.

Estela Suarez, Hans-Christian Hoppe
Juelich Supercomputing Centre
e.suarez@fz-juelich.de, h.hoppe@fz-juelich.de

MS45

Inspired by Nature: Dynamic Graphs and Their Applications

While graphs have always been an important area of mathematical study, recent work focuses on a range of creative variations on these simple structures. Dynamic graphs, i.e. graphs that change over time, capture graph structures in dynamical systems. These structures can provide insight into a given problem, whether from a theoretic perspective or for a practical application. Recently, dynamic graphs in conjunction with machine learning and statistical models have been used to analyze a wide range of natural data, e.g. human behavior, embryos, animal cohorts, and organism connectomes. In this talk, we will uncover a few basic and important concepts in dynamic graphs and how

they can be used for understanding the world around us.

Devavrat V. Dabke
Princeton University
ddabke@princeton.edu

MS45

Spatio-Temporal K-Means

Spatio-temporal data is readily available due to technologies, such as GPS and IOT, that track the positions of moving objects of interest. Many spatio-temporal clustering methods exist to efficiently discover patterns in moving object behavior without human supervision. Our focus is the discovery of moving clusters, where clusters have a static identity, but their location and content can change over time. It is informative to know both which individuals associate at every time step, as well as in the long-term. We propose a two phase spatio-temporal clustering method, spatio-temporal k-means (STKM), to analyze the multi-scale relationships within spatio-temporal data. Phase 1 of STKM outputs the temporary associations between objects. The clustering objective function provides a unified formulation over space and time and less parameter tuning compared to existing methods. Phase 1 also tracks cluster paths without any post-processing, allowing it to identify long-term point behavior even in a dynamic environment. Phase 2 can be optionally applied to output the stable associations between objects. The combination of both phases allows us to analyze the interaction of short- and long-term point behavior. Unlike previous work, we evaluate STKM against baseline methods on a recently developed benchmark dataset and show that STKM performs competitively in extracting the multi-scale relationships of moving objects, particularly outperforming other methods in the low-data domain.

Olga Dorabiala, Nathan Kutz, Aleksandr Aravkin
University of Washington
olgad400@uw.edu, kutz@uw.edu, saravkin@uw.edu

Jennifer Webster
Pacific Northwest National Laboratory
jwebster42@gmail.com

MS45

Geomstats: Current State and Future Works

To be updated.

Luis Gomes-Pereira
Inria
luis.gomes-pereira@inria.fr

MS45

Geometries of Covariance and Correlation Matrices

Many data from neuroscience can be modeled by covariance matrices (EEG, MEG, MRI). This talk tackles two important challenges: decoupling the correlation part from the scaling part in the covariance matrix, and computing with singular covariance matrices. To do so, we present new geometries of the set of full-rank correlation matrices and we present the Bures-Wasserstein geometry of symmetric positive semi-definite matrices.

Yann Thanwerdas
INRIA

yann.thanwerdas@inria.fr

MS45

Models and Algorithms for Grouping in Trajectory Data

The collective motion of a set of moving entities like people, birds, or other animals, is characterized by groups arising, merging, splitting, and ending. Given the trajectories of these entities, we define and model a structure that captures all of such changes using the Reeb graph, a concept from topology. The trajectory grouping structure has three natural parameters that allow more global views of the data in group size, group duration, and entity inter-distance. In this talk, I will present on the latest models and algorithms for grouping in trajectory data.

Marc van Kreveld
Utrecht University
m.j.vankreveld@uu.nl

MS46

TNet: A Tikhonov Neural Network Approach to Deterministic and Bayesian Inverse Problems

Deep Learning (DL) by design is purely data-driven and in general does not require physics. This is the strength of DL but also one of its key limitations. DL methods in their original forms are not capable of respecting the underlying mathematical models or achieving desired accuracy even in big-data regimes. On the other hand, many data-driven science and engineering problems, such as inverse problems, typically have limited experimental or observational data, and DL would overfit the data in this case. Leveraging information encoded in the underlying mathematical models not only compensates missing information in low data regimes but also provides opportunities to equip DL methods with the underlying physics and hence obtaining higher accuracy. This talk introduces a Tikhonov Network (TNet) that is capable of learning Tikhonov regularized inverse problems. We rigorously show that our TNet approach can learn information encoded in the underlying mathematical models, and thus can produce consistent or equivalent inverse solutions, while naive purely data-based counterparts cannot. Furthermore, we theoretically study the error estimate between TNet and Tikhonov inverse solutions and under which conditions they are the same. Extension to statistical inverse problems will also be presented.

Hai Nguyen
The University of Texas at Austin
hainguyen@utexas.edu

Russell Phillely
University of Texas at Austin, U.S.
rsphillely@gmail.com

Tan Bui-Thanh
University of Texas
tanbui@oden.utexas.edu

MS46

Variable Shape Parameter Strategy for Radial Basis Function Approximation Using Neural Networks

The choice of the shape parameter highly effects the be-

haviour of radial basis function (RBF) approximations, as it needs to be selected to balance between ill-condition of the interpolation matrix and high accuracy. In this work, we demonstrate how to use neural networks to determine the shape parameters in RBFs. In particular, we construct a multilayer perceptron trained using an unsupervised learning strategy, and use it to predict shape parameters for inverse multiquadric and Gaussian kernels. We test the neural network approach in RBF interpolation tasks and in a RBF-finite difference method in one and two-space dimensions, demonstrating promising results.

Fatemeh Nassajian
University of Zurich
fatemeh.nassajian@math.uzh.ch

Maria Han Veiga
University of Michigan
mhanveig@umich.edu

Jan S. Hesthaven
EPFL
jan.hesthaven@epfl.ch

Philipp fner
Johannes Gutenberg-University, Mainz
poeffner@uni-mainz.de

MS46

Learning the Electrical Impedance Tomography Inversion Operator

Electrical impedance tomography (EIT) is a class of non-invasive imaging techniques with uses in medical and geophysical applications. EIT may be mathematically abstracted as the severely ill-posed nonlinear inverse problem of recovering the interior conductivity coefficient (the parameter, a function) of an elliptic partial differential equation from its Dirichlet-to-Neumann boundary map (the data, a linear operator). Existing methods for this problem are highly sensitive to noisy boundary measurements and tend to suffer from low accuracy. This work bypasses expensive iterative optimization or Bayesian inference solution approaches by instead directly learning the data-to-parameter solution operator of the inverse problem from training data. Theoretical analysis based on rigorous direct reconstruction algorithms for EIT establishes that the solution operator is well-approximated by a new class of neural operators suitable for operator-valued data. Numerical evidence suggests that incorporating theoretical and EIT-specific problem structure into these neural operator architectures can lead to highly accurate trained models.

Nicholas H. Nelsen
California Institute of Technology
nnelsen@caltech.edu

Maarten V. de Hoop
Rice University
mdehoop@rice.edu

Nikola Kovachki
California Institute of Technology
nkovachki@caltech.edu

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology

astuart@caltech.edu

MS46

Leveraging Multifidelity in Scientific Machine Learning for Inverse Design

Wide adoption of inverse design requires global surrogate models that can simulate desired real-world outcomes for all feasible inputs. With the universal approximation theorem and fast evaluations, neural networks (NN) are promising. However, by the curse of dimensionality, with many input variables, NNs become too data greedy for practical usage. For many problems in engineering, there exist data and/or models with multiple fidelity levels. This multifidelity knowledge can be exploited to make data-efficient models that are accurate enough for inverse design. I will present a recently developed physics-enhanced deep surrogate (PEDS) that combines a NN generator, trained end-to-end with a low-fidelity PDE solver, to match costly high-fidelity data. PEDS saves at least two orders of magnitude in data needed to inverse design. The solver is inaccurate but fast and enforces inductive biases that make PEDS designs always satisfy the PDE knowledge. When only multifidelity data is available, the low-fidelity solver can be replaced by a learned operator on the low-fidelity dataset. To finish, we will present inverse design results for thermoelectrics where we compare two optimization approaches: one relying on topology optimization of the continuous relaxation of the problem, another based on derivative-free genetic algorithm.

Raphael Pestourie

Massachusetts Institute of Technology
raphael.pestourie@gmail.com

Lu Lu

University of Pennsylvania
lulu1@seas.upenn.edu

Youssef Mroueh

IBM Research
mroueh@us.ibm.com

Giuseppe Romano, Christopher V. Rackauckas
Massachusetts Institute of Technology
romanog@mit.edu, accounts@chrisrackauckas.com

Das Payel

IBM Research
daspa@us.ibm.com

Steven Johnson

MIT
stevenj@math.mit.edu

MS46

Sampling Strategies for Training Machine Learning Emulators of Gravity Wave Momentum Transport

With the goal of developing a data-driven parameterization of gravity waves (GWP) for use in general circulation models, we train various machine learning architectures to emulate an existing GWP scheme. We diagnose the disparity between online and offline performance of the trained emulators by identifying a subspace of the phase space that is prone to large errors and sparse samples, and develop a sampling algorithm to treat biases that stem from underrepresentation. This strategy can be used for regression

tasks over long-tailed (and other imbalanced) distributions. We find that error-prone samples often have large shears in the wind profile this is corroborated with physical intuition as large shears indicate many breaking levels, which requires a more complex, nonlocal computation. To remedy this, we develop a sampling strategy that performs a parameterized histogram equalization. The sampling algorithm uses a linear mapping from the original histogram to the uniform histogram parameterized by $t \in [0, 1]$. Parameters t and “maximum repeat” assign each bin a new probability. The new probability is applied in two different implementations: 1) by sampling the bins to adjust the training set distribution; 2) by weighting the loss function to achieve the same effect in expectation. We find that this strategy improves the errors at the tail of the distribution except at the extreme end, while maintaining minimal loss of accuracy at the peak of the distribution.

Lucia M. Yang

New York University
minah.yang@nyu.edu

Edwin Gerber

NYU
epg2@nyu.edu

MS47

Q-Next: A Fast, Parallel, and Diagonalization-Free Self-Consistent Field Algorithm

As computer systems dedicated to scientific calculations become massively parallel, the poor parallel performance of the Fock matrix diagonalization becomes a major impediment to achieving larger molecular sizes in self-consistent field (SCF) calculations. In this Lecture I will present a novel, highly parallel, and diagonalization-free algorithm for the accelerated convergence of the SCF procedure is presented. The algorithm, called Q-Next, draws on the second-order SCF, quadratically convergent SCF, and direct inversion of the iterative subspace (DIIS) approaches to enable fast convergence while replacing the Fock matrix diagonalization SCF bottleneck with higher parallel efficiency matrix multiplications. Performance results on both parallel multicore CPU and GPU hardware for a variety of test molecules and basis sets are presented, showing that Q-Next achieves a convergence rate comparable to the DIIS method while being, on average, one order of magnitude faster.

Giuseppe Barca

The Australian National University
giuseppe.barca@anu.edu.au

MS47

Scalable Eigenvalue and Density Matrix Solutions in ELSI

The ELSI Infrastructure provides uniform, general interfaces to several frameworks for serial and parallel (including GPU) eigenvalue and density matrix solutions. The primary target application is electronic structure theory but any other application making use of parallel eigenvalue solutions can equally benefit from access to solvers such as Lapack, ELPA, SlePc, EigenExa and others through a common interface. We review recent progress in ELSI, especially related to mixed-precision solution and GPU solutions of large eigenvalue problems.

Volker Blum

Duke University
Department of Mechanical Engineering and Materials
Science
volker.blum@duke.edu

MS47

Latest Improvements of the Elpa Eigenvalue Solvers

In the electronic structure community, the solution of eigenvalue problems lies at the core of the numerical approaches and in most simulation packages the numerical eigenvalue solver can consume a substantial fraction of the complete runtime. Since more than a decade the ELPA library for dense, symmetric (hermitian) eigenvalue problems tackles this computationally severe challenge, by providing optimized and scalable algorithms and is used by researchers on high-performance computing (HPC) systems worldwide. Nowadays, with the transition to Exascale HPC systems being on its way and the enormous compute power provided by the upcoming HPC machines, researchers will be able to tackle problems in material sciences which were not possible before. In this presentation, we will discuss the path of the ELPA library to support exascale simulations, and we will discuss the latest improvements of the ELPA eigenvalue solver library. In particular, we will present and show some results for the extreme-scale GPU support of the ELPA library of AMD, Nvidia, and Intel GPUs, which will be available in the upcoming HPC systems.

Peter Karpov, Andreas Marek
Max Planck Computing and Data Facility
petr.karpov@mpcdf.mpg.de, andreas.marek@mpcdf.mpg.de

MS47

Economic Quasi-Newton Self Consistent Field Solver

We present an efficient Quasi-Newton solver for calculating Hartree-Fock (HF) and Kohn-Sham Density Functional Theory (KS-DFT) single determinant wavefunctions free of any matrix diagonalization (except possibly when generating the initial guess). The solver combines the L-BFGS hessian update with trust region step restriction to minimize the need for line search and thus greatly reduce the number of gradient (i.e., Fock matrix) evaluations. Solving the trust region problem leverages the low-rank structure of the L-BFGS inverse hessian. Robust performance of the solver has been assessed for standard benchmark sets of molecules. Particular utility of the solver for locating physically-reasonable (but non-aufbau) KS-DFT solutions in biopolymers will be illustrated.

Edward F. Valeev
Virginia Tech
efv@vt.edu

Samuel Slattery
Department of Chemistry
Virginia Tech
samwise@vt.edu

MS48

Towards Models of Blood Clotting on Moving Sur-

faces

Blood clotting is a complex phenomenon which involves interplay between biochemistry and fluid-structure interactions. For example, subclinical leaflet thrombosis is a potentially serious complication of aortic valve replacement in which clots form on the replacement valve. Clots in the left atrial appendage can dislodge during atrial fibrillation, which leads to an increased risk of strokes. Models of thrombosis in these kinds of systems must take into account the motion of both the structure and thrombus as well as the biochemistry to accurately describe formation of the thrombus. Herein, we describe recent approaches to model the clotting process within the context of fluid-structure interaction. We focus on two key components of thrombus modeling: deposition of material onto surfaces, and the growth of the thrombus into the lumen. While the methods presented are applied to clotting problems, these methods can also be used to model deposition and absorption along other moving boundaries, such as drug absorption in the gut or particulate deposition in the lungs.

Aaron Barrett
University of Utah
barrett@math.utah.edu

Boyce E. Griffith
University of North Carolina at Chapel Hill
boyceg@email.unc.edu

Aaron L. Fogelson
University of Utah
fogelson@math.utah.edu

MS48

Phase-Field Modelling of Evolving Adhesive Interfaces

Computational phase-field modelling is a powerful and versatile mathematical technique to describe and solve problems that involve evolving interfaces. In this contribution, we present a new phase-field model for adhesion and discuss its underlying foundations in terms of thermomechanical consistency. The interaction between adhesive interfaces is modelled by the introduction of a coupling term in the classical free energy. The resulting phase-field model is a stiff, higher-order, non-linear partial differential equation, in which adhesive interaction is regularized onto the moving interfaces. We characterize the singular limit of the adhesion problem. Under the assumption of the existence of a smooth limiting interface, we identify the limiting sharp-interface energy of our adhesion problem and demonstrate that its minimisers are geometric surfaces with a non-trivial adhered part. Lastly, we provide numerical examples of adhesion and show how unconditionally energy-stable results can be obtained by using the scalar auxiliary variable approach as time-discretization scheme. Furthermore, we discuss application of the presented model to problems in cell biology, which is vital in understanding numerous biological processes, including cell migration, structural integrity and signal transduction. We show how the model can be extended to account for the protein dynamics at a cell membrane, which govern the binding and unbinding of a cell to its extracellular environment.

Anne M. Boschman, Bindi S. Brook, Matteo Icardi, Kristoffer G. van der Zee
University of Nottingham
anne.boschman@nottingham.ac.uk,

bindi.brook@nottingham.ac.uk,
matteo.icardi@nottingham.ac.uk,
kg.vanderzee@nottingham.ac.uk

MS48

How Fluid Rheology Shapes Microorganism Swimming Gait in Viscoelastic Fluids

Many important biological functions depend on microorganisms' ability to move in viscoelastic fluids such as mucus and wet soil. The effects of fluid elasticity on motility remain poorly understood, partly because, the swimmer strokes depend on the properties of the fluid medium, which obfuscates the mechanisms responsible for observed behavioral changes. Gait adaptation is essential for biological function, but how organisms adapt their gait and the significance of these changes for biological function are not known. We use computational and mathematical modeling to (1) disentangle these effects by changing the gait and fluid rheology independently and to (2) examine how the swimmer gait emerges from the properties of the environment. We present a series of problems of increasing complexity: prescribed gait, prescribed motor activity, and fully emergent motion. We also discuss methodological advances to the Immersed Boundary method involved in performing these simulations.

Robert D. Guy
Mathematics Department
University of California Davis
guy@math.ucdavis.edu

Becca Thomases, Kathryn Link
University of California at Davis
thomases@math.ucdavis.edu, klink@math.ucdavis.edu

MS48

Effects of Transmural Pressure on Clot Structure and Occlusion Times

Hemostasis is the process by which a blood clot forms to prevent bleeding at the site of an injury. The formation time, size, and structure of a blood clot depends on the local hemodynamics, transmural pressure, and the nature and size of the injury. We have previously developed computational models to study intravascular clotting, a process confined to the interior of a vessel. Modeling extravascular injuries, where blood leaks from a vessel into extravascular space, requires a set of new computational tools for the complex geometries that simulate the injury. Similar to our previous intravascular model, our new model of extravascular clotting uses a continuum approach to track the advection, diffusion, and aggregation of platelet densities in a dynamic fluid environment. The transport of platelet densities into any spatial location is limited by the platelet fraction that already resides within that location, i.e., the densities satisfy a maximum packing constraint using a hindered transport coefficient. We used a finite volume method in a T shaped geometry where blood flows in through a top channel and can escape through a downward injury channel. Portions of the walls of the injury channel are reactive and can initiate platelet adhesion and aggregation. We varied transmural pressures across the injury channel and injury channel diameters and to study their effects on injury channel occlusion times and clot densities.

David Montgomery
Colorado School of Mines

dmontgomery@mines.edu

Karin Leiderman
Department of Applied Mathematics & Statistics
Colorado School of Mines
karinlg@unc.edu

MS48

Using IBFE to Simulate Muscle Driven Swimming of Blue Blubber Jellyfish

Jellyfish swimming is well-studied in the scientific community. Numerical simulations are commonly used to understand the mechanism of bell propulsion and the resulting fluid dynamics. These simulations use simplified jellyfish models, typically only simulating the jellyfish bell and neglecting the oral arms. However, a blue blubber jellyfish's oral arms are prominent; together, the oral arms are close to the diameter of the jelly bell. Neglecting the oral arms would result in overly idealized fluid dynamics, not considering the vortex shedding from the bell interacting with passive elastic structures. This work uses the immersed boundary finite element (IBFE) method to simulate a fully 3D swimming blue blubber jellyfish, including the oral arms. Additionally, swimming is driven using a muscle model rather than with prescribed motion, allowing for emergent behavior. The work is implemented in IBAMR, an adaptive, parallelizable implementation of the immersed boundary method with IBFE capabilities. The methodology and modeling will be discussed in addition to the challenges and need for improved methods.

Matea Santiago
University of California Merced
mateasantiago@arizona.edu

Alex Hoover
Cleveland State University
a.p.hoover@csuohio.edu

Citlali Ramirez, Laura Miller
University of Arizona
citlaliramirez@arizona.edu, lauram9@math.arizona.edu

MS50

Towards a Task-Based SPH Simulation Code for Exascale Computing

Future supercomputing architectures will exhibit unprecedented hardware parallelism that needs to be harvested on the software side, which will most likely involve radical changes to supercomputer software architectures. Hence, traditional global loop-based parallelism and notably the bulk-synchronous/BSP paradigm will become insufficient, and developers might have to go back to the drawing board and uncover alternatives to equip their code with higher concurrency. A candidate solution for this is task parallelism, which frames an algorithm in terms of tasks that can be executed concurrently by different processors. In this talk, I will introduce a newly developed, task-based Smoothed Particle Hydrodynamics (SPH) code integrated within the Peano/ExaHyPE framework. I will discuss the main features of the code such as the modelling of the compute steps and tasks in terms of the transversal of mesh hierarchies (dubbed spacetrees) as well as simulation results for standard SPH benchmarks. Furthermore, I will discuss the integration of a C++ extension with a memory-centric specification language into the SPH code which adds fea-

tures that can improve the performance of large-scale simulations, such as reordering data temporarily from an array of structures (AoS) into a structure of arrays (SoA) and the compression of floats with arbitrary precision.

Cristian Barrera Hinojosa
Durham University
United Kingdom
cristian.g.barrera@durham.ac.uk

MS50

Experiences with Running the Flash Multiphysics Simulation Code on Ookami, An A64FX Testbed Platform

We present lessons learned from porting and tuning FLASH, originally developed at the University of Chicago for multi-scale multi-physics applications, for Ookami, an HPE Apollo 80 machine featuring the A64FX processors developed by Fujitsu. We report performance and scaling results for different compilers, memory limitations, our efforts at effectively using the Scalar Vector Extensions and NUMA architecture, and use of Huge Pages to minimize translation lookaside buffer misses. We found that FLASH readily ran on this architecture, but obtaining good performance requires some effort. This research and the machine Ookami are supported by the US National Science Foundation (NSF) grant OAC 1927880. This research was also supported in part by the US Department of Energy (DOE) under grant DE-FG02-87ER40317.

David Carlson, Catherine Feldman
Stony Brook University
david.carlson@stonybrook.edu,
catherine.feldman@stonybrook.edu

Eva Siegmann
Institute for Advanced Computational Science
Stony Brook, USA
eva.siegmann@stonybrook.edu

Benjamin Michalowicz
Ohio State University
michalowicz.2@osu.edu

Anthony Curtis, John Dey
Stony Brook University
anthony.curtis@stonybrook.edu, john.dey@stonybrook.edu

Robert J. Harrison
Brookhaven National Laboratory
Stony Brook University
robert.harrison@stonybrook.edu

Alan Calder
Stony Brook University
alan.calder@stonybrook.edu

MS50

Extending Science Codes to Quantum Computers

There has been a growing recent interest in quantum computing as a way to solve scientific problems that are currently intractable through classical resources alone. Here I will discuss the work our classical supercomputing institution (the National Energy Research Scientific Computing Center, or NERSC) is doing to incorporate quantum resources into our traditional supercomputing workflows,

and which algorithmic areas we expect to benefit the most as quantum hardware matures. I will also discuss specific quantum algorithms we are developing for efficient eigenvalue calculation, applicable to studying chemical and materials systems.

Katherine Klymko
Lawrence Berkeley National Laboratory
Computational Chemistry, Materials and Climate Group
kklymko@lbl.gov

MS51

Iterative Solutions and Tuning Accuracy for Interpolatory Reduced Models in Inverse Problems

Medical image reconstruction presents huge computational challenges due to the quantity of data generated by modern equipment. In this talk, we focus on diffuse optical tomography (DOT) in medical imaging. In DOT, each stage of processing requires the solution of more than a thousand large, three-dimensional problems. Moreover, as rapid advances in technology allow for larger numbers of sources and detectors and using multiple frequencies, these problems become computationally prohibitively expensive. In a recent study, we combined two powerful methods, randomization and reduced order models (ROM) to drastically reduce this cost. While ROM provides a way to drastically reduce the cost of many expensive linear solves, randomization can drastically reduce the number of large linear solves needed for constructing the global ROM bases. Since the size of a realistic linear system in DOT is at least $O(10^6)$, we use iterative methods to handle the large linear systems. Hence, further computational savings are also possible by reducing the solver tolerance. In this talk, we provide a numerical study on how sensitive the quality of the reduced order model is to the chosen tolerance.

Selin Aslan
Argonne National Laboratory
selin@vt.edu

Eric De Sturler
Virginia Tech
sturler@vt.edu

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS51

New Robust-to-Stiffness Low-Rank Techniques for Sylvester-Like Equations

In this talk, we present a new technique for solving the dynamical low-rank approximation of stiff differential equations. It is specific to stiff Sylvester-like differential equations that often appear after discretising parabolic PDEs. The new method keeps all computations low-rank and is obtained from exponential integrators. An efficient technique is proposed for computing the order one and order two schemes in a large-scale setting. Theoretical guarantees are given with numerical experiments.

Benjamin Carrel
University of Geneva

benjamin.carrel@unige.ch

MS51

A Local Macroscopic Conservative Low Rank Tensor Method for the Vlasov Dynamics

In this talk, we introduce a novel Local Macroscopic Conservative (LoMaC) low rank tensor method for simulating the Vlasov-Poisson (VP) and Vlasov-Maxwell (VM) systems. The LoMaC property refers to the exact local conservation of macroscopic mass, momentum and energy at the discrete level. This is a follow-up work of our previous development of a conservative low rank tensor approach for Vlasov dynamics (arXiv:2201.10397). In that work, we applied a low rank tensor method with a conservative singular value decomposition (SVD) to the high dimensional VP/VM system to mitigate the curse of dimensionality, while maintaining the local conservation of mass and momentum. However, energy conservation is not guaranteed, which is a critical property to avoid unphysical plasma self-heating or cooling. The new ingredient in the LoMaC low rank tensor algorithm is that we simultaneously evolve the macroscopic conservation laws of mass, momentum and energy using a flux-difference form with kinetic flux vector splitting; then the LoMaC property is realized by projecting the low rank kinetic solution onto a subspace that shares the same macroscopic observables by a conservative orthogonal projection. The algorithm is extended to the high dimensional problems by hierarchical Tuck decomposition of solution tensors and a corresponding conservative projection algorithm. Extensive numerical tests on the VP and VM systems are showcased for the algorithm's efficacy.

Wei Guo

Texas Tech University
weimath.guo@ttu.edu

Jingmei Qiu

University of Delaware
jingqiu@udel.edu

MS51

Parallel Randomized Algorithms for Tucker Decompositions

The Tucker decomposition is a low-rank tensor decomposition generalizing the matrix SVD to higher dimensions. Traditional algorithms used to compute Tucker decompositions can be computationally expensive as they involve computing the SVD of mode unfolding, which can be quite large. Existing randomized and parallel algorithms have reduced this cost, but computational challenges remain. We propose new randomized parallel algorithms to address these challenges. Using randomized matrix techniques, we accelerate a distributed-memory implementation of the Tucker decomposition to obtain an efficient algorithm that avoids communication. Specifically, we employ a new sketching method that exploits Kronecker structure to accelerate a key computation. We also present probabilistic analysis of the error resulting from the algorithm, as well as numerical results demonstrating the computational benefits of this approach.

Rachel Minster

Wake Forest University
minsterr@wfu.edu

Zitong Li

University of California Irvine
zitongl5@uci.edu

Grey Ballard

Wake Forest University
ballard@wfu.edu

MS51

Dynamical Low-Rank Approximation for Second-Order Matrix Differential Equations

In this talk, we introduce a new dynamical low-rank integrator for second-order matrix differential equations

$$A'(t) = F(A(t)), \quad A(0) = A_0, \quad A'(0) = B_0, \quad t \in [0, T],$$

typically stemming from space discretizations of wave equations. The integrator is constructed by combining the projector-splitting integrator introduced in [2] with a Strang splitting ansatz. We also present a variant of the new integrator which is tailored to stiff second-order problems. The performance of the new schemes is illustrated by numerical experiments.

Stefan Schrammer

Karlsruhe Institute of Technology
stefan.schrammer@kit.edu

Marlis Hochbruck, Markus Neher

Institute for Applied and Numerical Mathematics
Karlsruhe Institute of Technology
marlis.hochbruck@kit.edu, markus.neher@kit.edu

MS52

Non-Linear Surrogate Models Solving Variational PDE Formulations with Sobolev Cubatures

We address the problem of solving Partial Differential Equations (PDEs) with general boundary conditions by proposing (weak) variational PDE formulations resulting in a soft-constrained optimization problem. We theoretically prove and empirically show that, due to the formulation, solutions of linear, non-linear, forward, and inverse PDE problems can be efficiently and numerically stable computed. Moreover, under the same formulation, optimal domain decomposition for approximating discontinuous shocks can be given. We investigate the properties and performance of two kinds of variational solvers: Physics Informed Neural Networks (PINNs) and Polynomial (spectral) Surrogate Models (PSMs). While already linear PDEs result in non-linear and non-convex optimization problems when addressed with PINNs, the corresponding PSM optimization problems maintain linear and convex. By computing the PDE losses on a generalization of classic Gauss-Legendre cubatures, termed Sobolev cubatures, we can show that PSMs models are as flexible as PINNs for general (inverse) non-linear problems, with several orders of magnitude of better accuracy and computational efficiency, opening up the discussion regarding When are PINNs necessary? How can we enhance existing methods with hybrid formulations? Our approach is intended to be a unifying framework of the classic numerical methods, together with the novel machine learning ones, exploiting the best of each field.

Juan E. Suarez

Helmholtz-Zentrum Dresden Rossendorf
juesuarezca@gmail.com

Michael Hecht

Center for Advanced Systems Understanding (CASUS)
m.hecht@hzdr.de

MS53

Stochastic Gradient Descent Alternating Least Square Method for High-Order Tensor Decomposition

High-order tensors have applications in many areas (biology, finance, engineering, etc.). One of the key issues in high-order tensor research is the optimal rank one decomposition of a tensor, which is comparable to the singular value decompositions (SVD) for matrices. Unlike SVD, the rank one decomposition problem for high-order tensors is NP-hard. The Stochastic Gradient Descent Alternating Least Squares (SALS) method is a generalization of the well-known Alternating Least Squares (ALS) method that approximates the canonical decomposition of averages of sampled random tensors. Its simplicity and efficient memory usage make the SALS algorithm an ideal tool for decomposing tensors in an online setting. This talk will discuss the convergence of the SALS as well as its application to scientific data analysis.

Yanzhao Cao

Department of Mathematics & Statistics
Auburn University
yzc0009@auburn.edu

MS53

Defining a Trend from a MultiScale Time Series

We propose criteria that define a trend for time series with inherent multi-scale features. We call this trend the tendency of a time series. The tendency can represent an executive summary of a complex time series, and as such can be viewed of a dimension-reduced representation of original signal. The tendency is defined empirically by a set of criteria and captures the large-scale temporal variability of the original signal as well as the most frequent events in its histogram. Among other properties, the tendency has a variance no larger than that of the original signal; the histogram of the difference between the original signal and the tendency is as symmetric as possible; and with reduced complexity, the tendency captures essential features of the signal.

Juan M. Restrepo

Oak Ridge National Lab
restrepojm@ornl.gov

MS53

Random Weight Factorization Improves the Training of Continuous Neural Representations

Continuous neural representations have recently emerged as a powerful and flexible alternative to classical discretized representations of signals. However, training them to capture fine details in multi-scale signals is difficult and computationally expensive. Here we propose random weight factorization as a simple drop-in replacement of conventional dense layers for accelerating and improving the training of coordinate-based multi-layer perceptrons (MLPs). We show that this factorization essentially alters the loss landscape and effectively enables each neuron in the network to learn its own self-adaptive learning rate. As a result, it not only helps with mitigating spectral bias, but also allows networks to quickly recover from poor initial-

izations and reach better local minima. We demonstrate how random weight factorization can be leveraged to accelerate and improve the training of neural representations on a variety of tasks, including image regression, shape representation, computed tomography, inverse rendering, solving partial differential equations, and learning operators between function spaces.

Sifan Wang, Paris Perdikaris

University of Pennsylvania
sifanw@sas.upenn.edu, pgp@seas.upenn.edu

MS53

Lagrange α -Exponential Synchronization of Non Identical Fractional Order Complex Valued Neural Networks

In this article, Lagrange α -exponential synchronization of non-identical fractional-order complex-valued neural networks (FOCVNNs) is studied. Numerous favorable conditions for achieving Lagrange α -exponential synchronization and α -exponential convergence of the descriptive networks are constructed using additional inequalities and the Lyapunov method. Furthermore, the structure of the α -exponential convergence ball, in which the rate of convergence is linked with the systems characteristics and order of differential, has also been demonstrated. These findings, which do not require consideration of the existence and uniqueness of equilibrium points, help to generalize and improve previous works and may be used to mono-stable and multi-stable of the FOCVNNs. The salient feature of the article is the graphical exhibition of the effectiveness of the proposed method by using numerical simulation for synchronization of a particular case of the considered fractional-order drive and response systems.

Vijay K. Yadav

Department of Mathematics, Amity School of Applied Sciences,
Amity University Haryana, Gurugram, Haryana, 122413, India
vijayky999@gmail.com

MS53

MetaNO: A Meta-Learnt Nonlocal Neural Operator Approach for Efficient Material Modeling

In real-world material modeling problems, the data acquisition is often very challenging and expensive, which makes learning the material model with a limited number of measurements critical. Herein, we propose a meta-learned approach for transfer-learning between neural operators, Meta-NO, based on the implicit Fourier neural operator (IFNO) approach. The overall goal is to efficiently provide accurate solution surrogates for new and unknown material-learning tasks (e.g., with different microstructure or mechanical parameters), from multiple training tasks where each task corresponds to different materials. The proposed sample-efficient meta-learning algorithm consists of two phases: (1) learning a common representation by sharing the same iterative integral layers from existing tasks; and (2) transferring the learned knowledge and rapidly learning surrogate operators for new and unseen tasks with a different material, where only a few test samples are required. The result demonstrates that the meta-learned representation would handle complex and nonlinear material response learning tasks, while greatly improving the sampling efficiency in new and unseen microstructures. Hence, it would substantially reduce the cost on lab testing

for new materials.

Lu Zhang
Lehigh University
luz319@lehigh.edu

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

MS54

Parallel-In-Time Solver for the All-At-Once Runge-Kutta Discretization

Time-dependent PDEs arise quite often in many scientific areas, such as mechanics, biology, economics, or chemistry. Of late, researchers have devoted their effort in devising parallel-in-time methods for the numerical solution of such PDEs, adding a new dimension of parallelism and allowing to speed-up the solution process on modern supercomputers. In this talk, we present a fully parallelizable preconditioner for the all-at-once linear system arising when employing a Runge-Kutta (RK) method in time. The resulting system is solved iteratively for the numerical solution and for the stages. The proposed preconditioner results in a block-diagonal solve for the stages, accelerated by a novel block-preconditioner based on the SVD of the RK coefficient matrix, and a Schur complement obtained by solving again systems for the stages. Preliminary parallel results on the heat equation show the robustness of the preconditioner with respect to the discretization parameters and to the number of stages, as well as very promising scalability and parallel efficiency indices.

Luca Bergamaschi
Universita di Padova
Italy
luca.bergamaschi@unipd.it

John W. Pearson, Santolo Leveque
University of Edinburgh
j.pearson@ed.ac.uk, s.leveque@sms.ed.ac.uk

Angeles Martinez
University of Trieste
amartinez@units.it

MS54

Parallel-in-Time Multilevel Krylov for Solving Hyperbolic Wave Equations

Parareal, introduced by Lions, has opened ways of computing solution of evolutionary differential equations in a non-traditional way: instead of sequentially marching in time from the initial time, Parareal algorithm allows computation of solutions in parallel over time subintervals independently. Connection between time subintervals is established through sequential solution using larger time step, which is used to define and improve the initial solution for each time subinterval. A few interpretations of Parareal have been derived, and one of them links Parareal with a special type of multigrid. The latter paves the way to the development of parallel-in-time multigrid methods, such as MGRIT. Borrowing a similar structure to multigrid, it is also possible to construct a parallel-in-time method based on Krylov methods with projection-like preconditioners. Out of several possible implementations, this talk focuses on shift operator, which, like multigrid, allows multilevel solution procedure in a stable way. The coarse time-level

system can still be modified so that parallel computation can be performed at the second level, resulting in a parallel-in-time two-level method. Numerical results will be presented for linear time-domain wave equations, discretized using a finite-difference method in combination with theta-scheme for time integration.

Yogi Erlangga
Zayed University
yogi.erlangga@zu.ac.ae

MS55

Plug and Play is a Convergent Regularization Method for Inverse Imaging Problems

A characteristic feature of image reconstruction problems is the non-uniqueness and instability against data perturbations. Therefore, it is necessary to develop regularization methods that aim to find an approximate but stable solutions. A regularization method consists of a family of stable reconstruction methods that converge, as the noise level tends to zero, to an exact solution of the inverse problem. The standard approach is variational regularization, which minimizes a data discrepancy term augmented by a regularizer. The actual numerical implementation applies iterative methods, often involving proximal maps of the regularizer. In recent years, Plug and Play (PnP) iterations have been developed as a new powerful generalization by replacing proximal mappings in variational regularization with general image denoisers. While PnP iterations achieve excellent results, neither stability nor convergence in the sense of regularization have been investigated so far. In this paper, we extend the idea of PnP by considering families of PnP iterations, each accompanied by its own denoiser. As main theoretical result, we show that such PnP methods leads to a stable and convergent regularization method for decreasing noise level. This demonstrates that the PnP method is mathematically equally justified as the variational model for convergent and stable image reconstruction.

Andrea Ebner
University of Innsbruck
andrea.ebner@uibk.ac.at

Markus Haltmeier
Universität Innsbruck
markus.haltmeier@uibk.ac.at

MS55

Hybrid Projection Methods for Solution Decomposition in Large-Scale Bayesian Inverse Problems

We develop hybrid projection methods for computing solutions to large-scale inverse problems, where the solution represents a sum of different stochastic components. Such scenarios arise in many imaging applications (e.g., anomaly detection in atmospheric emissions tomography) where the reconstructed solution can be represented as a combination of two or more components and each component contains different smoothness or stochastic properties. In a deterministic inversion or inverse modeling framework, these assumptions correspond to different regularization terms for each solution in the sum. Although various prior assumptions can be included in our framework, we focus on the scenario where the solution is a sum of a sparse solution and a smooth solution. For computing solution estimates, we develop hybrid projection methods for solution decomposition that are based on a combined flexible and gener-

alized Golub-Kahan processes. This approach integrates techniques from the generalized Golub-Kahan bidiagonalization and the flexible Krylov methods. The benefits of the proposed methods are that the decomposition of the solution can be done iteratively, and the regularization terms and regularization parameters are adaptively chosen at each iteration. Numerical results from photoacoustic tomography and atmospheric inverse modeling demonstrate the potential for these methods to be used for anomaly detection.

Jiahua Jiang
University of Birmingham
j.jiang.3@bham.ac.uk

Julianne Chung
Department of Mathematics
Emory University
julianne.mei-lynn.chung@emory.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Scot Miller
Johns Hopkins University
smill191@jhu.edu

MS55

Convergence Analysis of Critical Point Regularization with Non-Convex Regularizers

In recent years, several methods using regularizers defined by neural networks as penalty terms in variational methods have been developed. One of the key assumptions in the stability and convergence analysis of these methods is the ability of finding global minimizers. However, such an assumption is often not feasible when the regularizer is a black box or non-convex, making the search for global minimizers of the involved Tikhonov functional a challenging task. Instead, standard minimization schemes are applied which typically only guarantee that a critical point is found. To address this issue, we study stability and convergence properties of critical points of Tikhonov functionals with a possible non-convex regularizer. We show that even in such cases a full convergence analysis can be given and we provide numerical simulations which support our findings.

Daniel Obmann
University of Innsbruck
daniel.obmann@uibk.ac.at

Markus Haltmeier
Universität Innsbruck
markus.haltmeier@uibk.ac.at

MS55

Inverse Source Problem for Linearly Anisotropic Sources in Absorbing and Scattering Medium

We considered in a two dimensional absorbing and scattering medium, an inverse source problem in the stationary radiative transport, where the source is generated by a scalar and vector field. The medium has an anisotropic scattering property, the attenuating and scattering properties of the medium are assumed known and the unknown scalar and isotropic vector field sources are to be reconstructed. For

scattering kernels of finite Fourier content in the angular variable, we show how to simultaneously recover the scalar and vector field sources from boundary measurements. The approach is based on the Cauchy problem for a Beltrami-like equation associated with A -analytic maps in the sense of Bukhgeim.

David Omogbhe
University of Vienna
University of Mathematics
david.omogbhe@univie.ac.at

Kamran Sadiq
Radon Institute for Computational and Applied
Mathematics
kamran.sadiq@oeaw.ac.at

MS55

On Regularization in L^∞

Many inverse problems arise as parameter identification problems for PDE coefficients which are naturally assigned to L^∞ . However, this space is non-smooth, non-reflexive and non-separable which hinders the application of standard Banach space methods. In this talk we present a novel regularisation method which generates uniformly bounded iterates as approximate solutions to locally ill-posed equations and for which the regularisation property holds with respect to weak-* convergence. Numerical examples from Full Waveform Inversion will complete our analysis.

Lukas Pieronek
Karlsruhe Institute of Technology
lukas.pieronek@kit.edu

Andreas Rieder
University of Karlsruhe, Germany
andreas.rieder@kit.edu

MS56

Comparative Study of Profiling Tools on Fugaku Supercomputer

The overhead of several commonly used performance monitoring tools are measured when profiling a Fast Fourier transform based solver for the Klein Gordon equations. Final aim is to see if profiling can help provide an explanation for the better large scaling of FFTe as compared to 2DECOMP&FFT.

Samar A. Aseeri
Computational Scientist at KAUST
samar.aseeri@kaust.edu.sa

Judit Gimenez
Barcelona Supercomputing Center
Universitat Politècnica de Catalunya
judit@bsc.es

Sameer Shende
Dept. of Computer and Information Science
University of Oregon
sameer@cs.uoregon.edu

Benson Muite
Kichakato Kizito, Kenya
benson_muitemailplus.org

David E. Keyes
King Abdullah University of Science and Technology
(KAUST)
david.keyes@kaust.edu.sa

MS56

Updates on FFTX and NTTX

We present an update on the design of API of FFTX and introduce NTTX. The FFTX is developed as part of the DOE ExaScale effort by LBL, Carnegie Mellon University, and SpiralGen, Inc. We aim at translating the LAPACK/BLAS approach from the numerical linear algebra world to the spectral algorithm domain. FFTX is extending and updating FFTW for the exascale era and beyond while providing backwards compatibility. NTTX is a variant of FFTX aimed at the number theoretical transform (NTT). We apply the same design methodology across both packages. In both cases we utilize the SPIRAL code generation system as backend, and while initially we targeted offline code generation, we started investigating run time compilation.

Franz Franchetti
Carnegie Mellon University
franzf@andrew.cmu.edu

MS56

Implementation of Parallel Number-Theoretic Transform on Manycore Clusters

In this talk, we propose an implementation of parallel number-theoretic transform (NTT) on manycore clusters. The butterfly operation of the NTT can be performed using modular addition, subtraction, and multiplication. We show that a method known as the six-step fast Fourier transform (FFT) algorithm can be applied to the NTT. We parallelized the six-step NTT using MPI and OpenMP. Performance results of parallel NTTs on manycore clusters are reported.

Daisuke Takahashi
Center for Computational Sciences
University of Tsukuba
daisuke@cs.tsukuba.ac.jp

MS57

Machine Learning of Nonlocal and Fractional Models for Subsurface Transport in Heterogeneous Media

Anomalous behavior is ubiquitous in subsurface solute transport due to the presence of high degrees of heterogeneity at different scales in the media. Nonlocal and fractional models are the best candidates to capture such a multiscale behavior. In this talk, we propose a data-driven framework for the discovery of optimal nonlocal kernels or fractional orders on the basis of very small and sparse data sets in the context of anomalous subsurface transport. Using spatially sparse breakthrough curves recovered from fine-scale particle-density simulations, we learn the best coarse-scale nonlocal or fractional model using operator regression techniques. Several numerical tests illustrate the effectiveness of our approach.

Marta D'Elia
Sandia National Laboratories

marti.delia@gmail.com

MS57

Granular Media and the Effect of Grain Shape on Dynamics

We examine the effect of grain shape on the dynamics of granular media. Grain shape influences the dynamics of hopper flow, the interaction between vehicle tire and gravel road, avalanches, mudslides, etc. In this talk we discuss a model for mesoscopic modeling of the granular flow that is focused on grain shape and strength. The dynamics is represented by an ode in terms of nonlocal constitutive laws relating force to deformation. Here the length scale of nonlocality is taken on the particle length scale and is smaller than the mesoscopic length scale characterizing the flow. We provide a method for solving these odes and several simulations of road bed dynamics are provided. This is in collaboration with Debdeep Bhattacharya at LSU.

Robert P. Lipton
Department of Mathematics
Louisiana State University
lipton@math.lsu.edu

MS57

A Nonlocal Model for the Numerical Solution to Richards Equation

We present a nonlocal, derivative free model for transient flow in unsaturated, heterogeneous, and anisotropic soils based on the peridynamic formulation of continuum mechanics. In the proposed model, we consider a spectral method based on the implementation of Chebyshev polynomials to approximate numerically the solution of the model. We show the convergence of the method and present several simulations to study the properties of the solutions.

Sabrina F Pellegrino
University of Bari
pellegrino@lum.it

MS57

Physics-Guided Nonlocal Neural Operators

Neural operators, which emerge as implicit solution operators of hidden governing equations, have recently become popular tools for learning responses of complex real-world physical systems. Nevertheless, the majority of neural operator applications has thus far been data-driven, which neglects the intrinsic preservation of fundamental physical laws in data. In this paper, we introduce a novel integral neural operator architecture, to learn physical models with fundamental conservation laws automatically guaranteed. In particular, by replacing the frame-dependent position information with its invariant counterpart in the kernel space, the proposed neural operator is by design translation- and rotation-invariant, and consequently abides by the conservation laws of linear and angular momentums. As applications, we demonstrate the expressivity and efficacy of our model in learning complex material behaviors from both synthetic and experimental datasets, and show that, by automatically satisfying these essential physical laws, our learned neural operator is not only generalizable in handling translated and rotated datasets, but also achieves state-of-the-art accuracy and ef-

efficiency as compared to baseline neural operator models.

Yue Yu

Department of Mathematics, Lehigh University
yuy214@lehigh.edu

MS58

Analysis of the Transmission Eigenvalue Problem with Two Conductivity Parameters

In this talk, we show the existence and discreteness of the transmission eigenvalue problem with two conductivity parameters. In previous studies, this problem was analyzed with one conductive boundary parameter, while in this study we consider the case of two parameters. The underlying physical model is given by the scattering of a plane wave for an isotropic scatterer. We will study the dependence on the physical parameters and the monotonicity of the first transmission eigenvalue with respect to the parameters. Lastly, we will consider the limiting procedure as the second boundary parameter vanishes and present numerical results.

Rafael Ceja Ayala
Purdue University
rcejaaya@purdue.edu

MS58

Interior Transmission Eigenvalue Trajectories

Properties of complex-valued eigenvalue trajectories for the interior transmission problem parametrized by a constant index of refraction are investigated. At first, the unit disk is considered and several properties are derived such as that the only intersection points with the real axis are Dirichlet eigenvalues of the Laplacian. Then, for general sufficiently smooth scatterers also the only trajectorial limit points are shown to be Dirichlet eigenvalues of the Laplacian as the refractive index tends to infinity. Numerical results for several scatterers are presented which even give rise to an underlying one-to-one correspondence. Finally, a conjecture on the link between these two eigenvalue families is given.

Andreas Kleefeld
Forschungszentrum Jülich GmbH
Jülich Supercomputing Centre
a.kleefeld@fz-juelich.de

Lukas Pieronek
Karlsruhe Institute of Technology
lukas.pieronek@kit.edu

MS59

Synergies Between Neural Networks and Finite Elements

The direct numerical approximation of PDEs using neural networks is gaining lots of attraction, but many problems still need to be solved. One is the numerical integration of the resulting terms (which involves integrating the neural network and its derivatives). Monte Carlo is the most common approach but suffers a low convergence rate. Another problem is the required regularity of the neural network to end up with a differentiable functional. E.g., ReLU activation functions are not regular enough, even for weak formulations. In this presentation, we will discuss how to design effective adaptive quadratures for smooth enough

activation functions and compare the proposed approach with Monte Carlo methods. Alternatively, we explore using interpolated neural networks on finite element spaces to compute the PDE loss function instead of dealing with the neural network itself. Doing this, the integration is straightforward. However, this approach relies on a mesh. In this situation, we also discuss the ill-posedness of standard PDE losses and how to design well-posed preconditioned losses. We show how by combining standard finite element preconditioners and neural networks, we can substantially accelerate the training process, one of the most severe issues of neural network PDE approximations.

Santiago Badia, Alexandre Magueresse
Monash University
santiago.badia@monash.edu,
alexandre.magueresse@monash.edu

Alberto Martin
Australian National University
alberto.martin@monash.edu

Wei Li
Monash University
wei.li@monash.edu

MS59

CutFEM for Acoustic Shape Optimization: Features and Surprises

Fictitious-domain methods like CutFEM are nonstandard but attractive alternatives for shape optimization, as they bypass the need for mesh deformation or regeneration. The so-called boundary expressions for shape directional derivatives turn out to be exact for fictitious-domain methods. However, when mesh deformations are used, only the more cumbersome volume expressions turn out to be exact. Recent numerical experiences using CutFEM together with level-set geometry descriptions indicate less parameterization-induced artificial stiffness along boundaries compared to previously used methods. Rerunning previous optimization studies with the new method led to the discovery of unexpected but acoustically very beneficial subwavelength patterns when optimizing acoustic horns. In a 3D application, an analogous CutFEM approach was used to optimize the shape of the phase plug of a compression driver, the standard sound source of acoustic horns. The optimization used a base design that currently is assumed to be unpractical, due to the lack of simple design guidelines. However, after shape optimization, it turned out that this base design very well may be superior to the industry standard design.

Martin Berggren
Department of Computing Science, Umea University
martin.berggren@cs.umu.se

MS59

A Data-Oriented Programming Model for Massively Parallel Finite Element Computations

The development of MPI-based finite element (FE) applications is significantly more involved than for their sequential counterparts. Some bugs can be tracked with conventional tools in a single MPI rank, but genuine parallel errors often need to be debugged in parallel. Fixing an application that requires at least 100 ranks to crash can be very challenging, even with state-of-the-art parallel debuggers. We developed a novel data-oriented parallel programming

model that allows one to express parallel algorithms in a generic way, without explicitly relying on MPI communication directives. This makes possible to use different back-ends to run the generic parallel algorithms. With a sequential back-end, the data structures are logically parallel from the user perspective, but they are processed using a single rank. This makes possible to use serial debuggers for code development, which dramatically improves the user experience. Once the code has been debugged with the sequential back-end, it can be deployed in a supercomputer using a MPI back-end. Using the new model, we have implemented distributed vectors, sparse matrices, and other functionality needed in distributed-memory FE computations. They are implemented in Julia and made freely available in github [github.com/fverdugo/PartitionedArrays.jl]. With these tools, we were able to solve large FE computations with nearly optimal weak and strong scaling up to tens of thousands of CPU cores [doi.org/10.21105/joss.04157].

Francesc Verdugo

Department of Computer Science
Vrije Universiteit Amsterdam
f.verdugo.rojano@vu.nl

MS59

Neural Control of Discrete Weak Formulations: Galerkin, Least-Squares and Minimal-Residual Methods with Quasi-Optimal Weights

There is tremendous potential in using neural networks to optimize numerical methods. In this talk, I will introduce and analyse a framework [1] for the neural optimization of discrete weak formulations, suitable for finite element methods. The main idea of the framework is to include a neural-network function acting as a control variable in the weak form. Finding the neural control that (quasi-) minimizes a suitable cost (or loss) functional, then yields a numerical approximation with desirable attributes. In particular, the framework allows in a natural way the incorporation of known data of the exact solution, or the incorporation of stabilization mechanisms (e.g., to remove spurious oscillations). The main result of the analysis pertains to the well-posedness and convergence of the associated constrained-optimization problem. In particular, under certain conditions, discrete weak forms are stable, and quasi-minimizing neural controls exist, which converge quasi-optimally. The analysis results are specialized to Galerkin, least-squares and minimal-residual formulations, where the neural-network dependence appears in the form of suitable weights. Elementary numerical experiments support our findings and demonstrate the potential of the framework. [1] Brevis, Muga, Van der Zee, Neural Control of Discrete Weak Formulations: Galerkin, Least-Squares and Minimal-Residual Methods with Quasi-Optimal Weights, arXiv:2206.07475, (2022)

Kris van Der Zee

University of Nottingham
kg.vanderzee@nottingham.ac.uk

Ignacio Brevis

School of Mathematical Sciences
University of Nottingham
ignacio.brevis@nottingham.ac.uk

Ignacio Muga

Instituto de Matemáticas
Pontificia Universidad Católica de Valparaíso

ignacio.muga@pucv.cl

MS60

Statistical Shape Analysis of Shape Graphs with Applications to Retinal Blood-Vessel Networks

This paper provides theoretical and computational developments in statistical shape analysis of shape graphs and demonstrates them using analysis of data from complex retinal blood-vessel (RBV) networks. The shape graphs are represented by a set of nodes and a set of edges (planar articulated curves) connecting some of these nodes. The goals are to utilize shapes of edges and connectivities and locations of nodes to: (1) characterize full shapes, (2) quantify shape differences, and (3) model statistical variability. We develop a mathematical representation, elastic Riemannian shape metrics, and associated tools for such statistical analysis. Specifically, we derive tools for graph matching, shape geodesics, shape summaries, and shape modeling. One key challenge here is the registration of nodes across large graphs and we develop a novel multi-scale representation of shape graphs to handle this challenge. We utilize the concepts of (1) "effective resistance" to cluster nodes and (2) elastic shape averaging of edge curves, to change graph details while maintaining overall structures. Registration is then performed by bringing graphs to similar scales before matching. We demonstrate these ideas on Retinal Blood Vessel (RBV) networks taken from the STARE and DRIVE databases.

Aditi Basu Bal

Florida State University
ab18z@fsu.edu

Xiaoyang Guo

Meta
xiaoyangg@fb.com

Tom Needham, Anuj Srivastava

Florida State University
tneedham@fsu.edu, anuj@stat.fsu.edu

MS60

Data - Driven Continuum Modeling of Active Nematics via Sparse Identification of Nonlinear Dynamics

Data-driven modeling methods have recently shown great potential in determining accurate continuum models for complex systems directly from experimental measurements. One such complex system is the active nematic liquid crystal system consisting of microtubule-motor protein assemblies immersed in a fluid. This system exhibits rich non-equilibrium behavior, including spontaneous creation and annihilation of topological defects. Although several models have been proposed for the system, the governing equations remain under debate. We here present a model extracted directly from experimental image data via the "sparse identification of nonlinear dynamics" (SINDy) data-driven modeling technique. This model discovery process includes extracting appropriate data for a continuum model from experimental data, constructing a plausible library of model terms, and solving a sparse regression problem. A number of issues arise in this process, including strong correlations between library terms and instabilities in the discovered models. We present numerical simulations and the results of some modern statistical and classical pen-and-paper analyses which combine to clarify the discovery results. We then discuss the physical implica-

tions of the learned model, and compare the model with those proposed previously.

Anand Oza, Connor Robertson, Travis Askham
New Jersey Institute of Technology
anand.u.ozaj@njit.edu, cjr59@njit.edu, askham@njit.edu

MS60

Fluid-Solid Interaction Inspired by the Lotus Leaf

We present theory and develop a fast and accurate boundary integral method for the computation of incompressible Stokes flow over a superhydrophobic (SH) surface, such as occurs during rainwater shedding on a lotus leaf. Boundary integral methods have several advantages in SH flow computations. However, such problems exhibit flow singularities due to complex surface microstructure, which leads to mixed boundary conditions and geometric corners. Standard quadrature rules for smooth integrals result in a severe loss of accuracy. Adaptive mesh refinement mitigates the issue, however, the size of the discrete problem grows significantly with refinement level, and it can still be difficult to obtain satisfactory accuracy due to the ill-conditioning of the linear system. To resolve these issues, we combine the recently developed Recursively Compressed Inverse Preconditioning (RCIP) method with a scaling technique and the Fast Multipole Method to obtain a fast and accurate numerical scheme for SH flow computations. Several examples are presented to illustrate the performance of the method.

Michael Siegel
New Jersey Institute of Technology
misieg@njit.edu

Shidong Jiang
Flatiron Institute, Simons Foundation
Center for Computational Mathematics
sjiang@flatironinstitute.org

Kosuke Sugita
New Jersey Institute of Technology
ks823@njit.edu

MS60

Self-Organization of Microtubules Through Hydrodynamic Interactions Drives Robust and Oscillatory Rotational Flows

The piconewton forces generated by microtubules polymerizing against the cell cortex or other artificial boundaries are sufficient to deform long microtubules. When microtubules are sparse, their deformations are disordered, characterized by high-frequency buckling and inducing only localized flows. When many microtubules are present the nature of the microtubule deformations and induced flows can change dramatically, giving rise to long-range order and coherent flows. Using a combination of experiments, large-scale simulations of microtubules interacting hydrodynamically through a viscous fluid, and a coarse-grained theory for dense beds of filaments, we elucidate the mechanisms that underlie the self-organization of microtubule ensembles and their subsequent generation of rotation in artificially confined asters in *Xenopus Laevis* extract.

David Stein
Flatiron Institute
Simons Foundation
dstein@flatironinstitute.org

Abdullah Sami
Department of Molecular Biology
University of Wyoming
asami1@uwyo.edu

Gokberk Kabacaoglu
Bilkent University
gkabacaoglu@bilkent.edu.tr

Jesse Gatlin
Department of Molecular Biology
University of Wyoming
jgatlin@uwyo.edu

Michael Shelley
Group Leader in Biophysical Modeling Group,
Center for Computational Biology, Simons Foundation
mshelley@simonsfoundation.org

MS61

Rigid and Non-Rigid Multi-Way Point Cloud Matching via Late Fusion

Correspondences fuel a variety of applications from texture-transfer to structure from motion. However, simultaneous registration or alignment of multiple, rigid, articulated or non-rigid partial point clouds is a notoriously difficult challenge in 3D computer vision. With the advances in 3D sensing, solving this problem becomes even more crucial than ever as the observations for visual perception hardly ever come as a single image or scan. In this talk, I will present an unfinished quest in pursuit of generalizable, robust, scalable and flexible methods, designed to solve this problem. The talk is composed of two sections diving into (i) MultiBodySync, specialized in multi-body and articulated generalizable 3D motion segmentation as well as estimation, and (ii) SyNoRim, aiming at jointly matching multiple non-rigid shapes by relating learned functions defined on the point clouds. Both of these methods utilize a family of recently matured graph optimization techniques called synchronization as differentiable modules to ensure multi-scan / multi-view consistency in the late stages of deep architectures. Our methods can work on a diverse set of datasets and are general in the sense that they can solve a larger class of problems than the existing methods.

Tolga Birdal
Imperial College London
t.birdal@imperial.ac.uk

MS61

Hierarchical Bayesian 3D Synthetic Aperture Radar Reconstruction Using Joint Sparsity

Three-dimensional (3D) synthetic aperture radar (SAR) imaging is a burgeoning field of research with various applications in both military and civilian domains. Methods in compressive sensing and sparse reconstruction have proven to be effective in recovering highly resolved SAR images. These techniques provide a point estimate for the volumetric image. New research in the realm of Bayesian modeling has opened the door to quantifying uncertainty information in the image recovery. In this work, we introduce a Bayesian volumetric approach to reconstruct 3D SAR images from noisy SAR data as well as provide uncertainty quantification for two-dimensional projections of the volumetric image. We consider various metrics comparing our technique with similar point-estimation methods.

Distribution A: Cleared for Public Release; Distribution unlimited. PA Approval #: [AFRL-2022-4296]

Dylan Green
Dartmouth College
dylan.p.green.gr@dartmouth.edu

Jr Jamora
Air Force Research Lab
Wright-Patterson Air Force Base
jan_rainer.jamora@us.af.mil

Anne Gelb
Dartmouth College
Department of Mathematics
annegelb@math.dartmouth.edu

MS61

Refraction and Absorption for Underwater Shape Recovery

Underwater 3D reconstruction has been gaining more attention due to its profound applications in a wide range of fields including medicine (e.g., endoscopy), marine biology, oceanography, as well as general surveillance and navigation. The properties of water not only bring challenges, but are also part of the solution. The transmittance of water depends on both the distance and the wavelength, and water has significant absorption of light in the near-infrared wavelength range. We leverage this effect and propose two methods for passive underwater 3D sensing. The first method, Surface Normals and Shape From Water, estimates dense per-pixel surface normals and shape of underwater objects without any artificial constraint on one another (e.g., smoothness), enabling recovering intricate surface details that would otherwise be challenging for conventional methods. The second method, Non-rigid Shape From Water, reconstructs the holistic and consolidated shape of a dynamic, non-rigid object. The shape of an underwater object at different time instances are recovered, integrated and refined as it deforms and moves in water. In this talk, I will not only present the theory of these methods, but also successfully validate the effectiveness of the theory with the experimental results by applying the methods to a number of real-world objects in water.

Meng-Yu Kuo
University of Minnesota
kuo00052@umn.edu

MS61

Fast, Accurate and Memory-Efficient Multi-Image Matching

Previous partial permutation synchronization (PPS) algorithms, which are commonly used for multi-image matching, often involve computation-intensive and memory-demanding matrix operations. These operations become intractable for large scale structure-from-motion datasets. For pure permutation synchronization, the recent Cycle-Edge Message Passing (CEMP) framework suggests a memory-efficient and fast solution. Here we overcome the restriction of CEMP to compact groups and propose an improved algorithm, CEMP-Partial, for estimating the corruption levels of the observed partial permutations. It allows us to subsequently implement a nonconvex weighted projected power method without the need of spectral initialization. The resulting new PPS algorithm, MatchFAME (Fast, Accurate and Memory-Efficient Matching),

only involves sparse matrix operations, and thus enjoys lower time and space complexities in comparison to previous PPS algorithms. We prove that under adversarial corruption, though without additive noise and with certain assumptions, CEMP-Partial is able to exactly classify corrupted and clean partial permutations. We demonstrate the state-of-the-art accuracy, speed and memory efficiency of our method on both synthetic and real datasets.

Gilad Lerman
University of Minnesota, Minneapolis
School of Mathematics
lerman@umn.edu

MS61

Reconstruction of 3D Molecular Structure from 2D Covariance in Cryo-EM

In single-particle cryo-electron microscopy (EM), the 3D structure of a molecule needs to be determined from its noisy 2D projection images. Unfortunately, each of these projection images is taken in an unknown viewing direction. The high noise level makes it challenging to accurately estimate the viewing directions, ultimately affecting the entire reconstruction process. In this talk, we describe an approach for obtaining a 3D model using low-order statistics without directly estimating the viewing directions based on the non-uniformity of the distribution of viewing directions.

Nir Sharon
Tel Aviv University
nsharon@tauex.tau.ac.il

MS62

ML Accelerator Hardware: A New Model For Parallel Sparse Computations?

A major recent development in computer hardware was the rise of dedicated accelerator hardware for machine learning applications such as the Graphcore IPU and Cerebras WSE. These processors have evolved from the experimental state into market-ready products, and they have the potential to constitute the next major architectural shift after GPUs saw widespread adoption a decade ago. In this talk we will present the new hardware and discuss the programming techniques that are required to unlock their potential. We present implementations of basic graph and matrix algorithms and show early results on the attainable performance, as well as comparisons to other architectures. We follow up by discussing the wider implications of the architecture for algorithm design and programming, along with the wider implications of adopting such hardware.

Johannes Langguth
Simula Research Laboratory
langguth@simula.no

Luk Burchard
Simula Research Laboratory
University of Oslo
luk@simula.no

Xing Cai
Simula Research Laboratory
1325 Lysaker, Norway

xingca@simula.no

MS62

The Chain Rule of Differentiation Is Associative - So What?

The Chain Rule of Differentiation is Associative - So What? ... well, feasibility of backpropagation for training artificial neural networks by some variant of stochastic gradient descent is one rather obvious consequence. Adjoint algorithmic differentiation (AD) of numerical simulation programs enabling large-scale error control, uncertainty quantification or nonlinear optimization turns out to be the more general concept. Resulting challenges are discussed briefly. ... in fact, several combinatorial optimization problems arise when aiming to design a near-optimal AD algorithm for a given differentiable program. We discuss the main challenges and ideas behind AD mission planning.

Uwe Naumann
RWTH Aachen University
Software and Tools for Computational Engineering
naumann@stce.rwth-aachen.de

Simon Märtens
RWTH Aachen University, Germany
maertens@stce.rwth-aachen.de

MS62

Parallel Hdbscan* on GPUs

This talk presents efficient parallel algorithms for computing HDBSCAN* (hierarchical DBSCAN) on modern GPU and multicore architectures. HDBSCAN* is a density-based algorithm that finds clusters of arbitrary shapes, sizes, and densities within a single dataset. HDBSCAN* can compute both hierarchical and flat clusterings. The HDBSCAN* computation consists of two steps: 1) construction of a minimal spanning tree (MST), and 2) construction of single linkage hierarchy using the constructed MST. Optionally, a third step to produce flat clustering can be performed by using the stability scores on the hierarchy. For the MST computation, we present an efficient single-tree based algorithm that combines the simplicity of a tree traversal with several pruning optimizations to improve parallel efficiency. For the dendrogram computation, we present a new algorithm that exploits the dendrograms highly chained structure using a tree-contraction data structure that we call α -tree. The α -tree reduces the dendrogram computation to a set of list ranking, prefix sum, and segment sort kernels, all of which are can be efficiently implemented on GPUs and multicore architectures. Additionally, we show that the α -tree can also be used to speed up the computation of the flat clustering on parallel architectures.

Andrey Prokopenko, Piyush Sao, Damien Lebrun-Grandie
Oak Ridge National Laboratory
prokopenkoav@ornl.gov, saopk@ornl.gov, lebrun-grandie@ornl.gov

MS62

Towards Accelerating AI Using Fast and Feasible Matrix Multiplication

Training DNN's increasingly requires large resources. Significant time is spent on matrix multiplication, typically between 45%-95%. Most current math libraries (for CPU

and GPU) and all state-of-the-art hardware accelerators (such as Googles TPU and Intels Habana Labs Gaudi) are based on the cubic-time matrix multiplication algorithm, despite more than five decades of research on sub-cubic time algorithms. Why is that? Many of the sub-cubic time algorithms have large hidden constants in the arithmetic complexity. We obtain a high performance fast matrix multiplication that outperforms DGEMM of Intels MKL on feasible matrix size - up to nearly 2 speedup. Pans 1980-1982 algorithms have the lowest asymptotic complexity of all algorithms applicable to matrices of feasible dimensions. However, their large hidden constants make them impractical. We reduce these coefficients by 90% - 98%, often down to 2. We show that these are optimal or close to optimal. Fast matrix multiplication calls the cubic time algorithm on small sub-blocks. We obtain a new set of algorithms that outperform the cubic time one on small blocks, by trading multiplications for additions. E.g., we multiply 22 blocks using only four multiplications. This seemingly contradicts Winograd's 1971 lower bound. We obtain a lower bound proving our algorithm is optimal. Based on joint work with Gal Beniamini, Nathan Yan Cheng, Yoav Gross, Tor Hadas, Olga Holtz, Elaye Karstadt, and Noa Vaknin.

Oded Schwartz
Hebrew University of Jerusalem
odedsc@cs.huji.ac.il

MS62

High Performance Sparse Tensor Contractions

Tensors are widely encountered in several application domains such as scientific computing, machine learning and data analytics. Tensor contraction is a key algebraic operation in many applications involving multi-dimensional data. It is a higher dimensional analog of matrix-matrix multiplication. Sparse tensor contraction suffers from poor data-locality and irregular accesses, which poses a significant performance challenge. We explore the use of hashing-based methods to make sparse tensor contraction operation more efficient and performant on shared-memory systems.

Somesh Singh
ENS de Lyon
somesh.singh@ens-lyon.fr

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

MS63

Quasi Projective Synchronization of Complex Valued Cohen-Grossberg Neural Networks with Time Varying Delay and Mismatched Parameters

In this article the quasi-projective synchronization of time-varying delayed complex-valued Cohen Grossberg Neural Networks (CGNNs) with non-identical parameters has been studied. As complete projective synchronization is impossible due to parameters mismatches projective coefficient and controller, a drive has been taken to achieve quasi-projective synchronization of distinct complex-valued CGNNs. The purpose of this study is to find a criterion for quasi-projective synchronization of two non-identical CGNNs by constructing a suitable controller and by using direct method. The important contribution is to estimate the bound on the synchronization error. Some

sufficient criteria for synchronization between master and response systems are also established. The efficiency of the proposed method is justified through numerical simulation applied to a specific example.

Sapna Baluni

Department of Mathematical Sciences, Indian Institute of Technology (BHU), Varanasi-221005, India.
sapnabaluni.rs.mat19@itbhu.ac.in

MS63

Some Mathematical Aspects of PINN-VPINN Discretization of PDEs

We introduce a Petrov-Galerkin framework to derive rigorous a priori and a posteriori error estimates for VPINN discretizations of elliptic boundary-value problems. The analysis relies on an inf-sup condition between trial and test spaces; this allows us to control the error, in the energy norm, between the exact solution and a suitable high-order piecewise-polynomial interpolant of the computed neural network. The efficiency and accuracy of the computation depends, among other factors, on the way boundary conditions are imposed. We will discuss several options for enforcing Dirichlet or Neumann boundary conditions, and we will compare their behavior for both PINNs and VPINNs, and for various types of equations.

Claudio Canuto

Politecnico di Torino
claudio.canuto@polito.it

Stefano Berrone

Politecnico di Torino
Dipartimento di Scienze Matematiche
stefano.berrone@polito.it

Moreno Pintore

Politecnico di Torino
moreno.pintore@polito.it

MS63

Thermodynamics-Informed Neural Networks

In this talk, we review the recent advances in the development of thermodynamics-informed neural networks for learning physical phenomena. These networks work under a dynamical systems analogy and employ inductive biases so as to guarantee the fulfillment of the laws of thermodynamics. While for conservative dynamics a Hamiltonian structure could be a very convenient bias, in the case of dissipative phenomena an alternative description should be used. We employ a metriplectic description coined as GENERIC [Oettinger & Grmela, 1997] that ensure conservation of energy in closed systems and non-negative entropy production. In combination with a Graph Neural Network approach, these networks can even impose by construction translational invariance and other symmetries of the studied systems.

Quercus Hernandez, Beatriz Moya

Universidad de Zaragoza
quercus@unizar.es, beam@unizar.es

Alberto Badias

Universidad Politecnica de Madrid
alberto.badias@upm.es

David Gonzalez

Universidad de Zaragoza

gonzal

Francisco Chinesta

ENSAM Institute of Technology
unizar.es

Elias Cueto

I3A, University of Zaragoza, Spain
francisco.chinesta@ensam.eu

MS63

Two Approaches Using Deep Learning to Solve Partial Differential Equations

Many differential equations and partial differential equations (PDEs) are being studied to model physical phenomena in nature with mathematical expressions. Recently, new numerical approaches using machine learning and deep learning have been actively studied. There are two mainstream deep learning approaches to approximate solutions to the PDEs, i.e., using neural networks directly to parametrize the solution to the PDE and learning operators from the parameters of the PDEs to their solutions. As the first direction, Physics-Informed Neural Network was introduced in (Raissi, Perdikaris, and Karniadakis 2019), which learns the neural network parameters to minimize the PDE residuals in the least-squares sense. On the other side, operator learning using neural networks has been studied to approximate a PDE solution operator, which is nonlinear and complex in general. In this talk, I will introduce these two ways to approximate the solution of PDE and my research related to them.

Jaeyong Lee

KIAS
jaeyong@kias.re.kr

MS64

Discovering Reduced Order Models from Partial Measurements with Deep Delay Autoencoders

A central challenge in data-driven model discovery is the presence of hidden, or latent, variables that are not directly measured but are dynamically important. We develop a deep latent-model autoencoder that simultaneously optimizes for a nonlinear transformation that projects the measurement data on an effective coordinate system and discovers a sparse differential equation governing the latent variable. Our approach is motivated by Takens' theorem which provides conditions for when it is possible to augment partial measurements with time delayed information, resulting in an attractor that is diffeomorphic to that of the original full-state system. We show that the proposed architecture can discover a coordinate transformation from the delay embedded space into a new space where it is possible to represent the dynamics in a sparse, closed form, differential equation. Furthermore, we show how this architecture can be used for discovering reduced order models from given partial or full measurements of the system. We demonstrate this approach on the Lorenz, Rössler, and Lotka-Volterra systems, learning dynamics from a single measurement variable. As a challenging example, we learn a Lorenz analogue from a single scalar variable extracted from a video of a chaotic waterwheel experiment. Finally, we discuss the limitations and robustness of the method and propose future research directions.

Joseph Bakarji

University of Washington
jbakarji@uw.edu

Kathleen Champion
University of Washington, Department of Applied
Mathematics
kpchamp@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Steven Brunton
University of Washington
sbrunton@uw.edu

MS64

Non-Intrusive Physics-Informed Reduced Order Modelling of Multiscale Computational Mechanics Problems

Numerical methods for multiscale modelling and simulation require the solution of a micro-scale problem at every Gauss quadrature point of the macro-scale computational domain. This renders the computational costs associated with solving a non-linear forward problem to be exorbitant. Furthermore, while there exist intrusive and non-intrusive machine learning (ML) based hyper-reduction techniques, the corresponding offline and online training phase is inefficient. With state-of-the-art methods in mind, a numerical framework is needed that offers: (i) a significant improvement in the generalization accuracy for the unknown solution and the derived quantities when the full-order training data is limited, (ii) an efficient offline training phase in the sense that an improved sampling strategy is invoked to identify where to perform the full-order solve, and (iii) an offline- and online-efficient, physically consistent, non-intrusive affine decomposition and hyper-reduction technique. To this end, inspired from the reduced basis empirical quadrature procedure (RB-EQP) and scientific ML techniques, we develop a novel, non-intrusive numerical method, trained using pretraining and semi-supervised learning strategies, to efficiently and cheaply solve non-linear forward multiscale mechanics problems. Finally, we demonstrate the performance of the proposed method on a hyper elasticity and a coupled thermo-electro-mechanical problem in a multi-scale framework.

Harshit Bansal, Rahul Dhopeshwar, Yankun Hong
Eindhoven University of Technology
h.bansal@tue.nl, r.dhopeshwar@tue.nl, y.hong@tue.nl

Karen Veroy
Department of Mathematics and Computer Science
Eindhoven University of Technology
k.p.veroy@tue.nl

MS64

GPLaSDI: Identifying Autoencoders Latent Space Dynamics with Gaussian Processes for Non-Intrusive Reduced Order Models

Solving PDE numerically is challenging and computationally costly, which motivates the development of faster reduced order models (ROMs). In recent years, the advance in machine learning allows the development of novel non-linear projection methods, such as Latent Space Dy-

namics Identification (LaSDI) [Fries W., He X., Choi Y., LaSDI: Parametric Latent Space Dynamics Identification, Computer Methods in Applied Mechanics and Engineering (2022)]. LaSDI maps full-order PDE solutions into a latent space using auto-encoders and learns the system of ODEs governing the latent space dynamics. Interpolating and solving the system of ODEs allows for fast ROM predictions by feeding the predicted latent dynamics into the decoder. In this talk, we present Gaussian-process-LaSDI (GPLaSDI). GPLaSDI improves LaSDI in two ways: First, it uses Gaussian processes to interpolate the coefficients of the learnt latent dynamics, which allows for quantifying the uncertainty over the ROM predictions. Second, it uses the uncertainty to pick next parameter point where additional ROM data is collected to improve accuracy. This allows us to apply GPLaSDI to problems without any known underlying governing equations, e.g., noisy experimental data. Our method can achieve up to a few thousand time speed up with minimal reliance on the full order solver, provides meaningful confidence intervals over its prediction, and does not require to know the governing PDE or have access to its residual.

Christophe Bonneville
Cornell University
cpb97@cornell.edu

Youngsoo Choi
Lawrence Livermore National Laboratory
choi15@llnl.gov

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Livermore, CA
ghosh5@llnl.gov

MS64

Deep Learning-Based Reduced Order Modelling for the Real-Time Simulation of Steady-State Non-linear Dynamics in MemS

Micro-Electro-Mechanical-Systems (MEMS) represent essential components in high-end technology applications e.g., resonators and gyroscopes. The increasing performance demand in terms of lower power consumption and enhanced sensitivity requires performing nonlinear dynamic simulations that require adequate tools. MEMS designers can resort to standard full order simulation techniques such as the Finite Element (FE) Method, nevertheless, the high computational burden is limiting their practical usage. Reduced Order Models (ROMs) provide an appealing alternative to building reliable and efficient non-linear dynamical models. In particular Data-Driven methods allow building a ROM starting from data and without acting intrusively in the reduction procedure itself. Most of these approaches aim at modelling the dynamic behaviour of general systems, i.e. transient and steady-state regimes. Despite the generality of these approaches, some properties typical of the steady-state regime e.g., periodicity are difficult to reproduce. In this contribution, we propose a technique that, through deep learning autoencoders and harmonic decomposition, allows building efficient and accurate ROMs for the steady state regime for resonating MEMS.

Giorgio Gobat
DICA, Politecnico di Milano
giorgio.gobat@polimi.it

Alessia Baronchelli, Attilio Frangi
 Politecnico di Milano
 Dipartimento di Ingegneria Civile e Ambientale, DICA
 alessia.baronchelli@mail.polimi.it, attilio.frangi@polimi.it

MS65

Towards Learning to Rank and Decompose Rough Surfaces Based on Drag Penalty

Turbulent flows over irregular rough surfaces are ubiquitous in both nature and industry. Roughness increases the momentum transfer near the wall and the hydrodynamic drag on the wall. The increase of drag depends on various topographic features. Several studies have investigated the effects of statistical parameters of roughness topography on the increase of hydrodynamic drag. However, no generally applicable statistical model exists that can accurately predict drag on various rough surfaces (e.g., gaussian surfaces, positively and negatively skewed surfaces, etc.). Moreover, many models struggle to properly rank the rough surfaces based on drag penalty. In practice, ranking the drag penalties from rough surfaces is generally more critical than estimating the absolute value of drag itself. This presentation will introduce attempts to rank rough surfaces using various machine learning techniques based on their drag penalties. In addition, we will discuss distinctive roughness structures suspected to increase drag significantly by decomposing modes of feature maps in convolutional neural networks.

Zhaoyu Shi, Seyed Morteza Khorasanib
 KTH Royal Institute of Technology
 zhaoyus@kth.se, smhk2@mech.kth.se

Sangseung Lee
 Inha University
 sangseunglee@inha.ac.kr

Shervin Bagheri
 KTH Royal Institute of Technology
 shervin@mech.kth.se

MS65

Super-Resolving Turbulent Flows with Machine Learning: a Survey

Machine-learning-based super resolution has become a powerful tool for turbulent flows. Super resolution reconstructs fine-scale structures from their coarse input. This concept amounts not only to sparse reconstruction but can also be related to sub-grid scale modeling for turbulent flow simulations. We present a case study of machine-learning-based super-resolution analysis of turbulent flows to discuss its capabilities and extensions for a range of fluid mechanics problems. Convolutional neural network-based methods are used for the present survey of machine-learning-based super resolution. We find that embedding scale invariance is important in the construction of machine-learning model. Furthermore, physics-based cost function can greatly assist with the reconstruction in terms of accuracy and robustness against noisy low-resolution input. Towards the end of the presentation, we will also discuss the challenges and outlook of machine-learning-based super-resolution reconstruction with turbulence.

Kai Fukami
 University of California, Los Angeles
 kfukami1@g.ucla.edu

Koji Fukagata
 Keio University, Japan
 fukagata@mech.keio.ac.jp

Kunihiko Taira
 University of California, Los Angeles
 Mechanical and Aerospace Engineering
 ktaira@seas.ucla.edu

MS65

Submodular Optimization for Near-Optimal Sensor and Actuator Placement

Adequately choosing sensors to monitor a high-dimensional system is of utmost importance in numerous situations, most notably for feedback control applications. Yet, optimal sensor and actuator placement is a rapidly intractable combinatorial problem. Following an optimal design strategy, near-optimal selections can still be obtained using various convex relaxations of the otherwise combinatorial problem. One major drawback however is the limited scalability of the algorithms used to solve these convex problems in practice, or their long time to solution. Based on the theory of modular and submodular functions, this contribution will present different algorithms for sensor selection having extremely high computational efficiency along with near-optimal guarantees. We will focus in particular on D-optimality and compare the performances of a semi-definite programming relaxation of the problem against that of the corresponding submodular relaxation and SupSub approximation.

Jean-Christophe Loiseau
 ParisTech
 loiseau.jc@gmail.com

MS65

Interpretable Nonlinear Reduced-Order Modelling with Autoencoders

Autoencoders are machine-learning methods that enable a reduced-order representation of data. They consist of an encoder, which compresses the data in a latent space, and a decoder, which decompresses the data back to the original space. If only linear operations are performed during the encoding and decoding phases, an autoencoder can learn the principal components of the data. On the other hand, if nonlinear activation functions are employed, an autoencoder learns a nonlinear model of the data in the latent space. The interpretability of the latent space, however, is not yet fully established. In this work, we physically interpret the latent space with simple tools from differential geometry. The interpretation is employed on canonical turbulent flows, i.e., the Kolmogorov flow and the minimal flow unit. The results show that the autoencoder learns the optimal submanifold in which the reduced-order dynamics is well represented. The latent variables are exploited for reducing the model's complexity whilst keeping optimal accuracy on the spatiotemporal dynamics. This work opens opportunities for extracting physical insight from the latent space and for nonlinear model reduction.

Luca Magri
 Imperial College London
 l.magri@imperial.ac.uk

Nguyen Anh Khoa Doan
 TU Delft

n.a.k.doan@tudelft.nl

MS65

Reduced-Order Modeling with a Regression-Aware Autoencoder

The first step in reduced-order modeling (ROM) workflows is finding a low-dimensional representation of a highly-dimensional system. The second step of ROM often requires training a nonlinear regression model to predict physical quantities of interest from the reduced representation. Much of the research on training ROMs thus far has tackled those two steps separately. While they both come with their challenges, a good-quality low-dimensional system representation usually facilitates building a regression model. In this work, we leverage the link between dimensionality reduction and nonlinear regression. We propose an approach where dimensionality reduction and nonlinear regression are considered jointly within an autoencoder-like neural-network architecture. The dimensionality reduction (encoding) is affected by forcing accurate regression (decoding) of the quantities of interest. We show that such a joint architecture leads to improved low-dimensional representations as the two steps communicate with each other through backpropagation. We apply our regression-aware autoencoder on test cases coming from reacting and non-reacting flow systems. The relevant quantities of interest are the important state variables and highly nonlinear source terms required by the reduced model. The proposed approach can serve as an effective replacement of standalone dimensionality-reduction techniques whenever nonlinear regression is anticipated in the downstream use.

Kamila Zdybal, Riccardo Malpica Galassi
 Université Libre de Bruxelles
 kamilazdybal@gmail.com, riccardo.malpica.galassi@ulb.be

James Sutherland
 University of Utah
 james.sutherland@utah.edu

Ricardo Vinuesa
 KTH Royal Institute of Technology
 rvinuesa@mech.kth.se

Alessandro Parente
 Université Libre de Bruxelles
 alessandro.parente@ulb.be

MS66

Efficient Discretizations and Multigrid Techniques for Coupled Hyperbolic-Parabolic Systems

For coupled hyperbolic-parabolic systems, we study their numerical approximation and the efficient solution of the resulting algebraic systems. The models are prototype systems of dynamic poro- and thermoelasticity. Different formulations of the continuous model, including a first-order one in space and time, are studied. Well-posedness of solutions to the models in fitted functional analytical settings is reviewed. Tailored families of space-time finite element discretizations are analyzed for the representations of the continuous problem. For the solution of the algebraic equations, the efficiency of preconditioning GMRES iterations by V-cycle geometric multigrid methods is demonstrated numerically. Therein, higher order time discretizations, increasing the complexity of the block matrix structures, are

focused.

Markus Bause
 Helmut-Schmidt-University
 Department of Mechanical Engineering
 bause@hsu-hh.de

Mathias Anselmann
 Helmut-Schmidt-University Department of Mechanical Engineering
 anselmann@hsu-hh.de

MS66

Simulation of Mixing and Reactions in a Full Size Porous Column Using a Hybrid Eulerian-Lagrangian Approach

We perform the largest and most highly resolved simulation to date of solute transport through a laboratory scale porous medium column. Flow and transport are simulated using a finite volume Eulerian approach, while mixing and spreading processes are studied using a Lagrangian approach. This hybrid approach allows us to have an unprecedented look at the fine and larger scale processes that impact observations and interpretations from real column experiments. In particular we are able to study the influence of boundary effects, which are typically neglected in current interpretations and develop rigorous theories that enable a more accurate interpretation of what is actually happening with the end goal of developing improved upscaled models that account for the full range of complexities that any real experiment will have. Additionally we have developed novel approaches for interface tracking through the porous column that provide an unprecedented view of mixing processes that are key to understanding and ultimately predicting mixing-driven reactions where current predictions either overestimate the amount of reaction or artificially and unphysical adjust reaction parameters to match observations, but at the cost of loss of generalization to other settings. Our novel numerical approaches and theories overcome many previous limitations that will benefit the world of scientific computing and experimentalists alike.

Diogo Bolster
 University of Notre-Dame
 dbolster@nd.edu

MS66

HPC Multiphysics Simulation of CO₂ Geological Storage

Safe and efficient operation of CO₂ geological storage projects requires numerical simulation of a multi-physics problem in which compositional multiphase flow and transport are tightly coupled with the porous medium deformation. To simulate these processes, one needs to solve a set of coupled, nonlinear, time-dependent partial differential equations (PDEs) governing the conservation of mass of each component and linear momentum of the solid-fluid mixture. We present a scalable fully-implicit framework based on a displacement/pressure/component density formulation. The discrete displacement field is approximated using nodal interpolation. Pressure and component densities are assumed cell-wise constant, with interface pressure also present if a hybridization strategy is chosen. A Newton-Krylov approach is used to advance the solution in time, with a finely-tuned multigrid reduction (MGR) preconditioner introduced to accelerate con-

vergence. The solver is implemented in an open-source, exascale-compatible, research-oriented simulator for modeling fully coupled flow, transport and geomechanics in geological formations. Numerical results are presented to illustrate performance and robustness of the proposed solver on a variety of challenging test problems, including realistic field cases.

Nicola Castelletto

Lawrence Livermore National Laboratory
Livermore, CA, USA
castelletto1@llnl.gov

Matteo Cusini

Lawrence Livermore National Laboratory
cusini1@llnl.gov

François P. Hamon, Dick Kachuma

TotalEnergies E&P Research and Technology
francois.hamon@totalenergies.com,
dick.kachuma@totalenergies.com

Randolph R. Settgest, Joshua A. White

Lawrence Livermore National Laboratory
settgest1@llnl.gov, white230@llnl.gov

MS66

On a Stabilized Scheme for Biot's Model

Nonphysical oscillations might appear in the numerical solution of the fluid pressure for the quasi-static Biot's model for poroelasticity when standard finite element discretizations are considered for the discretization of the model. In this work, we address the issue related to the presence of non-physical oscillations in the pressure approximation for low permeabilities and/or small time steps, by proposing a new stabilization based in adding a term in the flow equation. We consider different finite element discretizations and illustrate how such a stabilized scheme provides numerical solutions that are free of non-physical oscillations. The new scheme allows us to iterate the fluid and mechanics problems in a fashion similar to the well-known fixed-stress split method. We also present numerical results illustrating the robust behavior of both the stabilization and iterative solver with respect to the physical and discretization parameters of the model.

Carmen Rodrigo, Ivaro Pé de la Riva, Francisco José

Gaspar
University of Zaragoza
carmenr@unizar.es, apedelariva@unizar.es,
fjgaspar@unizar.es

Xiaozhe Hu

Tufts University
xiaozhe.hu@tufts.edu

Ludmil Zikatanov

Pennsylvania State University
ludmil@psu.edu

MS66

A Deep ILearning Approach to Reduced Order Modeling of Partial Differential Equations for Porous Media

In the context of parametrized PDEs, Reduced Order Models (ROMs) allow for an efficient approximation of

the parameter-to-solution map, which is extremely useful whenever dealing with expensive many-query routines such as constrained optimization, sensitivity analysis and uncertainty quantification. Recently, motivated by the limitations of classical approaches such as the Reduced Basis method, many authors have been considering the use of neural networks and Deep Learning techniques for building non-intrusive ROMs. We tackle the development of a constructive approach based on deep neural networks for the efficient approximation of the parameter-to-solution map of PDEs. Our work is based on the use of deep autoencoders, which we employ for encoding and decoding a high-fidelity approximation of the solution manifold. We study the expressivity of these deep neural networks with respect to the solution manifold and we present some results that illustrate the connection between the expressivity and the complexity of the networks. Furthermore, we discuss its generalization to PDEs involving complex domains, leveraging on the idea of sparsifying the network architecture by means of information about the topology of the computational mesh used for the full order model. Finally, we present how these methodologies are applied to sensitivity analysis and uncertainty quantification of mathematical models for porous media such as Darcy flow in domains with fractures or channels.

Enrico Ballini

Department of Mathematics, Politecnico di Milano
enrico.ballini@polimi.it

Nicola Rares Franco, Stefania Fresca

Politecnico di Milano
nicolarares.franco@polimi.it, stefania.fresca@polimi.it

Luca Formaggia

MOX, Department of Mathematics
Politecnico di Milano, Italy
luca.formaggia@polimi.it

Alessio Fumagalli, Alessio Fumagalli, Andrea Manzoni

Politecnico di Milano
alessio.fumagalli@polimi.it, alessio.fumagalli@polimi.it,
andrea1.manzoni@polimi.it

Anna Scotti

anna.scotti@polimi.it
Politecnico di Milano, Italy

Piermario Vitullo

Department of Mathematics, Politecnico di Milano
piermario.vitullo@polimi.it

Paolo Zunino

Politecnico di Milano
Department of Mathematics
paolo.zunino@polimi.it

MS67

Novel Flexible Treatment Planning Framework for Proton Radiotherapy Using Julia and Automatic Differentiation

Radiotherapy aims to cure or improve the quality of life of cancer patients by depositing a precise amount of radiation dose within the tumour while depositing only a minimal dose to the surrounding healthy tissue. The dose distribution is shaped during treatment planning, which is currently a semi-automatic process requiring human planners with years of training and experience. State of the art plan-

ning systems provide tools to automate and optimise the treatment planning, but they are inflexible, which makes it hard to incorporate advances and insights from disciplines like machine learning. We develop a novel treatment planning framework in the Julia programming language. All our computations are performed on graphics processing units, reducing the planning time even further. This allows researchers to write codes that execute quickly, facilitating the transition from research to the clinic. The new planning framework is designed to optimise many quantities simultaneously thanks to the state-of-the-art automatic differentiation libraries available in the Julia ecosystem. We adopt a flexible loss-function approach that enables a precise formulation of the treatment goals. The loss function can trivially be extended to include new insights and requirements in a short time, rendering the framework versatile and future-proof.

Renato Bellotti, Bastien Golomer, Andreas Adelman,
Antony Lomax, Jan Hrbacek
Paul Scherrer Institut
renato.bellotti@psi.ch, bastien.golomer@psi.ch,
andreas.adelman@psi.ch, tony.lomax@psi.ch,
jan.hrbacek@psi.ch

MS67

VORONOI-FVM.JL - a Multiphysics Finite Volume Solver for Elliptic and Parabolic Systems

The two point flux Voronoi finite volume method on boundary conforming Delaunay meshes provides framework for developing discretizations for nonlinear elliptic and parabolic systems which conform to first principles of physics like mass conservation, maximum principles, positivity of concentrations, decay of discrete free energy. Closely oriented at the structure of this method, the VoronoiFVM.jl package allows to specify systems of partial differential equations on one-, two- and three-dimensional domains by providing nonlinear reaction, boundary, storage and flux functions, and a simplex grid representing the domain geometry. Newton's method in combination with various linear solvers is used to solve the discrete nonlinear problems. The package takes advantage of several features of the Julia language which will be highlighted during the talk:

- Automatic differentiation based on ForwardDiff.jl for the assembly of Jacobi matrices
- Composability with ODE and DAE solvers from DifferentialEquations.jl
- Dependency management via the Julia package manager

Among other topics, the package is used for numerically modeling semiconductor devices, electrochemical systems. Corresponding examples will conclude the talk.

Jürgen Fuhrmann
WIAS Berlin
fuhrmann@wias-berlin.de

MS67

Scientific Supercomputing with Julia - Insights from HPC Experts and Domain Scientists

We present an efficient approach for the development of nonlinear multi-physics applications written in Julia which are deployable on next generation xPU supercomputers. Powerful costless abstractions, metaprogramming and multiple dispatch enable writing a single code that is suitable

for both productive prototyping on a single CPU thread and large scale production runs on GPU or CPU supercomputers. High performance stencil computations can be expressed with math-close notation in hardware-agnostic kernels for which launch parameters are automatically derivable from the kernel arguments. Halo updates required for distributed memory parallelization can be triggered with a single function call, can fully and automatically overlap with computations and reach performance close to hardware limit. We outline the wide applicability of our approach by reporting on multiple multi-xPU solvers for geosciences as, e.g., 3D solvers for poro-visco-elastic two-phase flow and for reactive porosity waves. The solvers achieve high performance and scale close to ideally on thousands of NVIDIA Tesla P100 GPUs at the Swiss National Supercomputing Centre. Moreover, we have shown in recent contributions that our approach is naturally in no way limited to geosciences: we have showcased a computational cognitive neuroscience application modelling visual target selection and a quantum fluid dynamics solver using the nonlinear Gross-Pitaevski equation.

Samuel Omlin
Swiss National Supercomputing Centre
ETH Zurich
omlin@cscs.ch

Ludovic Räss
Lab of Hydraulics, Hydrology and Glaciology (VAW)
ETH Zurich
luraess@ethz.ch

MS67

The SciML Common Solver Interface

The SciML common interface is an abstraction to automating scientific machine learning, the connection of machine learning libraries to traditional scientific computing. In this talk we will walk through the SciML common interface and how its design connects to the intricacies of the numerical methods. We will dive into details of how automatic differentiation is mixed with analytical adjoints of equation solvers to achieve better performance than either technique alone, the automation of solver choices through open benchmarking platforms, and how abstract hierarchical interfaces are used to classify algorithms into new generic forms. We will detail how these algorithms within the SciML software suite which are leading to techniques, such as for the automated construction of differential equation systems from agent-based based systems through new discrete-stochastic automatic differentiation engines, along with advancements in ODE solver software which have come by integrating neural networks into the methods.

Christopher V. Rackauckas
Massachusetts Institute of Technology
accounts@chrisrackauckas.com

MS67

A Showcase of the New Functionalities of BifurcationKit.jl

In this talk, we present a panorama of the new functionalities of BifurcationKit.jl, a Julia package for performing automatic numerical bifurcation analysis of ODEs or large-dimensional equations (e.g., PDEs, nonlocal equations, etc). The Julia programming language gives access to a rich ecosystem (PDE discretizers, GPUs, automatic differentiation) which serves as a basis for designing

a performant code for numerical bifurcation. The package incorporates several continuation algorithms (pseudo-arclength, Moore-Penrose, asymptotic numerical method, deflated continuation) that can be used to perform fully automated computation of bifurcation diagrams of stationary states. The use of the Julia programming language has been essential to provide access to the huge variety of situations where the user can select the continuation method, the linear / eigen solver with little code change, mainly by relying on multiple dispatch which allows to write generic algorithms. This paradigm is also applied to the study of codimension 2 bifurcations and the computation of periodic orbits with many different methods (trapezoid, orthogonal collocation, parallel Standard / Poincar shooting). Additionally, we use automatic testing (> 85% code coverage) against analytics and the docs / tutorials are automatically generated for removing errors between code versions. We will provide applications highlighting the versatility of the package to the computation of waves and their bifurcations.

Romain J. Veltz
INRIA Sophia-Antipolis
romain.veltz@inria.fr

MS68

Decision Theoretical Uncertainty Quantification after Seeing the Data: Worst-Case Bayesian Priors and the Accuracy-Robustness Trade-off

The need for robust UQ has lead to the development of uncertainty quantification techniques that identify the worst case Bayesian prior by solving two player adversarial games. While these games were usually considered before seeing the data and were computationally intractable, we present a new formulation that solves the game after seeing the data. That way, (1) the game is simplified and efficient solutions are available and (2) the notion of risk is not averaged over all possible datasets. Additionally, the method comes with a parameter in $[0,1]$ that is able to optimally navigate the accuracy-tradeoff of UQ methods. Examples of applications of the method in the physical sciences will be presented

Pau Batlle
Caltech
pbatlle@caltech.edu

MS68

Designing Reliable AI-Based Critical Systems

Increasingly, modern domains like automotive, railway or aerospace depend on more and more number of critical functionalities in Embedded Critical Systems which rely on Deep Learning (DL) models, such as ground and on-board operations in avionics and decision-making functions in autonomous automotive systems. However, nowadays there are significant obstacles that must be addressed before Deep Learning solutions can be seamlessly adopted in Embedded Critical Systems. In this talk, we will go into the details of how to create a system based on DL components that is reliable; by modelling all types of the surrounding uncertainty.

Axel Brando
Barcelona Supercomputing Center

axel.brand@bsc.es

MS68

Hierarchical Ensemble Kalman Methods with Sparsity-Promoting Generalized Gamma Hyperpriors

In this talk, we introduce a computational framework to incorporate flexible regularization techniques in ensemble Kalman methods for nonlinear inverse problems. The proposed methodology approximates the maximum a posteriori (MAP) estimate of a hierarchical Bayesian model characterized by a conditionally Gaussian prior and generalized gamma hyperpriors. Suitable choices of hyperparameters yield sparsity-promoting regularization. We propose an iterative algorithm for MAP estimation, which alternates between updating the unknown with an ensemble Kalman method and updating the hyperparameters in the regularization to promote sparsity. The effectiveness of our methodology is demonstrated in several computed examples, including compressed sensing and subsurface flow inverse problems.

Hwanwoo Kim
University of Chicago
hwkim@uchicago.edu

MS68

Lowering The Computational Barrier : Selective Bayesian Uncertainty Quantification for Transparency in Medical Imaging AI

Deep Learning algorithms, including Convolutional Neural Networks (CNN), provide clinicians with precise insights based on patient data for high-stakes medical decision-making such as brain tumor segmentation. However, these algorithms are not transparent about the uncertainty in their predictions, giving clinicians a false sense of reliability which may lead to grave consequences for patient care. Growing calls for Transparent and Responsible AI have promoted Uncertainty Quantification (UQ) algorithms that help communicate confidence in the algorithmic predictions. However, traditional Bayesian UQ methods remain prohibitively costly for large, million-dimensional tumor segmentation CNNs such as U-Net. In this talk, we discuss a computationally efficient UQ approach via the partially Bayesian neural networks (pBNN). In pBNN, a single strategically chosen layer is used for targeted Bayesian inference while the rest of the network is trained with less-expensive deterministic methods. Sensitivity Analysis is employed to select the layer for Bayesian inference. We illustrate the benefits of our approach, including computational efficiency, and how practitioners and model developers can use this approach to understand the models uncertainty with lowered computational resources.

Snehal Prabhudesai, Jeremiah Hauth
University of Michigan
snehalbp@umich.edu, hauthj@umich.edu

Dingkun Guo
Robotics
Carnegie Mellon University
dingkung@andrew.cmu.edu

Arvind Rao, Nikola Banovic, Xun Huan
University of Michigan
ukarvind@med.umich.edu, nbanovic@umich.edu,

xhuan@umich.edu

MS69

Hyper-Differential Sensitivity Analysis with Respect to Model Discrepancy

Scientific machine learning has opened new avenues of research to enable efficient outer loop analysis through learned models of physical systems. However, these learned models are imperfect representations of complex physical processes. The discrepancy between such models and the underlying truth may be amplified by outer loop analysis such as optimization. We present a novel approach to compute the sensitivity of optimization problems with respect to model discrepancy and use this information to improve the solution obtained using learned models. By posing a Bayesian inverse problem to calibrate the discrepancy, we compute a posterior discrepancy distribution and then propagate it through post-optimality sensitivities to compute a posterior distribution on the optimal solution.

Joseph Hart, Bart G. Van Bloemen Waanders
Sandia National Laboratories
joshart@sandia.gov, bartv@sandia.gov

MS69

Learning High-Dimensional Parameters from Sparse Observations in Tsunami Forecasting

To improve tsunami preparedness, early-alert systems and real-time monitoring are essential. We use a novel approach for predictive tsunami modeling within the Bayesian inversion framework. This effort focuses on informing the immediate response to an occurring tsunami event using near-field data observation. Our forward model is based on a coupled acoustic-gravity model (e.g., Lotto and Dunham, *Comput Geosci* (2015) 19:327340). Similar to other tsunami models, our forward model relies on transient boundary data describing the location and magnitude of the seafloor deformation. In a real-time scenario, these parameter fields must be inferred from a variety of measurements, including observations from pressure gauges mounted on the seafloor. One particular difficulty of this inference problem lies in the accurate inversion from sparse pressure data recorded in the near-field where strong hydroacoustic waves propagate in the compressible ocean; these acoustic waves complicate the task of estimating the hydrostatic pressure changes related to the forming surface gravity wave. Furthermore, the forward model incurs a high computational complexity, since the pressure waves must be resolved in the 3D compressible ocean over a sufficiently long time span. Due to the infeasibility of rapidly solving the corresponding inverse problem for the fully discretized space-time operator, we explore options for using surrogate operators of the parameter-to-observable map.

Stefan Henneking
Oden Institute, UT Austin
stefan@oden.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS69

New Multifidelity Sobol Estimators for Large-Scale Multidisciplinary Systems: Application to the JW

Space Telescope

Sobol global sensitivity indices quantify how uncertainty in model inputs contributes to uncertainty in the model output. Such sensitivity indices allow inputs to be ranked in importance and are typically computed using Monte Carlo estimation, but the many samples required for Monte Carlo to be sufficiently accurate can make this computation intractable when the model is expensive. This talk will present a multifidelity Monte Carlo approach to estimating Sobol indices that combines samples from both cheaper lower-fidelity models (e.g., models learned from data) and expensive high-fidelity models to achieve computational acceleration with accuracy guarantees. We present new multifidelity Sobol index estimators based on rank statistics that can estimate Sobol indices for all inputs from a single set of independent samples. This significantly reduces the cost of Sobol analysis when the number of inputs is large. The approach accelerates the computation time required for sensitivity analysis of the JW Space Telescope thermal models from more than two months to less than two days. This demonstrates the power of the multifidelity approach to make sensitivity analysis tractable for large-scale engineering systems in the real world.

Elizabeth Qian
Georgia Institute of Technology
eqian@gatech.edu

MS69

Reinforcement Learning for Control of Spatially Developing Distributed Systems

Even though classical optimal control tools such as linear-quadratic ones remain widely used owing to a theoretically well grounded framework, they are primarily designed for deterministic evolutions, relying on the modelling of the dynamics, and often exhibiting low robustness with respect to noise estimation. The use of reinforcement learning techniques relying solely on input-output data has gained popularity in recent years with applications ranging from robotics to portfolio management, control of chaos and stabilization of fluid flows. However, the lack of uncertainty estimates and frequent failures to converge to an optimal control prevent its widespread use in practice. Our work focuses on leveraging optimal control theory concepts for Markov decision processes in view of analysis and design of deep reinforcement learning algorithms for the control of spatially distributed systems in an inherently stochastic setting. We use the complex Ginzburg-Landau equation as a test case for the stabilization of absolute instability and the control of local convective instability from noisy measurements. Limitations of the framework arising from the lack of observability and controllability and their effect on the convergence to efficient control strategies are discussed.

Amine Saibi, Onofrio Semeraro
LISN-CNRS, Univ. Paris-Saclay, France
amine.saibi@universite-paris-saclay.fr,
onofrio.semeraro@universite-paris-saclay.fr

Lionel Mathelin
LISN, CNRS, Univ. Paris-Saclay
lionel.mathelin@lisn.upsaclay.fr

MS69

Data-Driven Control Strategies for PDE Environ-

ments Using Reinforcement Learning

The process of identifying optimal control policies plays a key role in a wide range of scientific applications. While many complex systems can be modeled effectively with numerical methods, it is often difficult to determine the best strategy for interacting with such systems based on the forward model alone. In this poster, we explore the potential for leveraging deep reinforcement learning algorithms to automate the decision-making process in the context of scientific applications. In particular, we show that a proximal policy optimization algorithm can successfully identify a near-optimal strategy for positioning sinks within a simplified PDE model of contaminant flow. The generality of the underlying learning scheme, along with the promising results from initial experiments, suggest that reinforcement learning may provide an effective framework for optimizing the decision-making process in a wide range of scientific applications.

Nickolas Winovich

Sandia National Laboratory, U.S.
nwinovi@sandia.gov

Deepanshu Verma
Emory University
dverma4@emory.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS70

Implementing a SYCL Backend for Kokkos

Aurora, to be deployed at Argonne National Laboratory, is expected to be the next exascale computer in the United States. Just as with Frontier, the first exascale computer at Oak Ridge National Laboratory, Aurora's main performance will come from GPUs by a new vendor. Kokkos' (main) choice for supporting these GPUs developed by Intel is leveraging DPC++, Intel's implementation of the SYCL standard. In this talk, I discuss my experience while developing a SYCL backend for Kokkos and issues applications using this backend encounter. In particular, I will compare features that work better or worse in Kokkos than with the HIP or CUDA backend.

Daniel Arndt

Oak Ridge National Laboratory
arndtd@ornl.gov

MS70

A Summary of the Parallel Implementation Processes of Milc-Dslash Benchmark Using SYCL

Milc-Dslash is a benchmark derived from the Milc code which simulates four-dimension SU(3) lattice-gauge theory (or lattice quantum chromodynamics). The sequential version of Milc-Dslash basically consists of the execution of five nested `for` loop structures in order to calculate multiple matrix-vector product operations, whose entries are complex numbers of double-precision. The parallel Milc-Dslash implementation explores the loop-level parallelism

with and without dependencies between iterations. In this talk, we will describe the parallel implementation processes of Milc-Dslash using the SYCL programming model. We will show whether, and if so how, different parallel strategies and index orders can impact the SYCL performance of Milc-Dslash on the NVIDIA A100 GPU as well as its dependence on work-group size. Some examples of parallel strategies include: using atomic operations, local variable, synchronization barrier, unified shared memory with explicit and implicit data movement, buffers and accessors, and SYCL complex library. A comparative performance analysis between SYCL and CUDA programming models will also be provided for some cases, in addition to highlighting their main differences in code writing.

Amanda Dufek

NERSC
asdufek@lbl.gov

Brandon Cook
Lawrence Berkeley National Laboratory
bgcook@lbl.gov

Jack Deslippe
National Energy Research Scientific Computing Center
jrdeslippe@lbl.gov

MS70

Intel Developer Tools for Serious SYCL

Intel with its oneAPI developer suite brings in compatibility and ease of use on the table. In this session we shall explore how Intel is collaborating with open source communities from compilers to productivity tools, promoting SYCL to standardize interfaces over multiple HW vendors and vision for parallel programming.

Nisha Patel

Intel
p.nisha@gmail.com

MS70

Exploration of Performance-Portability in the EXCALIBUR Fusion Use Case

Project NEPTUNE approaches one of the grand challenge problems in fusion - modelling plasma in the edge region of a tokamak. In this edge region, hot plasma interfaces with the reactor wall and cold neutral gas resulting in a computationally expensive multi-scale problem that is considered an exascale challenge. For the fusion community it is highly desirable to implement the computational methods that simulate this multi-scale model in a performance-portable manner. The NESO (Neptune Exploratory Software) implementation is an intermediate project to evaluate the suitability and performance of proposed techniques and algorithms for NEPTUNE. As part of NESO we are evaluating SYCL as a performance-portable language to target multiple hardware architectures. In this talk we present how SYCL is used in the NESO project and our proposed approach for sustaining performance across different hardware architectures in the fusion use case.

Will R. Saunders

UKAEA

will.saunders@ukaea.uk

MS71

Learning High Order Taylor Series Surrogate Models via Randomized Hierarchical Tensor Compression

The most popular surrogate models for large-scale Bayesian inverse problems are based on local linearization and low rank approximation of the Jacobian. However, often non-linear effects are important. To capture non-linear effects, we present a new randomized method for approximating high order derivative tensors with hierarchical tensor networks, which we use to form high order Taylor series surrogate models. We demonstrate that the method can efficiently construct accurate high order Taylor series surrogate models at low cost. We use the resulting Taylor series surrogate models to accelerate MCMC sampling.

Nicholas Alger

The University of Texas at Austin
Oden Institute
nalger225@gmail.com

Blake Christerson

University of Texas at Austin
bechristerson@utexas.edu

Peng Chen, Omar Ghattas

The University of Texas at Austin
peng@cc.gatech.edu, omar@oden.utexas.edu

MS71

Machine Learning Surrogates for Time Dependent Fuel Degradation Processes in Nuclear Waste Repository Simulations

Performance assessment of a subsurface nuclear waste repository requires tracking the evolution and degradation of the spent uranium in thousands of waste packages over several hundred thousand years. While detailed process models for the uranium degradation in these waste packages are available, they are too computationally expensive to be used in a comprehensive full-repository simulation. In this work, we develop Machine Learning surrogate models for the fuel degradation process models and study their accuracy compared to the detailed model as a function of how much information from the internal fuel cask state is incorporated. Both k-Nearest Neighbor regression and Artificial Neural Networks are considered. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525

Bert J. Debusschere

Chemistry, Combustion and Materials Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

Caitlin Curry, Jacob Harvey

Sandia National Labs
cjcurry@sandia.gov, jharve@sandia.gov

Tom Seidl

Sandia National Lab
dtseidl@sandia.gov

Paul Mariner

Sandia National Laboratory

pmarine@sandia.gov

MS71

Bayesian Data-Driven Discovery of Coordinates, Governing Equations and Fundamental Constants

With the recent development of data science, there is a demanding need to transform rich data into well-characterized quantitative descriptions. In the era of big data, snapshots of a dynamical system are frequently captured with camera sensors in a high-dimensional space. Even if there is a high need for understanding high-dimensional video data, the study on data-driven discovery for coordinates and governing equations is still very limited on this data type due to its requirement in computation, difficulty in dimensionality reduction, and sensitivity to noisy observations. In this work, we propose a Bayesian sparse identification of nonlinear dynamics autoencoder (Bayesian SINDy autoencoder) to discover a coordinate transformation into a reduced space with a sparsely represented dynamical system. Bayesian SINDy autoencoder could not only enable the joint discovery of coordinates and governing equations with small sample sizes under noisy environments but could also quantify the estimation uncertainty for the equation discovery and trajectory prediction. We demonstrate the effectiveness of the Bayesian SINDy autoencoder with both synthetic data and real high-dimensional datasets with accelerated sparse inference. Especially in the real video dataset, Bayesian SINDy autoencoder performs GoPro physics discovery which correctly identifies the governing equation and provides an estimate for standard gravity constant g with only 390 video snapshots.

Mars L. Gao

University of Washington
marsgao@uw.edu

J. Nathan Kutz

University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS71

Development of Surrogate Models for Uncertainty Quantification in Manufacturing Engineering

Computer simulations are a widely used tool in computational science and engineering to, e.g., analyze the behavior of components or materials, to enhance manufacturing processes with fast and accurate a-priori forecasts, or even to control those processes during ongoing operation. With the help of these tools, we wish to make reliable assertions and predictions for one or more quantity of interests (QoI), also in the presence of uncertainty, e.g., in process conditions, material properties or similar. Thus, methods from the field of Uncertainty Quantification (UQ) can enhance the quality of processes and products by augmenting the results for the QoI with quantified probability measures. We consider sampling-based UQ methods that usually require a great number of model evaluations. Thus, employing high-fidelity models may easily exceed available resources. Here, the usage of surrogate models, which are computationally cheaper, can provide a remedy. Therefore, we first follow the path of Model Order Reduction (MOR). In particular, we explore the benefits of intrusive MOR techniques and use Proper Orthogonal Decomposition (POD) with a subsequent Galerkin projection of the operators onto the constructed subspaces. As an alternative, we also investi-

gate the advantages of Gaussian Process Regression (GPR) as a meta model. Finally, the integration of the resulting surrogate models into an UQ setting is demonstrated for applications coming from polymer processing.

Fabian Key, Stefanie Elgeti
Institute of Lightweight Design and Structural Biomechanics
TU Wien
fabian.key@tuwien.ac.at, stefanie.elgeti@tuwien.ac.at

MS71

Exploring Context-Aware Learning of Low-Fidelity Models for Multi-Fidelity Monte Carlo Sampling in Plasma Micro-Turbulence Analysis

Multi-fidelity Monte Carlo methods leverage hierarchies of high- and low-fidelity models for speeding up the computation of quantities of interest depending on high-fidelity models. However, the low-fidelity models are assumed to be fixed and therefore cannot be refined or improved, which is relevant, e.g. for data-driven models. To this end, the context-aware multi-fidelity Monte Carlo algorithm optimally balances the costs of training low-fidelity models with the sampling costs, which means that it takes into account the context in which the low-fidelity models will be used, i.e., for variance reduction in Monte Carlo estimation. In the original formulation of the algorithm, all low-fidelity models depend on all uncertain inputs. In this work, we use the fact that in many applications only a subset of the inputs is important for the underlying problem, and explore the potential of reduced-dimension low-fidelity models for speeding up computations further. With sensitivity information obtained from a first fully-dimensional context-aware low-fidelity model, we also construct reduced-dimension context-aware low-fidelity models, which only depend on subsets of selected, important uncertain inputs. We explore this idea in the highly relevant context of quantifying uncertainty in small-scale fluctuations in confined fusion plasmas, which is a problem typically characterized by a large number of uncertain inputs and is therefore challenging to address with conventional methods.

Julia Konrad
Technical University of Munich
julia.konrad@tum.de

Ionut-Gabriel Farcas
The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences
ionut.farcas@austin.utexas.edu

Tobias Neckel
Technical University of Munich
neckel@in.tum.de

Hans-Joachim Bungartz
Technical University of Munich, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

MS72

A Parallel Solver for Nonhydrostatic Ocean Models

The Model for Prediction Across Scales-Ocean (MPAS-Ocean) is an open-source, global ocean model and is one

component of a family of climate models within the MPAS framework, including atmosphere, sea-ice, and land-ice models. In this work, a new formulation for the ocean model is presented that solves the nonhydrostatic, incompressible Boussinesq equations on an unstructured, staggered, z-level grid. The introduction of this nonhydrostatic capability is necessary for the resolution of internal wave dynamics and large eddy simulations. Compared to the standard, hydrostatic formulation, a nonhydrostatic pressure solver and a vertical momentum equation are added, where the PETSc (Portable Extensible Toolkit for Scientific Computation) library is used for the inversion of a large sparse system for the nonhydrostatic pressure. Numerical results on a stratified seiche, internal solitary wave and nonhydrostatic overflow test cases are presented, and the parallel efficiency and accuracy of the time-stepper are evaluated.

Sara Calandrini, Darren Engwirda, Luke Van Roekel
Los Alamos National Laboratory
scalandrini@lanl.gov, dengwirda@lanl.gov, lvan-roekel@lanl.gov

MS72

Solver Improvements for Ice Sheets and Adaptive Mesh Refinement

The BISICLES ice sheet model (Cornford et al, 2013) employs adaptive mesh refinement (AMR) to efficiently resolve relatively fine dynamical features like grounding lines and ice streams. Like most ice sheet models, BISICLES spends the vast majority of its computational effort solving a coupled nonlinear viscous tensor equation for the ice velocity field. Solving these equations in the context of an ice sheet presents specific challenges which must be addressed for our model to be of practical use in understanding and projecting ice sheet response to climate changes. We will discuss strategies we have adopted to improve the efficiency and performance of our solvers, along with recent developments like the implementation of a parallel-in-time scheme for time integration.

Daniel Martin
Lawrence Berkeley National Laboratory
dfmartin@lbl.gov

MS72

Computational Challenges in Designing Robust Solvers for Ice Sheet Modeling

Modeling the dynamics of Greenland and Antarctic ice sheets is of paramount importance for accurately modeling climate and in particular for providing projections of sea level rise. This requires solving optimization problems constrained by partial differential equations for calibrating the models, running the model forward in time for decades or centuries, and performing analysis to quantify the uncertainty in the model projections. All these tasks rely on the repeated solution of nonlinear and linear solvers, which makes it critical to have robust large-scale solvers that can efficiently run on new architectures. In this talk we present the solver strategies we adopt in the ice-sheet code MALI and present result targeting forward and inverse ice-sheet problems.

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

Luca Bertagna, Kim Liegeois, Raymond Tuminaro, Jerry Watkins
 Sandia National Laboratories
 lbtag@sandia.gov, knliege@sandia.gov, rstu-
 min@sandia.gov, jwatkin@sandia.gov

MS72

Augmented Block-Arnoldi Recycling Solvers for Atmospheric Models

We present a one-reduce iterated Gauss-Seidel IGS-GMRES algorithm based on a hybrid MGS-CGS projector. The Krylov vectors are close to orthogonal and linearly independent in finite precision. The Arnoldi relative residual no longer stagnates above machine precision and the smallest singular value of V_k remains close to one. The new Arnoldi-QR algorithm can also be employed to compute eigenvalues with the Krylov-Schur algorithm and improves Krylov exponential integrators. An augmented-Arnoldi formulation is applied to Krylov subspace recycling. A block-Krylov algorithm is employed with an BMGS-LTS inter-ortho skeleton as in Carson et al (2022), and a block triangular solve. When combined with an oblique projection, similar to Embree et al (2017), the recycling algorithm leads to significant reductions in the number of solver iterations per linear system. The weight is also interpreted in terms of the angle between residual norms at restarts in LGMRES, Baker et al. (2006). We demonstrate a 10-fold reduction in the iteration count for warm bubble vertical convection governed by the non-hydrostatic Euler equations and pressure projection for the incompressible Navier-Stokes equations applied to wind-turbines. In many cases, the recycle subspace eigen-spectrum may substitute for a preconditioner.

Stephen Thomas

Advanced Micro Devices (AMD)
 stephethomas@gmail.com

Allison H. Baker

National Center for Atmospheric Research
 abaker@ucar.edu

Stéphane Gaudreault
 Environment Canada

stephane.gaudreault@ec.gc.ca

MS73

A Fully Parallelized and Budgeted Multilevel Monte Carlo (BMLMC) Method Applied to Wave Equations

Multilevel Monte Carlo methods (MLMC) have been successfully applied to many problems by numerous authors. In this talk, we will present a framework which can treat a wide range of PDEs with uncertain input data with various discretizations combined with non-intrusive statistical estimators. In particular, we want to present a new variant of the method - the budgeted (B)MLMC method with integrated load balancing. The resulting algorithm and its implementation is analyzed in combination with discontinuous Galerkin (DG) finite elements and space-time discretizations applied to the acoustic wave equation. Lastly, we will present benchmarked results on a high performance computing (HPC) clusters.

Niklas Baumgarten

Karlsruhe Institute of Technology
 niklas.baumgarten@kit.edu

Christian Wieners

IANM, Karlsruhe Institute of Technology, Germany
 christian.wieners@kit.edu

MS73

Ensemble Kalman Inversion for the Characterisation of Thermal Bridges in Buildings' Walls

I will discuss the application of Ensemble Kalman Inversion (EKI) for the probabilistic characterisation of the thermal performance of buildings' walls, in the presence of unknown thermal bridges which are known to be substantially detrimental to thermal efficiency. The aim of the framework is to infer thermal properties in a 3D heat transfer (forward) model of the wall given in-situ measurements. These thermal properties are parameterised using a level-set function that defines the region of anomalous thermal conductivity (from the potential presence of thermal bridges), as well as random fields that characterise variability in thermal properties within each region of the wall. Parameters are inferred via EKI in a Bayesian setting given data from multiple sources including surface temperatures from thermal camera images and thermocouples as well as surface heat flux measurements from heat flux meters. We show numerical examples with synthetic and real data in order to demonstrate the applicability of the proposed EKI framework in practical settings.

Marco Iglesias

University of Nottingham
 marco.iglesias@nottingham.ac.uk

MS73

Design and Implementation of UQ Software

In this talk I will discuss some of the challenges of applying advanced UQ methods to complex, large-scale forward models. I will also present UM-Bridge, an open source project that addresses some of these challenges. UM-Bridge provides a clean network-based interface to communicate between UQ and modeling software via HTTP. Containerization of models using UM-Bridge is trivial, allowing for unified, portable, fully reproducible and black-box models. As an example I will present a parallelized multi-level Markov chain Monte Carlo method, a state-of-the-art, algorithmically scalable UQ algorithm for Bayesian inverse problems. The integration between the UQ library MUQ and the PDE engine ExaHyPE facilitated by UM-Bridge allows for large-scale parallelism across forward model evaluations and the UQ algorithms themselves. The main scalability challenge presents itself in the form of strong data dependencies introduced by the MLMCMC method, prohibiting trivial parallelization. I will demonstrate the method by using it to infer the most probable locations for the initialisation of a tsunami from buoy data.

Anne Reinartz

Durham University
 anne.k.reinartz@durham.ac.uk

Linus Seelinger

Heidelberg University
 Institute for Applied Mathematics
 mail@linusseelinger.de

Robert Scheichl

Universitaet Heidelberg

r.scheichl@uni-heidelberg.de

MS73

Surrogate-Assisted Data-Free Inference with Summary Statistics for Predicting XE Diffusivity in UO₂ Nuclear Fuel

We present a data-free inference framework to estimate the parameters in an atomistic model for the diffusivity of Xe fission gas in UO₂ nuclear fuel. Our calibration strategy uses synthetic data sets to guarantee consistency between the model predictions and the summary statistics reported in historical gas release and thermodynamic experiments. In order to keep our approach computationally tractable, we replace evaluations of the forward model by evaluations of a prebuilt polynomial chaos surrogate model with a reduced number of parameters. This reduced set of parameters is identified by a two-step dimension reduction process, involving a global sensitivity analysis study as well as the construction of a likelihood-informed subspace. We discuss the performance of our calibration strategy, and consider a weighted construction that accounts for the different number of data points in each experimental data set. Finally, we investigate how the surrogate can be used to speed up the estimation of the parameters of the actual forward model.

Pieterjan Robbe
Sandia National Laboratories
pmrobbe@sandia.gov

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

Tiernan Casey, Khachik Sargsyan, Luc Bonnet
Sandia National Laboratories
tcasey@sandia.gov, ksargsy@sandia.gov, lbonet@sandia.gov

MS73

Combining the Sparse Grids Matlab Kit and UMBridge for Forward Uncertainty Quantification of a Naval Engineering Problem

In this work, we show how to couple the Sparse Grids Matlab Kit with a complex solver for computational fluid dynamics to perform a forward Uncertainty Quantification analysis with the classical sparse grids method. In details, the problem considered is the uncertainty quantification of the calm-water resistance of a destroyer-type vessel (the DTMB 5415 model) subject to uncertainty in the payload and advancing speed. The solver is written in Fortran and implements a free-surface potential flow solver for the water flow around the ship hull, coupled with the rigid-body equation of motions. The coupling between the sparse grids method and the fluid dynamics solver is handled through the UM-Bridge protocol.

Lorenzo Tamellini
Istituto di Matematica Applicata e Tecnologie Informatiche
tamellini@imati.cnr.it

Andrea Serani
Consiglio Nazionale delle Ricerche
Istituto di Ingegneria del Mare (CNR-INM)
andrea.serani@cnr.it

Chiara Piazzola
Consiglio Nazionale delle Ricerche - Istituto di Matematica Applicata e Tecnologie Informatiche (CNR - IMATI)
chiara.piazzola@imati.cnr.it

Matteo Diez
Consiglio Nazionale delle Ricerche
Istituto di Ingegneria del Mare (CNR-INM)
matteo.diez@cnr.it

MS74

Multiword Matrix Multiplication with Tensor Cores

In multiword arithmetic, a matrix is represented as the unevaluated sum of two or more lower-precision matrices, and a matrix product is formed by multiplying the constituents in low precision. We investigate the use of multiword arithmetic for improving the performance-accuracy tradeoff of matrix multiplication with mixed precision block fused multiplyadd (FMA) hardware, focusing especially on the tensor cores available on NVIDIA GPUs. Building on a general block FMA framework, we develop a comprehensive error analysis of multiword matrix multiplication. After confirming the theoretical error bounds experimentally by simulating low precision in software, we use the cuBLAS and CUTLASS libraries to implement a number of matrix multiplication algorithms using double-binary16 arithmetic. When running the algorithms on NVIDIA V100 and A100 GPUs, we find that double-binary16 is not as accurate as binary32 arithmetic despite satisfying the same worst-case error bound. Using probabilistic error analysis, we explain why this issue is likely to be caused by the rounding mode used by the NVIDIA tensor cores, and propose a parameterized blocked summation algorithm that alleviates the problem and significantly improves the performance-accuracy tradeoff.

Massimiliano Fasi
Durham University
massimiliano.fasi@durham.ac.uk

Nicholas J. Higham
School of Mathematics
The University of Manchester
nick.higham@manchester.ac.uk

Florent Lopez
Livermore Software Technology
florent.lopez@ansys.com

Theo Mary
Sorbonne Universite, CNRS, LIP6
theo.mary@lip6.fr

Mantas Mikaitis
The University of Manchester
m.mikaitis@leeds.ac.uk

MS74

Block preconditioners for the Marker and Cell discretization of the Stokes-Darcy equations

We develop block preconditioners for solving the Stokes-Darcy equations, discretized by the Marker and Cell (MAC) finite difference method. The discretization leads to a mildly nonsymmetric double saddle-point linear sys-

tem. We identify numerical properties and exploit the sparsity structure of the matrix, for the purpose of developing a fast preconditioned iterative solution procedure. The proposed preconditioners are based on approximations of two Schur complements that arise in decompositional relations associated with the double saddle-point matrix and its blocks. We analyze the eigenvalue distribution of the preconditioned matrices with respect to the physical parameters of the problem and show that the eigenvalues are strongly clustered. Consequently, preconditioned GMRES appears to be relatively insensitive to the mesh size and the physical parameters involved. Numerical results validate our theoretical observations.

Chen Greif

Department of Computer Science
The University of British Columbia
greif@cs.ubc.ca

Yunhui He

The University of British Columbia
yunhui.he@ubc.ca

MS74

Preconditioners for Multiple Saddle-Point Systems from PDE-Constrained Optimization

Optimization problems subject to PDE constraints form a mathematical tool that can be applied to a wide range of scientific processes, including fluid flow control, medical imaging, biological and chemical processes, and many others. These problems involve minimizing some function arising from a physical objective, while obeying a system of PDEs which describe the process. Of key interest is the numerical solution of the discretized linear systems arising from such problems, and in this talk we focus on preconditioned iterative methods for these systems. In particular, we describe recent advances in the preconditioning of multiple saddle-point systems, specifically positive definite preconditioners which can be applied within MINRES, which may find considerable utility for solving these optimization problems as well as other applications. In particular, these preconditioners lead to a guaranteed convergence rate for MINRES, and often demonstrate superior convergence as opposed to widely-used block diagonal preconditioners. Further, we briefly discuss an inexact active-set method for large-scale nonlinear PDE-constrained optimization problems, coupled with block diagonal and block triangular preconditioners for multiple saddle-point systems which utilize suitable approximations for the relevant Schur complements. This is joint work with Andreas Potschka (TU Clausthal).

John W. Pearson

University of Edinburgh
j.pearson@ed.ac.uk

MS74

Parallel Efficiency of Coarsest Grid Solvers in Multigrid

Multigrid methods find their way into many applications in scientific computing. Simulations in theoretical studies of the inner workings of nuclei, in particular, are bound by the use of supercomputers to speed them up. Multigrid methods have brought a way of pushing the computational boundaries in such large-scale simulations, and they open the possibility of scalably simulating at the exascale. Unfortunately, under certain extreme situations such as the

use of many processes and/or very high condition of the linear systems to be solved, the time spent at the coarsest grid ends up representing most of the execution time in those multigrid solves. If the solver employed at the coarsest level is e.g. GMRES, in which case we see the appearance of many dot products, scalability is at risk. We discuss here different ways in which we can improve the scalability of multigrid solvers by focusing on the coarsest level. All of our implementations and tests are performed within our solver library DD-alphaAMG, a solver for large and sparse matrices emerging in lattice quantum chromodynamics, but these techniques are useful in any application where multigrid fails to scale due to coarsest-level restrictions.

Gustavo Ramirez-Hidalgo

Applied Mathematics
University of Wuppertal
g.ramirez@math.uni-wuppertal.de

MS74

Structure-Exploiting Preconditioners for Data Assimilation

Correlation operators are used in data assimilation algorithms to weight the contribution of prior and observation information. Efficient implementation of these operators is therefore crucial for operational implementations. Diffusion-based correlation operators are popular in ocean data assimilation, but can require a large number of serial matrix-vector product. A parallel-in-time formulation removes this requirement, and offers the opportunity to exploit modern computer architectures. High quality preconditioners for the parallel-in-time approach are well-known, but impossible to apply in practice for the high-dimensional problems that occur in oceanography. In this talk we consider a nested preconditioning approach which retains many of the beneficial properties of the ideal analytic preconditioner while remaining affordable in terms of memory and computational resource.

Jemima M. Tabcart

Mathematical Institute
University of Oxford
jemima.tabcart@maths.ox.ac.uk

Selime Gurol

CERFACS, France
selime.gurol@cerfacs.fr

John Pearson

University of Edinburgh
john.pearson@ed.ac.uk

Anthony Weaver

CERFACS
anthony.weaver@cerfacs.fr

MS75

Asymptotic Formulation of Droplet Pinch-Off

Droplets are seen in applications like emulsion, spraying, ink-jet printing, atomization and entrainment, and many more. Understanding the droplet dynamics is crucial to improve the efficacy of these processes. Hybrid rockets present a promising alternative to purely liquid-fueled engines as common hybrid fuels are environmentally friendly, abundant, and easy to process and store. Paraffin wax is a prominent candidate among high regression rate fuels for

hybrid rocket engines [M. Karabeyoglu and B. J. Cantwell, Combustion of liquefying hybrid propellants: Part 2, stability of liquid films, *Journal of Propulsion and Power*, vol. 18, no. 3, pp. 621630, 2002]. The atomization of the paraffin wax that begins with the droplet formation and pinch-off, enables rapid burning and generates a much more specific thrust than other fuels. We present a one-dimensional model to simulate the droplet pinch-off in a quiescent environment using front tracking [D. K. Nathawani and M. G. Knepley, Droplet formation simulation using mixed finite elements, *Physics of Fluids*, vol. 34, no. 6, p. 064105, 2022]. The model is verified using the method of manufactured solution and validated against laboratory experiments. We propose a self-consistent algorithm with adaptively refined mesh. Furthermore, we expand the model for droplets in a continuous flow to incorporate the effects of shear force on droplet evolution.

Darsh Nathawani
University at Buffalo
darshkir@buffalo.edu

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

MS75

Numerical Simulation and Analysis of Multiscale Interface Coupling Between a Poroelastic Medium and a Lumped Hydraulic Circuit: Comparison Between Functional Iteration and Operator Splitting Methods

We consider a multiscale problem modeling the flow of a fluid through a deformable porous medium, described by a system of partial differential equations (PDEs), connected with a lumped hydraulic circuit, described by a system of ordinary differential equations (ODEs). This PDE/ODE coupled problem includes interface conditions enforcing the continuity of mass and the balance of stresses across models at different scales. In the present work, we address questions related to the solution methods of the PDE/ODE coupled problem via staggered algorithms, focusing on a detailed comparison between functional iterations and an energy-based operator splitting method and how they handle the interface conditions. We provide sufficient conditions for the convergence of functional iterations and prove that the energy-based operator splitting method is unconditionally stable with respect to the size of the time discretization step.

Daniele Prada
CNR
prada@imati.cnr.it

Lorena Bociu
NC State University
Raleigh, NC
lvbociu@ncsu.edu

Giovanna Guidoboni
Department of Electrical Engineering and Computer Science,
University of Missouri, Columbia, MO, USA
guidobonig@missouri.edu

Riccardo Sacco
Politecnico di Milano

Dipartimento di Matematica
riccardo.sacco@polimi.it

MS75

Keynote: Physical Insight + Mathematical Modeling + Numerical Methods = The Winning Synergy for the 21st Century Computational Science

The conjunction among climate change and overwhelming socio-economic urgencies is pushing human relationships to unprecedented levels of complexity, demanding for a new vision of the future. Mathematics can help cope with such a challenge by devising reliable and user-friendly tools to interpret phenomena, predict multiple scenarios and suggest optimal solutions. The success of this endeavor must rely upon the combination of sound physically-based models with robust and efficient computational techniques, the final aim being to produce an harmonized representation of mechanistic behavior and data uncertainty. In this lecture, the above vision of Mathematics will be applied to the study of Ophthalmology and Electronics, two branches of Science that, unexpectedly, share analogous microscopic functions and a wide range of spatial and temporal scales. Simulations of realistic problems will be illustrated to demonstrate the decisive importance of a cooperative synergy among physics, models and methods in the development of the 21st Century Computational Science.

Riccardo Sacco
Politecnico di Milano
Dipartimento di Matematica
riccardo.sacco@polimi.it

MS76

One Shot Learning of Stochastic Differential Equations with Gaussian Processes

We consider the problem of learning the drift and diffusion functions of a Stochastic Differential Equation of the form $dX_t = f(X_t)dt + \sigma(X_t)dW_t$ from one sample trajectory. This problem is challenging because one sample trajectory only provides indirect information on the unknown functions f , σ , and stochastic process dW_t representing the drift, the diffusion, and the stochastic forcing terms, respectively. In this talk we present a simple kernel-based solution that decomposes the problem as follows: 1) Express the increments of the SDE as a computational graph. 2) Recover the unknown functions through a Maximum a Posteriori Estimation (MAP) estimator with Gaussian Process (GP) priors. 3) Optimize the kernels of the GPs using randomized cross validation. We illustrate the efficacy, robustness, and scope of our method through numerical examples. Co-authors: Boumediene Hamzi, Giulia Livieri, Houman Owhadi, Peyman Tavallali.

Matthieu Darcy
Caltech
mdarcy@caltech.edu

Houman Owhadi
California Institute of Technology
owhadi@caltech.edu

Boumediene Hamzi
Johns Hopkins University
boumediene.hamzi@gmail.com

Giulia Livieri
Scuola Normale Superiore

giulia.livieri@sns.it

Peyman Tavallali
Jet Propulsion Lab, NASA
tavallali@gmail.com

MS76

On Uncertainty Quantification of Eigenpairs with Higher Multiplicity

We consider generalized variational eigenvalue problems with random perturbations in the bilinear forms. This setting is motivated by Galerkin discretizations of the Helmholtz equation or Maxwell's equations with random material laws, for example. The considered eigenpairs can be of higher but finite multiplicity. We investigate stochastic quantities of interest of the eigenspaces and discuss why, for multiplicity greater than 1, only the stochastic properties of the eigenspaces are meaningful, but not of individual eigenpairs. To that end, we characterize the Frchet derivatives of the eigenpairs with respect to the perturbation and provide a new linear characterization for eigenpairs of higher multiplicity. Based on the Frchet derivatives of the eigenpair we discuss a meaningful sampling strategy for multiple eigenvalues and develop an uncertainty quantification perturbation approach. In the latter, the arising tensor equations for the covariance can be efficiently solved by a low-rank method, although alternate approaches such as sparse grids are also feasible. Finally, we discuss performance gains compared to sampling based methods.

David Ebert, Jürgen Dölz
University of Bonn
ebert@ins.uni-bonn.de, doelz@ins.uni-bonn.de

MS76

Auto-Differentiable Ensemble Kalman Filters

Data assimilation is concerned with sequentially estimating a temporally-evolving state. This task, which arises in a wide range of scientific and engineering applications, is particularly challenging when the state is high-dimensional and the state-space dynamics are unknown. In this talk I will introduce a machine learning framework for learning dynamical systems in data assimilation. Our auto-differentiable ensemble Kalman filters (AD-EnKFs) blend ensemble Kalman filters for state recovery with machine learning tools for learning the dynamics. In doing so, AD-EnKFs leverage the ability of ensemble Kalman filters to scale to high-dimensional states and the power of automatic differentiation to train high-dimensional surrogate models for the dynamics. Numerical results using the Lorenz-96 model show that AD-EnKFs outperform existing methods that use expectation-maximization or particle filters to merge data assimilation and machine learning. In addition, AD-EnKFs are easy to implement and require minimal tuning.

Yuming Chen, Daniel Sanz-Alonso, Rebecca Willett
University of Chicago
ymchen@uchicago.edu, sanzalonso@uchicago.edu, willett@uchicago.edu

MS76

Parameter Estimation for the McKean-Vlasov Stochastic Differential Equation

McKean-Vlasov SDEs arise naturally in many applications

as the mean field limit of systems of interacting particles. They are important in many fields, from mathematical biology (e.g., neuroscience and population dynamics) to the social sciences (e.g., opinion dynamics and cooperative behaviours), to high-dimensional Bayesian inference. In this talk, we will discuss online parameter estimation for a McKean-Vlasov SDE and the associated system of weakly interacting particles. We propose a new online estimator, which evolves according to a continuous-time stochastic gradient descent algorithm on the asymptotic log-likelihood of the interacting particle system. We obtain various convergence results for this estimator, under assumptions which guarantee ergodicity and uniform-in-time propagation chaos. Our theoretical results are supported via several numerical examples, including a toy linear mean field model, a stochastic Kuramoto model, and a stochastic opinion dynamics model.

Louis Sharrock
Lancaster University
l.sharrock@lancaster.ac.uk

Nikolas Kantas, Panos Parpas, Greg Pavliotis
Imperial College London
n.kantas@imperial.ac.uk, panos.parpas@imperial.ac.uk,
g.pavliotis@imperial.ac.uk

MS76

Structure Preserving Neural Networks for Identifying and Solving Dynamical Systems

We propose the symplectic neural network (SympNet) that are composed of unit lower/upper triangular matrix-like maps for solving and identifying conservative dynamical systems from data. We theoretically proved that SympNets are universal approximators within the class of symplectic maps. We apply SympNet to learn the evolution map from the molecular dynamics (MD) simulation results across multiple steps, which reduces the computational cost of traditional MD solver by 20% on GPU while being able to keeping energy constant in the NVE ensemble. We apply SympNet to solve the high dimensional path planning problems. We experimentally showed that our neural-network-based method SympNets can overcome the curse of dimensionality by solving the path planning problem of 256 closely interacted drones in 1.5 hours. As a generalization of SympNet, we proposed the GENERIC formalism informed neural network (GFINN), which exactly satisfies the first and second law of thermodynamics, to facilitate prediction of the evolution of complex processes that are dissipative in nature. The architecture is inspired by the spectral decomposition of skew-symmetric and positive definite matrices. Based on rigorous approximation theory, we prove that our model approximates any system that can be described by the GENERIC formalism in a thermodynamics-consistent way. We apply GFINN to predict the motion of thermoelastic double-pendulum governed by Gough-Joule effect.

Zhen Zhang
Brown University
zhen.zhang1@brown.edu

MS77

Heterogeneity and HPC: The Messy Price for Performance

HPC platforms continue to embrace diverse and novel hardware, all the while requiring that core software stacks

remain functional, e.g., an implementation of the Message Passing Interface (MPI) or BLAS and LAPACK. This talk discusses some of the challenges that continued heterogeneity introduces both as a technical challenge for system administration and as an additional layer of complexity that users must become aware of. Unfortunately, this complexity spans all levels of the system, from administration to research and development and on to the user experience. In this talk, we will give some practical instances where challenges arise as well as discuss how co-design can be used to mitigate some of these risks and challenges.

James Elliot, Clayton Hughes
Sandia National Laboratories
jjellio@sandia.gov, chughes@sandia.gov

MS77

Accelerated Computing Using FleCSI

The Flexible Computational Science Infrastructure (FleCSI) framework is a compile time configurable runtime library that is designed to support multi-physics applications development. FleCSI's abstraction layer provides a single-source programming interface for shared-memory and distributed-memory parallelism through task and kernel execution, respectively. The shared memory parallelism interface in FleCSI supports portability across heterogeneous computing architectures, utilizing the underlying accelerator runtimes OpenMP, CUDA, and HIP. In this talk, we present the performance results for two FleCSI-based applications: MPAS (shallow water core), and a basic iterative solver for elliptic PDEs. Results were obtained on heterogeneous computing architectures that demonstrate the performance, portability, and productivity capabilities of FleCSI.

Sumathi Lakshmiranganatha, Andrew Reisner
Los Alamos National Laboratory
sumathil@lanl.gov, areisner@lanl.gov

MS77

Supporting Applications on Emerging Heterogeneous Supercomputers

The introduction of heterogeneous computing via GPUs represented a significant shift in direction for HPC centers, and therefore required significant preparation. Science applications faced many challenges and required rethinking and re-factoring of their algorithms, system software, and tools. Furthermore, partnerships with system vendors were necessary to co-design architectural features with software and applications. In this talk, we present key lessons learned from a concentrated effort at Lawrence Livermore National Laboratory to prepare applications, system software, and tools for existing and upcoming heterogeneous supercomputers. We share the process we applied at the laboratory with the hope that others will be able to learn from both our successes and intermediate setbacks. We describe best practices for algorithms and source code, system configuration and software stack, tools, and application performance. Finally, we present early results as we prepare for the arrival of our exascale supercomputer El Capitan.

David F. Richards
Lawrence Livermore National Laboratory

richards12@llnl.gov

MS77

Enhancing Productivity on Heterogeneous Supercomputers with Task-Based Programming Model

Heterogeneous supercomputers with GPUs are one of the best candidates to build Exascale machines. However, porting scientific applications with huge number of code lines is challenging. Data transfers/locality and exposing enough parallelism determine the maximum achievable performance on such systems. Porting efforts impose developers to rewrite parts of the application which is tedious and time-consuming and does not guarantee performances in all the cases. Being able to detect which parts can be expected to deliver performance gains on GPUs is therefore a major asset for developers. Moreover, task parallel programming model is a promising alternative to expose enough parallelism while allowing asynchronous execution between CPU and GPU. OpenMP 4.5 introduces the target directive to offload computation on GPU in a portable way. Target constructions are considered as explicit OpenMP task as for CPU but executed on GPU. In this work, we propose a methodology to detect the most profitable loops of an application that can be ported on GPU. While we have applied the detection part on several mini applications (LULESH, miniFE, XSBench and Quicksilver), we experimented the full methodology on LULESH through MPI+OpenMP task programming model with target directives. It relies on runtime modifications to overlap of data transfers and kernel execution through tasks. This work has been integrated into the MPC framework, and has been validated on distributed heterogeneous system.

Adrien Roussel

CEA
The French Alternative Energies and Atomic Energy
Commission
adrien.roussel@cea.fr

Romain Pereira, Mickael Boichot, Manuel Ferat
CEA, DAM, DIF
romain.pereira@cea.fr, mickael.boichot@cea.fr,
manuel.ferat@cea.fr

MS79

Disentangling Nuisance Factors During Geophysical Inversion Using Symmetric Autoencoders

We present an autoencoder architecture that assumes scale separation to represent the measurements of a given physical system using two separate latent codes: one correlated with the information that varies on a slow scale and another one on a much faster scale. The key idea is that, as opposed to using the underlying physical model, the architecture uses symmetry and stochastic regularization to achieve a data-driven disentanglement. We show that this disentanglement is invaluable when analyzing geophysical systems. Here, the variations in the measurements occur across two separate scales attached to different physical phenomena, with one of the phenomena being task-irrelevant and therefore resulting in nuisance parameters. Some examples include passive time-lapse seismic monitoring using measurements corrupted with nuisance variations due to uncontrollable sources; learning earthquake signatures in the presence of nuisances attached to multiple scattering in the subsurface; leveraging symmetry under reordering the pixels to disentangle the influence of the

nuisance atmosphere effects from hyperspectral imagery.

Pawan Bharadwaj
Department of Mathematics, MIT
pawbz@mit.edu

Matthew Li
Massachusetts Institute of Technology
mtcli@mit.edu

Laurent Demanet
Department of Mathematics, MIT
laurent@math.mit.edu

MS79

Full Waveform Inversion via Reduced Order Modeling

Full waveform inversion (FWI) is a seismic imaging method that aims at using the entire content of seismic traces to extract physical parameters of the medium being probed by seismic waves. Conventional FWI is carried out by nonlinear least-squares data fitting with some appropriate regularization and it can be interpreted as a Maximum A Posteriori in a Bayesian framework. It turns out that the FWI misfit function is high-dimensional and non-convex with many local minima. A novel approach to FWI based on a data driven reduced order model (ROM) of the wave equation operator is introduced and it is shown that the minimization of ROM misfit function performs much better. The talk is based on a joint work with L. Borcea (Univ. Michigan), A. Mamonov (Univ. Houston), J. Zimmerling (Uppsala Univ.).

Josselin Garnier
Ecole Polytechnique
josselin.garnier@polytechnique.edu

Liliana Borcea
University of Michigan
borcea@umich.edu

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

Jörn Zimmerling
University of Michigan
jtzimmerling@gmail.com

MS79

Augmented Pseudo-Marginal Metropolis-Hastings for Partially Observed Diffusion Processes

We consider the problem of parameter inference for partially observed Markov process models using data at discrete times that may be incomplete and subject to measurement error. We integrate over uncertainty in the latent process between observation times via a state-of-the-art correlated pseudo-marginal Metropolis-Hastings algorithm, that aims to improve mixing of the parameter chains by inducing positive correlation between successive estimates of the intractable observed data likelihood. However, unless the measurement error or the dimension of the latent process is small, correlation can be eroded by the resampling steps in the particle filter. We therefore propose a novel augmentation scheme, that allows for conditioning on values of the latent process at the observation

times, completely avoiding the need for resampling steps. We integrate over the uncertainty at the observation times with an additional Gibbs step. We illustrate the resulting methodology in the context of nonlinear multivariate diffusion processes and find that our approach offers substantial increases in overall efficiency, compared to some competing methods.

Andrew Golightly
Durham University
andrew.golightly@durham.ac.uk

Chris Sherlock
Lancaster University
c.sherlock@lancaster.ac.uk

MS79

Geometry of Molecular Conformations in Cryo-EM

Cryo-Electron Microscopy (cryo-EM) is an imaging technology that is revolutionizing structural biology. Cryo-electron microscopes produce many very noisy two-dimensional projection images of individual frozen molecules; unlike related methods, such as computed tomography (CT), the viewing direction of each particle image is unknown. The unknown directions and extreme noise make the determination of the structure of molecules challenging. While other methods for structure determination, such as x-ray crystallography and NMR, measure ensembles of molecules, cryo-electron microscopes produce images of individual particles. Therefore, cryo-EM could potentially be used to study mixtures of conformations of molecules. We will discuss a range of recent methods for analyzing the geometry of molecular conformations using cryo-EM data.

Roy R. Lederman
Yale University
roy.lederman@yale.edu

MS79

Optimal Transport for Elastic Source Inversion

Full waveform inversion (FWI) is a state-of-the-art method for imaging Earth's subsurface. However, FWI is notorious for local-minimum trapping, or "cycle skipping," and thus requires an accurate initial model ([Mtivier et al., Optimal transport for mitigating cycle skipping in full-waveform inversion]). Cycle skipping is caused by the nonconvex nature of the misfit optimization landscape in a least-squares formulation. In contrast, the Wasserstein-2 distance, denoted W_2 , is convex with respect to shifts and dilations, which occur naturally in seismic data ([Engquist et al. Optimal transport for seismic full waveform inversion]). Therefore, we propose using this optimal-transport metric as our misfit for source inversion. Previous work using optimal transport for source inversion has shown promise ([Chen et al., The quadratic Wasserstein distance for earthquake location]). However, this work uses the acoustic wave equation, which is less accurate than an elastic model. In this talk, we extend these results to elastic source inversion in two spatial dimensions and show that they translate well to the elastic model. We will show a comparison of the optimization landscapes with respect to source location for the L^2 misfit and W_2 distance for the IASP-91 model. We then show inversion results using USArray data from the western United States for various earthquake sources,

assuming known Lam parameters.

Tyler Masthay
University of Texas at Austin
tyler@oden.utexas.edu

Björn Engquist
Department of Mathematics
the University of Texas at Austin
engquist@oden.utexas.edu

MS80

Lbm Bloodflow and Endovascular Coiling-Simulations in 3D Aneurysm-Geometries

Cerebral aneurysms pose a big and creeping danger to patients due to their ubiquitous risk of rupture. In order to quantify the rupture-risk as well as to assess different treatment strategies in-silico, detailed numerical simulations of the hemodynamics within the aneurysm as well as the interaction with treatment devices such as coils are necessary. In this talk, we present research and simulation results based on the Lattice-Boltzmann method to simulate the bloodflow within real, 3D and patient specific vessel-geometries obtained from medical scans. Realistic coil geometries are obtained via a one-dimensional spring-wire insertion simulation that takes into account bending, wall-contact, self-intersection as well as a manufacturer-predefined coil micro-structure that allows to fill the aneurysms space to a large extend, eventually occluding it from further bloodflow. In order to lower computational costs, we also aim to replace and compare such, computationally costly, fully resolved coiling-bloodflow-simulations by means of porous medium surrogate models, where parameters such as porosity and permeability are obtained via homogenization from different, realistic coils as mentioned before. It is our goal to eventually provide a simulation tool that will allow medical doctors to judge and select based on quantities of interest such as the wall-shear-stress on the aneurysm walls different treatment methods and devices in advance, based on the specific patients data.

Markus Muhr, Natalia Nebulishvili, Barbara Wohlmuth
Technical University Munich
muhr@cit.tum.de, nebulish@ma.tum.de,
wohlmuth@ma.tum.de

MS80

Solving Problems in Complex Multiscale Geometries for Porous Media, Vegetation and Vascular Applications

We present a strategy for the approximation of multiscale problems which feature a very complex geometry and overlapping continua models. We start with digitized images of realistic media digitized into voxels. The challenges include the coupled nature of the governing PDEs of different types as well as the fact that the geometry may be changing, but it is not feasible to obtain the geometry snapshots in time. The goals are to eventually upscale the solutions. By design, we employ therefore the lowest order mixed finite element methods on rectangular grids, and we embed the domains in realistically simple fictitious domain. This strategy makes tracking the behavior of the entire system feasible while the pre-processing and post-processing efforts are manageable and allow to focus on the deliverables of interest. This strategy is motivated by applications in human tissue, soil-vegetation systems, and multiphase flow

coupled to energy conservation in porous media.

Malgorzata Peszynska
Oregon State University
mpesz@math.oregonstate.edu

Azhar Alhammali
Imam Abdulrahman Bin Faisal University
aalhammali@iau.edu.sa

Tyler Fara
Department of Mathematics, Oregon State University
farat@oregonstate.edu

Choah Shin
Oregon State University
Department of Mathematics
shinc@oregonstate.edu

Nachuan Zhang
Department of Mathematics, Oregon State University
zhangnac@oregonstate.edu

MS80

Numerical Modeling of Permafrost Thermal Dynamics in Arctic Lake-Soil Systems

Permafrost degradation is observed to increase rapidly in Arctic coastal regions due to permafrost warming. Increases in air temperature is the primary direct driver of permafrost warming, but frequent coastal flooding has been recognized as another important driver indirectly contributing to permafrost warming. However, the impact of coastal flooding on permafrost has not been tested and remains poorly understood. Our goals are to 1) understand permafrost hydro-thermal dynamics under floods and 2) develop new capability to simulate arctic lake and permafrost thermal dynamics. We consider a coupled one-dimensional model of snow-lake-soil system with variable water body depth. The temperature distribution in the lake (with or without ice) is calculated by means of a heat diffusion equation in the scaled coordinate system to account for variability in the depth. The soil is a coupled model of heat equation for the temperature and Richards equation for the water content. Finally, we use the surface energy balance equation to model the snow layer. The numerical discretization and coupling of the physical processes are performed using the Amanzi-ATS software infrastructure. For model verification we compare the results of Amanzi-ATS simulations for three different lakes in Alaska with the seasonal observation data for the surface temperature and the depth of the snow and ice layers.

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

David Moulton
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
moulton@lanl.gov

Konstantin Lipnikov
Los Alamos National Laboratory

lipnikov@lanl.gov

MS80

Finite Elements for Thermo-Hydro-Mechanical Coupling in Modeling Permafrost Thaw

We present an algorithm and its analysis for the finite element approximation to thermo-hydro-mechanical (TpHM) phenomena associated with permafrost thaw. Our work involves a new mixed finite element formulation for the phase change problem; this model works very well with the appropriate heterogeneous non-linearities, both at the microscopic scale as well as at the Darcy scale. We demonstrate robustness of our algorithm for thermo-hydro properties when handling irregular data such as the temperature-enthalpy relationship calibrated using data from Alaska. For the hydro-mechanical portion of the model, we consider the three-field formulation of Biot's system of poroelasticity, which we demonstrate is robust in physical scenarios involving heterogeneity and non-linear elastic and hydrological constitutive relationships and parameters. We further discuss our progress towards the coupled TpHM system, particularly the challenges of a monolithic scheme and our progress on sequential coupling.

Naren Vohra, Malgorzata Peszynska
Oregon State University
vohran@oregonstate.edu, mpsz@math.oregonstate.edu

MS81

Rational Approximations in Boundary Value Problems with Corner Singularities

Rational functions are well suited to approximate a wide range of functions. In comparison to polynomial approximations, rational functions can converge at exponential rates (root-exponential to be precise) even for functions with singularities. In this talk we focus on branch point singularities, which may arise in applications as singularities at corners or edges of a computational domain. We show that poles in a typical rational approximation for such functions cluster towards the singularity, not unlike graded meshes do in adaptive numerical simulation methods. The class of so-called lightning methods for boundary value problems, developed by Trefethen and collaborators, explicitly uses rational approximations and exploits this property by choosing clustered poles in advance. This turns the rational approximation problem into a linear one, since rational functions depend non-linearly on the location of the poles but linearly on the associated residues. We show that a judicious choice of the distribution of poles has a significant impact on the rate of convergence. Moreover, the chosen representation of the rational functions, i.e., the basis of the approximation space, strongly affects the numerical stability of these classes of methods.

Astrid Herremans, Daan Huybrechs
Department of Computer Science
KU Leuven
astrid.herremans@kuleuven.be,
daan.huybrechs@kuleuven.be

MS81

Constructing Summation-by-Parts Operators on Point Clouds

Computational scientists and engineers use summation-by-parts (SBP) operators to construct PDE discretizations

with desirable properties, such as nonlinear stability and superconvergent functionals. SBP operators are typically defined on a geometrically simple reference domain—for example, a quadrilateral or triangular domain—and then mapped to elements in an unstructured mesh or blocks in a multi-block grid. Here we consider a more flexible construction of SBP operators that is suitable for degrees of freedom over point clouds. The construction is a generalization of the discontinuous Galerkin difference (DGD) method. Unlike DGD, the proposed method does not pair degrees of freedom with elements in a one-to-one manner; instead, the element mesh is used only to perform integration and define the support of the basis functions. We will present results that verify the method and demonstrate its flexibility; in particular, we will present how r -adaption can be applied to the method without modifying the mesh.

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

Ge Yan
Rensselaer Polytechnic Institute
Department of Mechanical, Aerospace, and Nuclear Engineering
yang@rpi.edu

MS81

Efficient Tensor-Product Spectral-Element Methods with the Summation-by-Parts Property on Triangles and Tetrahedra

We present an extension of the summation-by-parts approach to nodal and modal spectral-element methods employing tensor-product operators in collapsed coordinates, enabling the construction of provably stable and conservative schemes of arbitrary order which combine the geometric flexibility of unstructured triangular and tetrahedral grids with the efficiency of matrix-free sum-factorization algorithms. Employing such a methodology within the context of a skew-symmetric discontinuous Galerkin formulation, we develop discretizations on curvilinear meshes which are energy conservative for the linear advection equation in the case of a central numerical flux and energy dissipative for an upwind numerical flux. We then describe the implementation of efficient algorithms for evaluating the proposed tensor-product operators within an open-source Julia code. Finally, the stability and conservation properties of the proposed methods are verified numerically, and comparisons with existing methods on the basis of accuracy and computational cost are presented, demonstrating that for higher polynomial degrees, the proposed schemes are more efficient than comparable multidimensional (i.e. non-tensor-product) formulations on triangles and tetrahedra.

Tristan Montoya
University of Toronto Institute for Aerospace Studies
tristan.montoya@mail.utoronto.ca

David Zingg
University of Toronto Institute for Aerospace Studies
dwz@utias.utoronto.ca

MS81

Nonlinear Boundary Conditions for Energy and Entropy Stable Initial Boundary Value Problems

in Computational Fluid Dynamics

We derive new boundary conditions and implementation procedures for nonlinear initial boundary value problems that lead to energy and entropy bounded solutions. A step-by-step procedure for general nonlinear hyperbolic problems on skew-symmetric form is presented. That procedure is subsequently applied to the three most important equations in computational fluid dynamics: the shallow water equations and the incompressible and compressible Euler equations. Both strong and weak imposition of the nonlinear boundary conditions are discussed. Based on the continuous analysis, we show that the new nonlinear boundary procedure lead to energy and entropy stable discrete approximations if the scheme is formulated on summation-by-parts form in combination with a weak implementation of the boundary conditions.

Jan Nordström

Department of Mathematics
Linköping University
jan.nordstrom@liu.se

MS81

Stability of IMEX Schemes for Convection-Diffusion Equations Discretized by Upwind SBP Schemes

In the context of convection-diffusion equations, IMEX approaches generally discretize the convection terms explicitly and treat the diffusion terms implicitly in order to avoid parabolic time step restrictions based on the stiff diffusion terms. Recent research has shown that specific combinations of space and time discretization even yield unconditional stability of the resulting scheme with grid-independent time step restrictions only based on the convection and diffusion constants. For DG space discretizations using upwind numerical fluxes, previous analysis via energy estimates has shown that the choice of diffusion fluxes influences this stabilization mechanism. For instance, whereas LDG fluxes allow for grid-independent time step sizes when IMEX time stepping is applied, the classical BR1 fluxes do not. In this talk, we focus on the spatial discretization by upwind SBP schemes which introduce an upwind mechanism similar to numerical fluxes within the DG method. Using energy equations, we investigate the stability of upwind-SBP-discretized convection-diffusion problems using IMEX time stepping. This analysis is quite general and only requires a compatibility of the first and second derivative SBP operators. In particular, the upwind SBP framework allows to show that replacing upwind advection fluxes by central fluxes in fact suffices to guarantee grid-independent time step restrictions in case of BR1 diffusion fluxes.

Sigrun Ortleb

Universität Kassel
ortleb@mathematik.uni-kassel.de

MS82

Bioeconomic Model of Three Marine Populations Including Tide Effects

The main objective of this work is the study of the effects of high tides and low tides on fishing effort, catches as well as profits in a bioeconomic model of populations of *Sardina pilchardus*, *Engraulis encrasicolus* and *Xiphias gladius* in Moroccan areas. To achieve this objective, we studied the stability of the equilibrium points of our biological model

then we added in our model the effect of the tides in the fishing effort which maximizes the profits of the fishermen under the constraint of the conservation of the biodiversity of these marine species using the generalized Nash equilibrium in the resolution of the bioeconomic model. As results, we were able to give the best fishing times according to the tides of each month of the whole year which will allow us to achieve better yields. Hence the importance of introducing the effect of high and low tides in bioeconomic models.

Noussaiba Baba

Hassan II university
noussaibababa1@gmail.com

Imane Agmour, Youssef El Foutayeni, Naceur Achtaich

Hassan II University of Casablanca

agmour.imane@gmail.com, foutayeni@gmail.com, nachtaich@gmail.com

MS82

Deep Autoencoder-Based Approaches for Plume Transport Predictions

Wind-driven spatial patterns are of great interest in the climate and geophysical sciences. Examples include features in Aeolian dunes and the deposition of volcanic ash, wildfire smoke, and air pollution plumes. Machine learning models are attractive because they can make accurate and fast predictions when the physical processes are not fully understood, experiments are unavailable, or computer simulations are too costly. We explore deep convolutional neural network-based autoencoders exploiting relationships in geophysical spatial patterns to reduce their dimensionality. Reducing the dimension size with an encoder allows us to train regression models linking geographic and meteorological scalar input quantities to the encoded space. Once this is achieved, full predictive spatial patterns are reconstructed using the decoder. We also designed a novel spatiotemporal prediction model, ST-GasNet, inspired by the partial differential equations used to model plume evolution. The ST-GasNet model learns the spatiotemporal dependencies from limited binary training data and accurately predicts the evolution of spatial distributions of contaminant plumes later in time (15-time steps or 9 minutes) given the initial time sequences (5-time steps or 3 minutes) as the input. Finally, we are currently working on a 3D CNN Autoencoder Architecture to extend our previous work to temporal predictions of the magnitude of the quantities of interest instead of just binary data.

M. Giselle Fernandez-Godino

Lawrence Livermore National Laboratory
fernandez48@llnl.gov

Wai Tong Chung

Stanford University
wtchung@stanford.edu

Yinan Wang

Rensselaer Polytechnic Institute
wangy88@rpi.edu

Donald D. Lucas

Lawrence Livermore National Laboratory

lucas26@llnl.gov

MS82

Learning Gas Transport Through Fracture Networks from Multi-Fidelity Data

Modeling gas transport through fractures of subsurface rock is particularly difficult because of the heterogeneous nature of the material. Structural information such as fracture size, geometry, and other material properties plays a key role in modeling the dominant physics for flow and transport. We explore the application of machine learning to create surrogate reduced-order models coupled with graph theoretic techniques, to characterize flow through fracture networks. We utilize training data from computationally intensive high-fidelity discrete fracture network simulations, augmented by easily generated low-fidelity (i.e. noisy) graph flow models, in a multi-fidelity approach. Specifically, we take advantage of Bayesian methods to predict the high-fidelity model output, along with uncertainty bounds. In this way, we aim to generate rapid and accurate predictions of particle breakthrough times, orders of magnitude faster than the simulations themselves, with associated uncertainty quantification.

John Kath, Donald Berry, An Ly, Yingqi Shi, Allon Percus
Claremont Graduate University
john.kath@cgu.edu, donald.berry@cgu.edu,
an.ly@cgu.edu, yingqi.shi@cgu.edu, allon.percus@cgu.edu

Jeffrey Hyman
Los Alamos National Laboratory
Earth and Environmental Sciences Division
jhyman@lanl.gov

Kelly Moran, Philip Stauffer, Justin Strait
Los Alamos National Laboratory
kmoran@lanl.gov, stauffer@lanl.gov, jstrait@lanl.gov

Matthew R. Sweeney
Earth and Environmental Sciences Division
Los Alamos National Laboratory
msweeney2796@lanl.gov

Hari Viswanathan
Los Alamos National Laboratory
Earth and Environmental Sciences Division
viswana@lanl.gov

MS82

Global Sensitivity Analysis Using the Ultra-Low Resolution Energy Exascale Earth System Model to Investigate Parametric Uncertainty in Arctic Climate

Accurate seasonal predictions of sea ice minimum extent and long-term estimates of timing for a seasonally ice-free Arctic depend on a better understanding of the factors influencing sea ice dynamics in the strongly coupled Earth system. As a first step in quantifying uncertainties in simulated Arctic climate response, we performed a variance-based global sensitivity analysis (GSA) using a fully coupled ultra-low resolution (ULR) configuration of version 1 of the US Department of Energy's Energy Exascale Earth System Model (E3SMv1). We quantify the sensitivity of six Arctic quantities of interest (QOI), which characterize changes in Arctic climate over a 75 year period, to un-

certainties in nine model parameters spanning the sea ice, atmosphere, and ocean components of E3SMv1. Sensitivity indices for each QOI were computed with a Gaussian process emulator using 139 random realizations of the random parameters and fixed preindustrial forcing. The parameter variations showed significant impact on the Arctic climate state with the largest impact coming from atmospheric parameters in the Cloud Layers Unified by Binormals (CLUBB) scheme. Our results demonstrate the importance of conducting sensitivity analyses with fully coupled climate models and motivate follow-on investigations using the ULR model.

Irina K. Tezaur, Kara J. Peterson, Amy Powell
Sandia National Laboratories
ikalash@sandia.gov, kjpeter@sandia.gov,
ajpowel@sandia.gov

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Erika Roesler
Sandia National Laboratories
elroesl@sandia.gov

MS83

Optimal Control of Automated Vehicles for Flow Smoothing in Mixed Autonomy Traffic

The introduction and control of automated vehicles (AVs) in predominantly human-driven traffic has the potential to improve traffic conditions at both the local and the system levels. In this work, we study the optimal control of AVs in mixed-autonomy traffic with the objective of smoothing unstable traffic flows. We first model the dynamics of a mixed-autonomy system where the acceleration of human-driven vehicles is governed by a car-following model, and the acceleration of AVs is to be controlled. This is realized as a system of first-order ODEs. We then formulate the optimal control problem and show the well-posedness of the system dynamics for a reasonable set of admissible controls and establish the existence of minimizers. Given this problem setup, we devise an approach for solving the optimal control problem based on Pontryagin's maximum principle (PMP) and the adjoint method. Further, we propose to approximate the optimal controller using a set of parametrized basis functions. Their parameters are tuned by optimizing an objective function that involves the original objective of the optimal control problem over a distribution of trajectories. Through this approach, we aim to develop a feedback control scheme based on our open-loop optimal control problem. We demonstrate the effectiveness of both approaches on real-world trajectories data collected from the I-24 in Tennessee, United States.

Arwa AlAnqary
UC Berkeley
arwa@berkeley.edu

Alexander Keimer
ITS, UC Berkeley
keimer@berkeley.edu

Xiaoqian Gong
Arizona State University
U.S.
xgong14@asu.edu

Alexandre Bayen
University of California, Berkeley
bayen@berkeley.edu

MS83

Detecting Congestion and Boundary Conditions: How Machine Learning Techniques can Improve Differential Traffic Models

Vehicular traffic has traditionally been described and forecast by means of mathematical models, either differential and nondifferential. In recent years, machine learning came into play in order to analyze traffic data, trying to discover patterns and substitute, when possible, traditional models. Hybrid approaches like the one based on physics-informed neural networks were also proposed. In this talk we focus on a road segment where traffic flux is continuously counted by means of some fixed sensors. The idea is to couple machine learning and differential models, putting the former at the service of the latter. More precisely, we use a neural network to predict sensor data at the inflow boundary of the road, where the model expects information about the next-to-come traffic flow. Moreover, we use a neural network to detect congestion formations near the sensors. This information is used to recover the car density by inverting the (noninjective) fundamental diagram. Some examples motivated by real scenarios will be discussed. Real data are provided by the Italian company Autovie Venete S.p.A.

Emiliano Cristiani

Istituto per le Applicazioni del Calcolo
Consiglio Nazionale delle Ricerche
e.cristiani@iac.cnr.it

MS83

Aspects of Nonlocal Traffic Flow Modeling

The technological progress in autonomous driving brings new challenges for the modeling of traffic flow and extensions of classical approaches, such as the LWR or ARZ model, are needed in order to manage these. As autonomous cars possess more data about the surrounding traffic, there is a need that traffic models incorporate those. Therefore, nonlocal traffic flow models can be considered. They include more information in a certain nonlocal range about the traffic of the road. In case of non-autonomous drivers, the nonlocal range can stand for the sight of a human driver. Moreover, this nonlocal range can be interpreted as the connection radius of autonomous cars. In this talk we will give an overview on nonlocal traffic flow models and discuss the challenges which arise when comparing these models to real traffic data.

Jan Friedrich

RWTH Aachen University
friedrich@igpm.rwth-aachen.de

MS83

Macroscopic Traffic Models and Autonomous Vehicles

We present two models for controlling car traffic through special vehicles, like autonomous ones. Car traffic is described by macroscopic models, while the dynamics of the special vehicles is described by a microscopic one. More precisely, either the Lighthill-Whitham-Richards (LWR) model [?, ?] or, alternatively, the Colombo-Marcellini-Rascle

(CMR) model [?] describes the evolution of traffic in a road. Assume that a vehicle, whose position is described by the function $y(t)$, aims at controlling the behavior of traffic. The evolution of such a vehicle is described by the ODE $\dot{y}(t) = u(t)$ where $u(t)$ is a control function, which selects the desired speed. Following the ideas proposed in [?], we consider two control models, one based on the LWR model and one on the CMR model, where the control acts on the autonomous vehicle. We discuss about the concepts of solutions for the two systems and we show that, given a control function $u = u(t)$ with finite total variation, a solution exists for both systems. The proofs are based on the wave-front tracking technique. These are joint works with P. Goatin, T. Liard, F. Marcellini, and B. Piccoli.

Mauro Garavello

University of Milan
mauro.garavello@unimib.it

MS83

Uncertainty Quantification for Microscale Traffic-Related Pollutant Dispersion CFD Simulations

Uncertainty quantification (UQ) in urban computational fluid dynamics (CFD) is challenging due to significant computational costs, which complicate the efficient construction of response surfaces. Moreover, exploration and accuracy in the UQ space are all the more difficult to reach when more than 2 uncertain variables and functional outputs are considered, as is the case for instance when constructing atmospheric pollution maps. Traditionally, such UQ analyses are carried out using proper orthogonal decomposition with gaussian processes to emulate the coefficients of a truncated Karhunen-Loeve decomposition (POD-GP). This method is used in this work to carry out a global sensitivity analysis for a realistic high-dimensional traffic-related atmospheric pollutant dispersion CFD study with up to 5 uncertain variables. The pollutant dispersion is modelled with the open-source OpenLB lattice-Boltzmann CFD code, coupled with the SUMO microscale traffic simulator and a physical engine emissions model. We also compare the POD-GP method with another metamodeling strategy based on an anchored-ANOVA decomposition of the quantity of interest, which was designed to reduce the number of required samples in the UQ space for a given accuracy.

Mathis J. Pasquier

IFP Energies Nouvelles
IFP Energies Nouvelles
mathis.pasquier@ifpen.fr

Delphine Sinoquet

IFP Energies nouvelles
delphine.sinoquet@ifpen.fr

Stephane Jay

IFP Energies Nouvelles
Stephane.Jay@ifpen.fr

Jérôme Jacob, Pierre Sagaut

Aix-Marseille Université
jerome.jacob@univ-amu.fr, pierre.sagaut@univ-amu.fr

MS84

Fast Finite Element Solver for Focused Ultrasound Applications

Focused ultrasound (FUS) application is a therapeutic

medical technology that uses ultrasound for tissue heating and ablation. In recent years, interest in FUS application has been growing, especially for the treatment of some cancers. In contrast to conventional cancer therapy, FUS application is non-invasive and as such makes it an attractive treatment method. However, effective patient-specific FUS treatment would be supported by accurate simulations of treatment. This presents a computational challenge as the ratio of the domain size to the wavelength is typically large, and therefore computationally demanding. We have developed a fast, high-order matrix-free finite element solvers to simulate a 3D FUS application. We show that by employing a high-order methods and fast assembly techniques, we simulate with high accuracy and acceptable computation time. We show that we can use methods with the lowest known computational cost complexity and achieve a very significant fraction of the peak performance of modern CPU architectures, and with excellent parallel scalability. Overall, we show that our solver is fast, accurate, scalable, and is suitable for large-scale time-domain nonlinear acoustics simulations.

Adeeb Arif Kor
University of Cambridge
maabm4@cam.ac.uk

Igor Baratta, Garth Wells
Department of Engineering
University of Cambridge
ia397@cam.ac.uk, gnw20@cam.ac.uk

MS84 Adapting Exascale Techniques Through Ginkgo

This talk presents modern, exascale-ready techniques within the Ginkgo high-performance linear algebra library, focused on manycore systems. The application area is the implicit or semi-implicit time stepping of hyperbolic-parabolic partial differential equations discretized with discontinuous Galerkin methods. Within this area, strategies such as batched solver, or mixed-precision are discussed. The talk will highlight the necessary abstractions and implementations on GPU hardware, and present the performance on sample problems.

Marcel Koch
Karlsruhe Institute of Technology
marcel.koch@kit.edu

MS84 Large Eddy Simulation of Aircraft at Affordable Cost: a Milestone in Computational Fluid Dynamics

While there have been numerous applications of large eddy simulations (LES) to complex flows, their application to practical engineering configurations, such as full aircraft models, have been limited to date. Recently, however, advances in rapid, high quality mesh generation, low-dissipation numerical schemes and physics-based subgrid-scale and wall models have led to, for the first time, accurate simulations of a realistic aircraft in landing configuration (the Japanese Aerospace Exploration Agency Standard Model) in less than a day of turnaround time with modest resource requirements. In this paper, a systematic study of the predictive capability of LES across a range of angles of attack (including maximum lift and post-stall regimes), the robustness of the predictions to grid resolution and the incorporation of wind tunnel effects is carried

out. Integrated engineering quantities of interest, such as lift, drag and pitching moment will be compared with experimental data, while sectional pressure forces will be used to corroborate the accuracy of the integrated quantities.

Oriol Lehmkuhl
BSC
oriol.lehmkuhl@bsc.es

MS84 Towards Exascale for Reactor and Wind Energy Simulations

As part of DOE's ECP Center for Efficient Exascale Discretizations, we have developed Nek5000/RS, a highly-performant open-source thermal-fluids code based on high-order spectral element discretizations that is targeting GPU-based platforms. Recent developments and optimizations have led to an unprecedented milestone in the ability to perform large-eddy simulation of a full reactor core. Careful tuning of the high-performance kernels, the multilevel solvers, and timestepping algorithms led to a 4X speed-up in performance when running on all GPUs of Summit (4608 nodes=27648 GPUs) at Oak Ridge Leadership Computing Facility. As a result, it is now possible to simulate flow through a 352000-pebble bed reactor (98 million spectral elements of order $N=8$, $n=51$ billion gridpoints) in just six hours of wall clock time. We also explore large-eddy-simulation modeling approaches and computational performance for the simulation of atmospheric boundary layer (ABL) flows that are of direct relevance to wind energy production. We present the performance of NekRS on Summit and Crusher, the testbed for the Frontier exascale system, using 18 to 384 Graphics Compute Dies on AMD MI250X GPUs. We compare strong- and weak-scaling capabilities, linear solver performance, and time to solution of NekRS with AMR-Wind. We also identify leading inhibitors to parallel scaling.

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

MS84 Performance-Portable p -Multigrid for Nonlinear Solid Mechanics in Nearly Incompressible Regime

We introduce a matrix-free p -multigrid finite element method with Newton-Krylov and quasi-Newton iterative solvers for hyperelasticity problems. Contrary to standard formulations of finite strain constitutive models, which lead to loss of several digits of relative accuracy for small deformations, we introduce a stable formulation that is accurate for all-deformation elasticity. We investigate reliability, efficiency, and accuracy on a collection of multiscale compressible and incompressible materials up to billions of degrees of freedom running on CPU and GPU architectures. We find that high order methods deliver significant costs savings even for loose error tolerances and physical singularities in complex geometric models.

Rezgar Shakeri, Jeremy Thompson, Jed Brown
University of Colorado Boulder
rezgar.shakeri@colorado.edu,

jeremy.thompson@colorado.edu, jed@jedbrown.org

MS85

Forward-Mode Enzyme in Developing Constitutive Models with Ratel

Ratel is a new, open-source package built on libCEED and PETSc capable of solving complex solid mechanics problems without sacrificing computational performance. Computing derivatives is essential to the algorithms employed within Ratel. However, deriving and implementing derivatives of some constitutive models could be cumbersome hence exploiting automatic differentiation (AD) tools could simplify the implementation or provide verification results for hand-coded derivatives. Enzyme is a new LLVM plugin with GPU support that provides split forward and reverse mode AD on LLVM intermediate representation (IR). We explore the applicability and performance of Enzyme in computing tensor gradients in some of our elasticity experiments in Ratel.

Leila Ghaffari

University of Colorado Boulder
leila.ghaffari@colorado.edu

MS85

Efficient Batched Forward-Mode Derivatives with Compiler-Based Automatic Differentiation

Derivatives are becoming ever more central to modern approaches in the computational sciences. From gradient-based optimization over uncertainty quantification to the incorporation of machine learning approaches into our simulation workflows, the ability to obtain gradients of our simulation code determines our ability to access this exceedingly large toolbox. But for a large number of our simulations the rewriting of entire simulations in differentiable domain-specific languages such as JAX or PyTorch is simply infeasible. Compiler-based automatic differentiation with Enzyme enables the synthesization of gradients even for these simulations in any language which utilizes LLVM in their compiler such as e.g. C/C++, Julia, Fortran, Rust, and Swift. This tight integration into the compiler enables Enzyme to synthesize much more efficient gradients by operating on representations already optimized by the compiler. In this talk we present key extensions to Enzyme which enable the auto-batching of operations inside of Enzyme to convert scalar functions into vectorized functions for more efficient gradients in downstream applications.

Tim Gymnich

Technical University of Munich
tim.gymnich@tum.de

MS85

The Pitfalls of Backpropagation Some Perspectives and Alternatives

The ubiquity of neural networks has pushed the boundaries of human understanding in many areas such as science, healthcare, finance. The key reason for this popularity is the availability of GPU infrastructure to enable fast matrix calculations necessary to perform back-propagation the workhorse behind the design and development of neural network. While backpropagation has been proven effective, it is not without its pitfalls. In this talk, we will discuss the many challenges that arise during the design

and development of neural networks through backpropagation especially in science application. We will then provide alternatives to backpropagation that are both theoretically sound and can be easily amalgamated into the current GPU infrastructure for fast and efficient training of neural networks. We will show that, our alternatives can provide performance equivalent to backpropagation when gradients are available with the added advantage of being able to address scenarios where the gradients may be inaccurate. We will demonstrate the efficiency of these alternatives in GPU infrastructures and in AI accelerators and derives perspectives for future research.

Krishnan Raghavan

Mathematics and Computer Science
Argonne National Laboratory
kraghavan@anl.gov

MS85

Randomized Dimensionality Reduction Techniques for Automatic Differentiation

Recent research has shown the power of randomized dimensionality reduction techniques in creating accuracy preserving reduction of general nonlinear models. The reduction is in the form of active subspaces, described by linear transformations of the model parameters and responses. Theoretical results prove that the errors resulting from the reduction can be upper bounded with high probability. This talk shows how the resulting active subspaces can be leveraged to reduce the computational burden associated with Automatic Differentiation by confining the differentiation to a number of pseudo parameters and responses as defined by the dimensionality reduction algorithm.

Arvind Sundaram

Purdue
sundara4@purdue.edu

Hany S. Abdel-Khalik

North Carolina State Univ.
abdelkhalik@ncsu.edu

MS85

Parallel Adjoint Taping Using MPI

Automatic Differentiation (AD) allows to efficiently and accurately calculate derivatives of expressions stated as computer code. The adjoint data flow reversal of long evolutionary calculations (e.g. loops), where each iteration depends on a set of parameters and the output of the previous iterate, is a common occurrence in computational engineering (e.g. computational fluid dynamics simulation), physics (e.g. molecular dynamics) and computational finance (e.g. Monte Carlo paths). For the extreme case of a scalar state, the execution as well as adjoint control flow reversal are inherently serial operations, as there is no spatial dimension to parallelize. We propose a method, focusing on programs with such a structure, that exploits the run time difference typically exhibited by AD tools between pure function evaluation and evaluation with additional calculation of local derivative information (a process frequently called pre-accumulation), which is then later used by the data flow reversal. Additional parallelism is introduced into the computation by distributing the aforementioned calculation of local derivatives onto multiple processes. A reference C++ implementation using MPI is presented, which allows us to reverse an OpenFOAM simulation. The proposed method is most beneficial for operator overloading AD tools, how-

ever the concepts are also applicable to source-to-source transformation and handwritten adjoints, or a hybrid of all approaches.

Markus Towara

STCE, RWTH Aachen University
towara@stce.rwth-aachen.de

Uwe Naumann

RWTH Aachen University
Software and Tools for Computational Engineering
naumann@stce.rwth-aachen.de

MS86

Geometric Mechanics Formulations of Momentum-Based Descriptions for Fluids Using Bundle-Valued Differential Forms

The use of differential forms to represent physical quantities enables a dimension and coordinate system independent description of continuum mechanics valid on arbitrary manifolds. Typically these are scalar-valued differential forms (SVDFs), however certain quantities such as momentum and stress require the use of (vector) bundle-valued differential forms (BVDFs). In this work we leverage the representation of momentum as a twisted covector-valued volume form to develop geometric mechanics formulations (variational, Hamiltonian, etc.) for momentum-based description of fluids, including a general conservation-type form. We consider a general single-velocity model based a semi-direct product theory (a special case of matched pair dynamics) with arbitrary advected quantities, including those with internal dynamics (ex. electromagnetic fields in electrohydrodynamics). This is sufficient to treat many models of interest, such as the thermal shallow water equations, compressible Euler equations, magnetohydrodynamics and electrohydrodynamics (including Euler-Poisson and Euler-Maxwell). If time permits, there will be discussion about the use of these new BVDF-based geometric mechanics formulations to develop structure-preserving discretizations for momentum-based descriptions of fluids, just as SVDFs are used for velocity-based descriptions.

Christopher Eldred

Sandia National Laboratories
celdred@sandia.gov

MS86

Structure Preserving Hybrid Finite Element - Finite Volume for MHD

In this talk we present a structure preserving hybrid Finite-Volume approach for solving the compressible MHD equations in all Mach number regimes. A splitting approach is designed so that we may take advantage of the conservation properties and robustness of the finite-volume (FV) schemes against the nonlinear advection, while relying on a structure-preserving discretization of the magneto-acoustic terms. Indeed, the nonlinear convective terms are treated via an explicit time-discretization, while the magneto- and acoustic- terms are solved via an implicit time-discretization. Thanks to this, the resulting CFL condition will depend formally only on the fluid velocity and not on the Alfvn- and/or sound-speeds that may become too stringent in the low Mach regimes. In this approach, the divergence free constraint of the magnetic field is always preserved up to machine precision, and the symmetry of the physical model is also reflected to the symmetry of the final algebraic nonlinear systems.

Thanks to the symmetry of the systems, the very-efficient matrix-free conjugate gradient method may be employed. First, we'll describe and show the main results obtained with a shock-capturing Finite-Volume/staggered Finite-Difference solver, which is shown to be very robust even against stringent MHD shock problems. Second, our most recent results about the extension of this new family of hybrid schemes to compatible Finite-Elements for weakly compressible flows will be presented.

Francesco Fambri

Max-Planck Institute of Plasma Physics
Garching
francesco.fambri@ipp.mpg.de

Eric Sonnendrucker

Max-Planck Institute of Plasmas Physics
Garching
eric.sonnendrucker@ipp.mpg.de

MS86

Variational and Thermodynamically Consistent Discretization for Heat Conducting Fluids

We construct a structure-preserving and thermodynamically consistent finite element method and time-stepping scheme for heat conducting viscous fluids. The method is deduced by discretizing a variational formulation for nonequilibrium thermodynamics that extends Hamilton's principle for fluids to systems with irreversible processes. The resulting scheme preserves the balance of energy and mass to machine precision, as well as the second law of thermodynamics, both at the spatially and temporally discrete levels.

François J. Gay-Balmaz

CNRS Ecole Normale Supérieure
francois.gay-balmaz@lmd.ens.fr

MS86

Structure-Preserving Discretisations of Incompressible Fluids on the Sphere

Geometric mechanics provides a powerful theoretical framework for formulating models in geophysical fluid dynamics. The discretisation of such models poses significant computational challenges, both in representing discretely the structure of the original problem and the corresponding computational cost. In this work we will present recent developments in geometric integration for fluid dynamics in two dimensions. In particular, we will illustrate a recently developed efficient and scalable Lie-Poisson integrator for flows on the sphere. We will show that it is possible to design geometric integrators which conserve the Casimirs of the system at modest computational costs. The construction of such schemes, the main numerical algorithms and their parallelisation on modern supercomputing facilities will be discussed. An application to the spectrum of homogeneous two-dimensional turbulence and quasi-geostrophic turbulence will be illustrated.

Erwin Luesink

University of Twente
e.luesink@utwente.nl

MS87

A Space-Time Adaptive Low-Rank Method for High-Dimensional Parabolic Partial Differential

Equations

In this talk we present the construction and analysis of a space-time adaptive method for parabolic partial differential equations that combines sparse wavelet expansions in time with adaptive low-rank approximations in the spatial variables. Similar to the existing adaptive low-rank method for elliptic problems, the method is based on a perturbed Richardson iteration, in the present case applied to a standard space-time variational formulation of the parabolic initial-boundary value problem. The analysis of the method requires a new approximation class for the temporal operator, taking into account the interaction between hierarchical tensor formats of different time indices. Since the parabolic operator is an isomorphism with respect to spaces not endowed with a cross norm, we devise a new low-rank preconditioning scheme based on exponential sum approximations that is adapted to the parabolic case. The method is shown to converge and satisfy similar complexity bounds as the existing adaptive low-rank methods for elliptic problems, establishing its suitability for parabolic problems on high-dimensional spatial domains and does not suffer from the curse of dimensionality. The construction also yields computable rigorous a posteriori error bounds for the total error depending on the activated basis functions and ranks in the approximation.

Markus Bachmayr, Manfred Faldum
 RWTH Aachen University
 bachmayr@igpm.rwth-aachen.de, faldum@igpm.rwth-aachen.de

MS87

Rank-Adaptive Time Integration of Tree Tensor Networks

The numerical approximation of time dependent high dimensional problems is a challenging task: The storage required by the numerical approximation of the solution often exceeds standard memory resources. In this setting, time-dependent model order reduction and low-rank approximation techniques are of interest. In the present contribution, dynamical low-rank approximation for matrices together with recent developments on the topic of rank-adaptivity is briefly recapitulated. Then, appropriate notations and definitions for hierarchical tensors in Tucker format - also known as Tree Tensor Networks (TTNs) - are presented, and a robust rank-adaptive numerical integrator for the dynamical low-rank approximation of TTNs is introduced. The proposed algorithm recursively moves from the leaves to the root, updating the bases and the connection tensors via a Galerkin projection in an augmented subspace followed by a rank truncation up to a given error tolerance. It is shown that the memory requirements are linear in the order of the tensor and linear in the maximal mode dimension. Furthermore, the integrator inherits from the matrix setting property of interests such as robustness to small singular values of matricizations of the connection tensors. Up to the rank truncation error, the proposed integrator preserves norm and energy for Schroedinger equations, and it dissipates the energy in gradient systems.

Gianluca Ceruti
 EPFL
 gianluca.ceruti@epfl.ch

Dominik Sulz
 University of Tuebingen
 dominik.sulz@uni-tuebingen.de

Christian Lubich
 Mathematisches Institut
 Universität Tübingen
 lubich@na.uni-tuebingen.de

MS87

Streaming Tensor Train Approximation

Tensor trains are a versatile tool to compress and work with high-dimensional data and functions. In this talk, we will introduce the Streaming Tensor Train Approximation (STTA), a new class of algorithms for approximating a given tensor T in the tensor train format. STTA accesses T exclusively via two-sided random sketches of the original data, making it streamable and easy to implement in parallel unlike existing deterministic and randomized tensor train approximations. This property also allows STTA to conveniently leverage structure in T , such as sparsity and various low-rank tensor formats, as well as linear combinations thereof. When Gaussian random matrices are used for sketching, STTA is admissible to an analysis that builds and extends upon existing results on the generalized Nystrom approximation for matrices. Our results show that STTA can be expected to attain a nearly optimal approximation error if the sizes of the sketches are suitably chosen. A range of numerical experiments illustrates the performance of STTA compared to existing deterministic and randomized approaches.

Daniel Kressner
 EPFL Lausanne
 Mathicse
 daniel.kressner@epfl.ch

Bart Vandereycken, Rik Voorhaar
 University of Geneva
 bart.vandereycken@unige.ch, rik.voorhaar@unige.ch

MS87

Structure Preserving Hyper-Reduction of Parametric Hamiltonian Systems

Many conservative processes in a wide range of physical and engineering applications can be modelled by Hamiltonian systems. The dynamics of such systems is described by the gradient of a scalar function, the Hamiltonian, which is preserved over time. Since the number of degrees of freedom required to provide accurate numerical solutions can be very large, model order reduction and hyper-reduction methods are typically used to achieve computational speedups, in particular in many-query simulations. However, traditional hyper-reduction techniques, such as discrete empirical interpolation (DEIM), fail to preserve the gradient structure of the flow velocity, which might lead to unphysical and unstable solutions. Furthermore, conservative phenomena do not exhibit global reducibility properties, in the sense that large approximation spaces are usually needed to accurately approximate the full order solution. To overcome these issues, we construct adaptive hyper-reduced models that retain the gradient structure and whose complexity is independent of the dimension of the full order problem. Finally, in order to exploit the local low-rank nature of conservative phenomena, both the reduced space and the DEIM basis are updated in time. In this talk, we will discuss this adaptive approach and present several numerical results.

Federico Vismara
 Eindhoven University of Technology

f.vismara@tue.nl

Cecilia Pagliantini
TU/e Eindhoven University of Technology
c.pagliantini@tue.nl

MS88

SPH-EXA: A Framework for Scalable, Flexible, and Extensible Astrophysical and Cosmological Simulations

SPH-EXA is a highly scalable and extendable simulation framework for astrophysical and cosmological simulations. It is codesigned by computational (astrophysicists and cosmologists) and computer scientists (high-performance computing) to achieve scientific advances in astrophysics, cosmology, and high-performance computing, for highly scalable simulations using Smoothed Particle Hydrodynamics. SPH-EXA includes highly optimized and parallelized off-the-shelf hydrodynamics and gravity solvers. It supports new particle types and particle data fields that can be combined with custom-made simulation-derived observable properties and in-situ data analysis. It requires minimal software dependencies, provides scalable parallelization and communication support, and portability with optimizations for recent CPUs and GPUs. This design relieves potential users from architectural specifics and performance concerns for scaling up simulations with SPH-EXA. This talk will describe the SPH-EXA framework, its approach to domain decomposition, parallelization, gravity and hydrodynamics solvers, and their flexible use as pluggable components. A scalable initial conditions generator, test cases and simulations will also be presented. SPH-EXA is extendable with additional physical effects through the use of propagators, whereby the pluggable framework components are easily customized or implemented from scratch using the provided building blocks in an abstract, efficient and scalable way.

Florina M. Ciorba
University of Basel, Switzerland
florina.ciorba@unibas.ch

Ruben Cabezon
University of Basel
ruben.cabezon@unibas.ch

Osman Simsek
University of Basel, Switzerland
osman.simsek@unibas.ch

Lucio Mayer
University of Zurich, Switzerland
lmayer@physik.uzh.ch

Sebastian Keller, Jean-Guillaume Piccinali
Swiss National Supercomputing Centre, Switzerland
keller@cscs.ch, jgp@cscs.ch

Axel Sanz
Universitat Politècnica de Catalunya
axel.sanz@unibas.ch

Joseph Touzet
École Normale Supérieure Paris-Saclay
joseph.touzet@ens-paris-saclay.fr

Noah Kubli

University of Zurich, Switzerland
noah.kubli@uzh.ch

Lukas Schmidt
University of Basel, Switzerland
luke.schmidt@stud.unibas.ch

Darren Reed
University of Zurich, Switzerland
darren.reed@uzh.ch

Jose Escartin
Max Planck Institute for Extraterrestrial Physics
jescarti@mpe.mpg.de

MS88

Fully Kinetic Simulations of Particle Acceleration in Plasma Turbulence

Turbulence is often invoked to explain the origin of non-thermal particles in space and astrophysical plasmas. By means of 3D fully kinetic particle-in-cell simulations, we demonstrate that turbulence in low-beta plasmas (beta is the ratio of plasma pressure to magnetic pressure) accelerates ions and electrons into a nonthermal energy distribution with a power-law energy range. The ion spectrum is harder than the electron one, and both distributions get steeper for higher beta. We show that the energization of electrons is accompanied by a significant energy-dependent pitch-angle anisotropy, with most electrons moving parallel to the local magnetic field, while ions stay roughly isotropic. We demonstrate that particle injection from the thermal pool occurs in regions of high current density. Parallel electric fields associated with magnetic reconnection are responsible for the initial energy gain of electrons, whereas perpendicular electric fields control the overall energization of ions. Our findings have important implications for the origin of nonthermal particles in space and astrophysical plasmas.

Luca Comisso
Columbia University, NY
USA
luca.comisso@columbia.edu

Lorenzo Sironi
Columbia University
lsironi@astro.columbia.edu

MS88

An Event-Based Approach to Monte Carlo Radiation Transport

Radiation transport modeling forms the backbone of predicting the electromagnetic emission from high energy astrophysical events, including supernova explosions and neutron star mergers. Monte Carlo (MC) methods are a favorable algorithmic choice in these situations, given the inherent multi-dimensional and multi-physics nature of such phenomena. MC transport methods have traditionally used a history-based approach, where individual particle histories are tracked until census. With the recent computational shift to GPUs, an alternative "event-based" approach is becoming more attractive given their potential speedups compared to their history-based counterparts. In this talk, we present a CUDA implementation of an event-based MC method in the astrophysical radiative transfer code Sedona. Through a series of benchmark

test problems, we examine the computational efficiency relative to a history-based CPU approach including strong and weak scaling measurements. We also demonstrate the drawbacks of event-based approaches for memory-bound situations that may appear in more realistic science production runs, and suggest potential workarounds in such cases.

David Khatami, Benny Tsang
University of California, Berkeley, USA
david_khatami@berkeley.edu, benny.tsang@berkeley.edu

MS88

Simulations of Magnetic Reconnection with a Pseudo-Spectral Maxwell Solver

Magnetized high-temperature plasmas are found in astrophysical settings like gamma-ray bursts, magnetars, and stellar coronae. These plasmas can undergo magnetic reconnection, which converts the potential energy of the magnetic fields to particle kinetic energy. This mechanism is an important source of particle acceleration and heating in astrophysical plasmas. Using the particle-in-cell code WarpX, we perform large first-principles 2D simulations of relativistic reconnection plasmas representative of those found in these environments. These simulations produce particle energies and plasmoid structures in agreement with results previously reported in the literature. Using this robust baseline case, we compare the accuracy and computational performance of three methods of solving Maxwell's equations, including the commonly-used second-order finite difference time domain (FDTD) method and the more advanced ultra-high-order pseudo-spectral analytical time domain (PSATD) method. This is the first time PSATD has been used in simulations of astrophysical reconnection. We find that for the reconnection problem, FDTD and PSATD are comparably accurate, but that time-to-solution with PSATD is half that of FDTD. These performance gains will make our 3D simulations of reconnection more tractable, and complement other efforts to improve simulation efficiency, such as the use of mesh refinement.

Hannah Klion, Revathi Jambunathan, Michael Rowan
Lawrence Berkeley National Laboratory
klion@lbl.gov, rjambunathan@lbl.gov, mrowan@lbl.gov

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

Rémi Lehe
Lawrence Berkeley National Laboratory
ATAP
rlehe@lbl.gov

Jean-Luc Vay
Lawrence Berkeley National Laboratory
Berkeley, CA
JLVay@lbl.gov

Wei-qun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering

weiqunzhang@lbl.gov

MS90

Modular Data-Driven Elastoplasticity

The development of highly accurate constitutive models for materials that undergo path-dependent processes continues to be a complex challenge in computational solid mechanics. Challenges arise both in considering the appropriate model assumptions and from the viewpoint of data availability, verification, and validation. Recently, data-driven modeling approaches have been proposed that aim to establish stress-evolution laws that avoid user-chosen functional forms by relying on machine learning algorithms. However, these approaches not only require a significant amount of data but also need data that probes the full stress space and also probes a variety of complex loading paths. In this work, we discuss a hybrid framework that can work on a variable amount of data by relying on the modularity of the elastoplasticity formulation where each component of the model can be chosen to be either a classical phenomenological or a data-driven model depending on the amount of available information. The method is tested on synthetic uniaxial data coming from simulations as well as cyclic experimental data for structural materials. The discovered material models are found to not only interpolate well but also for accurate extrapolation in a thermodynamically consistent manner far outside the domain of the training data.

Jan Niklas Fuhg
Cornell University
414 Upson Hall, 124 Hoy Rd, Ithaca, NY 14850
jf853@cornell.edu

Craig M. Hamel, Kyle Johnson, Reese Jones
Sandia National Laboratories
chamel@sandia.gov, kyljohn@sandia.gov,
rjones@sandia.gov

Nikolaos Bouklas
Cornell University
nb589@cornell.edu

MS90

Exploring the Low-Data Regime with Model-Data-Driven Yield Functions

Constitutive modeling of path-dependent processes continues to be a complex challenge in computational solid mechanics. Recently, data-driven modeling approaches have been proposed that aim to establish stress-evolution laws that avoid user-chosen functional forms by relying on machine learning algorithms. However, these approaches not only require a significant amount of data but also need data of the full stress space and hence are in particular not suitable for small-data regimes. In this talk, we discuss a hybrid framework that can work on a variable amount of data by relying on the modularity of the elastoplasticity formulation where each component of the model can be chosen to be either a classical phenomenological or a data-driven model depending on the amount of available information. The method is tested on synthetic uniaxial data coming from simulations as well as cyclic experimental data of steels. The obtained material models are found to not only interpolate well but also allow to extrapolate in a thermodynamically consistent manner far outside the domain of the training data. Training aspects and details about the implementation of these models into Finite-Element

simulations are discussed and analyzed.

Jan Niklas Fuhg
Cornell University
414 Upson Hall, 124 Hoy Rd, Ithaca, NY 14850
jf853@cornell.edu

MS90

Physics Informed Neural Networks for the Solution of Phase-Field Models of Brittle Fracture

We contrast the performance of a variety of formulations from the family of physics informed neural networks (PINNs) to the solution of phase-field models of crack propagation. Comparison of classical PINNs, variational PINNs (VPINNs) and variational energy based PINNs (VE-PINNs) to solve a simple 1D problem leads us to conclude that the variational approaches may be more accurate than the standard PINNs, whilst VPINNs is less straightforward to implement than VE-PINNs: leading to a preference for the latter. This is then deployed for the solution of more challenging phase-field problems in two dimensions and the performance is analysed in terms of both accuracy and robustness. Our conclusions suggest that, whilst accurate solutions are certainly achievable for this class of problem, further research is required to ensure the robustness and the generalization of the approach.

Peter K. Jimack
School of Computing
University of Leeds
p.k.jimack@leeds.ac.uk

Yousef Ghaffari Motlagh
University of Leeds
y.ghaffarimotlagh@leeds.ac.uk

Rene de Borst
University of Sheffield
r.deborst@sheffield.ac.uk

MS90

Machine Learning Constitutive Models of Inelastic Materials with Microstructure

Traditional simulations of complex physical processes, such as material deformation, are both crucial technologically and expensive computationally. Furthermore the development of physical models via traditional methods is particularly time-consuming in human terms. Developing comparably accurate models directly from data can enable rapid development of accurate models as well as more robust design, uncertainty quantification, and exhaustive structure-property exploration. We have been developing neural network models that are guided by traditional constitutive theory, such as tensor function representation theorems to embed symmetries, and also exploiting deep learning to infer intrinsic microstructural features. Neural networks are flexible since sub-components of their graph-like structure can be arranged to suit particular tasks, such as image processing and time integration, and represent the mechanistic flow of information. This talk will describe the architectures and demonstrate the efficacy of neural networks designed to model the response of complex history-dependent materials with pores, inclusions or grains based solely on observable data.

Reese Jones, Cosmin Safta, Kyle Johnson
Sandia National Laboratories

rjones@sandia.gov,
kyljohn@sandia.gov

csafta@sandia.gov,

MS90

Integrated Characterization and Calibration of Elastoplastic Constitutive Models

Calibration of constitutive models in solid mechanics remains both challenging and costly, despite the widespread availability of high-quality full-field deformation measurement techniques (e.g. digital image/volume correlation). Furthermore, the process of collecting data from a mechanical test and its ultimate utilization within a computational model calibration workflow are often loosely coupled. To address these issues, we propose an Interlaced Characterization and Calibration (ICC) paradigm that leverages concepts from Bayesian optimal experimental design (OED) to actively drive an experiment that is tailored to the accurate calibration of a given material. In this talk we will provide an overview of the ICC concept and its application to elastoplastic constitutive models and share technical progress on a few of its key components. These topics will include the construction of surrogate model graphs that replace expensive computational simulations of the experiment and the “step selection” OED problem, i.e. given a set of choices for the next material state and current posterior distribution, which one should we choose?

Daniel T. Seidl, Denielle Ricciardi, Brian Lester, Amanda Jones, Elizabeth Jones
Sandia National Laboratories
dtseidl@sandia.gov, dericci@sandia.gov,
btleste@sandia.gov, ajones1@sandia.gov,
emjones@sandia.gov

MS91

Learning the Evolution of Unknown Systems via Deep Neural Networks

Many phenomena in science and engineering are observable but not yet explainable. That is, we can observe solution data generated from many physical systems, but the actual physics, e.g. an ordinary or partial differential equation model, are unknown. In this case, developing a deep neural network based model that replicates the systems behavior is desirable. Hence in this talk, we will explore how to learn the time evolution of unknown ODE and PDE systems from their solution data using deep neural networks. The specific neural network architectures used are grounded in numerical methods for solving ODEs and PDEs. We also consider the case of partially observing the solution vector, where a time history of the observed variables are required.

Victor Churchill
The Ohio State University
churchill.77@osu.edu

Dongbin Xiu
Ohio State University
xiu.16@osu.edu

MS91

Machine Learning Methods for Particle-Based Plasma Computations

The development of accurate and efficient methods for transport computations is critical in magnetically confined plasmas for controlled nuclear fusion. In this presentation

we discuss how machine learning methods can help alleviate two aspects of this problem that are a bottle neck in current particle-based computations. The first one is magnetic field interpolation needed for particle tracking. We present a neural network-based divergence-free (NN-DivFree) interpolation method for arbitrary magnetic field training data points. The method is mesh-free and based on a data-driven reconstruction of the magnetic field using a feedforward neural network. The NN-DivFree method exhibits a significant reduction of the computational complexity compared to local splines while maintaining a small error in the divergence. The second problem is the estimation of the particle density, given the 6-dimensional coordinates of an ensemble of computed particle orbits. To address this problem, we propose the use of Normalizing Flows (NF) which are a family of generative models where both sampling and density evaluation can be efficient and exact. Our method is based on a pseudo-reversible neural network architecture, which improves the estimation by relaxing the invertibility constraint of the nonlinear transformation. Although NF can be computationally inefficient in very high-dimensional problems, numerical experiments show that this is not the case for the 6-D problem of interest.

Diego Del-Castillo-Negrete, Minglei Yang, Guannan Zhang
Oak Ridge National Laboratory
delcastillod@ornl.gov, yanggm@ornl.gov, zhangg@ornl.gov

MS91

Effective Deep Neural Network Architectures for Learning High-Dimensional Banach-Valued Functions from Limited Data

In the past few decades the problem of reconstructing high-dimensional functions taking values in abstract spaces from limited samples has received increasing attention, largely due to its relevance to uncertainty quantification (UQ) for computational science and engineering. These UQ problems are often posed in terms of parameterized partial differential equations whose solutions take values in Hilbert or Banach spaces. Impressive results have been achieved on such problems with deep learning (DL), i.e. machine learning with deep neural networks (DNN). This work focuses on approximating high-dimensional smooth functions taking values in reflexive and typically infinite-dimensional Banach spaces. Our novel approach to this problem is fully algorithmic, combining DL, compressed sensing (CS), orthogonal polynomials, and finite element discretization. We present a full theoretical analysis for DNN approximation with explicit guarantees on the error and sample complexity, and a clear accounting of all sources of error. We also provide numerical experiments showing that DNNs can produce accurate approximations on challenging Banach-valued benchmark problems.

Nick Dexter
Florida State University, U.S.
nick.dexter@fsu.edu

Ben Adcock
Simon Fraser University
ben_adcock@sfu.ca

Simone Brugiapaglia
Concordia University
simone.brugiapaglia@concordia.ca

Sabastian Moraga
Department of Mathematics
Simon Fraser University
smoragas@sfu.ca

MS91

Level Set Learning with Pseudo-Reversible Neural Networks for Nonlinear Dimension Reduction in Function Approximation

Due to the curse of dimensionality and the limitation on training data, approximating high-dimensional functions is a very challenging task even for powerful deep neural networks. Inspired by the Nonlinear Level set Learning (NLL) method that uses the reversible residual network (RevNet), in this paper we propose a new method of Dimension Reduction via Learning Level Sets (DRiLLS) for function approximation. Our method contains two major components: one is the pseudo-reversible neural network (PRNN) module that effectively transforms high-dimensional input variables to low-dimensional active variables, and the other is the synthesized regression module for approximating function values based on the transformed data in the low-dimensional space. The PRNN not only relaxes the invertibility constraint of the nonlinear transformation present in the NLL method due to the use of RevNet, but also adaptively weights the influence of each sample and controls the sensitivity of the function to the learned active variables. The synthesized regression uses Euclidean distance in the input space to select neighboring samples, whose projections on the space of active variables are used to perform local least-squares polynomial fitting. This helps to resolve numerical oscillation issues present in traditional local and global regressions.

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Yuankai Teng
University of South Carolina
yteng@email.sc.edu

Zhu Wang
Department of Mathematics
University of South Carolina
wangzhu@math.sc.edu

Anthony Gruber
Sandia National Laboratories
adgrube@sandia.gov

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

MS91

Reinforcement Learning for Discretization-Aware LES Models

Over the past few years, increasing efforts have been devoted to leveraging the recent advances in machine learning for the field of turbulence modeling. Most approaches in this field were based on Supervised Learning, for which artificial neural networks (ANN) are trained by means of a precomputed training dataset. However, for large eddy simulations (LES), this approach can cause instabilities

in practical simulations, since the dynamics of the discretized equations are not captured by the training process. An approach which avoids this pitfall is the Reinforcement Learning (RL) paradigm, which, in contrast to Supervised Learning, trains ANN by interacting directly with the discretized dynamical system. We demonstrate how the RL paradigm can be applied for turbulence modeling by presenting the novel RL framework Relexi, which couples flow solvers with the machine learning library TensorFlow, while leveraging modern high-performance computing resources. Relexi is applied to canonical flows, for which ANN are trained to adapt the coefficients of analytical LES models dynamically in space and time. We show that these data-driven LES models provide stable simulations, while outperforming traditional LES models in terms of accuracy. Thus, the proposed framework can provide a novel class of discretization-aware LES models, which can incorporate complex LES filter formulations and discretization effects by design.

Marius Kurz, Andrea D. Beck
Institute of Aerodynamics and Gas Dynamics
University of Stuttgart
marius.kurz@iag.uni-stuttgart.de, andrea.beck@iag.uni-stuttgart.de

MS92

Particle Filtering Method for Estimating Cell Membrane Permeability from Surface PH Data

The transport of gases through cell membranes is one of the fundamental processes for oxygen-based life forms. Measuring the membrane permeability directly is challenging, and one possible way to estimate it is to consider another process, the pH regulation of cells by carbon dioxide transport, and use the membrane surface pH as a proxy and to estimate the permeability by solving an inverse problem. In this talk, we discuss the computational and modeling challenges of this approach, and show how particle filtering can be used to solve the inverse problem.

Alberto Bocchinfuso
Case Western Reserve University
axb1094@case.edu

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and Statistic
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS92

Chemical Thermodynamics and the Mathematical Integration of Reaction Kinetics

Key advances in numerical methods often stem from translating physical requirements into concrete numerical guidelines, such as discrete entropy inequalities in compressible flow models. In the present work, we present nonlinear analysis tools that draw from Chemical Thermodynamics. Through Gibbs formalism, chemical thermodynamics provides a well-known theoretical expression for a reaction's chemical equilibrium constant in terms of reduced chemical potentials. A less-known yet extremely valuable result is that if this expression is implemented, mass action kinetic

models are consistent with the dynamical prescriptions of the 2nd law. For fixed-temperature chemical ODEs, this incurs a decreasing Helmholtz free energy. If the temperature is allowed to vary in accordance with conservation of energy (1st law), this incurs the popular statement of increasing entropy. These *nonlinear* prescriptions can and should be used to further develop temporal integration techniques for reaction kinetics. We demonstrate that these prescriptions hold even the equilibrium constants are approximated from data, by constructing the free energy and entropy implicitly defined by the data. With this structure established, one can examine the consistency of time integration schemes with these prescriptions. Using chemical potentials, one can compute the respective contributions of the kinetics model and of the temporal scheme to free energy/entropy variations, and develop more robust schemes.

Ayoub Gouasmi
University of Michigan
Department of Aerospace Engineering
gouasmia@umich.edu

Scott Murman
NASA Ames Research Center
scott.murman@nasa.gov

MS92

PeleSurface: A Low Mach Number Heterogeneous Chemistry Solver for Catalytic Combustion Simulations

Catalyst-assisted heterogeneous reacting flows are a ubiquitous and integral component of several chemical industries. As such, a thorough understanding of the interactions between the reacting flow and catalytic media is of significant value to both academia as well as industry. In this work, to enable high fidelity numerical explorations of catalytic combustion processes, we build upon the adaptive mesh refinement enabled, low Mach number reacting flow solver, PeleLM (<https://amrex-combustion.github.io/PeleLM/>), and present a new catalytic reacting flow solver called PeleSurface. The required pore-scale heterogeneous chemical kinetics and surface-species specific thermodynamic information adhere to the SURFACE-CHEMKIN standard and are implemented within the PelePhysics library (<https://github.com/AMReX-Combustion/PelePhysics>). Since the framework employed is suited for high performance exascale compute architectures, the proposed solver is expected to facilitate large scale numerical workloads of catalytic reacting flows with detailed reaction chemistry for problems of theoretical and practical relevance.

Anirudh Jonnalagadda, Pavan Kotakonda, Konduri Aditya
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
anirudhj@iisc.ac.in, venkatapkk@iisc.ac.in, konduriadi@iisc.ac.in

Konduri Tiwari, Piet Moeleker
Shell India Markets Private Limited, Bengaluru, India
shashank.s.tiwari@shell.com, piet.moeleker@shell.com

MS92

A Multiscale Framework for Chemical Kinetic Modeling via Reaxff and Ab-Initio Molecular Dy-

namics: Application to Hydrogen Combustion

Chemical kinetics (CK) are an integral part of hierarchical multiscale modelling of combustion systems which provides reaction mechanisms and rate constants, extracted from theoretical approaches such as transition state theory and molecular dynamics (MD), to enable continuum scale simulations. To extend CK models for extreme conditions (e.g., thermochemical non-equilibrium), MD simulations based on an empirical reactive force field (reaxff-based MD: RMD) or a quantum-mechanical description (density functional theory MD: DFT-MD) to capture the atomistic interactions, has a potential to easily incorporate the necessary complexities (e.g., imposing vibrational non-equilibrium). In this work, we propose a multiscale modelling workflow using reactive MD, and validate its performance based on simple hydrogen-oxygen (HO) combustion at equilibrium conditions. For HO system, we extract the complete reaction mechanism and the statistically computed rate constants. Arrhenius parameters are then fitted for each elementary reaction and a detailed CK model is generated for hydrogen oxidation. The parametrization obtained from both the MD approaches are evaluated by computing the flame speed of a one-dimensional premixed flame in continuum scale simulations and compared against standard results. A good agreement in flames speed is obtained at different equivalence ratios, indicating the robustness of the information extracted from the MD simulations.

Raj Maddipati

Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
rajmaddipati@iisc.ac.in

Phani S. Motamarri

Indian Institute of Science Bangalore
phanim@iisc.ac.in

Konduri Aditya

Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
konduriadi@iisc.ac.in

MS93**Nvidia Modulus: A Physics-ML Framework for Industry Scale Science and Engineering Problems**

High fidelity simulations are quintessential in the science and engineering domain. Like many applications, Machine Learning and AI can enhance the current ISV solvers and custom codes used by the enterprises and research groups. Specifically Physics based Machine learning (Physics-ML) can usher in AI based acceleration while meeting the demands of high fidelity simulations. In this session, we present Nvidia Modulus, an open Physics-ML Framework for the science and engineering community to use for their research and leverage the work of other researchers working on the development of Physics-ML approaches like Neural operators, PINNs etc. Nvidia Modulus provides a customizable training and inference pipeline that is abstracted away, so domain experts can use simple python based APIs for solving their problems. In this session, you will learn about the work we have done so far in enabling the community and the resources you can use to get started on applying AI and Physics ML approaches to your use cases.

Sanjay Choudhry
Nvidia

schoudhry@nvidia.com

MS93**Multi-Hierarchic Data-Driven Reduced Order Modeling Based on Mesh Simplification and Graph Convolutional Autoencoders**

Besides classic data-driven model order reduction (MOR) techniques like PCA, autoencoders (AEs) have been established as a nonlinear alternative. They are well-suited to find low-dimensional representations of high-dimensional systems and are a powerful tool in the creation of surrogate models but require an expensive offline training-phase. In addition, simulation models often consist of many degrees of freedom due to the modeling approach, but not due to the underlying requirements. Hence, they possess a higher resolution than necessary which leads to computational ballast. Accordingly, we propose to decrease the level of detail of the problem first to enable an efficient offline phase and hyperparameter tuning during training of AEs. Hereafter, a higher-detail surrogate model is built which is only refined in areas where necessary. For this, we use transfer-learning approaches to take advantage of already learned information. In detail, we work on down-sampled finite element models using mesh simplification. On a very coarse mesh, the first surrogate model is created using a graph convolutional autoencoder and a multi-layer perceptron to capture the model's parameter dependencies. Based on its approximation quality, some areas are refined by replacing up-sampling matrices with adaptive up-sampling networks. We demonstrate the capabilities of this approach on high-dimensional static as well as dynamic problems and compare it to classic AE-based MOR.

Jonas Kneifl

Institute of Engineering and Computational Mechanics
University of Stuttgart
jonas.kneifl@itm.uni-stuttgart.de

José Kutz

University of Washington
kutz@uw.edu

Jörg Fehr

Institute of Engineering and Computational Mechanics
University of Stuttgart
joerg.fehr@itm.uni-stuttgart.de

MS93**Inverse Design for Fluid-Structure Interactions Using Graph Network Simulators**

Designing physical artifacts that serve a purpose —such as tools and other functional structures —is central to engineering as well as everyday human behavior. Though automating design using machine learning has tremendous promise, existing methods are often limited by the task-dependent distributions they were exposed to during training. We showcase a task-agnostic approach to inverse design, by combining general-purpose graph network simulators with gradient-based design optimization. This constitutes a simple, fast, and reusable approach that solves high-dimensional problems with complex physical dynamics, including designing surfaces and tools to manipulate fluid flows and optimizing the shape of an airfoil to minimize drag. This framework produces high-quality designs by propagating gradients through trajectories of hundreds of steps, even when using models that were pre-trained for single-step predictions on data substantially different

from the design tasks. In our fluid manipulation tasks, the resulting designs outperformed those found by sampling-based optimization techniques. In airfoil design, they matched the quality of those obtained with a specialized solver. Our results suggest that despite some remaining challenges, machine learning-based simulators are maturing to the point where they can support general-purpose design optimization across a variety of fluid-structure interaction domains.

Tatiana Lopez Guevara
Deepmind
zepolitat@deepmind.com

MS93

Combining Autoencoders with DeepONet for Applications in Materials Processing and Climate

Phase-field modeling is an effective but computationally expensive method for capturing the mesoscale morphological and microstructure evolution in materials. Similarly in climate modeling, high-fidelity E3SM (Energy Exascale Earth System Model) simulations can resolve small flow structures but are very expensive. Hence, fast and generalizable surrogate models are needed to alleviate the cost of computationally taxing processes. The intrinsic discontinuous nature of the microstructures and the presence of high frequency features in climate prognostic variables make the training of the surrogate model cumbersome. In this study, we develop an algorithm that pairs a convolutional autoencoder with a deep neural operator (DeepONet) to learn the dynamic evolution of systems in a lower dimensional latent space. We demonstrate the effectiveness of the proposed hybrid framework through key examples in material science and climate modeling. For the first problem, we train the proposed model to learn the dynamic evolution of two-phase mixtures. The trained model is used to replace the high-fidelity simulation solver in interpolation tasks and is integrated with the numerical solver for accelerating extrapolation tasks. In the second problem, we train the model to predict corrections to the atmospheric states obtained by the coarse E3SM model. The trained model is able to forecast extreme events like hurricane and tropical cyclones and replaces the ‘nudging module in the E3SM.

Vivek Oommen, Aniruddha Bora, Khemraj Shukla
Brown University
vivek_oommen@brown.edu, aniruddha_bora@brown.edu, khemraj_shukla@brown.edu

Remi Dingreville
Sandia National Laboratories
rdingre@sandia.gov

Shixuan Zhang, Ruby Leung
Pacific Northwest National Laboratory
shixuan.zhang@pnnl.gov, ruby.leung@pnnl.gov

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS93

DeepOnet for Sub-Surface Flow Problems

Groundwater numerical models support the critical decision of issuing and reviewing permits such as water abstraction licences and enable forecasting the impact of land use

and climate change on groundwater resources. They typically use control volumes (cells) for solving the discretized groundwater flow equation. As high spatial and temporal resolution are generally required, the use in supporting planning and mitigation policies is often limited by computational cost. In this work, we explore the use of neural operator surrogates to mitigate this expense. We address this challenge by learning the distribution of hydraulic head in a confined aquifer with one or multiple pumping wells, which act as forcing terms on Darcy’s law. The heterogeneous aquifer is composed of various lenses of fine to medium sand, each with a different hydraulic conductivity. One or multiple wells are randomly placed in the domain and deep operator networks (DeepONet) learn the mapping from the hydraulic conductivity field and the forcing terms to the resulting groundwater head. Several application examples are deployed here with the aim of revealing how to best formulate the learning tasks for constructing a computationally inexpensive emulator for the steady-state and transient solutions of heterogeneous aquifers.

Maria Luisa Taccari
University of Leeds
cnmlt@leeds.ac.uk

He Wang
School of Computing
University of Leeds
h.e.wang@leeds.ac.uk

Jonathan Nuttall
Deltares
jonathan.nuttall@deltares.nl

Xiaohui Chen
School of Civil Engineering
University of Leeds
x.chen@leeds.ac.uk

Peter K. Jimack
School of Computing
University of Leeds
p.k.jimack@leeds.ac.uk

MS94

Tackling Advection-Dominated Problems Using Linear Optimal Transport

In model reduction of parametrized partial differential equations, Reduced Basis Methods are well-established for the calculation of solutions in low-dimensional linear spaces. However, advection-dominated problems are difficult to tackle this way.

In their work *Nonlinear model reduction on metric spaces*, Ehrlacher et al. (2020) consider PDE solutions as elements of the L^2 -Wasserstein space. Distributions are expressed as weighted means (in the sense of the Wasserstein distance) of a select few reference elements.

Using Linear Optimal Transportation (also known as Monge Embeddings), we extend this approach to more than one spatial dimension. In LOT, a distribution μ is mapped to the optimal transport map $T_{\sigma \rightarrow \mu}$ that takes a fixed reference distribution σ to μ . While the set of solutions $\{\mu_i\}_{i=1,2,\dots}$ typically exhibits slow n-width decay, that of the set of maps $\{T_{\sigma \rightarrow \mu_i}\}_{i=1,2,\dots}$ is much faster. This allows us to find a low-dimensional approximation of transport maps, which are elements of a linear space.

Since reconstruction of a solution via these reduced maps is not exact, we can improve the approximation by using

a reduced basis in the reference domain. This leads to a method in the spirit of snapshot remapping that is illustrated with numerical examples.

Tobias Blickhan

Max Planck Institute for Plasma Physics
tobias.blickhan@ipp.mpg.de

MS94

Machine Learning Enhanced a Posteriori Error Estimation for Model Reduction of Nonlinear Parametric Systems

In this work, we are interested in obtaining parametric reduced-order models (pROMs) for parametric nonlinear dynamical systems using snapshot-based model order reduction techniques. Residual-based *a posteriori* error estimators for pROMs are usually derived based on a given and known time discretization scheme for the spatially discretized full-order model. However, for many cases (such as stiff systems), the time integration is often performed using black-box solvers (such as `ode45` in MATLAB, `lsode` in Octave, etc.) that use automatic order and step-size control. Such systems pose a challenge for residual-based error estimation of their pROMs since the exact form of the time integration is unknown. To address this, we use machine learning to *learn* the mismatch between the true value of the residual and the one obtained by imposing a chosen time discretization scheme. Our approach aims at deriving a sharp *a posteriori* output error estimator for such cases. Using examples from process engineering (batch chromatography) and fluid dynamics (Burgers' equation), we show that our proposed approach is able to more accurately characterize the true output errors of pROMs for systems solved by black-box time integration solvers.

Sridhar Chellappa

Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg
chellappa@mpi-magdeburg.mpg.de

Lihong Feng

Max Planck Institute, Magdeburg, Germany
feng@mpi-magdeburg.mpg.de

Peter Benner

Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS94

Learning Nonlinear Reduced Models on Spectral Submanifolds from Data

We discuss a recent dynamical systems alternative to neural networks in the data-driven reduced-order modeling of nonlinear phenomena. Specifically, we show that the concept of spectral submanifolds (SSMs) provides very low-dimensional attractors in a large family of physical problems that range from wing oscillations to transitions in shear flows. A data-driven identification of the reduced dynamics on these SSMs gives a rigorous way to construct accurate and predictive reduced-order models in applications without the use of governing equations. We illustrate this on problems that include accelerated finite-element simulations of large structures, prediction of transitions in plane Couette flow, reduced-order modeling of fluid sloshing in a

tank, and model-predictive control of soft robots.

George Haller

ETH, Zurich
Switzerland
georgehaller@ethz.ch

MS94

Non-Linear Manifold Rom with Convolutional Autoencoders and Reduced Over-Collocation Method

Non-affine parametric dependencies, nonlinearities, and advection-dominated regimes of the model of interest can result in a slow Kolmogorov n -width decay, which precludes the realization of efficient reduced-order models based on Proper Orthogonal Decomposition. Among the possible solutions, there are purely data-driven methods that leverage nonlinear approximation techniques such as autoencoders and their variants to learn a latent representation of the dynamical system, and then evolve it in time with another architecture. Despite their success in many applications where standard linear techniques fail, more has to be done to increase the interpretability of the results, especially outside the training range and not in regimes characterized by an abundance of data. Not to mention that none of the knowledge on the physics of the model is exploited during the predictive phase. In this talk, in order to overcome these weaknesses, we implement a variant of the nonlinear manifold method introduced in previous works with hyper-reduction achieved through reduced over-collocation and teacher-student training of a reduced decoder. We test the methodology on a 2d nonlinear conservation law model and compare the results we would obtain with a purely data-driven method for which the dynamics is evolved in time with a long-short term memory network.

Francesco Romor

SISSA, International School for Advanced Studies
Trieste, Italy
fromor@sissa.it

Giovanni Stabile

SISSA mathLab, Trieste, Italy
giovanni.stabile@uniurb.it

Gianluigi Rozza

SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS94

Structure-Preserving Model Reduction for Transport-Dominated Port-Hamiltonian Systems

The framework of port-Hamiltonian (pH) systems provides a generalization of classical Hamiltonian systems which allows to describe internal energy dissipation as well as energy exchange with the environment. This system class comes with many desirable properties such as passivity and often also stability. Here, passivity refers, roughly speaking, to the inability of the system to internally generate energy out of nowhere. As a consequence, the pH framework may be used in the field of model reduction to ensure that the reduced-order models (ROMs) are stable and passive by employing structure-preserving methods. In this talk, we present a structure-preserving model reduction technique for port-Hamiltonian systems based on a recently proposed nonlinear model reduction approach. Especially, we consider a special nonlinear approximation

ansatz where the state of the full-order model is approximated by a linear combination of transformed modes. These transformations may be for instance achieved by suitable translation or shift operators and allow to obtain low-dimensional and accurate approximations of transport-dominated systems. In order to achieve structure preservation, we demonstrate that the ROMs may be constructed by minimizing the residual in a suitably weighted norm. In this context, we also discuss stability of the resulting ROMs. Finally, the new approach is illustrated by means of a numerical test case.

Philipp Schulze
Chemnitz University of Technology
philipp.schulze@math.tu-chemnitz.de

MS95

Combining Layer-Parallel Training with Multilevel Optimization Techniques

Training of deep neural networks is computationally costly due to slow convergence of optimizers like stochastic gradient descent, coupled with limited parallelism in terms of number of layers. The later point implies that there is an upper bound on the strong scaling of these methods. However, recent advances have pursued multilevel approaches that tackle each of these issues independently. Convergence has been accelerated using algorithms based on the Multigrid Optimization (MGOpt) framework. These approaches in combination with a trust-region based search direction can reduce the number of epochs required to train. On the other hand, a Layer-Parallel scheme that applies multigrid-in-time (MGRIT) approaches for both forward and back propagation has shown order of magnitude level parallel speedups for various neural network architectures. The proposed MGOpt algorithm and the MGRIT strategy both use an identical coarsening in time algorithm. This talk will consider leveraging this similarity to improve the parallel performance of the MGOpt approach using a Layer-Parallel decomposition. To demonstrate the performance of this methodology we will apply our approach to standard problems from the deep learning literature.

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratories
eccyr@sandia.gov

Nicholas Moore
Computational Mathematics Department
Sandia National Laboratories
nsmoore@sandia.gov

Jacob B. Schroder
Department of Mathematics and Statistics
University of New Mexico
jbschroder@unm.edu

MS95

Domain Decomposition Training Strategies for Physics-Informed Neural Networks

Physics-informed neural networks (PINNs) are a solution method for solving boundary value problems based on differential equations (PDEs). The key idea of PINNs is to incorporate the residual of the PDE as well as boundary conditions into the loss function of the neural network. This provides a simple and mesh-free approach for solving problems relating to PDEs. However, a key limita-

tion of PINNs is their lack of accuracy and efficiency when solving problems with larger domains and more complex, multi-scale solutions. In a more recent approach, Finite Basis Physics-Informed Neural Networks (FBPINNs), the authors use ideas from domain decomposition to accelerate the learning process of PINNs and improve their accuracy in this setting. In this talk, we show how Schwarz-like additive, multiplicative, and hybrid iteration methods for training FBPINNs can be developed. Furthermore, we will present numerical experiments on the influence on convergence and accuracy of these different variants.

Victorita Dolean Maini
University of Strathclyde, UK and Université Côte d'Azur
victorita.dolean@strath.ac.uk

Alexander Heinlein

Delft University of Technology
Delft Institute of Applied Mathematics
a.heinlein@tudelft.nl

Benjamin Moseley
ETH Zürich
benjamin.moseley@ai.ethz.ch

MS95

Towards Large-Scale Training of Deep Neural Networks Using Domain-Decomposition Methods

Deep neural networks (DNNs) are routinely used in a wide range of application areas and scientific fields, as they allow to efficiently predict the behavior of complex systems. However, before the DNNs can be effectively used for the prediction, their parameters have to be determined during the training process. Traditionally, the training process is associated with the minimization of a loss function, which is commonly performed using variants of the stochastic gradient (SGD) method. Although SGD and its variants have a low computational cost per iteration, their convergence properties tend to deteriorate with increasing network size. In this talk, we aim to alleviate the training cost of DNNs by leveraging the nonlinear domain decomposition strategies. The subdomains will be constructed by exploiting two complementary approaches, namely the decomposition of the data and the parameter space. This will give rise to two new classes of training methods, convergence properties of which will be analyzed using a series of numerical experiments. Comparison with state-of-the-art optimizers will be also performed, demonstrating the efficiency of our novel training methods. Moreover, we will discuss our parallelization strategy, its shortcomings, and possible future directions.

Alena Kopanicakova
Università della Svizzera italiana
Switzerland
alena.kopanicakova@brown.edu

Hardik Kothari
Università della Svizzera italiana
hardik.kothari@usi.ch

Rolf Krause
Università della Svizzera italiana, Switzerland
rolf.krause@usi.ch

MS95

Learning the Constraints in Adaptive Feti-Dp Do-

main Decomposition Methods

Domain decomposition methods (DDMs) are highly scalable, iterative solvers for the solution of large systems of linear equations, e.g., arising from the discretization of PDEs. The convergence rate of classic DDMs in general deteriorates severely for coefficient distributions with large contrasts in the coefficient function. To retain the robustness for such problems, the coarse space of the DDM can be enriched by additional coarse basis functions, often obtained by solving local generalized eigenvalue problems. However, the set-up and the solution of these eigenvalue problems typically takes up a significant part of the total time to solution. Additionally, for many realistic problems, only the solution of a small number of the eigenvalue problems is necessary to design a robust algorithm. In general, it is difficult to predict a priori which of the eigenvalue problems are needed. Using a neural network model we can predict the geometric location where eigenvalue problems have to be solved, often reducing its number significantly (joint work of the authors with Alexander Heinlein). To obtain such an a priori classification, we use a mesh-independent sampling strategy which is comparable to an image recognition problem. In a next step, we train a surrogate model which directly learns the necessary coarse basis functions themselves using again an image representation of the underlying coefficient function. Numerical results indicate the robustness of this approach.

Axel Klawonn
Universität zu Köln
axel.klawonn@uni-koeln.de

Martin Lanser
Universitaet zu Koeln
Mathematisches Institut
martin.lanser@uni-koeln.de

Janine Weber
Universität zu Köln
Department of Mathematics and Computer Science
janine.weber@uni-koeln.de

MS96

A Bayesian Approximation Error Approach for Design of Inverse Problems under Uncertainty

We consider optimal experimental design for infinite-dimensional Bayesian nonlinear inverse problems governed by PDEs. In addition to the inversion parameters, the governing PDEs typically have additional model parameters that are not known exactly and are uncertain. Inversion and design of experiments must account for these secondary model uncertainties. This talk considers methods for optimal experimental design for such inverse problems that are robust with respect to modeling uncertainties. Our approach builds on the Bayesian Approximation Error framework, which facilitates incorporating modeling uncertainties in the inverse problem. This is coupled with structure exploiting algorithms that enable computing optimal designs at a cost that does not grow with the dimension of the discretized inversion and secondary parameter dimensions.

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Noemi Petra
University of California, Merced

npetra@ucmerced.edu

Ruanui Nicholson
Department of Engineering Science
University of Auckland
ruanui.nicholson@auckland.ac.nz

MS96

Global Solution of Some Nonconvex Variational Problems and Some Optimal PDE Control Problems

This talk will present a computational scheme to approximate with guaranteed convergence global optimizers of some nonconvex integral variational problems and of some optimal control problems for semilinear elliptic partial differential equations (PDEs). The approach assumes the integral cost and any PDE constraints are described by polynomials, and is based on a discretize-then-relax strategy that leverages tools from polynomial optimization. First, the original variational or PDE control problem is discretized into a sparse polynomial optimization problem using a bounded finite-element scheme. This polynomial optimization problem is then relaxed into a hierarchy of convex and highly structured semidefinite programs (SDPs), which in principle can be solved using off-the-shelves software. We prove that solutions of such SDPs will converge to global optimizers of the original variational or PDE control problem at least when the latter are unique. Note that this assumption does not preclude the existence of other stationary points, including local minimizers. The potential of this computational strategy will be illustrated on a range of examples. Remaining challenges for both analysis and computation will be discussed if time permits.

Giovanni Fantuzzi
Aeronautics
Imperial College London
giovanni.fantuzzi@fau.de

Federico Fuentes
Institute for Mathematical and Computational
Engineering
Pontificia Universidad Católica de Chile, Santiago
federico.fuentes@uc.cl

MS96

Randomized Algorithms for Design of Experiments in High-Dimensions

We propose and analyze batch greedy heuristics for cardinality constrained maximization in the context of optimal experimental design for linear Bayesian inverse problems. Our approach shows particular promise for cardinality constraint experimental design in high dimensions. We demonstrate our theoretical findings on synthetic problems and on a real-world climate monitoring example.

Jayanth Jagalur-Mohan, Youssef M. Marzouk
Massachusetts Institute of Technology
jagalur@mit.edu, ymarz@mit.edu

MS96

Optimal Experimental Design for Bayesian Inverse Problems Using Transport Maps

For Bayesian inverse problems, the task of choosing designs or experimental conditions for 'optimal' inference of

unknown parameters requires optimizing some expected utility function or optimality criterion that assesses the effectiveness of any design. In contrast to Bayesian inverse problems governed by linear parameter-to-observable maps, a closed-form expression for the optimality criterion is typically not available for nonlinear inverse problems because the posterior distribution is generally non-Gaussian. Common approaches for dealing with this challenge rely on: (i) reformulation or approximation of the optimality criterion via linearization of the problem, or, (ii) sampling-based Monte Carlo estimates of the expected utility. We follow the latter approach and focus on designs maximizing the expected information gain from prior to posterior, quantified by the Kullback-Leibler divergence. Using transport maps, we present a computationally tractable approach for finding optimal designs for Bayesian inverse problems governed by nonlinear parameter-to-observable maps. More precisely, we approximate the joint density of the design, observable, and parameter random variables as the pushforward of a reference probability density, allowing real-time sampling from the (approximate) posterior distribution for any realization of data and design. Additionally, we touch upon possible extensions of our method to sequential optimal experimental design.

Karina Koval, Roland Herzog
Heidelberg University
karina.koval@iwr.uni-heidelberg.de,
roland.herzog@iwr.uni-heidelberg.de

Robert Scheichl
Universitaet Heidelberg
r.scheichl@uni-heidelberg.de

MS96

Coupling Optimal Experimental Design and Optimal Control

Optimal experimental design (OED) allows us to determine, a-priori, how to collect the most informative experimental data. Traditional OED approaches focus on optimizing the solution to the Bayesian inverse problem, which is often an intermediate step to the true objectives of making accurate model predictions and determining optimal control policies. In this work, we provide a novel formulation of a Bayesian OED problem that couples the goal-oriented experimental design and optimal control problems, thus minimizing uncertainty in the control objective directly. Through a numerical example, we illustrate how the resulting control-oriented OED approach leads to a better allocation of experimental resources that decreases uncertainty regarding the optimal control policy in comparison to classical OED strategies.

Bart G. Van Bloemen Waanders, Rebekah White
Sandia National Laboratories
bartv@sandia.gov, rebwhit@sandia.gov

MS97

Frequency-Domain Based Learning of Dynamical Systems from Purely Time-Domain Data

Classically, constructing reduced order models (ROMs) of large scale dynamical systems has required access to internal system matrices. When these quantities are not available, data driven approaches can be used that require access only to input-output data. Frequency-based data-driven methods have been very successful in creating high fidelity ROMs from data, but require access to values

(and sometimes derivatives) of the transfer function. These quantities can at times be costly or difficult to obtain, but one may have ample access to time-domain input-output data. In 2020, Burohman et al. introduced a framework to directly calculate transfer function values and derivatives using only time-domain data. We first discuss improvements to this method that allow for a more efficient and robust numerical implementation, which lead to a theoretical result on optimal diagonal scaling of a class of matrices. We then use these improvements to develop an algorithm that extends H2-optimal approximation to purely time-domain data. We also compare the H2-optimal algorithm's performance to other well established frequency based ROM techniques (such as Loewner framework and Vector Fitting) using only time-domain data.

Michael S. Ackermann
Virginia Tech
amike98@vt.edu

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS97

End-to-End Learning of Dynamical Systems with the Mori-Zwanzig Formalism

Deep learning approaches are extensively applied in order to capture the underlying dynamics of a physical phenomenon but often lack interpretability, performance, robustness guarantees and insights into the structure of the dynamics. In this work, we formulate the problem within the Mori-Zwanzig (MZ) formalism, a technique that arose in the field of statistical mechanics that allows to derive exact evolution equation for coarse-grained quantities of interest from high-dimensional systems. Not only this approach offers, by design, an explainable structure by separating the dynamics into a Markovian and a memory term but it allows a principled use of the available observational data by providing a physically consistent functional structure to learn. A neural network approximation format and efficient numerical integrator are coupled together, leading to a flexible and powerful representation model compliant with the MZ framework. The method is demonstrated on linear systems for which an analytical derivation of the MZ formalism is available allowing rigorous comparison; as well as on nonlinear, possibly chaotic, systems such as Van der Pol oscillator and the Kuramoto-Sivashinsky model, exhibiting a superior approximation performance in terms of data requirement and generalizability.

Thibault Monseil
LISN, CNRS
thibault.monseil@universite-paris-saclay.fr

Lionel Mathelin
LISN, CNRS, Univ. Paris-Saclay
lionel.mathelin@lisn.upsaclay.fr

Onofrio Semeraro
LISN-CNRS, Univ. Paris-Saclay, France
onofrio.semeraro@universite-paris-saclay.fr

Guillaume Charpiat
INRIA

guillaume.charpiat@inria.fr

MS97

Least Squares Rational Approximation: An Optimization-Based Bridge Between the Loewner and Astolfi Frameworks

Moment matching (MM) methods provide an efficient way to address the Model Order Reduction (MOR) problem (Antoulas, Beattie, Gugercin'21). The basic idea is to use (rational) interpolation theory to approximate the transfer function of a system with another transfer function of lower degree. The Loewner framework (Antoulas07) offers a classical solution to achieve MM: all tangential interpolation conditions are first imposed simultaneously; a reduced order model (ROM) is then obtained by truncation of the Loewner matrices via a Singular Value Decomposition (SVD). About a decade ago, an alternative approach has been developed by Astolfi (Astolfi'10) by identifying moments of a system with steady-state responses. In this context, (one-sided) interpolation conditions are imposed by ensuring that the steady-state response of the ROM coincides with that of the original system. Recently, this approach has been extended in (Padoan'21) by requiring that the interpolation conditions imposed by MM are satisfied only in a least squares sense. Building on these ideas, we discuss a new MOR method based on the solution of a separable nonlinear least squares problem with separable nonlinear equality constraints, which can be efficiently solved using, e.g., the Variable Projection (VARPRO) approach in (Golub, Pereira73). We also link the proposed approach with the Astolfi framework and the Loewner framework, thus providing an optimization viewpoint on the effect of SVD-based truncation.

Alberto Padoan
ETH Zurich
apadoan@ethz.ch

Ion Victor Gosea
Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

MS97

Data-Driven Balancing for Acoustical Systems

In engineering acoustics, one oftentimes encounters systems that are linear but have very complex dynamics. Although the underlying physics are usually well understood, adequately modelling these systems with constructive or physically motivated approaches is challenging in practice because e.g. material properties or domain geometries cannot be known at the required level of fidelity. Additionally, the model has to encompass the frequency span of human hearing - about three orders of magnitude. Hence, measurements have proven to be indispensable when dealing with real-world acoustical systems which makes high dimensional measurement data abundantly available. Classical realization methods, such as the Eigensystem Realization Algorithm, rely on the computation of orthogonal decompositions like SVD which is prohibitively expensive for the data at hand. As a remedy, these realization methods can be endowed with recent methods from randomized linear algebra for efficiently approximating orthogonal decompositions of large matrices. In this way, it is possible to construct accurate reduced order models from high dimensional data efficiently. In turn, this sheds new light onto many acoustical modelling challenges. The aim of this talk is to showcase and validate several contemporary exten-

sions and augmentations to established realization methods and analyze their aptitude and performance in practical applications that involve high dimensional measurements of real systems.

Art Pelling, Ennes Sarradj
TU Berlin
Department of Engineering Acoustics
a.pelling@tu-berlin.de, ennes.sarradj@tu-berlin.de

MS97

Data-Driven Balancing of Continuous-Time Systems via the Hankel Operator

Balanced truncation (BT) is one of the most successful and commonly employed system-theoretic model reduction methods. The foundations of BT lie in the Hankel operator and its singular values. For continuous-time linear dynamical systems, BT preserves system stability in the reduced model and provides a priori bounds for the model reduction error. Although BT is generally conceived as an intrusive method, recent developments in [Gosea/Gugercin/Beattie, SIAM J. Sci. Comput., 2022] have shown that it can be reformulated as a non-intrusive, data-driven process. In this work, we study the connection between data-driven balancing and the Hankel operator associated with a continuous-time linear dynamical system. We develop data-driven approximations of the Hankel operator based solely on input-output data and then show how a balanced reduced order realization can be constructed using this approximation. Based on these results, we propose an extension of the Eigensystem Realization Algorithm to a large class of continuous-time systems.

Igor Pontes Duff
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
pontes@mpi-magdeburg.mpg.de

Ion Victor Gosea
Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

Christopher A. Beattie
Virginia Polytechnic Institute and State University
beattie@vt.edu

MS98

An Efficient Iterative Solver for the GEM Model

Efficient and highly scalable numerical algorithms are increasingly needed to solve atmospheric fluid dynamics equations in weather forecasting models. In the Canadian GEM weather forecasting model, a system of nonlinear equations is linearized around a basic state and then reduced to an elliptic boundary value (EBV) problem. The large linear system arising from the discretization of the EBV problem is solved either by a parallel direct method or by a parallel iterative method based on a variant of Krylov subspace methods, namely the Flexible GMRES (FGMRES) algorithm. The operational GEM parallel solver account for a significant amount of the total model run-time, pointing to the need for a performance upgrade. Recently the efficiency and the parallel scalability of the GEM model

iterative solver have been improved by the implementation of a suitable preconditioner, based on the Restricted Additive Schwarz, leading to the FGMRES convergence rate acceleration, therefore to a reduction in the number of iterations. Furthermore, the use of a newly reformulated Arnoldi-Gram-Schmidt algorithm (Swirydowicz et al., 2019) resulted in a small number of global reductions and synchronizations (communication) required by the inner products.

Rabah Aider
Environment Canada
rabah.aider@ec.gc.ca

Abdessamad Qaddouri
Envi
abdessamad.qaddouri@ec.gc.ca

Stéphane Gaudreault
Environment Canada
stephane.gaudreault@ec.gc.ca

MS98

Anderson Acceleration on Emerging Architectures

Anderson Acceleration (AA) is a method to accelerate the convergence of fixed point iterations for nonlinear, algebraic systems of equations. Due to the requirement of solving a least squares problem (LSP) at each iteration, every iteration of AA requires some number of synchronization steps for global reductions. Performance of AA at scale on distributed systems is typically dependent upon the cost of solving this LSP or application of the underlying fixed point function. In this talk, we highlight approaches to low synchronization orthogonalization routines to reduce the overhead of the LSP solve and note that when performance is dependent upon the fixed point function evaluations, the reduction in synchronizations has little impact on overall performance. To reduce the overhead of costly function evaluations, there has been recent work done to modify the AA algorithm itself, such as alternating AA or non-stationary AA, methods which result in more rapid convergence than standard AA, thus reducing the number of function evaluations required for convergence. Overall, we explore several of these advances within the field of AA and their potential in developing scalable AA solvers on emerging supercomputer architectures.

Shelby Lockhart
University of Illinois at Urbana-Champaign
sll2@illinois.edu

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

Carol S. Woodward, David J. Gardner
Lawrence Livermore National Laboratory
woodward6@llnl.gov, gardner48@llnl.gov

Stephen Thomas
Advanced Micro Devices (AMD)
stephethomas@gmail.com

MS98

Integrating Ginkgos Distributed Linear GPU

Solvers for CFD Applications

To serve the ever-increasing demand for computing resources, many HPC systems are equipped with multiple highly parallel general-purpose GPUs. In most cases, the GPUs are integrated as discrete server-type GPUs that are attached to the nodes as coprocessors and provide the lions share of the theoretical compute performance. This concept allows to significantly increase the theoretical compute power, however, it poses a challenge to the scientific software developers that need to redesign their software such that it can benefit from the additional resources. An attractive strategy is to use the GPUs for the computationally most expensive part of an application. For computational fluid dynamics (CFD) simulations, a substantial part of the overall effort is spent on the solution of the system of linear equations arising from the discretized transport equations. This paper investigates how the currently ongoing effort within the open-source linear algebra library Ginkgo towards distributed GPU computing can be used as a backend for the CFD framework OpenFOAM to improve the simulation performance by offloading the linear algebra computations to GPUs.

Gregor Olenik
Steinbuch Centre for Computing (SCC)
Karlsruher Institut für Technologie (KIT)
gregor.olenik@kit.edu

MS98

GPU-Resident Sparse Factorizations

We present a set of preprocessing steps, symbolic and numerical factorization algorithms on GPU-resident matrices well-suited for the solution of sparse linear systems producing only small amounts of fill-in, where classical supernodal approaches fail to provide adequate performance. We present the performance characteristics of the heavily latency-bound factorization kernels on NVIDIA and AMD GPUs and suggest alternatives to established approaches where feasible.

Tobias Ribizel, Fritz Göbel
Karlsruhe Institute of Technology
tobias.ribizel@kit.edu, fritz.goebel@kit.edu

MS98

GPU Acceleration of the Pardiso 8.0 Sparse Direct Solver

Sparse direct solvers are a central component of many engineering analysis codes and often represent most of the full analysis applications run-time. An instance of this is in electronic design automation or in power grid optimization that requires a very large number of iteration steps such that simulation can take as long as several months. The matrices involved are characterized as small (100k-1M rows), extremely sparse, have little fill-in, and are poorly conditioned. Even though these matrices are considered low effort, their numerical factorization is the simulation bottleneck. This work pursues the GPU acceleration of sparse direct solvers tailored for such applications. As the size and sparsity of these matrices do not favor acceleration through BLAS-3 operations (they have very few supernodes large enough to show a performance benefit) an approach primarily leveraging the combination of batched operations, distribution across multiple GPUs, and leveraging a constant sparsity pattern over many solves, is followed. This is implemented in the PARDISO 8.0 solver with its use being demonstrated

in NXP's Mica analysis code and in the IPOPT optimization code. Details of the implementation will be discussed. Performance benchmarks will be presented along with comparisons to the performance of CPU-only versions of PAR-DISO 8.0 as well as other publicly-available solvers. Its use in Mica and associated impact on full-application performance will also be shown.

Olaf Schenk

Panua Technologies, Lugano
olaf.schenk@panua.ch

Steven C. Rennich
NVIDIA
srennich@nvidia.com

Kiran Gullapalli
NXP Semiconductors
kiran.gullapalli@nxp.com

Andreas Waechter
Panua Technologies, Lugano, Switzerland
andreas.waechter@northwestern.edu

MS99

Innovative Space-Time Domain Decomposition Method Solving Uncertainty Quantification Models

We present an innovative mathematical model and the related numerical algorithms, based on the simultaneous introduction of space-time decomposition on the PDEs governing the physical model and on the DA model. The core of our approach is that the DA model acts as coarse predictor operator solving the local PDE model, by providing the background values as initial conditions of the local PDE models. To enforce the matching of local solutions on overlapping regions, local problems are modified by adding a correction term keeping track of contributions of adjacent subdomains to overlapping regions. Such a correction term balances localization errors along overlapping regions, acting as a regularization constraint on local solutions. Moreover, in our approach both the coarse and the fine solvers run concurrently from the beginning. Consequently, the resulting algorithm requires only the exchange of boundary conditions between adjacent subdomains. Algorithm scalability is measured in terms of strong scaling and of weak scaling. A mechanism for dynamically balancing the loads has been included in the parallel framework for solving large-scale DA models. In particular, we refer to the introduction of a dynamic redefining of initial DD in order to deal with problems where the observations are non uniformly distributed and general sparse, a quite common issue in DA.

Luisa D'Amore

University of Naples Federico II Naples, ITALY
Mathematics and Applications
luisa.damore@unina.it

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS99

First Steps Toward a Generic Tool for Comparing Parallel Efficiency of Iterative Parallel-in-Time Al-

gorithms

To increase parallel concurrency for the simulation of time-dependent problems, many ideas for time parallelization algorithms have been developed the last two decades. Iterative parallel-in-time (PinT) methods, like Parareal, MGRIT, and PFAST, have received the most attention. Since those algorithms have different variants and can depend on many parameters, there is many ways for the scientific community to use them for PinT simulations. The optimal approach for one given problem is then quite difficult to determine. Parallel efficiency of iterative PinT algorithms relies on two aspects : the number of iterations required to get the PinT error below an acceptable level (convergence) and the computational cost for doing all those iterations, compared to a sequential simulation (speedup). Recently, a new analysis approach has been introduced to study and compare the convergence analysis of all iterative PinT methods in a common framework. On another side, recent research work have investigated how to compare the parallel performance of PinT methods using a common approach. In this talk, we present our recent work in trying to merge both analysis (convergence & speedup), in order to develop a generic tool that can be used to compare the parallel performance of several PinT methods on simple problems, and provide indications on optimal approach for a given type of problem.

Thibaut Lunet

Hamburg University of Technology
thibaut.lunet@tuhh.de

Jens Hahne
University of Wuppertal
jehahne@uni-wuppertal.de

MS99

Performance of Multigrid Reduction in Time (MGRIT) on GPU Machines

Multigrid reduction in time (MGRIT) is an optimal, scalable, multilevel parallel in time algorithm that has been shown to achieve significant speedup vs. sequential time stepping on massively parallel CPU machines. This speedup is attained through the use of additional computational resources and only occurs after a "cross-over point" where sufficient parallel resources are concurrently used. This talk examines how the move to GPU architectures impacts speedup, the crossover point, and the best choice of parallel decomposition of the space-time domain for MGRIT compared to CPU architectures.

Wayne Mitchell

Lawrence Livermore National Laboratory
mitchell82@llnl.gov

MS99

GParareal: (Towards) a Probabilistic Time-Parallel ODE Solver

Parareal is a well-studied (deterministic) time-parallel algorithm designed to integrate initial value problems (IVPs) by combining solutions from cheap (coarse) and expensive (fine) numerical integrators using a predictor-corrector (PC). We present *GParareal*, a Parareal-type algorithm which models the correction term in the PC using a Gaussian process emulator, trained on all previously collected fine and coarse solution information. For a number of test IVPs, we demonstrate that *GParareal* locates solutions in

faster wallclock time than Parareal and can locate solutions in situations where Parareal fails. We highlight that GParareal has the advantageous capability to re-use legacy solution data, e.g. solutions from prior simulations of the IVP for different initial conditions, to further accelerate convergence of the method - something that existing time-parallel methods cannot do.

Kamran Pentland, Massimiliano Tamborrino, Tim Sullivan

University of Warwick
kamran.pentland@warwick.ac.uk,
massimiliano.tamborrino@warwick.ac.uk,
t.j.sullivan@warwick.ac.uk

James Buchanan, Lynton Appel
Culham Centre for Fusion Energy
james.buchanan@ukaea.uk, lynton.appel@ukaea.uk

MS99

On the Combination Between Space-Time Model Order Reduction Methods and Deep Learning for Haemodynamic Problems

Deep learning (DL) techniques have nowadays become central in a number of research areas. However, a well-known downside of DL-based methods consists in the huge amount of data required, which often hinders their application. Recent works in the scientific machine learning area have demonstrated that the data greed bottleneck can be circumvented by enforcing physical awareness. Indeed, if DL models are enriched with knowledge of the physical laws underlying the phenomenon of interest, they can attain remarkable levels of accuracy even in a small data regime and, moreover, their generalization capabilities can be improved. In this talk, we present a physics-informed model - called ST-RB-DNN - that encodes scattered values of solutions to unsteady parametrized PDEs to the underlying parameter values, leveraging model order reduction techniques. In particular, we employed space-time reduced basis (ST-RB) methods, since they allow for dramatic lightening of the computational burden, by means of dimensionality reduction along both the spatial and the temporal dimensions. For the numerical tests, we consider blood flow - modeled via the Navier-Stokes equations - in different patient-specific 3D geometries. In all cases, the ST-RB-DNN model provides good estimates of the underlying parameter values. Further extensions of this work involve the implementation of RFSI models to account for the vessel wall compliance and of ad hoc boundary conditions to better model coronary flow.

Riccardo Tenderini, Simone Deparis
Ecole Polytechnique Federale de Lausanne
riccardo.tenderini@epfl.ch, simone.deparis@epfl.ch

MS100

Randomized Maximum Likelihood via High-Dimensional Bayesian Optimization

Posterior sampling for high-dimensional Bayesian inverse problems where the log-likelihood has a low-dimensional structure, known as an active subspace, is a problem commonly faced in real-world applications. Existing approaches assume that the computational budget is sufficient to estimate the active subspace, either via gradient-based methods or using a large number of simulator evaluations. Here we tackle the more challenging (and practically relevant) case where we do not have sufficient com-

putational budget to satisfactorily estimate the active subspace. We develop a high-dimensional Bayesian optimization approach to solve the Randomized Maximum Likelihood (RML) problem. RML is an approximate posterior sampling methodology based on multi-objective optimization, first developed for petroleum engineering applications. By sharing data between the different objective functions, we are able to implement RML at a greatly reduced computational cost compared to existing methods, allowing us to efficiently sample from the posterior distribution of the inverse problem. We demonstrate the benefits of this approach in comparison to alternative optimization methods on a variety of synthetic and real-world problems, including medical and fluid dynamics applications. Furthermore, we show that the samples produced by our method cover well the high-posterior density regions in all of the experiments.

Valentin Braez, Richard Wilkinson
University of Nottingham
valentin.braez@nottingham.ac.uk,
r.d.wilkinson@nottingham.ac.uk

MS100

Convergence of an Empirical Galerkin Method for Parametric PDEs

Weighted least squares methods have been examined thoroughly to obtain quasi-optimal convergence results for a chosen (polynomial) basis of a linear space. A focus in the analysis lies on the construction of optimal sampling measures and the derivation of a sufficient sample complexity for stable reconstructions. When considering holomorphic functions such as solutions of common parametric PDEs, the anisotropic sparsity they exhibit can be exploited to achieve improved results adapted to the considered problem. In particular, the sparsity of the data transfers to the solution sparsity in terms of polynomial chaos coefficients. When using nonlinear model classes, it turns out that the known results cannot be used directly. To obtain comparable a priori rates, we introduce a new weighted version of Stechkin's lemma. This enables to obtain optimal complexity results for a model class of low-rank tensor train networks. We also show that the solution sparsity results in sparse component tensors. Numerical examples illustrate the theoretical findings and demonstrate a remarkable performance of the derived RALS algorithm in comparison to existing state-of-the-art algorithms.

Martin Eigel
WIAS Berlin
martin.eigel@wias-berlin.de

Anthony Nouy, Philipp Trunschke
Ecole Centrale Nantes
anthony.nouy@ec-nantes.fr, philipp.trunschke@univ-nantes.fr

MS100

Generating Problem-Adapted Basis Functions Parallel in Time via Random Sampling

To tackle time-dependent partial differential equations (PDEs) with coefficients that are rough in both space and time, we construct reduced ansatz functions defined in space that can be combined with time stepping schemes, e.g., within model order reduction methods. As a key new contribution, we propose to construct these ansatz functions in an embarrassingly parallel and local manner in

time by selecting important points in time and only performing local computations on the corresponding local time intervals. In detail, we perform several simulations of the PDE for only few time steps in parallel, starting at different, randomly drawn start time points, prescribing random initial conditions. Applying a singular value decomposition to a subset of the so obtained snapshots yields the reduced ansatz functions. To select suitable start points in time, we suggest using data-driven sampling strategies from randomized numerical linear algebra such as leverage score sampling. By solving the PDE locally in time with random initial conditions, we construct local ansatz spaces in time that converge provably at a quasi-optimal rate and allow for local error control. Numerical experiments demonstrate that the proposed approach can outperform existing methods like the proper orthogonal decomposition even in a sequential setting and is well capable of approximating advection-dominated problems.

Julia Schleuß

University of Muenster
julia.schleuss@uni-muenster.de

Kathrin Smetana

Stevens Institute of Technology, U.S.
ksmetana@stevens.edu

Lukas ter Maat

University of Twente
lukas.tm@hotmail.nl

MS100

Randomized Local Model Order Reduction for Nonlinear Partial Differential Equations

Applications that require multiple simulation requests or a real-time simulation response are ubiquitous in science and engineering. Model order reduction methods in which the problem is (approximately) solved in a carefully chosen subspace of the high-dimensional discretization space have been developed to tackle such problems. However, for many large-scale problems, and especially problems that exhibit multiscale features, full order solves are not affordable in a reasonable time frame. Localized model order reduction methods decompose the global computational domain into subdomains, build local reduced models from solutions of the partial differential equation (PDE) on the subdomains, and use some coupling to compute a global reduced approximation. For the local model order reduction methods that we will present in this talk we can control the global approximation error even if the global computational domain is unknown when the local reduced models are constructed, facilitating a LEGO brick like assembly of the global computational domain, imperative for many applications such as digital twins. While there has been a significant progress in recent years for the construction of local reduced models for linear PDEs, very few results have been obtained so far for nonlinear PDEs. In this talk, we will show how randomized methods and their probabilistic numerical analysis can be exploited for the construction and numerical analysis of local reduced models for nonlinear PDEs.

Kathrin Smetana

Stevens Institute of Technology, U.S.
ksmetana@stevens.edu

Tommaso Taddei

INRIA Bordeaux South-West

tommaso.taddei@inria.fr

MS101

A Stochastic ADMM Algorithm for Large-Scale Image Ptychography with Weighted Difference of Anisotropic and Isotropic Total Variation

Ptychography is an imaging technique that has various scientific applications, ranging from biology to optics. The method scans the object of interest in a series of overlapping positions, thereby generating a set of multiple Fourier magnitude measurements that are potentially corrupted by noise. From these measurements, a high-quality image of the object can be reconstructed depending on how the related inverse problem is formulated and solved. In this paper, we propose a class of variational models that incorporate the weighted anisotropic-isotropic total variation (AITV), an effective regularizer for image recovery. This class of models is applicable to measurements corrupted with either Gaussian or Poisson noise. In order to have the models applicable for large number of ptychographic scans, we design an efficient stochastic alternating direction method of multipliers (ADMM) algorithm and establish its convergence. Numerical experiments demonstrate that from a large set of highly corrupted Fourier measurements, the proposed stochastic algorithm with AITV regularization can reconstruct complex-valued images with satisfactory quality, especially for the phase components.

Kevin Bui

UC Irvine
kevinb3@uci.edu

Zichao Di

Argonne National Lab
wendydi@mcs.anl.gov

MS101

Improving the Robustness of Deep Unrolling-based MRI Reconstruction by Learned Randomized Smoothing

Since the data collection in magnetic resonance imaging (MRI) is sequential and slow, many methods have been proposed to improve image reconstruction from accelerated scans involving limited or undersampled measurements. Recently, deep learning-based methods have gained much popularity for the reconstruction of medical images. However, many currently available deep learning methods for MRI reconstruction are not too robust: small adversarial changes to the input can result in very substantial changes in the output reconstruction. In this work, we propose a novel framework to improve the robustness of deep learning-based MRI reconstruction inspired by and extended from the celebrated randomized smoothing (RS) technique in robust learning. Yet, different from conventional RS-based robustification techniques for classification networks, we work with deep unrolling-based reconstruction networks (e.g., MoDL), which impose new challenges of how to integrate RS with the involved unrolling operations. In our proposal, we augment RS with a learnable reweighting network, which yields a new variant of RS that we call learned RS. We tested the learned RS by injecting worst-case adversarial attacks as well as random noise into the MRI measurements. Our proposed scheme demonstrates significantly better accuracy-robustness trade-offs at a variety of MRI data undersampling factors compared

to the conventional MoDL networks.

Shijun Liang, Jinghan Jia, Sijia Liu,
 Saiprasad Ravishankar
 Michigan State University
 liangs16@msu.edu, jiajingh@msu.edu, liusiji5@msu.edu,
 ravisha3@msu.edu

MS101

Uncertainty Estimates for Volume Predictions in Cardiac MRIs

The volume of the left ventricle (LV) is an important and well-studied parameter to diagnose heart failure. Segmenting the voxels belonging to the LV in magnetic resonance images (MRI) allows the most accurate estimation of the volume of the LV. While deep learning-based approaches for automatic segmentation are achieving good results, those methods lack a proper representation of the uncertainties linked to their predictions. However, when using automatic predictions especially in critical applications such as medical diagnosis an accurate description of the attached uncertainties is essential so that medical experts can interfere in case of unreliable predictions. In this work, we propose a method to add probabilities to the volume prediction made by any segmentation algorithm. The segmentation algorithm returns a segmentation for the LV for each slice throughout the 3D MRI of the heart. To add uncertainties, we model the volume along the depth of the heart by using a combination of a stochastic differential equation and a separate distribution for the particularly error-prone first and last slices which cut through the cardiac wall. We ensure that the predicted probabilities are bias-free with respect to the deterministic voxel counts of the segmentation algorithm. This approach inherently models asymmetries and incorporates geometric features into the uncertainty. Furthermore, this uncertainty lets us detect images that should be revisited by a medical expert.

Felix Terhag, Philipp Knechtges
 German Aerospace Center (DLR)
 felix.terhag@dlr.de, philipp.knechtges@dlr.de

Achim Basermann
 German Aerospace Center (DLR)
 Simulation and Software Technology
 Achim.Basermann@dlr.de

Raúl Tempone
 Chair of Mathematics for Uncertainty Quantification
 RWTH-Aachen University, Germany
 tempone@uq.rwth-aachen.de

MS102

Modeling of Perovskite Solar Cells with Structured and Planar SnO₂ Electron Transport Layers

Solar cell devices made using perovskites have reached the power conversion efficiency values (PCEs) as high as 25.7% for single-junction devices and 29.9% for silicon/perovskite tandem solar cell devices. This is a remarkable improvement in the performance of perovskite solar cells (PSCs) considering that the PCE of the first solar cell devices that employed methylammonium lead iodide (MAPbI₃) perovskite as a photo-absorber layer in perovskite dye-sensitized solar cells was only 3.9%. The performance of perovskite solar cells (PSCs) with nanorod-based SnO₂ electron transport layers (ETL, hereon referred to as ‘struc-

tured’) and MAPbI₃-based perovskite layer are investigated using computer simulation methods and contrasted with the performance of PSCs with planar SnO₂ (hereon referred to as ‘reference planar’) ETLs. The numerical simulations are conducted for different aspect ratios and density of SnO₂ nanorods, considering the optimum electronic parameters of the materials selected for the device functional layers. For each geometrical configuration of structured ETLs, the exact optical response of corresponding devices are calculated to account for the photonic effects originating from particular geometric dimensions of structured ETLs. The underlying discrepancies in the device performance of PSCs with reference planar and structured ETLs are explained through the analysis of electric field and current density distributions in devices.

Assylan Akhanuly, Iliyas T. Dossayev
 Nazarbayev University
 assylan.akhanuly@nu.edu.kz, iliyas.dossayev@nu.edu.kz

Erik Shalenov
 Satbayev University
 erik.shalenov@nu.edu.kz

Constantinos Valagiannopoulos
 Department of Physics, Nazarbayev University, Astana,
 010000
 konstantinos.valagiannopoulos@nu.edu.kz

Karlygash Dzhumagulova
 Institute of Experimental and Theoretical Physics,
 al-Farabi
 dzhumagulova.karlygash@gmail.com

Annie Ng
 Department of Electrical and Computer Engineering,
 Nazarbayev
 annie.ng@nu.edu.kz

Askhat Jumabekov
 Nazarbayev University
 askhat.jumabekov@nu.edu.kz

MS102

An Efficient 2-RDM Method for Electronic Structure Theory

We propose a semidefinite programming approach to determine the ground state energy of N -electrons using the reduced density matrix (RDM) formalism. While optimizing low-order RDMs enjoys a polynomial running time, in practice the size of the variables involved can scale as N^6 . To this end, we propose an efficient divide-and-conquer strategy where the system is partitioned into multiple fragments. Then local RDMs of the fragments are obtained and stitched together in a convex optimization framework.

Fabian M. Faulstich
 University of California, Berkeley
 f.m.faulstich@berkeley.edu

Yuehaw Khoo
 University of Chicago
 ykhoo@uchicago.edu

Michael Lindsey
 University of California, Berkeley

lindsey@berkeley.edu

MS102

Numerical Analysis of Linear Response in Time-Dependent Mean-Field Models of Quantum Mechanics

We study a non linear mean field model for quantum mechanics for a N -particle system (such as time-dependent density functional theory). We compute the linear response of the system when a perturbative time-dependent term is added to the Hamiltonian. To this end, we study the Liouillian L which dictates the evolution of the system at first order depending of the perturbation. The spectrum of the evolution operator e^{iLt} , defined on the manifold of admissible perturbation of orbitals, is of particular interest. We define rigorously the linear response function, and prove some of its properties (for instance, that the ionization threshold is given by the opposite of the highest occupied energy level).

Mi-Song Dupuy

Laboratoire Jacques-Louis Lions, Sorbonne Université,
Paris
mi-song.dupuy@sorbonne-universite.fr

Eloïse Letournel

CERMICS, Ecole des Ponts Paristech, Marne-La-Vallée,
France,
Inria Paris and Université Paris-Est
eloise.letournel@enpc.fr

Antoine Levitt

Université Paris-Saclay
antoine.levitt@universite-paris-saclay.fr

MS102

Momentum Space Method for Incommensurate 2D Bilayers Including Mechanical Relaxation Effects

Since their discovery by Geim, 2D heterostructures have become a hotbed of research due to their novel structure. Stacking varying materials on top of each other allows for a vast range of possible materials and corresponding properties. However, this stacking typically leads to an aperiodic, or incommensurate, material due to lattice mismatch or rotational misalignment. These incommensurate patterns increases the potential for tunability of electronic and mechanical properties, leading for example to the famous discovery of unconventional superconductivity of twisted bilayer graphene at the famous magic angle. In this work, we discuss configuration and momentum space methodologies to build algorithms for computing observables and quasi-band structure of incommensurate heterostructures for ab-initio tight-binding models. We exploit the ergodicity of the misalignment, and for appropriate materials (such as those with conic or parabolic bands) we use carefully selected perturbative expansions in momentum space with a thorough error analysis describing the effects of the large moire scale felt strongly through the mechanical relaxation patterns.

Daniel Massatt

Department of Mathematics
Louisiana State University
dmassatt@lsu.edu

Stephen Carr

Brown University

stephen_carr1@brown.edu

Mitchell Luskin

School of Mathematics
University of Minnesota
luskin@umn.edu

MS102

A Two-Level Domain Decomposition Method for Periodic Schrödinger Eigenstates in Anisotropically Expanding Domains

This talk presents a two-level domain decomposition method for the linear Schrödinger eigenvalue problem with periodic potentials for a non-uniform spatial expansion of the domain. In this framework, the numerical solution using inner-outer iterative eigenvalue algorithms suffers from a collapsing spectral gap that leads to deteriorating convergence rates. We solve the collapsing gap problem by a quasi-optimal spectral-shift strategy to uniformly bound the number of iterations with respect to different domain sizes. Since the algorithm requires the solutions of shifted linear systems, we further propose a two-level decomposition method that includes the already computed spectral asymptotics in a suitable coarse space. This domain decomposition preconditioner is easy to construct, leads to a bounded condition number for the shifted linear systems, and efficiently utilizes the geometric subdivision of the domain. We test the method's robustness and efficiency numerically. If time permits, we analyze the method's scalability within the theoretical framework of spectral coarse spaces.

Lambert Theisen, Benjamin Stamm

Chair of Numerical Mathematics of High Performance
Computing

University of Stuttgart

theisen@acom.rwth-aachen.de, benjamin.stamm@ians.uni-stuttgart.de

MS103

Mini-Batch Stochastic Three-Operator Splitting for Distributed Optimization

We consider a large class of convex optimization problems given by the sum of a convex and continuously differentiable function, and two, possibly non-smooth convex functions, one of which is composed with a linear operator. Such a structure is very general, and describes many applications that range from signal processing to machine learning to control. We propose a stochastic primal-dual algorithm which is provably convergent and apply it in its multi-agent formulation where each agent in a network performs local calculations and can only communicate with its neighbors. The challenging aspect of our study is that the smooth part of the private cost function is given as an expected value and agents only have access to this part of the problem formulation via a heavy-tailed stochastic oracle. To tackle such sampling-based optimization problems, we propose a stochastic extension of the triangular preconditioned primal-dual algorithm [Latafat, Freris, Patrinos, A new randomized block-coordinate primal-dual proximal algorithm for distributed optimization, IEEE Transactions on Automatic Control, 2019]. We demonstrate almost sure convergence of the scheme and validate the performance of the method via numerical experiments on the economic dispatch problem for power grids.

Barbara Franci

Department of Advanced Computing Sciences
Maastricht University
b.franci@maastrichtuniversity.nl

Mathias Staudigl
Maastricht University
m.staudigl@maastrichtuniversity.nl

MS103

A Condition-Based Maintenance Lease Contract with Continuously Monitored Equipment

This paper develops a mathematical model for a lease contract equipment with continuously monitoring through condition based maintenance (CBM). This equipment is used in dynamic environment and can be either use in normal stress (N) or severe/high stress (H) and the failure depends on the usage intensity of the equipment. The transitions from N to H and from H to N occur in a random manner. So we model the transitions by a two state continuous time Markov chain formulation. The usage intensity varies across the population of users and is modeled as a continuous random variable. Since the equipment degradation and failure depends on the usage intensity, this in turn has an impact on the expected lease contract cost. Finally we find numerically the optimal inspection time which minimizes the expected lease contract using non-cooperative game theory.

Hennie Husniah
LANGLANGBUANA UNIVERSITY
hennie.husniah@gmail.com

Bermawi Iskandar
Institut Teknologi Bandung
Indonesia
bermawi@mail.ti.itb.ac.id

Asep Supriatna
Padjadjaran University
a.k.supriatna@unpad.ac.id

MS103

Tight Lyapunov Function Existence Analysis for First-Order Methods

We present a unifying framework for establishing linear convergence rates for common first-order methods used to solve convex optimization problems. In particular, we consider i) classes of convex optimization problems of finite sum form with (possibly strongly) convex and (possibly) smooth functional components, and ii) first-order methods that can be written in so-called state-space form, i.e., as a linear system in feedback interconnection with the sub-differentials of the functional components of the objective function. The validity of a given target linear convergence rate is established by deriving a necessary and sufficient condition for verifying the existence of a quadratic Lyapunov function for the algorithm and problem class under consideration for the chosen rate, which amounts to the feasibility of a small-sized semidefinite program. This allows us to find the smallest linear convergence rate for which such a quadratic Lyapunov function exists, by bisection search, yielding a tight procedure. The approach is numerically exemplified on several algorithmic schemes. This is joint work with Sebastian Banert, Adrien Taylor and Pontus Giselsson.

Manu Upadhyaya

Lund University
manu.upadhyaya@control.lth.se

MS104

Portable, Scalable, and Safe Asynchronous Compute on GPUs with PETSc

To maximize efficiency on modern GPU hardware the programmer must make use of streams to execute tasks asynchronously from the host system. In addition to allowing the programmer to hide launch and memory transfer overhead, streams enable extracting additional parallelism by executing multiple tasks simultaneously. However, streams have proven difficult to integrate into libraries as they break the otherwise natural encapsulation offered by functions. This mini-symposium reports recent progress and developments on safely integrating a performant stream-aware programming model in PETSc.

Jacob Faibussowitsch
University of Illinois Urbana-Champaign
jacob.fai@gmail.com

MS104

Mixed Precision in the Solution of Sparse Linear Systems

Sparse systems of linear equations are a crucial part of many applications in today's science and engineering. This is often accomplished using iterative solvers like Krylov subspace methods. On modern, massively parallel GPU architectures, these iterative methods are memory bound operations, making carefully optimized memory access crucial for solver runtime. This talk discusses efforts taken in the Ginkgo library to identify areas in iterative solver algorithms and preconditioners which allow for the use of mixed precision, i.e. trading some of the accuracy used to store floating point numbers for reduced memory traffic when accessing said values. We will see results for conservative approaches preserving convergence behaviour as well as more drastic ones giving up some of the convergence speed in favor of much faster iterations thanks to drastically reduced memory access.

Fritz Goebel
Karlsruhe Institute of Technology
fritz.goebel@kit.edu

MS104

Development and Optimization of a Coupled Multi-GPU LBM-MHFEM Solver for Vapor Transport in the Boundary Layer over a Moist Soil

We present an efficient computational approach for simulating component transport within single-phase free flow over a soil layer. A numerical model based on this approach is validated using controlled experiments in a climate-controlled low-speed wind tunnel that is interfaced with a soil tank to study problems of heat and mass flux across the land-atmospheric interface. The numerical model is based on a combination of the lattice Boltzmann method (LBM) formulated for a weakly compressible fluid flow and the mixed-hybrid finite element method (MHFEM) for solving constituent transport. Both of these methods individually, as well as when coupled, are implemented for execution entirely on a GPU accelerator in order to utilize its computational power and avoid the hardware limitations caused by slow communication between the GPU and

CPU over the PCI-E bus. Furthermore, domain decomposition and optimization techniques are used to balance the load and communication costs when using multiple GPUs on an HPC cluster. We describe the numerical method and implementation of the solver, focusing primarily on the coupling mechanisms and performance optimizations. The performance of the solver is studied on Nvidia A-100 GPUs within the Karolina supercomputer managed by IT4Innovations, Czech Republic. Flow and transport simulation results are shown and compared with experimental data measured in the wind tunnel.

Jakub Klinkovsk
FNSPE CTU in Prague
jakub.klinkovsky@fjfi.cvut.cz

Andrew Trautz
US Army Engineer Research and Development Center
klinkjak@fjfi.cvut.cz

Radek Fucík
FNSPE CTU in Prague
klinkjak@fjfi.cvut.cz

Tissa Illangasekare
Colorado School of Mines
klinkjak@fjfi.cvut.cz

MS104

Fifteen Years of GPUs for Computational Science: Past, Present, and Future

General purpose computations on Graphics Processing Units (GPUs) have become widely available with the first release of the CUDA toolkit in summer 2007. In the 15 years that have followed, GPUs have become widely adopted for many general purpose computations in computational science. And yet, there is still an ongoing debate whether "GPUs are worth it" from a productivity point of view. This talk revisits various technological developments in the GPU landscape over these 15 years. Based on the lessons learned from the past, current developments and promises are evaluated in order to derive an outlook on where computational science using GPUs is headed.

Karl Rupp
TU Wien
me@karlrupp.net

MS104

High-Order Cell Methods for Maxwell Equations on Accelerators

We present recent advances of the linear operator algebra of the finite element package NGSolve. Matrix-free high order finite element discretizations, preconditioners, or explicit time stepping methods can be expressed by the composition of matrices of various types. Since in many cases it is cheap, the operators are built on the host system, and the whole tree of matrices is moved to the accelerator. As a case study we consider symplectic time stepping methods for Maxwell equations discretized by high order cell methods. The matrices from the differential operators are universal for all elements, and mass matrices are block-diagonal with very small, order independent block sizes. Time domain PML leads to coupled ordinary differential equations, which can all be expressed by operator composition. We present our in-house cuda-based implementation to dispatch the operator tree to GPUs, but similar con-

cepts are portable to other popular linear algebra libraries.

Joachim Schoeberl, Markus Wess, Thomas Lutonsky
TU Wien
joachim.schoeberl@tuwien.ac.at,
markus.wess@tuwien.ac.at, lutonsky.thomas@gmx.at

Bernard Kapidani
EPFL
bernard.kapidani@epfl.ch

Lorenzo Codecasa
Politecnico di Milano
lorenzo.codecasa@polimi.it

Matthias Hochsteger
CERBSim GmbH
matthias.hochsteger@cerbsim.com

MS105

On the Entropy Projection and the Robustness of High Order Entropy Stable Discontinuous Galerkin Schemes

High order entropy stable schemes provide improved robustness for computational simulations of fluid flows. However, additional stabilization and positivity preserving limiting can still be required for variable-density flows with under-resolved features. We demonstrate numerically that entropy stable DG methods which incorporate an "entropy projection" are less likely to require additional limiting to retain positivity for certain types of flows. We conclude by investigating potential explanations and applications to shock capturing.

Jesse Chan
Rice University
jesse.chan@rice.edu

Hendrik Ranocha
Applied Mathematics, University of Münster
mail@ranocha.de

Gregor Gassner, Andres Rueda-Ramirez
University of Cologne
ggassner@uni-koeln.de, aruedara@uni-koeln.de

Timothy Warburton
Virginia Tech
tim.warburton@vt.edu

MS105

High-Order State Redistribution Methods on Cut Cell Grids

In this talk, we present high-order state redistribution methods on embedded boundary grids. State redistribution relaxes the overly restrictive CFL condition that results from arbitrarily small cut cells when explicit time steppers are used. We show how to stabilize both finite volume and discontinuous Galerkin spatial discretizations and take time steps that are proportional to the size of Cartesian cells in the background grid. State redistribution works by post-processing the numerical solution after each stage or step of an explicit time stepping method. The advantage of this approach is that it does not require complex geometric manipulations. We solve a number of test prob-

lems with smooth and shocked solutions that demonstrate the encouraging potential of this technique for applications on embedded boundary grids.

Andrew Giuliani, Marsha Berger
Flatiron Institute
agiuliani@flatironinstitute.org,
mberger@flatironinstitute.org

MS105

Analysis of Hybrid Finite Element / Neural Network Methods for Solving Partial Differential Equations

We present hybrid discretisation schemes for partial differential equations that mix classical finite element methods with neural network approaches. We study methods that aim to enrich an efficient finite element on a coarse grid with fine scale fluctuations obtained using neural networks. The aim of this talk is to give a first numerical analysis for hybrid methods. The focus is on stability estimations as well as a priori and a posteriori error estimations. Here we consider common low-dimensional problems, e.g., the Laplace, Stokes or Navier-Stokes equations, and work purely on a model basis. We show theoretical results and numerical examples that demonstrate the potential of hybrid methods. Furthermore, we demonstrate how neural network predictions can also be treated theoretically using methods of classical finite element techniques.

Thomas Richter
University of Magdeburg
Germany
thomas.richter@ovgu.de

MS105

Energy Stable State Redistribution Cut-Cell DG Methods for Wave Propagation

Cut cell methods provide the ability to represent complex geometries while maintaining the simplicity of a Cartesian mesh wherever possible. However, cut cell meshes can result in extremely small and/or skewed cut elements that severely restrict the maximum stable time step of a simulation. Special discretizations may also be required to ensure that a cut-cell method is energy stable. In this work we prove the L2 stability of state redistribution, a technique for relaxing the CFL condition when small elements are present, and combine it with a provably energy stable high order discontinuous Galerkin formulation for wave propagation problems.

Christina G. Taylor, Jesse Chan
Rice University
cgt3@rice.edu, jesse.chan@rice.edu

MS105

Neural Network Based Structure Preserving Method for Transport Equation

Deep learning method has emerged as a competitive mesh-free method for solving partial differential equations (PDEs). The idea is to either represent solutions for solution map of PDEs by neural networks to take advantage of the rich expressiveness of neural networks representation. In this talk, we will explore the applicability of this powerful framework to the a large class of PDEs, with the emphasis in dealing with multiple scales and obtaining long

time stability.

Li Wang
University of Minnesota
wang8818@umn.edu

MS106

Learning Hyperparameters via Deep Neural Networks for Large-Scale Atmospheric Inverse Modeling

Inverse problems arise in a wide variety of applications including biomedicine, environmental sciences, astronomy, and more, but computing reliable solutions to these problems requires the inclusion of prior knowledge, in a process that is often referred to as regularization. Most regularization techniques require suitable choices of regularization parameters, and in this work, we describe new approaches that use deep neural networks (DNN) to obtain regularization parameters. We consider a supervised learning approach, where multiple networks are trained to approximate mappings from observation data to individual regularization parameters. Once the networks are trained, regularization parameters for newly obtained data can be computed by efficient forward propagation of the DNNs. The network-obtained regularization parameters can be computed more efficiently and may even lead to more accurate solutions compared to existing regularization parameter selection methods. Numerical results for atmospheric inverse modeling demonstrate the potential of using DNNs to learn regularization parameters.

Julianne Chung
Department of Mathematics
Emory University
julianne.mei-lynn.chung@emory.edu

MS106

Stochastic Inversion with Maximal Updated Densities for Storm Surge Wind Drag Parameter Estimation

Recent advances in a measure-theoretic framework for solving stochastic inverse problems (SIPs) has led to a novel method for solving Parameter Identification Problems (PIPs) where uncertainties in model output data are predominantly due to epistemic (i.e. reducible) uncertainties in model inputs. This approach utilizes a solution to a stochastic forward problem (SFP) to update an initial density only in the system input directions informed by the model output data, performing a selective regularization that distinguishes it from other common inverse problem methods. The solution to the PIP is defined by the Maximal Updated Density (MUD) point. Data assimilation for reducing variance in MUD points is provided functionally through data-constructed Quantity of Interest (QoI) maps. This work presents recent advancements in MUD point estimation methods and their application to the problem of estimating wind-drag parameters in the state-of-the-art ADvanced CIRCulation (ADCIRC) storm-surge model. A novel QoI map that uses Principal Component Analysis on data collected spatio-temporally is introduced. Furthermore, iterative methods are presented where data collected over successive time-windows is used to solve the PIP, using updated densities from prior iterations to inform successive iterations. These methods are used to estimate wind drag coefficients for a simulated extreme weather event near the Shinnecock Inlet located in the Outer Barrier of Long Is-

land, NY, USA.

Carlos del-Castillo-Negrete
Oden Institute for Computational Engineering and Sciences
University of Texas at Austin
cdelcastillo21@gmail.com

Michael Pilosov
University of Colorado: Denver
michael.pilosov@ucdenver.edu

Tian Yu Yen
Sandia National Laboratories
tyen@sandia.gov

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

Clint Dawson
Oden Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@oden.utexas.edu

MS107

Real-Time Modeling Through Neural Networks: Challenges and Perspectives

In the last years, the exponential diffusion of machine learning algorithms has contributed to improve already consolidated procedure, as well as allowing the creation of new frameworks, also in the field of data-driven and reduced order modeling. In particular, the neural networks, thanks to their approximation capability, have enabled the treatment of complex nonlinear models, maintaining a very limited computational cost during the inference of these models. In this contribution, we present some of the enhancements obtained by employing neural networks to overcome the limitations of standard technique like 1) an automatic snapshots shifting to solve the weakness of proper orthogonal decomposition in advection-dominated problems, 2) a multi-fidelity approach to improve the accuracy of a generic data-driven reduced order model and 3) the development of continuous convolutional filter to reduce the dimensionality with unstructured data. Numerical experiments are finally shown, with a particular focus about the comparison with state-of-the-art methods, to conclude then with the open challenges raised by the proposed techniques.

Dario Coscia, Laura Meneghetti
SISSA, Mathematics Area, mathLab
dcoscia@sissa.it, laura-meneghetti@sissa.it

Nicola Demo, Michele Girfoglio
SISSA, Mathematics Area, mathLab
International School for Advanced Studies
ndemo@sissa.it, mgirfogl@sissa.it

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS107

A Pod Based Reduced Order Model for Computational Cost Reduction of Industrial Packaging Sys-

tems Heat Exchange Simulations

This contribution presents an application of model order reduction to the simulation of the temperature field evolution during the sealing of cardboard beverage packages in an industrial machinery. In such process, the folding flaps of the poly laminated package are heated through a hot air jet, and then pressed together to obtain a watertight welding. A satisfactory sealing is obtained when the local poly laminated material temperatures are in the correct window at the moment when folding flaps are pressed one against each other. Simulations of the temperature field evolution within the layers of the package based on the solution of the Fourier equation, have been carried out to analyze how much the process is dependent on parameters such as the package dimension and layers materials properties, the heating air temperature and speed, and the heating time. In the framework of the full order model, the heat exchange govern equation is discretized making use of the Finite Element Method, implemented in a C++ software. A reduced order model has also been developed to allow for faster computations, compatible with real time control algorithms wired in the packaging machinery. Based on POD modal decomposition, both a non intrusive model based on the interpolation of the proper orthogonal decomposition modes, and a POD-Galerkin intrusive solver have been developed. The results obtained with the reduced order models will be presented and compared to characterize their performance.

Martina Cracco
SISSA, International School for Advanced Studies
gnoselli@si
mcracco@sissa.it

Giovanni Noselli
SISSA, International School for Advanced Studies
gnoselli@sissa.it

Andrea Mola
IMT School for Advanced Studies Lucca
andrea.mola@imtlucca.it

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS107

Automated Acceleration of Differentiable Computational Fluid Dynamics for Bayesian Optimization

With the ever more prevalent use of differentiable programming techniques throughout simulations, and machine learning, the desire to extend these techniques to classical fluid dynamics simulations for outer-loop applications such as Bayesian Uncertainty Quantification, and Bayesian Optimization has also become ever more pressing. While current approaches in fluid dynamics still view simulation as a black box which is then integrated with an outer-loop optimization framework, and potentially utilizes Gaussian Process surrogates to accelerate the sampling, modern differentiable programming approaches promise to heavily accelerate the optimization process through the acceleration of simulations, and the access to gradients of simulation trajectories. In this talk we build on recent advances in compiler-based automatic differentiation of parallelized C/C++, and Fortran-based computational fluid dynamics simulators to for the first time accelerate the production-

level simulations with machine learning techniques in an intrusive fashion to enable Bayesian optimization, where it was previously infeasible to apply Bayesian optimization to.

Ludger Paehler, Josef Winter, Nikolaus Adams
 Technical University of Munich
 ludger.paehler@tum.de, josef.winter@tum.de, niko-
 laus.adams@tum.de

MS107

Multi-Fidelity No-U-Turn Sampling

Hamilton Monte Carlo (HMC) and No-U-Turn Sampling (NUTS) are a family of Markov Chain Monte Carlo (MCMC) algorithms that incorporate first-order derivative to avoid random walk behavior which negatively impact the performance of many MCMC algorithms. However, calculation of derivative becomes impractical for computational expensive problems. We propose to build a multi-fidelity surrogate model using Gaussian Process and then use the derivative from the surrogate to generate proposals for NUTS. The multi-fidelity surrogate is built upon a hierarchy of models of decreasing approximation error and increasing computational cost where we delegate the majority of computational cost to the lower fidelity models. The samples are accepted/rejected according to the delayed acceptance criterion to ensure the detailed balance with respect to highest fidelity model. Then we test our problem on some test cases and compare it to existing methods.

Kislaya Ravi
 Technical University Munich
 kislayaravi@tum.de

Tobias Neckel
 Technical University of Munich
 neckel@in.tum.de

Hans Joachim Bungartz
 Technical University Munich
 bungartz@tum.de

MS108

CUQIpy: A New Python Platform for Computational Uncertainty Quantification in Inverse Problems

In this talk we present CUQIpy (pronounced cookie pie) a new computational modelling environment in Python that uses uncertainty quantification (UQ) to access and quantify the uncertainties in solutions to inverse problems. The overall goal of the software package is to allow both expert and non-expert (without deep knowledge of statistics and UQ) users to perform UQ related analysis of their inverse problem while focusing on the modelling aspects. To achieve this goal the package utilizes state-of-the-art tools and methods in statistics and scientific computing specifically tuned to the ill-posed and often large-scale nature of inverse problems to make UQ feasible. We showcase the software on problems relevant to imaging science such as computed tomography and partial differential equation-based inverse problems. CUQIpy is developed as part of the CUQI project at the Technical University of Denmark and is available at <https://github.com/CUQI-DTU/CUQIpy>.

Nicolai A B Riis

Technical University of Denmark
 nabr@dtu.dk

Amal Alghamdi
 Department of Applied Mathematics and Computer Science
 Technical University of Denmark
 amaal@dtu.dk

Felipe Uribe
 Technical University of Denmark
 Department of Applied Mathematics and Computer Science
 felipe.uribe@lut.fi

Silja W. Christensen, Afkham Babak
 Technical University of Denmark
 swech@dtu.dk, bmaaf@dtu.dk

Per Christian Hansen
 Technical University of Denmark
 DTU Compute
 pcha@dtu.dk

Jakob Jørgensen
 Technical University of Denmark
 jakj@dtu.dk

MS108

Comparing Monte Carlo Sampling Algorithms Using a Benchmark Bayesian Inverse Problem

We present the details of a prototypical benchmark for Bayesian inverse problems in which we seek to identify the posterior probability distribution of a discretized coefficient in an elliptic partial differential equation. We then use this benchmark to compare a number of simple and not-so-simple Monte Carlo sampling algorithms to assess how they perform using this benchmark as a test case that has many of the same features as real-world inverse problems.

Wolfgang Bangerth
 Colorado State University
 bangerth@colostate.edu

MS109

Convergence Properties for Data Assimilation Based on Gauss-Newton Iteration

Data assimilation is the problem of finding a state of a dynamical system such that the difference between true state and an approximated one is small in a properly defined sense. The widely-used data-assimilation methods are variational methods. They aim at finding an optimal initial condition of the dynamical model such that the distance to the observations is minimized. The problem is formulated as a minimization of a nonlinear least-square problem with respect to initial condition, and it is usually solved using a Gauss-Newton method. Motivated by the variational approach, we consider a data-assimilation method that minimizes a cost function under assumption of model error. We are seeking a solution over a time window at once. We prove the method converges to the true state in the case of noise-free observations and provide error bound in the case of noisy observations. Furthermore, we extend the method to parameter estimation. We confirm our theoretical results with numerical experiments using

the Lorenz models.

Nazanin Abedini, Svetlana Dubinkina
VU Amsterdam
n.abedini@vu.nl, s.b.dubinkina@vu.nl

MS109

A Dynamics-Aware Measure Transport Algorithm for Bayesian Filtering in Uniformly Hyperbolic Dynamics

In Bayesian filtering, we recursively update the probability measures of the state of a dynamical system conditioned on past observations. Numerical methods for the filtering recursion based on Kalman filter variants make the Gaussian assumption, and particle filters are computationally prohibitive in high dimensions. Transport map-based methods are computationally tractable and applicable to non-Gaussian problems but do not make use of the underlying dynamics present in the filtering setting. Our goal is to construct a dynamics-aware computationally tractable algorithm to sample from filtering distributions in uniformly hyperbolic systems (idealized chaotic systems). To this end, we express the filtering recursion by involving the scores (gradients of logarithms) of the filtering distribution sequence and show that these scores can be estimated efficiently by exploiting uniform hyperbolicity. We propose an ansatz for the filtering map: composition of forecast and analysis steps on the phase space. While the filtering map is not, in general, unique, our ansatz allows an interpretation of the filtering map as a conjugacy between two dynamical systems and imposes uniqueness. Such an interpretation along with the score equation constraint gives rise to a novel algorithm to perform filtering via measure transport. This algorithm is derived and demonstrated on low-dimensional hyperbolic dynamics.

Nisha Chandramoorthy, Youssef M. Marzouk
Massachusetts Institute of Technology
nishac@mit.edu, ymarz@mit.edu

MS109

A Non-Intrusive Solution to the Ill-Conditioned Gaussian Kernel Covariance Matrix for Gradient-Enhanced Gaussian Processes

Gaussian processes (GPs) provide a probabilistic method of constructing surrogates that is used for various applications such as uncertainty quantification and Bayesian optimization. A common problem for GPs is the ill-conditioning of their covariance matrix, particularly when the popular Gaussian kernel is used. For gradient-free GPs, a modest nugget can be added to the diagonal of the covariance matrix to ensure its condition number is below a user-set threshold. Unfortunately, it is problematic to apply this method on its own for gradient-enhanced GPs since it can require a large nugget, which is detrimental to the accuracy of the GP. A novel method has been developed that guarantees that the condition number of the gradient-enhanced Gaussian kernel covariance matrix remains below a user-set threshold. This is achieved by using non-isotropic rescaling for the data and a modest nugget. This method is straightforward to implement, non-intrusive, applicable to problems of any dimension, and it allows all data to be kept. To demonstrate the effectiveness of this method, a Bayesian optimizer is used with and without this method. With this novel method, the optimizer does not encounter any ill-conditioned covariance matrices and the optimality is converged several orders of magnitude deeper than the

base case.

Andre Marchildon, David Zingg
University of Toronto Institute for Aerospace Studies
andre.marchildon@mail.utoronto.ca,
dwz@utias.utoronto.ca

MS109

Criteria for Uniform in Time Numerical Approximations of Stochastic Differential Equations

Complicated models, for which a detailed analysis is too far out of reach, are routinely approximated via a variety of procedures; this is the case when we use multiscale methods, when we take many particle limits and obtained a simplified, coarse-grained dynamics, or, simply, when we use numerical methods. While approximating, we make an error which is small over small time-intervals but it typically compounds over longer time-horizons. Hence, in general, the approximation error grows in time so that the results of our “predictions” are less reliable when we look at longer time-horizons. However this is not necessarily the case and one may be able to find dynamics and corresponding approximation procedures for which the error remains bounded, uniformly in time. We will discuss a very general approach to understand when this is possible. I will show how the approach we take is very broad and show how it can be used for all of the approximation procedures mentioned above with a particular focus on numerical approximations for SDEs. This is based on a series of joint works with L. Angeli, D. Crisan, P. Dobson, I. Souttar.

Michela Ottobre
Heriot Watt University
mo3@hw.ac.uk

MS109

On Ensemble Size in a Particle Method for Subsidence Estimation

We use a particle method in two different experiments with models of different complexity. The first model calculates subsidence for a single observation point due to a single source and considers independent and uncorrelated parameters and observations. The second model calculates the observed subsidence as a summation of subsidence contributions from multiple sources. In the latter model, the spatial response that a single subsidence source causes at the surface will result in correlated observations. The correlation in the resulting observations may trigger weight collapse in the particle method. The information loss related to the spatial correlation can be quantified with mutual information. Using the quantification of information loss, we illustrate how this loss of information is reflected in the log-likelihood of the estimation problem, and how this depends on the number of parameters of the model. Based on the results of these experiments, we propose criteria to evaluate the required ensemble size for the assimilation of spatially correlated subsidence observations.

Samantha Kim
Delft University of Technology
s.s.r.kim@tudelft.nl

Femke Vossepel
Delft University of Technology Dept. of Geoscience and Engin

f.c.vossepoel@tudelft.nl

MS110

Low and Mixed Precision Recycling on GPUs

Recycling selected subspaces can substantially improve the rate of convergence of iterative methods as well as provide good initial guesses. However, in GPU-based architectures, where memory is often limited, only a modest number of recycling basis vectors can be stored locally. On the other hand computations involving local data are very fast. We discuss a number of low or mixed precision algorithmic extensions as well as memory saving strategies to make recycling more effective on GPU-based architectures.

Eric De Sturler

Virginia Tech
sturler@vt.edu

MS110

Augmented Lagrangian Block Preconditioners for Incompressible Resistive Magnetohydrodynamics

The equations of magnetohydrodynamics are generally known to be difficult to solve numerically. They are highly nonlinear and exhibit strong coupling between the electromagnetic and hydrodynamic variables, especially for high Reynolds and coupling numbers. In this work, we present a scalable augmented Lagrangian preconditioner for a finite element discretization of the single-fluid incompressible viscoresistive MHD equations. For stationary problems, our solver achieves robust performance with respect to the Reynolds and coupling numbers in two dimensions and good results in three dimensions. We extend our method to fully implicit methods for time-dependent problems which we solve robustly in both two and three dimensions. Our approach relies on specialized parameter-robust multigrid methods for the hydrodynamic and electromagnetic blocks. The scheme ensures exactly divergence-free approximations of both the velocity and the magnetic field up to solver tolerances. We confirm the robustness of our solver by numerical experiments in which we consider fluid and magnetic Reynolds numbers and coupling numbers up to 10,000 for stationary problems and up to 100,000 for transient problems in two and three dimensions.

Patrick E. Farrell

Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

MS110

Robust Preconditioners for a High Order Finite Element Discretization of Planar Linear Elasticity

The discretization of the primal form of planar linear elasticity by high order continuous finite elements is known to be robust and locking-free in the incompressible limit. Recently, we showed that the stresses computed from the finite element solution are also robust and locking-free thanks to the uniform (in mesh size and polynomial degree) inf-sup stability of the Scott-Vogelius elements. Unfortunately, the condition number of the resulting linear system blows up in the incompressible limit, as the mesh is refined, and/or as the polynomial degree is increased. We present a family of preconditioners that control the growth of the condition number uniformly in the Lam parameters, the mesh size, and the polynomial degree. The

preconditioners are illustrated by applying them to the solution of a number of representative test problems. As a by-product of the analysis, we also obtain a preconditioner for the Morgan-Scott C^1 -conforming finite element space that is uniform in the mesh size and polynomial degree.

Charles Parker

Brown University
charles.parker@maths.ox.ac.uk

Mark Ainsworth

Division of Applied Mathematics
Brown University
Mark.Ainsworth@brown.edu

MS110

Krylov Subspace Recycling for the Application of Matrix Functions

We derive an augmented Krylov subspace method with subspace recycling for computing a sequence of matrix function applications on a set of vectors. The matrix is either fixed or changes as the sequence progresses. We assume consecutive matrices are closely related, but make no assumptions on the relationship between the vectors. This setting presents some unique challenges with regard to augmentation of Krylov subspaces. We demonstrate the effectiveness of the method using a range of numerical experiments with a selection of functions and matrices.

Kirk Soodhalter

Trinity College Dublin
ksoodha@maths.tcd.ie

MS110

Hybridization of Machine Learning and Numerical Linear Algebra Techniques for Scientific Computing: Learned Minimum Residual Solvers for the Helmholtz Equations

In the recent years, research on scientific machine learning based on deep learning solvers has been increasingly applied to scientific computing and computational engineering. However, even though these newly data-driven deep learning solvers could be very effective as soon as they have been properly trained, they can (in general) provide us with a solution of limited accuracy. Furthermore, the computational cost in the training phase can be extremely expensive. In this talk, we present some ways of hybridizing the newly emerging deep learning solvers and the more traditional numerical linear algebra techniques to let them benefit from each other. In the context of solving a 2D heterogeneous Helmholtz equation, we first focus on introducing some mathematical ingredients from classical iterative solver into the training phase of a recently proposed deep neural network solver. The main benefit is a significant improvement in the training phase that is more robust and faster, which turns out to be applicable to the testing process as well. Furthermore, once the network solvers have been properly trained, their inferences can be applied as a nonlinear preconditioner in the traditional flexible GMRES and flexible FOM methods. This part demonstrates that these hybrid variants have clear advantages over both the newly emerging deep neural network approach and the classical iterative Krylov solver in terms of both computational cost and accuracy of the computed solution.

Yan-Fei Xiang, Luc Giraud

Inria

yanfei.xiang@inria.fr, luc.giraud@inria.fr

Paul Mycek, Carola Kruse
Cerfacs
mycek@cerfacs.fr, carola.kruse@cerfacs.fr

MS111

Placental Haemodynamics: Transport Effects at the Organ Scale

The placenta provides nutrients and oxygen to a developing fetus and is therefore vital to fetal development. It brings maternal and fetal blood close together, allowing nutrients to diffuse across thin separating barriers in the fetal villous tree structure. Structurally, the placenta is divided into placentones: compartments that are partially separated by septal walls. An approach to modelling maternal blood flow is to treat the fetal villous tree as a porous medium; several authors have utilised this on representative placentone geometries, mainly focusing on arterial supply. However, whilst these simulations are a useful indicator of organ-level behaviour, they fail to describe the effects of the blood flux between neighbouring placentones, as well as neglecting the importance of maternal venous return. In 2020, a new phenomenon was discovered called the placental contraction, which is yet to be mathematically modelled. I will present some in-silico organ-scale maternal blood flow simulations, using Navier-Stokes and Brinkman equations to model blood flow on representative placental geometries, coupled to a reaction-advection-diffusion equation to model nutrient transport. I will show notable blood flux passing between placentones, the importance of considering maternal venous return paths on the uniformity of oxygen exchange, and that the recently-observed placental contraction phenomenon could be vital in redistributing blood and encouraging venous return.

Adam M. Blakey
University of Nottingham
adam.blakey@nott.ac.uk

Penny Gowland
University of Nottingham
penny.gowland@nottingham.ac.uk

Paul Houston
School of Mathematical Sciences
University of Nottingham, UK
Paul.Houston@nottingham.ac.uk

Matthew Hubbard, George Hutchinson, Lopa Leach, Reuben O'Dea
University of Nottingham
matthew.hubbard@nottingham.ac.uk,
george.hutchinson1@nottingham.ac.uk,
lopa.leach@nottingham.ac.uk,
reuben.odea@nottingham.ac.uk

MS111

A Spatially Distributed Model of Brain Metabolism

The different active roles of neurons and astrocytes during neuronal activation are associated with the metabolic processes necessary to meet the associated energetic needs. Metabolism, in turn, relies on the delivery of metabolites and removal of toxic byproducts through diffusion processes and the cerebral blood flow. A comprehensive

mathematical model of brain metabolism, should account not only for the biochemical processes and the interaction of neurons and astrocytes, but also the diffusion of metabolites. We present a computational methodology based on a multidomain modeling view of the brain tissue using a homogenization argument to account for the diffusion processes. In our spatially distributed compartment model, communication between compartments occur both through local transport fluxes, as is the case within local astrocyte-neuron complexes, and through diffusion of some substances in some of the compartments. Our model assumes that diffusion takes place in the extracellular space (ECS) and in the astrocyte compartment, where its strength depends on the strength of the gap-junctions. The diffusion process is implemented numerically by means of a finite element method (FEM) based spatial discretization, and robust stiff solvers are used to time integrate the resulting large systems. Computed experiments investigate how the effects of ECS tortuosity, gap junction strength and spatial anisotropies in the astrocyte network affect the brain energy metabolism.

Gideon O. Idumah
Case Western Reserve University
goi@case.edu

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and Statistic
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS111

Model Order Reduction for Complex Ocular Simulations Inside the Human Eyeball

Complex ocular simulations are devoted to modeling the interplay between tissue perfusion, biomechanics, fluid dynamics, and heat transfer within the eye. These different aspects of the same physical problem have to be properly connected and every step has to be verified and validated in the interest of a medical application. The models require the knowledge of various parameters and some may be important factors in the development of pathologies. However, despite recent significant advances in medical data acquisition, only some parameters and their variability are known, but others cannot be directly measured. To identify the main factors that influence the biomechanical behavior of the eye, we, therefore, need to study the influence of these parameters through an uncertainty quantification (UQ) process which requires many evaluations of the models. Since 3D models are not amenable directly to UQ, a reduction step is needed to mitigate the computational cost. We previously conducted a complete UQ analysis on a 0D reduced model. In the present talk, we propose a strategy to carry out further analysis using reduced order methods and in particular the reduced basis method. We present the application of these methods in the context of both continuous Galerkin and Hybridized Discontinuous Galerkin formulations. We discuss then the implementation with the library Feel++. Finally, we apply the methodology to some advanced ocular models and report our findings.

Thomas Saigre
IRMA, Université de Strasbourg
saigre@math.unistra.fr

Marcela Szopos
 Université Paris Cité
 MAP5
 marcela.szopos@parisdescartes.fr

Christophe Prud'homme
 University of Strasbourg
 France
 prudhomme@unistra.fr

MS111

A Computational Framework for the Simulation of Micro-Swimming

We are currently developing a computational framework for swimming simulation in order to address problems where rigid or deformable bodies move in Newtonian and complex fluids. The objective is to investigate the interactions of biological micro-organisms and micro-robots with their environment, in order to propose efficient driving strategies and obtain insights into propulsion at this scale. In this talk, we present the state of the art of our micro-swimming computational framework, capable of simulating sperm cells and biologically inspired micro-organisms as well as handling contact in some of the arising fluid-structure problems. The mathematical formulation and computational aspects of this fluid-structure interaction will be presented, ranging from the governing equations and gait models that are currently supported to the numerical solution of the swimming problem. A selection of applications and results will be presented as well.

Céline Van Landeghem
 IRMA, Université de Strasbourg
 celine.van-landeghem@etu.unistra.fr

Christophe Prud'homme
 University of Strasbourg
 France
 prudhomme@unistra.fr

Luca Berti
 IRMA, Université de Strasbourg
 berti@math.unistra.fr

Vincent Chabannes
 CEMOSIS, Université de Strasbourg
 vincent.chabannes@cemosis.fr

Laetitia Giraldi
 Inria Sophia Antipolis-Méditerranée
 laetitia.giraldi@inria.fr

Yannick Hoarau
 Icube, Université de Strasbourg
 hoarau@unistra.fr

Thibaut Metivet
 Inria Grenoble Rhône-Alpes
 thibaut.metivet@inria.fr

MS112

Ensemble Kalman Methods: A Mean Field Perspective

The ensemble Kalman methodology is an innovative and flexible set of tools which can be used for both state estimation in dynamical systems and parameter estimation

for generic inverse problems. It has primarily been developed by practitioners in the geophysical sciences, with notable impact on the fields of oceanography, oil reservoir simulation and weather forecasting. Despite its wide adoption in fields of application, firm theoretical foundations are only now starting to emerge. In this talk we consider the problem from the perspectives of both control theory and probability, and provide a unifying approach to algorithms that rests on transport of measures and mean field stochastic dynamical systems.

Edoardo Calvello
 California Institute of Technology
 e.calvello@caltech.edu

Sebastian Reich
 Universität Potsdam, Germany
 sereich@uni-potsdam.de

Andrew Stuart
 Computing + Mathematical Sciences
 California Institute of Technology
 astuart@caltech.edu

MS112

Quantum Mechanics for Closure of Dynamical Systems

When modelling chaotic dynamical systems, notably in climate prediction models, a common obstacle is the issue of approximating the dynamics of dimensions of the state space for which dynamical laws are unknown. Common approaches to solving this issue include parametrization schemes which employ stochasticity or which are exclusively dependent on the internal state of the system at a given time. We propose a novel parametrization scheme based on the mathematical framework of quantum mechanics and Koopman operator theory. Given a system in which some components of the state are unknown, this method involves defining the surrogate system as being in a time-dependent quantum-state which influences the choice of the unknown component of the classical state at each timestep. The quantum state is an operator on the space of classical observables and evolves over time under an action by the Koopman operator. The quantum state also updates with new data-points according to a quantum Bayes law, and evolves under the action of data-driven operators. Kernel functions are utilized to allow the quantum Bayes law to be implemented numerically. We analyze the results of two different modalities of this methodology applied to the Lorenz 63 and Lorenz 96 systems, and show how this approach preserves important statistical and qualitative properties of the underlying chaotic systems.

David Freeman, Dimitrios Giannakis
 Dartmouth College
 david.c.freeman.gr@dartmouth.edu,
 Dimitrios.Giannakis@dartmouth.edu

Joanna M. Slawinska
 Finnish Center for Artificial Intelligence
 Department of Computer Science, University of Helsinki
 joanna.m.slawinska@dartmouth.edu

MS112

Sparse Kernel Flows for Learning Dynamical Systems

Regressing the vector field of a dynamical system from a

finite number of observed states is a natural way to learn surrogate models for such systems. We introduce a new variant of the method of kernel flows that we call sparse kernel flows to learn the 131 chaotic dynamical systems proposed in <https://arxiv.org/abs/2110.05266>.

Boumediene Hamzi

Johns Hopkins University
boumediene.hamzi@gmail.com

Houman Owhadi

California Institute of Technology
owhadi@caltech.edu

Yang Lu

Imperial College London
Nanjing University of Aeronautics and Astronautics
l.yang22@imperial.ac.uk

MS112

Universal Inversion: a Framework for Infusing Expert Knowledge in Bayesian Inverse Problems

In the natural sciences, one is often faced with the problem of reconstructing some unknown function from indirect data. Such problems are broadly known as inverse problems. Solving these generally requires some form of regularization and Bayesian techniques such as Gaussian process (GP) priors have proved useful to that end. Nevertheless, in most fields of applications, there exists prior knowledge about the unknown phenomenon that cannot be directly included in the specification of the GP prior and would be better modelled by a set of trend functions. In this presentation, we will show how the traditional, GP-based Bayesian inversion framework can be extended to include (partially) known trends. These trends are modelled as linear combinations of basis functions with a multivariate Gaussian prior on the trend coefficients. In essence, this is an extension of the usual universal kriging approach to inverse problems, which we thus deem "universal inversion". We will demonstrate our universal inversion techniques on a large-scale gravimetric inverse problem based on data collected on Stromboli island, demonstrating how field knowledge can help improve the inversion. The introduction of trends in the inversion process fosters new questions pertaining to model selection. To tackle these, we extend recent results about fast k-fold cross-validation to the inverse problem and present an overview of the new venues of inquiry in model selection opened by fast cross-validation.

Cédric Travelletti

University of Bern
cedric.travelletti@stat.unibe.ch

MS112

Unbiased Estimators Applied to the Ensemble Kalman-Bucy Filter

We consider the development of an unbiased estimator for the ensemble Kalman-Bucy filter (EnKBF). The EnKBF is a continuous-time filtering methodology which can be viewed as a continuous-time analogue of the famous discrete-time ensemble Kalman filter. Our unbiased estimators will be motivated from recent work which introduces randomization as a means to produce unbiased and finite variance estimators. Our estimator will be specific to linear and Gaussian settings, where we know that the EnKBF is consistent, in the particle limit $N \rightarrow \infty$, with

the KBF. We highlight this for two particular variants of the EnKBF, i.e. the deterministic and vanilla variants, and demonstrate this on a linear Ornstein-Uhlenbeck process. We compare this with the EnKBF and the multilevel (MLEnKBF), for experiments with varying dimension size.

Miguel . lvarez B.

King Abdullah University of Science and Technology (KAUST)
miguelangel.alvarezballesteros@kaust.edu.sa

MS113

Data Science and Analysis with Earth Imaging Satellites

Maxar Technologies operates a constellation of high-resolution earth imaging satellites with varying spectral capabilities. We have a vast imagery archive with 125 petabytes of imagery, and the satellite images are loaded with content. This presentation will survey some data science problems that are key to unlocking the value in our archive. These problems range from the analysis of pixel data to feature and object detection. We use a variety of techniques to extract information and insight from these images to see a better world.

Brett Bader

Maxar Technologies
brett.bader@maxar.com

MS113

Computational Science and Engineering in Capital Markets

The rapid change of markets to electronic operation has been rapidly opening opportunities for massive data analytics and computation to play a key role. We will discuss the overall needs for computation and show that the industry is in great need of advances in algorithms, software and hardware all the way from bare metal to cloud computing. Thus, tremendous opportunities are open for computational scientists and engineers to impact a supercritical aspect of the world's economy as manifested in the capital markets.

Costas Bekas

Citadel Securities
bekas.costas@gmail.com

MS113

Mathematics to Transform the Everyday Applied Mathematics in Industry

In nearly all industries, computational algorithms are the innovation backbone of the 21st century. More and more innovations are based on digital solutions and powerful algorithms. This makes industrial applications a great exploration ground for mathematical innovations and novel applications. In this talk, I will review my own industrial career from consultancy and program management to becoming a technical leader and major innovator at Siemens. I will do this along specific examples highlighting how mathematics made the difference. My ambition is to get you excited about career opportunities in industry allowing you to make a positive impact on our everyday lives via mathematics.

Dirk Hartmann

Siemens AG
Corporate Technology
hartmann.dirk@siemens.com

MS113

**(Math + Life + Engineering + Statistical + Data)⁵
x Sciences + Industry : Great Careers and More**

Mathematical and Data Sciences are being applied across a broad range of industries, generating human knowledge (and, thus, informing decisions) by distilling data into information across disparate domains, scales of time and space, and sources. Brief examples will be given from a broad array of healthcare sectors such as pharmaceutical, biomedical equipment (instrumentation and prostheses), digital health, and insurance. Many also consider agricultural applications including plants and animals to be critical life science applications as well, and, thus, naturally embracing also quantitative ecological and climatic considerations. Ideas on the type of background valuable for a job search will also be discussed. One consideration in thinking about this background is that a very broad range of mathematical and computational sciences show up: classical applied math such as differential and partial differential equations, discrete mathematics, visualization, network inference, (stochastic) optimization, control theory, probability, dynamical systems, inverse problems, image processing, pattern recognition/AI, natural language processing, topology, and, of course, numerical analysis and simulation.

Jeffrey R. Sachs
Merck & Co., Inc.
Rahway, NJ, USA
jeff.sachs@merck.com

MS114

Application and System Co-Design for Extreme Scale AI Workloads

The true potential of AI rests on super-human learning capacity, and on the ability to selectively draw on that learning. Both properties—scale and selectivity—challenge the design of AI computers and the tools used to program them. This talk will discuss these challenges, and how Graphcore, a company that builds accelerators for AI, has been and will be breaking through them by co-designing new applications, software and hardware infrastructures to achieve super-human learning.

Mario Krell
GRAPHCORE
mariok@graphcore.ai

MS114

Hardware, Software, and Application Co-Design on the Frontier and Lumi Systems

The first Exascale systems are based on heterogeneous node designs. These systems are intended to run a range of applications across a wide set of domains scientific and machine learning applications. Ensuring these applications can leverage all the unique architectural features of the heterogeneous node design requires deep application co-design between hardware, software, and application teams. This talk discusses some of the challenges, and solutions, on preparing applications for the AMD-based OLCF Frontier and CSC LUMI-G systems. This talk will highlight key

hardware capabilities, such as matrix cores, packed Fused-multiply-add operations, large HBM capacities, coherent CPU and GPU memory, and how to leverage them in software. Supporting applications on multiple architectures via performance portable compilers such as OpenMP and HIP will also be discussed.

George Markomanolis
AMD
georgios.markomanolis@amd.com

MS114

Machine Learning Assisted Simulation Using Highly Coupled GPU, CPU and Deep Learning Accelerators

High performance computing workloads have started to leverage the benefits of AI both in software and hardware, namely merging traditional, first principles workloads with deep learning models. In multiphysics problems, intersection of machine learning with first principles simulation, sometimes referred to as cognitive simulation, has given rise to new paradigms in computational modeling and simulation where benefits are yet to be recognized. In this talk, we will focus on a specific subclass of ML assisted simulation, where tens, hundreds or potentially thousands of surrogate models representing different physiochemical phenomena, material properties, and regimes, are used to replace sections of the computational code historically solved through first principles methods. The simulation will run on innovative node-level heterogeneous design, consisting of Reconfigurable Dataflow Unit (RDU), GPU, and CPU. We will show for cases where strong coupling between the surrogate models and the main computational loop is required (i.e., calling the surrogates within a computational loop for each step of the simulation), and for problems that are latency-bound with low number of inference calls to each model (i.e., small batch sizes) SambaNovas Reconfigurable Dataflow Architecture (RDA) serves particularly well. An example use case will be presented and the improvements will be discussed.

Martin Mueller
SAMBANOVA
martin.mueller@sambanovasystems.com

MS114

Supporting Applications and Software for the Design of Post-Exascale Technology

Now that Exascale has been reached, how can we achieve another 1000x improvement to reach a ZettaFlop, at least for a simple benchmark? If we assume the size of computing centers remains roughly constant and we cannot increase the energy envelope, we fall back to a performance/density problem for all levels of a computer system. A computing unit today reaches about 100 TF/s in 1 kW(dp64); to reach 100 PF/s in a few years, that translates to 1 Summit supercomputer (ORNL) in a socket. If we consider that applications and their associated test cases can theoretically scale by simple oversampling (a factor of 10 in the 3 dimensions) and assuming that the numerical schemes and the physics can remain stable, then the problem comes back to the conservation of the ratios of bandwidth, floating point (and other data types), latencies for scalability of applications, and reduction of the energy needed to move a bit. It is difficult to predict if the current model of building HPC supercomputers can be scaled without new breakthroughs, including for programming mod-

els, or if limits will be reached. In this presentation, I will briefly discuss these points as well as the upcoming initiatives and investments in Europe to put applications and software at the center of the problem when designing new technologies.

Philippe Thierry
Intel Corp
philippe.thierry@intel.com

MS114

Newton's Method to Compute Taylor Series in Multiple Double Precision Accelerated by Graphics Processing Units

The problem is to investigate the scalability of a new path tracker (SISC 42(6), A3610-A3637, 2020) to solve systems of polynomial equations in many variables. The many in this context is about a thousand. Path trackers repeatedly run Newton's method, evaluating and differentiating polynomials, followed by the solution of a linear system. On Taylor series, the matrices are block Toeplitz lower triangular, obtained by convolutions. Proximity to singularities requires multiple double precision arithmetic, which causes a cost overhead to be offset by acceleration with Graphics Processing Units (GPUs). In particular, GPUs capable of teraflop performance compensate for the overhead caused by quad double arithmetic. Multiple double precision is necessary to adjust the parameter representation in case the radius of convergence of the Taylor series is too small. Singularities are located efficiently via the quadratically converging Newton's method at regular points and with extrapolation on the series. While the current implementation takes polynomials on input, viewing polynomials as truncated series extends its application to analytic systems, systems of functions with well defined Taylor series developments. The software written for this investigation is licensed under GPL-3.0, available at <https://github.com/janvershelde/PHCpack>. This research is supported by the U.S. National Science Foundation under grant 1854513 of the CDS&E-MSS program.

Jan Vershelde
University of Illinois at Chicago
Chicago, Illinois USA
janv@uic.edu

MS116

An Adaptive and Flexible Software Framework for Two-Scale Coupled Problems

For many challenging applications in simulation technology, micro-scale phenomena can dominate macro-scale behavior. Examples in this setting would be reactive porous-media flow and biomechanical models of human organs. We present a novel software framework to couple existing micro-scale and macro-scale simulations in a black-box fashion. We develop a software component which manages a set of micro simulations, and couples them to the macro simulation through the coupling library preCICE. The broad aim is to develop macro-micro coupling methods and software that are independent not only of the concrete multiscale application but also of the used macro-scale and micro-scale software. While reusing key coupling implementations of preCICE (e.g., parallel communication and fixed-point acceleration schemes), we present a new software component called Micro Manager. The Micro Manager calls all micro-scale simulations as libraries in an adaptive manner and is itself coupled to the macro-scale simu-

lation using preCICE. The working of the Micro Manager is demonstrated using a two-scale heat conduction scenario in porous-media. In this case, each micro simulation consists of a grain structure which evolves depending on the temperature at the corresponding macro location. The effective thermal conductivity and material amounts at each macro point is computed by resolving a micro simulation. The micro simulations are run in an adaptive manner based on a similarity criterion.

Ishaan Desai
University of Stuttgart
ishaan.desai@ipvs.uni-stuttgart.de

Carina Bringedal
University of Stuttgart
SimTech, IWS, Pfaffenwaldring 61, 70569 Stuttgart
carina.bringedal@hvl.no

Benjamin Uekermann
University of Stuttgart
benjamin.uekermann@ipvs.uni-stuttgart.de

MS116

Determining Obstacles to Generate Flow Structures

One of the ways to understand a fluid flow is to analyze the structures which govern their behavior. One such set of structures, called Lagrangian Coherent Structures (LCS), are of particular interest. LCSs serve as a "skeleton" of a fluid flow, as they describe the regions of maximal attraction, repulsion, and shearing of a flow. Typically these structures are used to qualitatively analyze flow behavior. However, one may wish for these structures to appear in specific regions of a flow, as a means of flow control. Specifically, we are interested in the following problem: can we introduce a flow obstruction to guarantee a specified LCS. In this talk, we will discuss the feasibility of this problem, whether or not solutions to this problem exist, and how we may go about solving it both analytically and numerically.

Derek A. Drumm
Worcester Polytechnic Institute
dadrumm@wpi.edu

MS116

Calibration of Model Parameters for 1D-0D Coupled Blood Flow Models

In this talk, we present a model for simulating blood flow within the largest arteries branching out of the heart. In order to decrease the computational effort one-dimensional (1D) Navier-Stokes equations are used to describe the flow field within this network. Since only large arteries are modeled by means of the 1D Navier-Stokes equations appropriate boundary conditions have to be imposed at the outlets of the considered network. One way to account for the missing vessels is to use systems of ordinary differential equations incorporating the resistance and capacity of the omitted vessels. Since they exhibit no space variable they are indicated as 0D models. All in all, this results in a 1D-0D coupled blood flow model. Parameters of the 0D models have to be chosen in a careful way to produce realistic simulation data. Using measurements of blood pressure curves, these parameters are calibrated formulating appropriate minimization problems. To avoid expensive evaluations of the respective cost functionals, physics informed neural networks are considered. For the solution of

the minimization problems, we apply two different types of approaches: The first one is a classical gradient based method, while the second one uses features of quantum computers. In case of the second approach, we outline how data are transferred from a classical to a quantum computer and discuss new algorithmic aspects. The performance of both algorithms is illustrated by several numerical tests.

Tobias Koepl
Hasselt University
tobias.koepl@uhasselt.be

Nikolay Tcholtchev
Faunhofer FOKUS Institute, System Quality Center (SQC), Berl
nikolay.tcholtchev@fokus.fraunhofer.de

Benedikt Hoock
Department of Physics, Humboldt University Berlin
benedikt.hoock@physik.hu-berlin.de

Ilie-Daniel Gheorghe-Pop
Faunhofer FOKUS Institute, System Quality Center (SQC), Berl
ilie-daniel.gheorghe-pop@fokus.fraunhofer.de

Luigi Iapichino
Leibniz Supercomputing Centre (LRZ), Garching
luigi.iapichino@lrz.de

MS116

A Decoupled Splitting Scheme Combined with a Discontinuous Galerkin Spatial Discretization for Solving the Cahn-Hilliard-Navier-Stokes Equations

We formulate and theoretically analyse a fully decoupled scheme for solving the Cahn-Hilliard-Navier-Stokes equations. This scheme combines a pressure correction approach with interior penalty discontinuous Galerkin spatial discretizations. We prove unique solvability and mass conservation. Under a CFL condition, we show that the order parameter is stable in the L^∞ norm and that the scheme is energy dissipative. Optimal a priori error rates in the broken H^1 and L^2 norms are proved. Our analysis is novel and robust in the sense that no artificial stabilization or regularization of the potential function are needed. We present numerical experiments which validate our theoretical results.

Chen Liu
Department of Mathematics, Purdue University, West Lafayette
liu3373@purdue.edu

Rami Masri
Simula
rami@simula.no

Beatrice Riviere
Rice University
Houston, Texas, USA
riviere@caam.rice.edu

MS116

Autonomous Finite Elements in Clinical Practice -

Coupling FEA with AI

Finite element analysis requires a qualified analyst to generate the necessary input data, verify the output and post-process the analysis results for a meaningful conclusion. The required expertise and labor efforts precluded the use of FEA in daily medical practice for example. Recent scientific advancements such as low-dose CT scans, machine learning, and high order FEA which allows an inherent verification methodology of the numerical accuracy, make it possible to provide a fully autonomous process for assessing bone strength and fracture risk. This autonomous process, which we refer to as autonomous streamline, named autonomous finite element (AFE) analysis, introduces a paradigm shift in the use of FEA. This talk addresses a novel AFE [Yosibash et al, *CMAME*, 80(11), p 2417, 2020] for patient-specific analysis of human femurs used in clinical practice: it involves an automatic segmentation of femurs from CT-scans by U-Net, an automatic mesh generation and application of boundary conditions based on anatomical points, a high-order FE analysis with numerical error control, and an automatic report with a clear assessment of bone fracture risk. Two specific applications of AFE are presented: a) Determination of the risk of fracture for patients with tumors of the femur [Sternheim et al, *Bone*, 110, p215, 2018, & *BJJ*, 102-B, p638, 2020]. b) Identifying patients with high hip fracture risk [Rotman et al, *BJJ*, 103-B, p1497, 2021].

Zohar Yosibash
Tel Aviv University
yosibash@tauex.tau.ac.il

Nir Trabelsi
Shamoon College of Engineering
nirtr@sce.ac.il

Kent Myers
PerSimiO Ltd
kent.myers@persimio.com

MS117

Extension of Optimized Grid Metrics to Non-Polynomial Curvilinear Coordinate Transformations for Summation-by-Parts Discretizations

We extend the optimization-based approach of Crean et al. (JCP 2018) for calculating grid metric terms to arbitrary non-polynomial mappings. This extension allows the use of summation-by-parts discretizations in curvilinear coordinates with non-polynomial geometry representations, including non-uniform rational basis splines, and general nodal distributions while still satisfying the discrete metric invariants. Furthermore, we provide a thorough comparison of this optimized approach to various standard approaches of calculating different geometric terms, including the grid metric terms, metric Jacobians, and boundary normals. In the context of diagonal-norm summation-by-parts discretizations, we demonstrate through both theoretical arguments and numerical examples involving the 3D Euler equations that the optimized approach can be beneficial in certain cases. In particular, when nodal distributions that do not include boundary nodes are used, only the optimized approach results in conservative discretizations that satisfy the metric invariants and superconvergent functionals in the presence of flow tangency boundary conditions.

Alex Bercik
University of Toronto Institute for Aerospace Studies

alex.bercik@mail.utoronto.ca

David Craig Penner
University of Toronto Institute for Aerospace Studies
david.craigpenner@mail.utoronto.ca

David W. Zingg
University of Toronto Institute for Aerospace Studies
dwz@oddjob.utas.utoronto.edu

MS117

A Stable Cut DG Method for Hyperbolic Conservation Laws Interpreted as an Sbp Method

Recently we developed a cut-DG method for hyperbolic conservation laws, where severe time-step restrictions are avoided by adding ghost penalty stabilization. The correspondence between standard DG methods and summation-by-parts finite difference methods is by now well known. In this talk we will explore how a cut-DG method for a hyperbolic conservation law can be interpreted as an immersed boundary finite difference method.

Gunilla Kreiss
Division of Scientific Computing
Uppsala University
gunilla.kreiss@it.uu.se

MS117

A Positivity Preserving Strategy for Entropy Stable Discontinuous Galerkin Discretizations of the Compressible Euler and Navier-Stokes Equations

High-order entropy-stable discontinuous Galerkin methods for the compressible Euler and Navier-Stokes equations require the positivity of thermodynamic quantities in order to guarantee their well-posedness. In this work, we introduce a positivity limiting strategy for entropy-stable discontinuous Galerkin discretizations based on high order DG (discontinuous Galerkin) positivity-preserving limiter. The key ingredient in the limiting procedure is a low order positivity-preserving discretization based on graph viscosity terms. The proposed limiting strategy is both positivity preserving and discretely entropy-stable for the compressible Euler and Navier-Stokes equations. We illustrate the behavior of the method using both analytical solutions and model problems such as the Sedov blast wave and Daru-Tenaud shock tube.

Yimin Lin
Rice University
yiminlin@rice.edu

MS117

A Discretization of Elliptic Terms with Improved Convergence Properties Using Summation by Parts Operators

Nishikawa (2007) proposed to reformulate the classical Poisson equation as a steady-state problem for a linear hyperbolic system. This enables a unified discretization based on hyperbolic PDE solvers, e.g., in the context of coupled elliptic-hyperbolic systems such as the Euler equations with self-gravity studied by Schlottke-Lakemper, Winters, Ranocha, and Gassner (2021). It also results in optimal error estimates for the solution of the elliptic equation and its gradients, which are of primary interest in self-gravity. However, it prevents the application of well-known solvers

for elliptic problems such as the (preconditioned) conjugate gradient method. We show connections to a discontinuous Galerkin (DG) method analyzed by Castillo, Cockburn, Perugia, and Shtzau (2000) that is very difficult to implement in general. Next, we demonstrate how this method can be implemented efficiently using summation by parts (SBP) operators, in particular in the context of SBP DG methods. The resulting scheme combines nice properties of both the hyperbolic and the elliptic point of view, in particular a high order of convergence of the gradients, which is one order higher than what one would usually expect from DG methods for elliptic problems.

Hendrik Ranocha
Applied Mathematics, University of Münster
mail@ranocha.de

MS118

Subsurface Flow and Transport with Differentiable Programming and Quantum Computing

Many subsurface flow applications involve components where physical laws are well understood and other components where the physical laws are either poorly understood or not applicable. Numerical modeling excels at the former whereas interpolating data with machine learning (ML) excels at the latter, but neither approach can tackle these components simultaneously. Existing ML approaches (often called physics-informed ML, or PIML) to handling these types of components simultaneously are minor tweaks to standard ML methods (e.g., PIML might use physics data to train or a loss function that encourages the ML to obey an equation without any accuracy guarantees). Tweaking black-box ML models is fundamentally limited because "big data does not interpret itself" meaningful, interpretable structure in models is a necessity to improve predictability, enable human understanding, and maximize the impact of small data. We show how Differentiable Programming (DP) enables us to meld trustworthy numerical modeling with trainable ML to enhance workflows for physical model development, inverse analysis, and machine learning. This talk will also consider applications of quantum computing to subsurface flow.

Daniel O'Malley
Los Alamos National Laboratory
omalld@gmail.com

MS118

Combining Imaged Data and Deep Learning for Prediction of Transport in Porous Media

Significant scientific opportunities can be realized by developing methods that study transport behavior within individual images (volumes) of rock samples (framework popularly called digital rock physics) and upscaling observations from many different samples and across different scales. The barrier to scientific progress is not acquiring data, but automation that would allow more efficient upscaling across the spatial and temporal scales. We show new machine (deep) learning algorithms that, in combination with data hosted in Digital Rocks Portal, creates the environment that directly links data, high performance computing simulation, and deep-learning-enabled upscaling of velocity field, concentration fields and electrical potential. The concept is applicable beyond rocks to estimating properties of other complex/porous structures such as those in materials science (foams), chemical engineering (batteries or fuel cells), and medicine (micro-vascular networks or bones).

This is joint work with Bernard Chang (UT Austin), Javier E. Santos (Los Alamos National Laboratory, Agnese Marcato and Gianluca Boccardo (Politecnico Torino).

Masa Prodanovic

The University of Texas at Austin
masha@utexas.edu

MS118

Modeling Geologic Disposal of Co2 Using Machine Learning

Fractured systems in the subsurface play a role in many natural and engineered applications such as geologic carbon sequestration, hydraulic fracturing and underground nuclear test detection. Structural information (fracture size, orientation, etc.) plays a key role in governing the dominant physics for these systems but can only be known statistically. Traditional modeling approaches either ignore or idealize structural information at these larger scales because we lack a computational framework that utilizes it in its entirety. The work presented here integrates computational physics, machine learning and graph theory to make a paradigm shift from computationally intensive high-fidelity models to coarse-scale graphs without loss of critical structural information. We exploit the underlying discrete structure of fracture networks in systems considering flow through fractures and fracture propagation. We demonstrate that compact graph representations require significantly fewer degrees of freedom (dof) to capture micro-fracture information and further accelerate these models with Machine Learning.

Gowri Srinivasan

T-5, Theoretical Division, Los Alamos National Laboratory
gowri@lanl.gov

MS118

Fast Modeling CO₂ Migration Dynamics System with Deep Learning Methods

Carbon capture and storage (CCS) is a viable and promising technology to reduce the CO₂ emissions and mitigate climate change. In CCS, the captured anthropogenic CO₂ is injected into the geological media such as depleted reservoirs and deep saline aquifers for long-term sequestration. Numerical simulation is an essential tool for managing the subsurface flow, but suffers prohibitively high computational cost due to the nonlinearity of the multi-physics nature. Surrogate models, constructing approximations with greater efficiency for computational expensive numerical models, provide alternative approaches to predict the CO₂ sequestration process. A deep-learning-based surrogate model is developed in this work for predicting the temporal-spatial evolution of CO₂ plume migration in heterogeneous geological formations with high computational efficiency. The surrogate model is trained to predict the CO₂ saturation and solution CO₂-brine ratio in the saline aquifer, which are two secure parameters to calculate the amount of CO₂ trapped by dissolution. The deep learning workflow not only provides high predictive fidelity across temporal and spatial scales, but also offers a speedup of 5 orders of magnitude compared to traditional numerical simulation, and thus will be a significant predictive tool for engineers to manage the long-term process of CCS.

Mengjie Zhao

TU Delft

m.zhao-2@tudelft.nl

MS119

Estimate of Traffic Quantities Through Multiscale Second Order Models with Heterogeneous Data

We present a multiscale traffic model, based on the family of Second Order Generic Models, which integrates multiple trajectory data into the velocity function. It is specifically designed to take advantage of the availability of heterogeneous data. By heterogeneous data, we mainly mean data from different scales of observation and different monitoring modes. We refer in particular to Lagrangian data, which provide information on the trajectories followed by vehicles, and Eulerian data, which measure the transit of cars from fixed positions. We show how this combination of a second-order macroscopic model with microscopic information makes it possible to reproduce significant variations in speed and acceleration that strongly influence traffic emissions. Accurate approximations are obtained even with few trajectory data. The proposed approach is therefore a computationally efficient and accurate tool for calculating macroscopic traffic quantities and estimating emissions.

Maya Briani

CNR, Rome
maya.briani@gmail.com

MS119

Inverse Design in Conservation Laws

Inverse Design in Conservation Laws R.M.Colombo (In collaboration with Vincent Perrollaz and Abraham Sylla) Consider a scalar Conservation Law $\partial_t u + \partial_x f(x, u) = 0$ where the flux f possibly depends also on the scalar space variable x . Fix an arbitrary function $w = w(x)$ of x and an arbitrary positive time T . Does there exist an initial datum $u_o = u_o(x)$ that, provided to the conservation law, yields a solution $u = u(t, x)$ that evolves into the function w at time T , so that $u(T, x) = w(x)$? In addressing the answer to this question, the relations between Conservation Laws and Hamilton-Jacobi equations are first rigorously set in framework unique to both classes of equations and, then, exploited. Applications to the management of vehicular traffic are also considered.

Rinaldo M. Colombo

University of Brescia
INdAM Unit
rinaldo.colombo@unibs.it

MS119

Global Sensitivity Analysis of Pollutant Dispersion Uncertainty Quantification Problem

Air pollution is a very important issue challenging modern society from both public health and environmental perspectives. For example, the World Health Organisation (WHO) estimate that 4.2 million deaths annually are linked to ambient air pollution. The use of computational modelling enables us study pollutant dispersion models over a range of parameter values in order to gain some insight into the factors that strongly influence pollution concentration in the air. We study a parametric advection-diffusion pollution dispersion model with parameters sampled from probability distribution functions to investigate how the pollutant concentration varies with respect to randomness

in the model parameters. We investigate uncertainty quantification in terms of the randomness in the pollutant concentration and also carry out global sensitivity analysis in order to assess how the variability in the pollutant concentration can be attributed to the variability in the model parameters

Tasnia T. Shahid
University of Strathclyde
tasnia.shahid@strath.ac.uk

MS119

Reconstruction of Traffic Speed Distributions from Kinetic Models with Uncertainties

In this talk, we discuss the ability of Boltzmann-type kinetic equations for traffic dynamics to predict speed distributions observable from rough traffic data. We adopt, in particular, the formalism of uncertainty quantification, since driver reactions are uncertain and related to different types of driver behaviour or different classes of vehicles present in the traffic stream. Therefore, the calibration of the theoretical speed distribution has to face the reconstruction of the distribution of the uncertainty. We rely on experimental microscopic measurements of vehicle speeds (recorded on a German motorway) whose statistical distribution shows a multimodal trend. The calibration is performed by extrapolating the uncertain parameters of the kinetic distribution via a constrained optimisation procedure.

Michael Herty
RWTH Aachen University
herty@igpm.rwth-aachen.de

Andrea Tosin
Department of Mathematics
Politecnico di Torino
andrea.tosin@polito.it

Giuseppe Visconti
Department of Mathematics "G. Castelnuovo"
Sapienza University of Rome
giuseppe.visconti@uniroma1.it

Mattia Zanella
Department of Mathematics
University of Pavia
mattia.zanella@unipv.it

MS119

Combining Physics Models and Gaussian Processes for Traffic Prediction

We propose a Bayesian framework for traffic state prediction by macroscopic traffic flow models. Due to limited access to both trajectory and loop detector traffic data, we perform our analysis on synthetic data generated by numerical simulations. Classically, macroscopic traffic flow models are calibrated by fitting the so-called fundamental diagram i.e., the density-flow mapping described by the model flux function. However, data noise and congested traffic situations make the parameter identification process more difficult to deal with. Thus we consider an alternative approach based on [M. Kennedy, A. O'Hagan. *Bayesian calibration of computer models*. Journal of the Royal Statistical Society, 2001] which introduces a bias term to better account for possible discrepancies between the mathematical model and reality; this bias term is modeled by

a Gaussian process (GP). Once the calibration parameters are obtained, our analysis distinguishes between travel time estimation and prediction where the former is related to realized traffic scenarios. For the second one, we apply a GP to predict future traffic conditions at loop detector locations and sparse time points. These serve as initial data to simulate the traffic conditions at a finer scale, which enables us to do travel time prediction. Finally we compare the travel times between the ground truth and simulated data observing that the usage of the physics model on top of the GP improves the prediction accuracy.

Alexandra Wuerth
INRIA Sophia Antipolis
aalexandra.wuerth@inria.fr

Paola Goatin
Team ACUMES
INRIA
paola.goatin@inria.fr

Mickael Binois
Inria
mickael.binois@inria.fr

MS120

Variable Projection Networks in Autonomous Driving

Data representation has a crucial role in machine learning (ML), since it influences heavily the performance of these methods. Accordingly, the first step of traditional ML techniques includes data preprocessing, information extraction, and dimension reduction. This can be done manually via feature engineering or automatically via representation learning. There is a trade-off between these two approaches, which is to say, utilizing domain knowledge and promoting handcrafted features results in suboptimal performance, whereas learning representations directly from the data increases the performance, but also reduces the interpretability of the ML algorithms. In this talk, we propose VPNet, a novel model-driven neural network architecture by combining variable projections (VPs) and neural networks (NNs). VP is a classical method for solving separable nonlinear least squares problems (SNLLSs) by which we extract features automatically in our learning framework. Note that many inverse problems can be viewed as SNLLS data fitting problems including a small set of adjustable nonlinear parameters with direct physical interpretations. Thus, VPNet retains the powerful learning ability of ML methods, and the representation abilities of VPs in which both the trainable parameters and the extracted features are interpretable. As a case study, we show the potential of VPNet in road surface classification, which is an important environmental recognition problem in autonomous driving.

Gergo Bognár, Péter Kovács
Department of Numerical Analysis
Eötvös Loránd University
bognargergo@staff.elte.hu, kovika@inf.elte.hu

MS120

Interpretability in Multimodal Sensing-Based AI Methods for Evidence-Based Personalized Health Support

Enhancement of medical diagnosis using machine learning models is highly effective in health analytics. It al-

lows personalized health support with improved assessment and diagnosis. Multimodal sensing (MS) and deep learning (DL) models are considered to provide evidence-based health support. However, interpretability of the generated MS-DL models is challenging and require evidence-based result for better assessment. In this work, we propose a learning framework for assessment of stress in drivers using MS and regularized deep kernel learning with better interpretability. We acquire electrocardiography (ECG), electrodermal activity, photoplethysmography (PPG), and respiration rate from healthy drivers (N=10) in a controlled environment. Four hours of recordings are obtained after the removal of unstable segments. Multimodal features are extracted and applied to regularized deep kernel learning for fused biomarkers. Two classifiers, support vector machine (SVM) and random forest (RF), are used to differentiate stress states. Results show that the proposed approach can discriminate the drivers stress state. The combination of PPG, ECG, and RF yields the highest F1-score of 0.97. PPG and RF yield a maximum F1-score of 0.90 to classify the stress state of the drivers. ECG and PPG signals proved to be more reliable in classifying the stress state of drivers. Thus, the proposed framework could be extended to real-time stress state assessment for evidence-based health support.

Nagarajan Ganapathy
Indian Institute of Technology Hyderabad
nagarajan@bme.iith.ac.in

Vishal Singh Roha
Indian Institute of Technology Patna, Patna, India
vishal.roha95@gmail.com

Bhavnoor Singh Marok
Indian Institute of Technology Bombay, Mumbai, India
bhavnoormarok@cse.iitb.ac.in

MS120

Deep Unfolding for Data Estimation in Wireless Communication Systems

Traditionally, data estimation on the receiver side of a wireless digital communication system is accomplished with model-based methods. These methods are based on well-established physical and statistical models. Consequently, they are well-interpretable and performance bounds can often be derived. However, modeling errors, oversimplifications, wrong statistical assumptions, or insufficient model knowledge may severely degrade the performance of model-based approaches, and incorporating empirical statistics of possibly available data is usually difficult. Data-driven approaches can resolve some of the aforementioned issues. However, they usually suffer from huge data hunger, and they typically lack interpretability. Some of these problems may be tackled by incorporating model knowledge into data-driven methods, which is a major challenge with lots of open research questions. In this talk, we will present model-inspired neural networks (NNs) for data estimation, which are derived by using deep unfolding. These NNs are designed by unfolding the iterations of iterative model-based algorithms to layers of NNs. We will highlight similarities between model-based methods and model-inspired NNs. For example, we will show that conducting preconditioning, which is known to improve iterative model-based methods, can boost the performance of NNs that are derived by deep unfolding. We will compare these NNs to traditional model-based methods, and highlight their pros

and cons.

Mario Huemer, Stefan Baumgartner
Institute of Signal Processing
Johannes Kepler University Linz
mario.huemer@jku.at, stefan.baumgartner@jku.at

MS120

Examination of a Deep Neural Network for Biosignal Classification Reveals Features Similar to Clinical Criteria

Computer-aided diagnosis of biomedical signals culminated recently in deep learning (DL) methods, often outperforming human experts. However, these methods remain unadopted in clinical practice which is partially due to their lack in explainability. In this talk, we address this issue using electrocardiography (ECG) as use-case. We apply a pre-trained DL network for classification of 12-lead ECG signals that has been trained on more than 2 million signals and classify signals from a publicly-available ECG dataset. Additionally, we apply the method 'Integrated Gradients' which assigns a positive or negative relevance to each sample of a classified signals, pointing towards or against the classification; thereby allowing to understand the relationship between model prediction and learned features. We compare the relevance values of patients suffering from atrial fibrillation (AF) or left bundle branch block (LBBB) to healthy controls. Results show that the network learned cardiology textbook knowledge: Visible P-waves and concordant T-waves both result in a distinct negative relevance in AF and LBBB classification, respectively. Moreover, the difference in relevance between healthy controls and patients w.r.t. ECG leads is in line with clinical recommendations. In summary, our analysis suggests that the network learned features similar to clinical experience. Integrating this analysis in an application could be useful for quality control of diagnosis or teaching.

Nicolai Spicher
Department of Medical Informatics
University Medical Center Göttingen
nicolai.spicher@med.uni-goettingen.de

MS121

Efficient Algorithms for the Simulation of Incompressible Fluid Flow with High-Order Methods

The talk will present algorithms and implementations for the efficient and robust simulation of incompressible flow phenomena. The core algorithm are high-order discontinuous Galerkin discretizations, which are assessed in a holistic manner in terms of discretization methods in time and space, iterative solvers and preconditioners, and implementation techniques, with the goal to optimize overall computational efficiency. The key ingredient for practical simulations of marginally resolved flows at moderate to high Reynolds numbers is the energy stability, which can be fulfilled either by L2-conforming discontinuous Galerkin discretizations with techniques to weakly enforce normal continuity and pointwise divergence-free condition, or via H(div)-conforming Raviart-Thomas finite element pairs. The algorithms are assessed on several challenging test problems. The implementations are available through the ExaDG project.

Niklas Fehn
Technical University of Munich

niklas.fehn@tum.de

MS121

Solver Advances for High-Order Discretizations and Exascale Applications

Nek5000/RS, a highly-performant open-source spectral element code, has recently achieved an unprecedented milestone in the simulation of nuclear reactors: the first full core computational fluid dynamics simulations of reactor cores, including pebble beds with 352,625 pebbles and 98M spectral elements (51 billion gridpoints), advanced in less than 0.25 seconds per Navier-Stokes timestep. This talk will present performance and optimization considerations necessary to achieve this milestone when running on all of Summit. These optimizations led to a four-fold reduction in time-to-solution, making it possible to perform high-fidelity simulations of a single flow-through time in less than six hours for a full reactor core under prototypical conditions.

Paul Fischer

University of Illinois at Urbana-Champaign
fischerp@illinois.edu

MS121

Computational Contact Mechanics Using Fenicsx

Phenomena involving two objects coming into contact are ubiquitous in science and engineering. At the same time algorithms for computational contact mechanics are notoriously difficult to implement and only very few results regarding scalability of the algorithms in parallel are available. Our goal is to develop a scalable and robust implementation for frictional contact problems as an open-source extension to the FEniCSx framework. FEniCSx is the most recent version of FEniCS, an open source software for solving partial differential equations based on finite element methods. One of the most attractive features of FEniCS is the so-called Unified Form Language (ufl) which allows the user to express the PDE in a notation very close to mathematical notation. One major challenge is that integrals on the contact surfaces typically cannot be expressed as a ufl-form. In our implementation, we combine custom integration kernels on the contact surfaces with automatically generated kernels based on ufl-forms for all remaining integrals. Currently, we use Nitsche's method to enforce the contact constraints, but in principle other methods, such as Lagrange multiplier methods, can be implemented in a similar fashion.

Sarah Roggenendorf

University of Cambridge, United Kingdom
University of Cambridge, United Kingdom
sr957@cam.ac.uk

Joergen Dokken

Simula Research Laboratory
dokken@simula.no

Chris Richardson

University of Cambridge
cnr12@cam.ac.uk

Garth Wells

Department of Engineering
University of Cambridge

gnw20@cam.ac.uk

MS121

Performance-Portable Implicit Scale-Resolving Compressible Flow Using libCEED

Unstructured mesh CFD solvers with second order accurate spatial discretizations are flexible and ubiquitous in industry. Scale-resolving simulation techniques are viewed as the future for predictive fluid simulations, but standard technology is inefficient when repurposed for scale-resolving simulation. We show that greater efficiency can be gained using a stabilized-continuous-Galerkin finite element method with high-order basis functions on GPUs, using implicit time-stepping schemes. We explore trade-offs in the discretization and solver using a GPU implementation based on libCEED and PETSc, which enables portable performance from a single source code. We demonstrate accuracy, stability, and efficiency of our solver on turbulent compressible boundary layer flows.

James Wright, Jed Brown

University of Colorado Boulder
james.wrightiii@colorado.edu, jed@jedbrown.org

Kenneth Jansen

University of Colorado at Boulder
kenneth.jansen@colorado.edu

Leila Ghaffari

University of Colorado Boulder
leila.ghaffari@colorado.edu

MS122

Dual Polynomial Spaces in Finite Element Methods

In this presentation we will introduce a sequence of finite dimensional function spaces for high order methods, which form a de Rham sequence. Vector operations like the gradient, curl and divergence in this representation take the particularly simple form of incidence matrices, which are very sparse matrices containing only 1 and -1 as non-zero entries. The dual vector operations, in general, will not be sparse, topological matrices unless a de Rham sequence is constructed for the dual space. In this talk these dual spaces will be presented, yielding also highly sparse matrix representations for the dual vector operations. The mass matrices of the primal spaces serve as Riesz maps to switch between the primal and dual spaces. This construction immediately gives a representation for the trace spaces. These ideas are presented in [Jain, Zhang, Palha & Gerritsma, Construction and Application of Algebraic Dual Polynomial Representations for Finite Element Methods on Quadrilateral and Hexahedral Meshes, Computers and Mathematics with Applications, 95, 101-142, 2021]. Time permitting, some examples of the approach will be given.

Mark Gerritsma, Artur Palha

Delft University of Technology
m.i.gerritsma@tudelft.nl, a.palha@esciencecenter.nl

Varun Jain

TU Delft
v.jain@tudelft.nl

Yi Zhang

Sun Yat-Sen University

zhangyi_aero@hotmail.com

MS122

The Use of the Mimetic Method for Two-Phase Flows

In an industrial setting the separation of gas bubbles from the liquid phase can be a slow and costly process. It is therefore desirable to study the separation process in order to understand the mechanisms involved and thereby optimise the separation process. In alkaline hydrogen PtX plants the separation of hydrogen from the lye requires tanks, which are costly, due to the large nickel content in the tank material, and also due to the large size of the tanks that is required for fast gas separation. A dynamic two-phase model is presented using the mimetic discretisation method. The gas-liquid flow is modelled by a Navier-Stokes system of equations in an Eulerian representation. The motion of gas is modelled by a separate continuity equation.

Bo Gervang

Technical University of Denmark
boge@dtu.dk

MS122

High-Order Accurate Unified Discontinuous Galerkin Methods for Fluid and Solid Mechanics

A novel high-order accurate computational framework for a unified treatment of fluid and solid mechanic problems is presented. The framework is based on the expression of the governing equations of either fluids or solids as the sum of three differential operators, namely an operator involving only the time derivatives of the unknown solution fields, an advection operator involving the first-order spatial derivatives and an elliptic operator involving the first- and second-order spatial derivatives. A single formulation for both solids and fluids is obtained by a special expression of the elliptic operator. High-order accuracy in space is achieved by discretizing these operators via suitably-introduced discontinuous Galerkin methods, which lead to block-structured mass matrices, enable a natural way to treat generally-shaped mesh elements, and are amenable to massive parallelization. Numerical examples are provided to model fluids within the incompressibility regime, generally anisotropic solids and fluid-solid interaction problems.

Vincenzo Gulizzi

LBNL
vincenzo.gulizzi@unipa.it

MS122

Geometric Discontinuous Galerkin Methods for Fluids and Plasmas

Most conservative problems in fluid dynamics, plasma physics as well as many other branches of science and engineering have the form of hyperbolic conservation laws that inherit a Lagrangian and/or Hamiltonian structure. That is their dynamical equations can be obtained from an action principle or a Poisson bracket and a Hamiltonian functional, typically the total energy of the system. Non-conservative problems are usually composed of a conservative (Lagrangian or Hamiltonian) and a dissipative part. In both cases, it is important to preserve the structure of the conservative part in the course of discretisation in order to obtain stable numerical schemes that deliver accu-

rate and reliable simulation results. We discuss Hamiltonian structure-preserving discretisation approaches based on high-order Discontinuous Galerkin Spectral Element Methods (DGSEM). We show how these approaches relate to and generalise known energy-stable schemes based on split-forms and summation-by-parts properties. The inviscid Burgers equation and the compressible Euler equations serve as main examples. Generalisations to other important fluid and plasma systems are sketched. A remarkable property of the proposed approach is that exact mass, momentum and energy conservation can be achieved even if the system of equations is not cast in conservative form and momentum and energy do not explicitly appear as variables.

Michael Kraus

Technical University Munich
Germany
michael.kraus@ipp.mpg.de

MS123

Density Estimation in RKHS with Application to Korobov Spaces in High Dimensions

In this talk, we will consider a kernel method for estimating a probability density function (pdf) from an i.i.d. sample drawn from such density. Our estimator is a linear combination of kernel functions, the coefficients of which are determined by a linear equation. We will present an error analysis for the mean integrated squared error in a general reproducing kernel Hilbert space setting. Then, we will discuss how this theory can be applied to estimate pdfs for circular data. Under a suitable smoothness assumption, our method attains a rate arbitrarily close to the optimal rate.

Yoshihito Kazashi

University of Strathclyde
y.kazashi@uni-heidelberg.de

Fabio Nobile

EPFL, Switzerland
fabio.nobile@epfl.ch

MS123

Constructive Sparsification of Finite Frames with Application in Optimal Function Recovery

We present a new constructive subsampling technique for finite frames to extract minimal plain subsystems which preserve a good lower frame bound. The technique is based on a greedy type selection of frame elements to positively influence the spectrum of rank one updates of a matrix. It is a modification of the 2009 algorithm by Batson, Spielman, Srivastava and produces an optimal size subsystem without additional weights. It moreover achieves this in polynomial time and avoids the Weaver subsampling (based on the Kadison-Singer theorem) which has been applied in earlier work, yielding rather bad oversampling constants. In the second part of the talk we give applications for multivariate function recovery. Here we consider the particular problem of L_2 and L_∞ recovery from sample values. In this context, the presented subsampling technique allows to determine optimal (in cardinality) node sets even suitable for plain least squares recovery. It can be applied, for instance, to reconstruct functions in dominating mixed-smoothness Sobolev spaces, where we are able to discretize trigonometric polynomials with frequencies from a hyperbolic cross with nodes coming from

an implementable subsampling procedure. In addition we may apply this to subspaces coming from hyperbolic cross wavelet subspaces. Numerical experiments illustrate the theoretical findings. Joint work with: Felix Bartel (Chemnitz), Martin Schaefer (Chemnitz)

Tino Ullrich

University of Technology Chemnitz
tino.ullrich@math.tu-chemnitz.de

MS124

Entity: a Novel General Coordinate Particle-in-Cell Code for Plasmas Around Neutron Stars and Black Holes

Particle-in-cell has been the go-to approach for modeling plasmas in the environments of compact astrophysical objects for the last decade. Yet, there is no single publicly available code that includes all relevant radiation-plasma coupling processes and is capable of modeling global systems. In this talk I will describe the development of a new-generation PIC code for extreme astrophysical plasmas, Entity. The code is based on the Kokkos framework, which enables efficient implicit multi-architecture portability including GPUs. The code features algorithms for various radiation-plasma coupling processes, such as Compton scattering, production of electron-positron pairs and their annihilation. In its core, the code is designed in general coordinate system, defined by the metric functions; this enables the Entity to also efficiently tackle the global (full-system) models of the magnetospheres of compact objects, which require algorithms on non-cartesian (spherical, cubed sphere) non-uniform grids, and even full general relativity.

Hayk Hakobyan

Princeton Plasma Physics Laboratory
hayk.hakobyan@gmail.com

Sasha Philippov

University of Maryland
sashaph@umd.edu

MS124

Fully Kinetic Simulations of Pulsar Magnetospheres Using WarpX

Pulsars are rapidly rotating neutron stars immersed in strong electromagnetic fields that emit twin beams of electromagnetic radiation. However, the plasma composition and structure in the region surrounding pulsars, called magnetospheres, and the physical processes that drive particle acceleration leading to the observed spectra are not well understood. Global pulsar magnetosphere simulations are required to answer these questions. However, resolving the current sheet skin-depth which is $O(10^6)$ smaller than the pulsar radius for realistic systems, is intractable even on large supercomputers. Thus, the magnetic field strength in global PIC simulations is typically scaled-down, restricting the maximum energy of the charged particles. We will present the effect of scaling down the magnetic field on particle acceleration, energy dissipation, and Poynting flux. We use WarpX, a highly scalable, electromagnetic PIC code with advanced algorithms to mitigate numerical artifacts in mesh-refinement simulations. We will also present 3D simulations to study the effect of pulsar obliquity and plasma injection rate on the plasma structure and Poynting flux. Additionally, we explore the use of ultra-high-order spectral methods (PSATD) to perform pul-

sar magnetosphere simulations and compare their accuracy and performance with traditional finite-difference methods.

Revathi Jambunathan, Hannah Klion
Lawrence Berkeley National Laboratory
rjambunathan@lbl.gov, klion@lbl.gov

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Jean-Luc Vay
Lawrence Berkeley National Laboratory
Berkeley, CA
JLVay@lbl.gov

Michael Rowan, Ann S. Almgren
Lawrence Berkeley National Laboratory
mrowan@lbl.gov, asalmgren@lbl.gov

Rémi Lehe
Lawrence Berkeley National Laboratory
ATAP
rlehe@lbl.gov

Axel Huebl
Lawrence Berkeley National Laboratory (LBNL), USA
axelhuebl@lbl.gov

Axel Huebl
Lawrence Berkeley National Laboratory (LBNL), USA
Helmholtz-Zentrum Dresden-Rossendorf (HZDR),
Germany
axelhuebl@lbl.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

MS124

Simulating the Life and Death of Massive Stars with Efficient Radiation Transport

During the formation of the densest and most massive star clusters, the intense radiation field from young stars carries significant momentum that can dominate the dynamics and efficiency of the star-forming process. Following the detailed transport of radiation in 3D hydrodynamical simulations can reveal multi-dimensional effects that would otherwise be lost in models assuming simple geometries. As massive stars approach the end of their lives and explode as core-collapse supernovae, radiation transport through the stellar ejecta also leaves informative imprints of the progenitor stars' properties on the observed light curves and spectra. Reliably capturing the relevant radiative processes is instrumental to proper inference of the progenitors' stellar properties. In this talk, I will discuss the application of Monte Carlo methods in numerical radiation transport, including acceleration schemes optimized to deliver efficient calculations in high-density media that shape the life and

death of massive stars.

Benny Tsang

University of California, Berkeley, USA
benny.tsang@berkeley.edu

MS125

Inferring the Basal Sliding Coefficient Field for the Stokes Ice Sheet Model under Rheological Uncertainty

Modeling of the dynamics of polar ice sheets is critical to enable relevant projections of future sea levels. One of the parameters that influences the predictions is the basal sliding coefficient. We consider the problem of inferring the basal sliding coefficient of an ice sheet from noisy surface velocity measurements. The dynamics of the ice are modeled by a nonlinear Stokes equation. Inference is further complicated by the fact that the rheological parameters are unknown and uncertain. The standard approach would be to jointly infer the basal sliding coefficient and rheological parameters. However, the rheological parameters are not of primary interest and are distributed throughout the domain, thus greatly increasing the computational costs and ill-posedness associated with the problem. To avoid carrying out joint inference, we employ the Bayesian approximation error (BAE) approach which allows for offline premarginalization over the rheological parameters. We show numerical examples which demonstrate that incorrectly fixing the rheological parameters may result in unphysical artifacts in the reconstructed basal sliding coefficient, and posterior infeasibility (ie. true parameter is not supported by the posterior model) due to underestimation of uncertainty in the reconstructions.

Olalekan Babaniyi

University of California, Merced
obsma@rit.edu

Ruanui Nicholson

Department of Engineering Science
University of Auckland
ruanui.nicholson@auckland.ac.nz

Umberto Villa

Oden Institute for Computational Engineering & Sciences
uvilla@oden.utexas.edu

Noemi Petra

University of California, Merced
npetra@ucmerced.edu

MS125

Machine Learning Solvers of PDEs with Uncertainty Quantification

We propose a physics-constrained neural network (NN) approach to solve partial differential equations (PDEs) without labels. We express the loss function of these NNs in terms of the residual of PDEs obtained through a discrete, convolution operator-based, and vectorized implementation. We explore an encoder-decoder NN structure for both deterministic and probabilistic models, with Bayesian NNs (BNNs) for the latter. For BNNs, the discretized residual is used to construct the likelihood function. Deterministic and probabilistic convolutional layers are used to learn the applied boundary conditions (BCs) and to detect the problem domain. Dirichlet and Neumann BCs are specified as inputs to NNs, and we explore whether a single

NN can solve for similar physics; i.e., the same PDE, but with different BCs and on a number of problem domains. The trained BNN PDE solvers demonstrate a degree of success at extrapolated predictions for BCs that they were not exposed to during training. We demonstrate the capacity and performance of the proposed framework by applying it to different steady-state and equilibrium boundary value problems of diffusion, linear and nonlinear elasticity. Such NN solution frameworks are important for high-throughput solutions of PDEs with different boundary conditions and on varying domains in support of design and decision-making.

Krishna Garikipati, Xiaoxuan Zhang

Mechanical Engineering
University of Michigan
krishna@umich.edu, xxzh@umich.edu

MS125

Incorporating Full-Field Diagnostic Data into Constitutive Model Discovery via Physics Informed Neural Networks

The Virtual Fields Method (VFM) and Finite Element Model Updating (FEMU) have shown success in calibrating constitutive model parameters for nonlinear materials. However, VFM faces difficulties where plane-stress assumptions break down, and FEMU remains challenging when calibrating complex models where computational inefficiencies of the inverse problem are intractable. In this presentation a novel physics informed neural network (PINN) architecture for calibration and discovery of hyperelastic constitutive models will be proposed which alleviates some of the weaknesses of VFM and FEMU due to the interpolation power of neural networks. Our architecture consists of radial basis function neural networks and input convex neural networks as a constitutive model surrogate trained with a combined least-squares and gradient descent update. This approach ensures that conservation of momentum is obeyed to higher precision than typical PINNs while searching for constitutive models that obey convexity and thus material stability. Exemplars utilizing synthetic experimental data will be shown that are illustrative of typical mechanical test specimens with non-homogenous strain fields. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Craig M. Hamel

Sandia National Laboratories
chamel@sandia.gov

MS125

A Machine Learning Enabled Scale Bridging Framework for Phase Transformations in Materials Physics

We propose a machine learning-enabled scale bridging framework that combines density functional theory (DFT) calculations at the atomistic scale with phase-field modeling at the continuum scale to understand the impact of phase stability on microstructure evolution. We accomplish scale bridging by incorporating traditional statistical mechanics methods with integrable deep neural networks, allowing formation energies for specific atomic configurations to be coarse-grained and incorporated in a neural

network description of the free energy of the material. The resulting realistic free energy functions enable atomistically informed phase-field simulations. Here, we use Li_xCoO_2 (LCO) as a model system to benchmark our scale-bridging framework. Li_xTMO_2 (TM=Ni, Co, Mn) are promising cathodes for Li-ion batteries, whose electrochemical cycling performance is strongly governed by crystal structure and phase stability as a function of Li content at the atomistic scale. Our computational results allow us to make connections to experimental work on LCO cathode degradation as a function of temperature, morphology, and particle size.

Mostafa Shojaei
University of Michigan
mfsh@umich.edu

Krishna Garikipati
Mechanical Engineering
University of Michigan
krishna@umich.edu

MS126

Equipping Neural PDE Surrogates with Better Geometric Priors

Partial differential equations (PDEs) see widespread use in sciences and engineering to describe simulation of physical processes as scalar and vector fields interacting and co-evolving over time. Due to the computationally expensive nature of their standard solution methods, neural PDE surrogates have become an active research topic to accelerate these simulations. However, the practical utility of training such surrogates is contingent on their ability to model complex multi-scale spatio-temporal phenomena. In this talk, we address two challenges of neural network PDE surrogates: (i) How to take into account the relationship between different fields and their internal components, which are often correlated? We therefore view the time evolution of correlated fields through the lens of multivector fields allows which allows us to introduce new operations that are grounded on the algebraic properties of Clifford algebras. (ii) How to model multi-scale phenomena and generalize across timescales and equations? We therefore analyze design considerations of Fourier and U-Net based PDE surrogate models, showing promising results on generalization to different PDE parameters and time-scales with a single surrogate model. We conclude by given an outlook how the proposed methods directly help to tackle imminent challenges in neural PDE surrogate modeling.

Johannes Brandstetter
Microsoft Research - AI4Science
johannesb@microsoft.com

MS126

Fusion Plasma Turbulence Simulation with Neural Network Surrogate Models

Accurate predictive modelling of tokamak turbulence is a key component of multiphysics simulation of fusion reactors. Many-query applications such as scenario optimization, experimental design, and controller design, demands both fast and accurate modelling, infeasible with direct numerical simulation. Surrogate models generated with machine learning methods circumvent the conflicting constraints of accuracy and tractability. A key enabling step is the development of reduced-order models, validated by direct numerical simulation. Reduced-order model calculation time is then sufficient for constructing extensive

databases of model input-output mappings using HPC resources. These are used as training sets for neural network regression. A key aspect is the physics-informed customization of regression variables and optimization cost functions, to capture known features of the system. The resultant surrogate models accurately reproduce the original reduced-order turbulence model with significant speedup, providing near-realtime capability, 1 trillion times faster than the anchoring direct numerical simulations. We summarize the state-of-the-art in tokamak turbulent transport neural network surrogate development, ranging from practical considerations on label generation, model training, and demonstration applications for scenario optimization at the JET tokamak and extrapolations to the next-generation ITER performance.

Jonathan Citrin
DIFFER - Dutch Institute for Fundamental Energy
Research
j.citrin@diffier.nl

MS126

Using Deep Neural Networks to Learn Memory Effects in Integro-Differential Equations

Memory effects are ubiquitous in a wide variety of complex physical phenomena, ranging from climate models and metamaterials to glassy dynamics. The Generalised Langevin Equation formalism provides a rigorous way to describe all memory effects via the so-called memory kernel in an integro-differential equation. However, the memory kernel is often unknown, and accurately predicting or measuring it via e.g. a numerical inverse Laplace transform remains a herculean task. In this talk we will first review the importance and difficulties of measuring realistic memory kernels, and describe our recent progress using deep neural networks (DNNs) to tackle this problem. We demonstrate that DNNs are remarkably robust against noisy data and can therefore successfully circumvent the bottleneck that plagues conventional inverse Laplace transform approaches. Furthermore we show that DNNs are capable of tackling the notoriously long-lived memory effects of glassy systems. Our work will provide a general purpose tool for extracting memory kernels from glasses and supercooled liquids, and, with sufficient data, from a broad range of other non-Markovian systems.

Max Kerr Winter, Liesbeth Janssen
Eindhoven University of Technology
m.j.kerr.winter@tue.nl, l.m.c.janssen@tue.nl

MS126

A PDE Structure-Preserving Deep Learning Approach to Predict Turbulent Reacting Flow Dynamics

High-fidelity simulations of complex chemically reacting flows are very expensive due to high temporal and spatial resolution requirements. Recently, there has been a significant progress in physics-informed deep learning (DL) approaches, which embed prior physics and domain knowledge into DL models, reducing the computational cost while retaining satisfactory solution accuracy. We utilize a PDE-preserved Neural Network (PPNN) approach. This architecture consists of a NN solver based on Convolutional ResNet, where the known governing PDEs are discretized and solved in a low-resolution grid, and high-resolution predictions are obtained with encoding-decoding convolution operations with trainable filters. We extend the

capability of the PPNN framework to turbulent reacting flows, demonstrated in the context of a methane tribrachial flame. The data are generated by solving the unsteady NS equations along with finite rate chemistry (FRC). In classical simulators, the addition of FRC at times requires even higher temporal resolution due to the inherent stiffness of the chemical system. It is shown that the PPNN framework offers promises in predicting the flow variables, temperature and species mass fractions within reasonable accuracy at different Reynolds numbers. Furthermore, the multi-resolution information passing approach in PPNN shows superior training efficiency and long-term prediction accuracy compared to the baseline Convolutional ResNet architecture.

Pushan Sharma
Stanford University
ps729@stanford.edu

MS127

Enable Predictive Reduced-Order Models for Turbulent Reacting Flows Using Adaptive Model Order Reduction

Even with exascale computing capabilities, high-fidelity, full-scale simulations of turbulent combustion in realistic applications like rocket combustion remain computationally expensive and inaccessible for many-query applications. Projection-based model order reduction (PMOR) methods have shown promise in greatly improving computational efficiency. However, classical MOR methods that seek reduced solutions in low-dimensional subspaces fail for realistic turbulent combustion problems because reacting flows feature extreme stiffness, sharp gradients, and multi-scale transport, which pose great challenges in deriving effective low-order representations and developing predictive reduced-order models (ROMs). In this talk, we introduce an adaptive reduced-order modeling technique which updates the low-dimensional space, thus circumventing representation barriers faced by static reduced-dimensional spaces. The method leverages model-form preserving least-squares projections with variable transformation (MP-LSVT) for improved robustness of ROM and adapts the low-dimensional subspaces based on the evaluated dynamics during online calculations to enable predictive ROMs for turbulent reacting flows.

Cheng Huang
University of Kansas
Department of Aerospace Engineering
chenghuang@ku.edu

MS127

Fully Flexible Gaussians for Laser Induced Quantum Dynamics

Fully flexible Gaussians for laser induced quantum dynamics Laser induced quantum dynamics pose oscillatory dynamics, typically in high dimensions coupling various temporal and spatial scales. The talk presents joint work with Adamowicz, Kvaal, Lubich, and Pedersen on sums of Gaussian ansatz functions for this challenging application. The Gaussians have a nonlinear parametrization with complex width matrices and phase space centers that evolve according to the time-dependent variational principle and related approximation methods.

Caroline Lasser
Technische Universität München

classer@ma.tum.de

MS127

Time-Marching Schemes for Dynamical Low Rank Approximations

The Dynamical Low Rank (DLR) method is a time-dependent reduced basis approach attractive for the quantification of uncertainties of time-dependent random PDEs for which the optimal low-rank subspace may vary significantly over time. The true solution is approximated by an expansion of deterministic modes weighted by stochastic coefficients, all time-dependent, mimicking a truncated Karhunen-Loewen expansion; evolution equations are available for the deterministic and stochastic modes. Time-integration of the DLR equations may be challenging as the DLR structure may lead to additional stiffness in the system. Recently, we have shown that a projector-splitting scheme, combined with standard first-order time integration of the operator, leads to a fully discrete DLR solution of a random parabolic PDE that verifies the same norm bound properties as a standard (non-DLR) discrete parabolic system, and is not affected by the additional stiffness in the equations. In this talk, we propose extensions of the projector-splitting scheme, including a practical Strang scheme as well as other time-marching schemes which improve on the accuracy of the approximation and its numerical stability compared to the explicit integrators while not being overly expensive.

Yoshihito Kazashi
University of Strathclyde
y.kazashi@uni-heidelberg.de

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Thomas Trigo Trindade
EPFL
thomas.trigotrindade@epfl.ch

Eva Vidlickova
EPFL, Lausanne
eva.vidlickova@epfl.ch

MS127

Adaptive Sampling for Efficiently Training Models of Nonlinear Latent Dynamics

A typical approach for training nonlinear parametrizations such as deep neural networks to approximate solutions of partial differential equations is to minimize the residual at collocation points. However, if the spatial domain is high-dimensional or the dynamics are transport dominated with local features such as waves, then just uniformly sampling collocation points becomes intractable because each collocation point entails at least one residual evaluation in each iteration of the training. In this work, we propose an active learning scheme that adapts collocation points to follow the dynamics described by the partial differential equations. This means that the solution dynamics are leveraged to adapt collocation points to where the residual is high. In experiments with high-dimensional evolution equations and equations that transport local features such as waves, the proposed active scheme learns accurate approximations of the solutions from few samples whereas traditional collocation methods based on static sampling

fail to provide meaningful predictions.

Yuxiao Wen, Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
yw3210@nyu.edu, pehersto@cims.nyu.edu

MS128

Uncertainty Quantification for Electromagnetic Scattering Problems via First-Order Sparse Boundary Element Approximation

Electromagnetic scattering problem under small magnitude shape randomness is an important field of research for robust design and high precision industries. Due to the limited and probability measure dependent convergence rate of Monte-Carlo-based simulations, we follow the first-order shape boundary method (FOSB), whose theoretical aspects for electromagnetic scattering were introduced in [Jerez-Hanckes and Schwab, IMA Journal of Numerical Analysis (2017)]. FOSB is a deterministic method based upon shape calculus, boundary integral equations and sparse tensor approximation, allowing for a computation of the statistical moments with a poly-logarithmic complexity for prescribed accuracy. We conduct extensive numerical experiments regarding accuracy and performance. The latter validates our claims and shows broad applicability of the technique. Hence, further improvements are suggested.

Paul Escapil-Inchauspe
Pontificia Universidad Católica de Chile,
Santiago, Chile
paul.escapil@edu.uai.cl

Carlos Jerez-Hanckes
Universidad Adolfo Ibañez, Santiago, Chile
carlos.jerez@uai.cl

MS128

Multilevel Domain Uncertainty Quantification In Computational Electromagnetics

We study the numerical approximation of time-harmonic electromagnetic fields for the Maxwell lossy cavity problem for uncertain geometries. We adopt the same affine-parametric shape parametrization framework, mapping the physical domains to a nominal polygonal domain with piecewise smooth maps. The regularity of the pullback solutions on the nominal domain is characterized in piecewise Sobolev spaces. We prove error convergence rates and optimize the algorithmic steering of parameters for edge-element discretizations in the nominal domain combined with: (a) multilevel Monte Carlo sampling, and (b) multilevel, sparse-grid quadrature for computing the expectation of the solutions with respect to uncertain domain ensembles. In addition, we analyze sparse-grid interpolation to compute surrogates of the domain-to-solution mappings. All calculations are performed on the polyhedral nominal domain, which enables the use of standard simplicial finite element meshes. We provide a rigorous fully discrete error analysis and show, in all cases, that dimension-independent algebraic convergence is achieved. For the multilevel sparse-grid quadrature methods, we prove higher order convergence rates which are free from the so-called curse of dimensionality, i.e. independent of the number of parameters used to parametrize the admissible shapes. Numerical experiments confirm our theoretical results and

verify the superiority of the sparse-grid methods.

Rubén Aylwin, Carlos Jerez-Hanckes
Universidad Adolfo Ibañez, Santiago, Chile
rbayl1993@gmail.com, carlos.jerez@uai.cl

Christoph Schwab
ETH-Zurich
schwab@sam.math.ethz.ch

Jakob Zech
Heidelberg University
jakob.zech@uni-heidelberg.de

MS128

Efficient Adaptive Stochastic Collocation Strategies for Advection-Diffusion Problems with Uncertain Inputs

Physical models with uncertain inputs are commonly represented as parametric partial differential equations (PDEs). That is, PDEs with inputs that are expressed as functions of parameters with an associated probability distribution. Developing efficient and accurate solution strategies that account for errors on the space, time and parameter domains simultaneously is highly challenging. Indeed, it is well known that standard polynomial-based approximations on the parameter domain can incur errors that grow in time. In this work, we focus on advection-diffusion problems with parameter-dependent wind fields. A novel adaptive solution strategy is proposed that allows users to combine stochastic collocation on the parameter domain with off-the-shelf adaptive timestepping algorithms with local error control. This is a non-intrusive strategy that builds a polynomial-based surrogate that is adapted sequentially in time. The algorithm is driven by a so-called hierarchical estimator for the parametric error and balances this against an estimate for the global time-stepping error which is derived from a scaling argument.

Benjamin Kent
The University of Manchester
benjamin.kent@manchester.ac.uk

Catherine Powell
University of Manchester
catherine.powell@manchester.ac.uk

David Silvester
School of Mathematics
University of Manchester
david.silvester@manchester.ac.uk

Malgorzata Zimon
IBM Research UK
Daresbury Laboratory
malgorzata.zimon@uk.ibm.com

MS128

An Adaptive Combination Technique for High-Dimensional Problems in Uncertainty Quantification

We present an adaptive multilevel approximation for the computation of quantities of interest involving the solution a stochastic elliptic PDE where the diffusion coefficient is parametrized by means of a Karhunen-Loève expansion. One common approach is to approximate the equivalent

parametric problem by a restriction of the countably infinite dimensional parameter space to a finite-dimensional parameter set and subsequently apply a spatial discretization and an approximation in the parametric variables. In order to reduce the computational effort, we combine these different approximation steps using the sparse grid combination technique. By this approach, no knowledge about the decay of the Karhunen-Love coefficients and about the regularity of the solution is required. In a further step, we consider, additionally to the adaptivity in the sparse grid approach, an adaptive spatial discretization which allows also to adjust to singularities in the spatial variable.

Uta Seidler

University of Bonn
Institute for Numerical Simulation
seidler@ins.uni-bonn.de

Michael Griebel

Institut für Numerische Simulation, Universität Bonn & and Fraunhofer-Institut für Algorithmen und Wissenschaftlich
griebel@ins.uni-bonn.de

MS129

TBA

no text

Axel Klawonn

Universität zu Köln
axel.klawonn@uni-koeln.de

Martin Lanser

Universitaet zu Koeln
Mathematisches Institut
martin.lanser@uni-koeln.de

Janine Weber

Universität zu Köln
Department of Mathematics and Computer Science
janine.weber@uni-koeln.de

MS129

A Nonlinear Preconditioning Strategy Based on Residual Learning for PDEs

We talk about a nonlinearly preconditioned inexact Newton method for solving highly nonlinear system of algebraic equations from the discretization of PDEs. From a large number of numerical experiments, we observe that when the inexact Newton stagnates or fails to converge, the space of residuals often contains a subspace that is difficult to reach by the usual Newton direction. We introduce a learning technique to identify this subspace and then speedup the convergence.

Li Luo

Faculty of Science and Technology
University of Macau
lilu@um.edu.mo

Xiao-Chuan Cai

University of Colorado, Boulder
Dept. of Computer Science

xccai@um.edu.mo

MS130

Fast and Scalable Algorithms for Bayesian Optimal Experimental Design

Bayesian optimal experimental design (OED) is a principled framework for maximizing information gained from limited data in Bayesian inverse problems. Unfortunately, conventional methods for OED are prohibitive when applied to expensive models with high-dimensional parameters. In this talk, I will present fast and scalable computational methods for large-scale Bayesian OED with infinite-dimensional parameters, including data-informed low-rank approximation, efficient offline-online decomposition, projected neural network approximation, and a new swapping greedy algorithm for combinatorial optimization.

Keyi Wu

Meta
keyiwu.w@gmail.com

Thomas O'Leary-Roseberry

The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Peng Chen, Omar Ghattas

The University of Texas at Austin
peng@cc.gatech.edu, omar@oden.utexas.edu

MS130

Source Detection on Graphs

The source detection problem consists of a weighted connected graph and a signal spreading through it with the following properties. The signal is sent with a constant velocity from a unique unknown source node. For each node there is the possibility to measure the time at which the signal reaches the node. The goal is to find the source node with as few measurements as possible without knowledge of the starting time or the velocity of the signal. This talk takes two cases into account. In the deterministic offline case the combinatorial concept of the spread dimension is introduced. Assuming exact computation and no measurement errors, the objective is to determine as few as possible measurement nodes to uniquely locate the source node, no matter which node actually turns out to be the source. In the stochastic online case it is discussed how to find a probable source with an iterative algorithm using parameter estimation and experimental design. This iterative approach requires a repetitive nature of the signal so that in each iteration the information gained in the previous iterations can be used.

Sarah Feldmann

OVGU
sarah.feldmann@ovgu.de

Sebastian Sager, Volker Kaibel

Otto-von-Guericke Universität Magdeburg
sager@ovgu.de, kaibel@ovgu.de

MS130

Maximin-Efficient Experimental Design for Or-

thogonally Invariant Information Criteria

Optimal selection of measurement locations for parameter estimation is usually focused on maximizing an optimality criterion defined on the Fisher information matrix (FIM) associated with the estimated parameters. But different optimality criteria may yield different optimal designs. In consequence, strong interest is generated by compromise solutions which would produce decent values for a broadest possible class of design criteria. In the search for such universally suboptimal solutions, an approach is proposed to compute designs maximizing the minimal efficiency with respect to the class of orthogonally invariant information criteria. This class is broad enough to include all optimum design criteria encountered in practice. It turns out that the minimal efficiency with respect to this class equals that with respect to a finite set of criteria which generalize the well-known E-optimum design criterion. An extremely simple and fast algorithm is proposed to numerically construct such designs, which is based on the appropriate semi-infinite programming problem formulation. Its idea is to alternate between solving the eigenproblem for the current FIM and finding a solution to a linear-programming problem. In this way the nondifferentiability of the optimality criterion is circumvented and the algorithm is guaranteed to converge in a finite number of steps.

Dariusz Ucinski

Institute of Control and Computation Engineering
University of Zielona Góra
d.ucinski@issi.uz.zgora.pl

MS130**Advances and Challenges in Solving High-Dimensional HJB Equations Arising in Optimal Control**

This talk presents recent advances in neural network approaches for approximating the value function of high-dimensional control problems. A core challenge of the training process is that the value function estimate and the relevant parts of the state space (those likely to be visited by optimal policies) need to be discovered. We show how insights from optimal control theory can be leveraged to achieve these goals. To focus the sampling on relevant states during neural network training, we use the Pontryagin maximum principle (PMP) to obtain the optimal controls for the current value function estimate. Our approaches can handle both stochastic and deterministic control problems. Our training loss consists of a weighted sum of the objective functional of the control problem and penalty terms that enforce the HJB equations along the sampled trajectories. Importantly, training is self-supervised, in that, it does not require solutions of the control problem. We will present several numerical experiments for deterministic and stochastic problems with state dimensions of about 100 and compare our methods to existing approaches.

Deepanshu Verma

Emory University
dverma4@emory.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University

lruthotto@emory.edu

MS131**Dissipativity-Preserving Reduced-Order Modeling from Data**

Reduced-order modeling from data with dissipativity preservation is discussed in this talk. Employing the data informativity framework, the dissipativity of all systems consistent with noisy data can be characterized by a data-based linear matrix inequality (LMI). Furthermore, semi-definite programming duality helps us to prove the existence of minimal and maximal solutions to the LMI. As in the classical bounded-real and positive-real balanced truncation, these extremal solutions play a role in the computation of well-approximating reduced-order models carrying the dissipativity property. As an additional advantage of using this balancing-type method, a priori error bounds of the reduced-order models are available.

Azka Muji Burohman, Bart Besselink

University of Groningen
a.m.burohman@rug.nl, b.besselink@rug.nl

Jacquelin M. Scherpen
Rijksuniversiteit Groningen
j.m.a.scherpen@rug.nl

Kanat Camlibel
University of Groningen
Institute of Mathematics and Computer Science
m.k.camlibel@rug.nl

MS131**Energy Matching in Reduced Passive and Port-Hamiltonian Systems**

Conventional structure-preserving *model order reduction* (MOR) methods for *port-Hamiltonian* (pH) systems focus on approximating the input-output dynamics by (approximately) minimizing classical system norms, such as the Hardy \mathcal{H}_2 norm. Nevertheless, the definition of a pH system consists of two objects: the input-output dynamics and an energy function that is typically called the Hamiltonian. If we thus measure the approximation quality only with respect to the input-output dynamics, then the approximation of the Hamiltonian is not reflected at all. This is particularly relevant since recent results [Breiten and Unger, *Automatica* 142, 2022] demonstrate that modifying the Hamiltonian of the *full order model* (FOM) may yield better *reduced order models* (ROMs), at least if only the input-output dynamics are analyzed. In this talk, we make a first step towards the dual-objective minimization problem of optimal approximation of the input-output dynamics and the Hamiltonian by noticing that in the pH ROM, we can modify the Hamiltonian without changing the system dynamics. Thus, for a given pH ROM, we can search for the reduced Hamiltonian that best approximates the FOM Hamiltonian. We prove that the resulting optimization problem is strictly convex and uniquely solvable. Moreover, we propose a numerical algorithm to solve the minimization problem and demonstrate its applicability with numerical examples.

Jonas Nicodemus

University of Stuttgart
Stuttgart Center for Simulation Science (SC SimTech)
jonas.nicodemus@simtech.uni-stuttgart.de

Benjamin Unger
University of Stuttgart, Germany
benjamin.unger@simtech.uni-stuttgart.de

Paul Schwerdtner
TU Berlin
schwerdt@math.tu-berlin.de

MS131

Data-Driven Parametric Reduced-Order Model for Aircraft Flutter Monitoring

This talk deals with the aircraft flutter analysis and detection. Such a phenomenon, responsible for stability loss and structure destruction, is commonly explained by the so-called flutter equation aggregating the structural dynamics, the aerodynamic forces and the actuators forces. Such a coupling leads to the following flexible aircraft model description:

$$(s^2 M + sD + K)x = Q(s_i, \rho_j)x + Bu \text{ and } y = Cx,$$

blending a structural dynamics (M, D, K) and the so-called generalized aeroelastic forces $(Q(s_i, \rho_j))$. More in details, M, D, K and B, C are real-valued matrices while $Q(s_i, \rho_j)$ is a complex-valued matrix known at frozen complex s_i values (usually on the imaginary axis) and real parameter ρ_j values (function of the flight configuration). Some challenges are to evaluate the flutter instability occurrence via the frequency/parameter couple, and second, to prevent it by monitoring the critical damping value. To address both objectives, SISO and MIMO parametric reduced order models are first constructed using the Loewner framework. These latter are then used to design a flutter detection mechanism involving nonlinear observers.

Charles Poussot-Vassal

Onera
charles.poussot-vassal@onera.fr

Tea Vojkovic
TU Delft
t.vojkovic@tudelft.nl

Jesus Toledo, Alex dos Reis de Souza
Onera
jesus.toledo_zucco@onera.fr,
alex.dos_reis_de_souza@onera.fr

David Quero
DLR
david.queromartin@dlr.de

Pierre Vuillemin
Onera
pierre.vuillemin@onera.fr

MS131

Structured Optimization-Based Reduction and Identification of Dynamical Systems

We discuss the benefits of optimization-based algorithms for the model order reduction (MOR) and system identification of structured dynamical systems. For this, we recall our MOR method [P. Schwerdtner and M. Voigt, SOB MOR: Structured Optimization-Based Model Order Reduction, Preprint arXiv:2011.07567, 2020] and present recent extensions for the approximation of structured systems with an additional parameter dependency [P. Schw-

erdtner and M. Schaller, Structured Optimization-Based Model Order Reduction for Parametric Systems, Preprint arXiv:2209.05101, 2022] and differential-algebraic equations [P. Schwerdtner, T. Moser, V. Mehrmann, and M. Voigt, Structure-Preserving Model Order Reduction for Index One Port-Hamiltonian Descriptor Systems, Preprint arXiv:2206.01608, 2022]. After that, we show how our method can be modified to perform frequency domain system identification to generate structured systems from potentially noisy data. We demonstrate the effectiveness of our method by several numerical experiments on well-accepted benchmark examples.

Paul Schwerdtner

TU Berlin
schwerdt@math.tu-berlin.de

Tim Moser
Technische Universität München
tim.moser@tum.de

Manuel Schaller
Technische Universität Ilmenau, Germany
manuel.schaller@tu-ilmenau.de

Matthias Voigt
UniDistance Suisse
matthias.voigt@fernuni.ch

Volker Mehrmann
Technische Universität Berlin
mehrmann@math.tu-berlin.de

MS132

Preconditioning a GPU-Enabled Compressible CFD Solver

Oak Ridge National Laboratory is developing the Spinnaker code, an implicit time-dependent compressible-flow computational fluid dynamics solver targeting a variety of fluid flow applications. Spinnaker is built using finite element assembly, linear algebra, and automatic differentiation capabilities from the Trilinos code library. The Kokkos performance portability layer allows kernels to be compiled for several different computing architectures, including multithreaded CPUs, Nvidia GPUs, and AMD GPUs. Spinnaker uses an implicit time-stepping approach with Newton's method as a nonlinear solver within each time step. GMRES (potentially with restarting) is typically used as the linear solver. The preconditioner selection has primarily focused on overlapping additive Schwarz type preconditioners with a sparse direct solver used for the local domain blocks. Preconditioners based on incomplete factorizations (also within an additive Schwarz approach) and algebraic multigrid methods have also been evaluated. This talk will provide an evaluation of the linear solvers and preconditioners used within Spinnaker for several compressible 2D and 3D fluid flow problems. A variety of CPU-only and GPU-enabled preconditioners are evaluated, including native options in Trilinos as well as the SuperLU and Nvidia cuSOLVER libraries. Results will include parallel scaling studies and profiling to illustrate current performance bottlenecks.

Steven Hamilton

ORNL

hamiltonsp@ornl.gov

MS132

Preparing Multigrid Solvers in Hypre for Exascale Computers

The emerging exascale computers provide opportunities to perform much larger scale simulations to obtain more accurate solutions than ever before. The increasing complexities of heterogeneous accelerators on such platforms have made the development of sparse linear solvers challenging to achieve high performance. In this talk, we will discuss the porting strategies, new developments and performance optimizations of the multigrid solvers in hypre in preparation for the exascale computers with the results from real application codes.

Ruipeng Li

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
li50@llnl.gov

MS132

Efficient Fine-Grain Parallel Solution Strategies for Extremely Ill-Conditioned Linear Systems

Design computations for complex engineering systems typically generate extremely sparse and ill-conditioned linear systems. Their nature makes it challenging to solve them efficiently in a fine-grain parallel fashion. Iterative linear solvers, which can be parallelized effectively, are not performing well on such ill-conditioned problems. Traditional supernodal and multifrontal parallel strategies used in direct linear solvers are not effective with extremely sparse systems, because dense blocks created by these methods are too small to take advantage of fast dense linear algebra. These challenges impede deployment of design computations to heterogeneous compute platforms, which use hardware accelerators such as GPUs. We propose linear solver strategy tailored to fine-grain parallel design computations. We take advantage of design computations creating a sequence of linear systems with the same sparsity pattern to perform first factorization on the host and then move sparse factors to the hardware accelerator. Each subsequent linear solve is reusing the same ordering and symbolic factorization and performs numerical (re)factorization using a partial pivoting method. This strategy requires advanced iterative refinement to recover solution of required quality. We present our results and show performance improvement over current state of the art. We discuss future research directions to develop robust sparse linear solvers that perform well on hardware accelerators.

Katarzyna Swirydowicz

Pacific Northwest National Laboratory
kasia.swirydowicz@pnnl.gov

Slaven Peles

Oak Ridge National Laboratory
peless@ornl.gov

MS132

Hykkt: A Hybrid Direct and Iterative Method for Solving Kkt Linear Systems

We propose a solution strategy for the large indefinite linear systems arising in interior methods for nonlinear optimization. The method is suitable for implementation on

hardware accelerators such as graphical processing units (GPUs). The current gold standard for sparse indefinite systems is the LBL factorization where L is a lower triangular matrix and B is 1X1 or 2X2 block diagonal. However, this requires pivoting, which substantially increases communication cost and degrades performance on GPUs. Our approach solves a large indefinite system by solving multiple smaller positive definite systems, using an iterative solver on the Schur complement and an inner direct solve (via Cholesky factorization) within each iteration. Cholesky is stable without pivoting, thereby reducing communication and allowing reuse of the symbolic factorization. We demonstrate the practicality of our approach on large optimal power flow problems and show that it can efficiently utilize GPUs and outperform LBL factorization of the full system.

Shaked Regev

Stanford University
Pacific Northwest National Lab
sregev@stanford.edu

Nai-Yuan Chiang

Lawrence Livermore National Laboratory
7000 East Avenue, Livermore, CA 94550
chiang7@llnl.gov

Eric F. Darve

Institute for Computational and Mathematical
Engineering
Stanford University
darve@stanford.edu

Cosmin G. Petra

Lawrence Livermore National
Center for Advanced Scientific Computing
petra1@llnl.gov

Michael A. Saunders

Stanford University
saunders@stanford.edu

Katarzyna Swirydowicz

Pacific Northwest National Laboratory
kasia.swirydowicz@pnnl.gov

Slaven Peles

Oak Ridge National Laboratory
peless@ornl.gov

MS132

Stopping Criteria for Coarsest-Grid Solver in Multigrid V-Cycle Method

Multigrid methods are frequently used when solving systems of linear equations, applied either as standalone solvers or as preconditioners for iterative methods. Within each cycle, the approximation is computed using smoothing on fine levels and solving on the coarsest level. With growth of the size of the problems that are being solved, the size of the problems on the coarsest grid is also growing and their solution can become a computational bottleneck. In practice the problems on the coarsest-grid are often solved approximately, for example by Krylov subspace methods or direct methods based on low rank approximation; see, e.g., [M. Huber, Massively parallel and fault-tolerant multigrid solvers on peta-scale systems, Ph.D. Thesis, Technical University of Munich, Germany, 2019], [Buttari et al., Numerical Linear Algebra with Applications (2021)]. The

accuracy of the coarsest-grid solver is typically determined experimentally in order to balance the cost of the solves and the total number of multigrid cycles required for convergence. In this talk, we present an approach to analyzing the effect of approximate coarsest-grid solves in the multigrid V-cycle method for symmetric positive definite problems. We discuss several stopping criteria derived based on the analysis and suggest a strategy for utilizing them in practice. The results are illustrated through numerical experiments.

Petr Vacek
Charles University
vacek@karlin.mff.cuni.cz

Erin Claire C. Carson
Charles University MFF
carson@karlin.mff.cuni.cz

Kirk M. Soodhalter
Trinity College Dublin
ksoodha@maths.tcd.ie

MS133

Stability Analysis of the Asynchronous Discontinuous-Galerkin Method

The asynchronous computing approach significantly improves the scalability of time-dependent partial differential equation (PDE) solvers on massively parallel supercomputers. The method relaxes data synchronization between processing elements (PEs) at a mathematical level, allowing computations to proceed regardless of communications status. Recently, we developed an asynchronous discontinuous-Galerkin (ADG) method, which can provide accurate solutions in the absence of data communication and/or synchronization. As the arrival time of messages between PEs is random, the latest available time level at the buffer nodes near PE boundaries is also random. This makes the update equations at the boundary elements of the PDE solver stochastic, posing a challenge in analyzing the stability properties of the underlying schemes using standard methods. In this work, we develop a novel Fourier stability analysis coupled with the eigen-decomposition of the amplification matrix to analyze the properties of schemes based on the ADG method. We investigate errors incurred in terms of dispersive and dissipative nature. Unlike the standard DG method, where the stability properties depend on the schemes as well as the simulation parameters like grid spacing and time step, we show that the properties of the ADG method can also depend on the number of PEs used and the amount of delay in communication. Our analysis results show that the ADG schemes are stable but with stricter CFL conditions.

Shubham K. Goswami, Konduri Aditya
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
shubhamkg@iisc.ac.in, konduriadi@iisc.ac.in

MS133

Application of ParaDiag to Finite Element Atmospheric Models

Numerical weather prediction involves solving very large simulations on an operational schedule. This requires vast amounts of computing power, so highly scalable algorithms are essential on massively parallel modern hardware. The ParaDiag method achieves this by exposing time paral-

lelism, in addition to the spatial parallelism exposed by traditional domain decomposition. We present recent progress on the application of ParaDiag to finite element discretizations of PDEs for atmospheric flow, which are particularly challenging for time-parallel methods due to their hyperbolic nature. These solvers are implemented as an open source general library using Firedrake, an automated code generation framework for the solution of finite element methods. Various ParaDiag formulations are explored, including quasi-Newton iterations and the preconditioned Newton-Krylov method, and their parallel scaling is compared. The performance of ParaDiag is compared against theoretical convergence estimates. We describe our approaches to the solution of the block systems within the ParaDiag matrix and show how these can be realised within the library.

Joshua Hope-Collins
Imperial College London
joshua.hope-collins13@imperial.ac.uk

Colin J. Cotter
Imperial College London
Department of Mathematics
colin.cotter@imperial.ac.uk

MS133

Parallel-in-Time for Weather and Climate Simulations

Weather and climate simulations face new challenges due to changes in computer architectures caused by physical limitations. From a pure computing perspective, algorithms are required to cope with stagnating or even decreasing per-core speed and increasing on-chip parallelism. These and other showstopping trends will continue and already led to research on partly disruptive mathematical and algorithmic reformulations of dynamic cores, e.g., using (additional) parallelism along the time dimension. This presentation provides an overview and introduction of a selection of the following promising newly developed and evaluated time integration methods for equations related to prototypical dynamical cores, all aimed at improving the ratio of wall clock time vs. error: Exponential Integration, Rational Approximation of Exponential Integration (REXI), Parallel Full Approximation Scheme in Space and Time (PFASST) and Semi-Lagrangian methods combined with Parareal. We get improved time-vs.-error rates, but sometimes with additional challenges on the way which need to be further overcome. Overall, our results motivate further investigation and combination of these methods for operational weather/climate systems. I gratefully acknowledge collaborators related to this presentation Jed Brown, Finn Capelle, Francois Hamon, Richard Loft, Michael Minion, Nathanael Schaeffer, Andreas Schmitt, Pedro S. Peixoto, Raphael Schilling

Martin Schreiber
Université Grenoble Alpes
martin.schreiber@univ-grenoble-alpes.fr

MS133

Towards Parallel in Time Methods for Compatible Finite Element Discretizations of the Shallow Water Equations

Compatible finite element methods are of interest to the weather and climate modelling community due to their conservation and wave propagation properties on non-

orthogonal meshes such as the cubed-sphere. These non-orthogonal meshes allow for better scaling from spatial domain decomposition than meshes based on the latitude-longitude grid which have grid points clustered at the poles. However, parallel scaling is still limited by the sequential timestepping schemes commonly used in these applications. Here I will present recent progress towards investigating various parallel in time algorithms for compatible finite element discretisations of the shallow water equations. I will describe an exponential integrator that can be computed in parallel using a rational approximation and show how this can be used to construct a coarse propagator for the parareal algorithm. I will also outline how this can be extended to equation sets used to model vertically stratified flows.

Jemma Shipton
University of Exeter
j.shipton@exeter.ac.uk

Colin J. Cotter
Imperial College London
Department of Mathematics
colin.cotter@imperial.ac.uk

Beth Wingate
University of Exeter
B.Wingate@exeter.ac.uk

MS133

Towards Exponential Semi-Lagrangian Parallel-in-Time Methods for the Shallow Water Equations on the Rotating Sphere

The hyperbolic nature and the development of small-scale spatial features in mathematical models for atmospheric circulation constitute major challenges for parallel-in-time (PinT) numerical simulations in climate modeling and numerical weather prediction. Therefore, it is of great importance to develop coarse time stepping schemes that are able to improve the stability and convergence of PinT methods. In this work, we explore the application of semi-Lagrangian exponential methods for the two- and multilevel parallel-in-time simulation of the shallow water equations (SWE) on the rotating sphere, using Parareal and MGRIT. As a main feature, exponential methods (e.g. ETDRK) solve exactly the linear terms of the governing equations, and their recently proposed semi-Lagrangian version (SL-ETDRK) was shown to be more stable in the serial integration of the SWE on the plane. We then consider these methods on the sphere and in the PinT context. We perform stability studies on a linearized ODE which indicate that Parareal and MGRIT using ETDRK and SL-ETDRK as coarse schemes have improved stability properties compared to the use of the well-known SL-SI-SETTLS method, which turns out to be highly unstable in the PinT framework despite of its successful application in serial simulations of atmospheric circulation. Numerical tests of standard benchmarks confirm that exponential schemes, notably in the semi-Lagrangian framework, are promising choices for parallel-in-time integration.

Joao Steintraesser
Universidade de São Paulo
joao.steintraesser@usp.br

Pedro Peixoto
Instituto de Matemática e Estatística
Universidade de São Paulo, Brazil

pedrosp@ime.usp.br

Martin Schreiber
Université Grenoble Alpes
martin.schreiber@univ-grenoble-alpes.fr

MS134

A Quantitative Justice Primer: The #MetaMath Project and Other Examples

The application of techniques, tools and topics from various quantitative sciences (e.g., applied mathematics, data science, computer science, etc.) in subject domains that are derived from the social sciences (e.g., political science, law, economics, history, etc.) with the goal of promoting social justice has been dubbed “quantitative justice.” This talk shall provide a primer on quantitative justice (QJ) for attendees of CSE23 and include several examples of work that can be described under the QJ moniker. The #MetaMath project is one of several initiatives to emerge from the ICERM Data Science and Social Justice program held in Summer 2022 and continuing in Summer 2023. The primary idea animating #MetaMath is the application of quantitative justice techniques to the mathematics community itself. Some things that fall under the #metaMath tag include meditations on the many definitions of the term “mathematician” and their implications and significance, suggestions for how to address the paucity of robust demographic data on important aspects of the mathematics community, and outlines for future projects that could fit under the aegis of #MetaMath.

Ron Buckmire
Occidental College
ron@oxy.edu

MS134

Fast Summation on a Sphere

Fast summation techniques have proven to be of great importance in a variety of fields. In this talk, I will present a new technique for performing fast summations on spheres, which is suitable for sums coming from spherical convolutions. I present applications of this technique to problems coming from atmospheric and oceanic modeling.

Anthony Chen
University of Michigan
cygnari@umich.edu

MS134

Interfacing in-Silico and in-Situ Experiments of Organismal Fluid Pumping

Far from the surface, the ocean’s midwater present a rich frontier of biodiversity that is not well understood. Part of this gap in our knowledge is the great expense involved in collecting data with remotely operated vehicles. In this presentation, we will discuss the pipeline of developing in-silico computational experiments in concert with in-situ experimental data. Using a combination of particle image velocimetry data, optical scans, and confocal microscopy, we will discuss the creation of fluid-structure interaction models for organismal pumping and fluid transport, with the goal of developing an intuition on the physical mechanisms that drive their success. Using a combination of simplified geometries and scanned body meshes, we will employ the immersed boundary/finite element (IB/FE) method to

simulate chambered, valveless pumping mechanism generated by the pelagic tunicate known as a larvacean. Additionally, we will use the same modeling methodology to explore the metachronal motion and fluid transport of the parapodial paddles of the pelagic, midwater polychaete known as tomopterids. All motion described in these systems will not be prescribed and will emerge from the interaction of active muscular tension, passive elastic recoil, and the local fluid environment.

Alexander P. Hoover
Cleveland State University
a.p.hoover@csuohio.edu

MS134

Exact Divisibility by the Powers of Some Binary Number Sequences

We obtain exact divisibility results for the powers of the balancing and Lucas balancing numbers. This gives all the results which are analogous to those for Fibonacci and Lucas numbers from 1970 to 2019. For example, Hoggatt and Bicknell-Johnson (1977) and Benjamin and Rouse (2009) proved that if $F_n^k | m$, then $F_n^{k+1} | F_{nm}$, which was later generalized by Pongsriiarn (2014) to include the exact divisibility such as $F_n^{k+1} | F_{nm}$ provided that $F_n^k | m$, $n = 3$, and $n \not\equiv 3 \pmod{6}$. Here F_n is the n th Fibonacci number. For the balancing numbers B_n , we show that $B_n^k | m$ if and only if $B_n^{k+1} | B_{nm}$ for all $k = 1$ and $m, n = 2$. The corresponding results for the Lucas-balancing numbers are also given

Asim Patra
National Institute of Technology Rourkela
INDIA
asimp1993@gmail.com

MS135

Mixed-Precision Randomized Solution of Least Squares Problems

We consider the preconditioned solution of least squares problems, where a randomized preconditioner is computed in lower precision. The idea is to reduce the amount of sampling by demoting the randomization to lower precision. For a tall and skinny $m \times n$ matrix, $\mathcal{O}(n)$ samples produce a randomized preconditioner that is highly ill-conditioned in double precision, but much better conditioned after demotion to single precision. From a deterministic perturbation perspective, we present lower bounds for the smallest singular values that explain the singular value increase in lower precision. These are genuine lower bounds, in contrast to existing expressions which hold only to first and second order. From a probabilistic perspective, we model perturbations as independent, bounded random variables, and present expressions for the expectation of the small singular values. Numerical experiments corroborate the effectiveness of the preconditioners.

Christos Boutsikas
Department of Computer Engineering and Informatics
University of Patras, Greece
boutsikas@ceid.upatras.gr

Petros Drineas
Purdue University
Computer Science
pdrineas@purdue.edu

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

MS135

Randomized Algorithms for Rounding in the Tensor-Train Format

The Tensor-Train (TT) format is a highly compact low-rank representation for high-dimensional tensors. The fundamental operation used to maintain feasible memory and computational time is called *rounding*, which truncates the internal ranks of a tensor already in TT format. We propose several randomized algorithms for this task that are generalizations of randomized low-rank matrix approximation algorithms and provide significant reduction in computation compared to deterministic TT-rounding algorithms. Randomization is particularly effective in the case of rounding a sum of TT-tensors (where we observe $20\times$ speedup), which is the bottleneck computation in the adaptation of GMRES to vectors in TT format. We present the randomized algorithms and compare their empirical accuracy and computational time with deterministic alternatives.

Hussam Al Daas
STFC, Rutherford Appleton Laboratory
Computational Mathematics Group
hussam.al-daas@stfc.ac.uk

Grey Ballard
Wake Forest University
ballard@wfu.edu

Paul Cazeaux
Department of Mathematics
University of Kansas
pcazeaux@ku.edu

Eric Hallman
North Carolina State University
erhallma@ncsu.edu

Agnieszka Miedlar
Virginia Tech
amiedlar@vt.edu

Mirjeta Pasha
Tufts University
mirjeta.pasha@tufts.edu

Timothy Reid, Arvind Saibaba
North Carolina State University
twreid@ncsu.edu, saibab@ncsu.edu

MS135

Randomized Algorithms for Tikhonov Regularization in Linear Least Squares

We describe three randomized algorithms to efficiently solve regularized linear least squares systems based on sketching. The algorithms compute preconditioners for $\min \|Ax - b\|_2^2 + \lambda \|x\|_2^2$, where $A \in \mathbb{R}^{m \times n}$ and $\lambda > 0$ is a regularization parameter, such that LSQR converges in $\mathcal{O}(\log(1/\epsilon))$ iterations for ϵ accuracy. We focus on the context where the optimal regularization parameter is unknown, and the system must be solved for a number of parameters λ . Our algorithms are applicable in both the

underdetermined $m \ll n$ and the overdetermined $m \gg n$ setting. Our algorithms efficiently update preconditioners for new regularization parameters. We introduce an algorithm specifically for an approximately low-rank setting, in which the matrix A has rapidly decreasing singular values and such the problem is of low statistical dimension. The scheme we propose exploits the low statistical dimension while not requiring the computation of the Gram matrix, resulting in a more stable scheme than existing algorithms in this context.

Maike Meier, Yuji Nakatsukasa
University of Oxford
maiike.meier@maths.ox.ac.uk,
nakatsukasa@maths.ox.ac.uk

nakat-

MS135

On the Parallelization of Sketching Algorithms for the Tensor-Train Decomposition

In this talk, we propose TT-Sketching, a new parallelizable tensor-train decomposition algorithm for streaming tensor data. We introduce a couple of variants of this algorithm for computation and storage efficiency. For these variants, we provide theoretical guarantees of accuracy, parallel implementation details using message passing interface (MPI), and scaling analysis. Strong scaling results on different tensors suggest that TT-Sketching is better than its serial counterpart, and scales well with the number of computing cores.

Tianyi Shi, Maximilian Ruth, Alex J. Townsend
Cornell University
tianyishi@lbl.gov, mer335@cornell.edu,
townsend@cornell.edu

MS136

Inferring Functions with Uncertain Regularity

This work describes a Bayesian framework for reconstructing functions with uncertain regularity, i.e. roughness vs. smoothness. The regularity of functions carries crucial information in many inverse problem applications, e.g., in medical imaging for identifying malignant tissues or in the analysis of electroencephalogram for epileptic patients. We characterize the regularity of a function with its fractional differentiability. We propose a hierarchical Bayesian method which estimates a function and its regularity, simultaneously. In addition, we quantify the uncertainties in the estimates. Numerical results suggest that the proposed method is a reliable approach for estimating functions in different types of inverse problems. Furthermore, this is a robust method under various noise types, noise levels, and incomplete

Babak Maboudi Afkham, Nicolai Riis
Technical University of Denmark
bmaaf@dtu.dk, nabr@dtu.dk

Yiqiu Dong
Technical University of Denmark
Department of Applied Mathematics and Computer Science
yido@dtu.dk

Per Christian Hansen
Technical University of Denmark
DTU Compute

pcha@dtu.dk

MS136

Variable Projection Methods for Solving Separable Nonlinear Inverse Problems

Variable projection methods are among the classical and efficient methods to solve separable nonlinear least squares problems. In this talk, I will introduce the variable projection method, its use to solve large-scale blind deconvolution problems, and some new variants that preserve the edges in the solution.

Malena Espanol
Arizona State University, U.S.
mespanol@asu.edu

MS136

Computational Imaging with Low Precision Arithmetic

Although some work has been done to exploit mixed precision computations for inverse problems, most previous work focuses on particular applications, such as image registration and image reconstruction. Extended (high) precision has been proposed for inverse problems to control significant digits to avoid the influence of rounding errors. In this talk we consider a different perspective: because we cannot expect to precisely know data, we develop and analyze solvers for general inverse problems that can take advantage of low precision speed of modern computer architectures, and which can be used in a variety of imaging applications.

James Nagy
Emory University
jnagy@emory.edu

MS136

Deterministic and Bayesian Spatio-Temporal Methods with Edge-Preserving Priors for Inverse Problems

Inverse problems are ubiquitous in many fields of science such as engineering, biology, medical imaging, atmospheric science, and geophysics. Three emerging challenges on obtaining relevant solutions to large-scale and data-intensive inverse problems are ill-posedness of the problem, large dimensionality of the parameters, and the complexity of the model constraints. In this talk we discuss efficient methods for computing solutions to dynamic inverse problems, where both the quantities of interest and the forward operator may change at different time instances. We consider large-scale ill-posed problems that are made more challenging by their dynamic nature and, possibly, by the limited amount of available data per measurement step. In the first part of the talk, to remedy these difficulties, we apply efficient regularization methods that enforce simultaneous regularization in space and time (such as edge enhancement at each time instant and proximity at consecutive time instants) and achieve this with low computational cost and enhanced accuracy. In the remainder of the talk, we focus on designing spatio-temporal Bayesian Besov priors for computing the MAP estimate and quantifying uncertainties in large-scale and dynamic inverse problems. Numerical examples from a wide range of applications, such as tomographic reconstruction, image deblurring, and multichannel dynamic tomography are used to illustrate the

effectiveness of the described approaches.

Mirjeta Pasha
Tufts University
mirjeta.pasha@tufts.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Silvia Gazzola
University of Bath
s.gazzola@bath.ac.uk

Malena Espanol
Arizona State University
malena.espanol@asu.edu

Eric De Sturler
Virginia Tech
sturler@vt.edu

Shiwei Lan, Shuyi Li
Arizona State University
slan7@asu.edu, shuyili3@asu.edu

MS136

The Horseshoe Prior for Bayesian Inverse Problems in Imaging

In inverse problems arising in imaging science characterization of sharp edges in the solution is desired. Within the Bayesian approach to these problems, edge-preservation is often achieved using Markov random field priors based on heavy-tailed distributions. Another strategy, popular in sparse statistics, is the application of hierarchical shrinkage priors. An advantage of this formulation lies in expressing the prior as a conditionally Gaussian distribution depending on heavy-tailed distributed hyperparameters. In this presentation, we revisit the shrinkage horseshoe prior and introduce its formulation for edge-preserving settings. Moreover, we discuss a Gibbs sampling framework to solve the Bayesian inverse problem. Image deblurring and computed tomography applications show that our computational procedure is able to compute sharp edge-preserving posterior point estimates with reduced uncertainty.

Felipe Uribe, Yiqiu Dong
Technical University of Denmark
Department of Applied Mathematics and Computer Science
felipe.uribe@lut.fi, yido@dtu.dk

Per Christian Hansen
Technical University of Denmark
DTU Compute
pcha@dtu.dk

MS137

Can the Hartree-Fock Kinetic Energy Exceed the Exact Kinetic Energy?

The Hartree-Fock (HF) approximation has been an important tool for quantum-chemical calculations since its earliest appearance in the late 1920s, and remains the starting point of most single-reference methods in use today. Intuition suggests that the HF kinetic energy should not exceed the exact kinetic energy, but no proof of this con-

jecture exists, despite a near century of development. Beginning from a generalized virial theorem derived from scaling considerations, we derive a general expression for the kinetic energy difference that applies to all systems. For any atom or ion this trivially reduces to the well-known result that the total energy is the negative of the kinetic energy and since correlation energies are never positive, proves the conjecture in this case. Similar considerations apply to molecules at their equilibrium bond lengths. We use highly precise calculations on Hooke's atom (two electrons in a parabolic well) to test the conjecture in a non-trivial case, and to parameterize the difference between density-functional and HF quantities, but find no violations of the conjecture.

Steven Crisostomo

Department of Physics and Astronomy
University of California, Irvine
crisosts@uci.edu

MS137

Topological Index and Homotopy in Coupled Cluster Theory

We propose a comprehensive mathematical framework for Coupled-Cluster-type methods based on topological degree theory. This allows us to establish more general existence results and deduce local information about the solutions of the CC equations. The idea of constructing a homotopy for CC theory is not new, and has been extensively studied in the past. We consider the more recent Kowalski-Piecuch (KP) homotopy from a mathematical point of view and use it as a theoretical tool to prove the existence of a truncated CC solution. This follows from a more general result guaranteeing the existence of a whole solution curve of the KP homotopy.

Mihaly A. Csirik
Hylleraas Centre for Quantum Molecular Sciences
Department of Chemistry, University of Oslo
m.a.csirik@kjemi.uio.no

Andre Laestadius
Hylleraas Centre for Quantum Molecular Sciences
Department o
andre.laestadius@kjemi.uio.no

MS137

Domain Wall Dynamics in Cubic Magnetostrictive Materials Subject to Rashba Effect and Nonlinear Dissipation

In this talk, we present an analytical investigation of domain wall motion occurring along the major axis of a thin magnetostrictive nanostrip perfectly arranged on the top of a thick piezoelectric actuator. The motion is driven by magnetic fields, spin-polarized currents, and spin-orbit torque effects and takes place in cubic magnetostrictive materials characterized by a nonlinear dissipation. The main aim is to describe how magnetoelasticity, Rashba field, dry-friction, chemical composition, and crystal symmetry affect the steady and precessional dynamics of magnetic domain walls. In detail, it is here analytically inspected how the key features (threshold, breakdown, domain wall mobility, and propagation direction) can be effectively manipulated by the above contributions. Finally, the theoretical results are numerically illustrated for realistic materials, revealing a satisfying qualitative agreement with experimental

observations.

Sumit Maity

National Institute of Technology Andhra Pradesh
sumitmaity1304@gmail.com

MS137

On Solutions to the Hartree-Fock Equations with a Self-Generated Magnetic Field

In this talk we will explore the existence and regularity of solutions to the Hartree-Fock equations coupled with a self-generated magnetic field. In the second part of the talk we will explore some methods to numerically approximate said solutions.

Rafael Reyes

Department of Mathematics
University of California, Santa Barbara
raan.lare@math.ucsb.edu

MS137

Modified-Operator Method for the Calculation of Band Diagrams of Crystalline Materials

In solid state physics, electronic properties of crystalline materials are often described by the spectrum of periodic Schrödinger operators. As a consequence of Bloch's theorem, the numerical computation of electronic quantities of interest involves computing derivatives or integrals over the Brillouin zone of so-called energy bands, which are piecewise smooth, Lipschitz continuous periodic functions obtained by solving a parametrized elliptic eigenvalue problem on a Hilbert space of periodic functions. Classic discretization strategies for resolving these eigenvalue problems produce approximate energy bands that are either non-periodic or discontinuous, both of which cause difficulty when computing numerical derivatives or employing numerical quadrature. We present here an alternative discretization strategy based on an ad hoc operator modification approach. While specific instances of this approach have been proposed in the physics literature, we describe a systematic formulation of this operator modification approach. We derive a priori error estimates for the resulting energy bands and we show that these bands are periodic and can be made arbitrarily smooth (away from band crossings) by adjusting suitable parameters in the operator modification approach. We also present numerical experiments involving a toy model in 1D, graphene in 2D, and silicon in 3D to validate our theoretical results and showcase the efficiency of the operator modification approach.

Eric Cancès

Ecole des Ponts Paristech
eric.cances@enpc.fr

Muhammad Hassan

Post-doctoral Researcher
Laboratoire Jacques-Louis Lions, Sorbonne Université
muhammad.hassan@sorbonne-universite.fr

Laurent Vidal

CERMICS
École Nationale des Ponts et Chaussées (ENPC)

laurent.vidal@enpc.fr

MS138

On the Distributed Computation of Mixed-Strategy Nash Equilibria

We address the problem of distributedly computing mixed-strategy Nash equilibria for repeated (generalized mixed-integer) games, specifically, a class with monotone and Lipschitz continuous pseudo-gradient mapping and jointly convex feasible set for fixed integer variables. We design efficient distributed algorithms that exploit the particular structure of the problem, namely the simplicial feasible sets of the mixed strategies. To that end, we exploit an operator theoretic approach and propose an extension of the forward-reflected-backward splitting method that incorporates the Bregman distance of a Legendre function. Then, we use the negative entropy function to obtain a closed-form formula for updating the mixed strategies, resulting in fast and cheap computations compared to standard Euclidean-projection-based methods. We illustrate the performance and effectiveness of our algorithms via numerical simulations on some engineering applications.

Wicak Ananduta

TU Delft
w.ananduta@tudelft.nl

Sergio Grammatico

ETH
grammatico.sergio@gmail.com

MS138

Evolution of Wardrop Equilibria and Braess's Paradox As Demand Varies

In this talk we consider routing games with affine cost functions and analyze the evolution of Wardrop equilibria (WE) and the Braess's paradox as the demand increases. Regarding Wardrop equilibria, we show that the computation of such an evolution essentially boils down to solving a finite number of variational inequalities. Such a computation also provides insights on how the travel cost under WE changes as demand increases. In particular, we derive the direction in which this cost will change as the demand increases unboundedly, and leverage this result to find a set of roads which will carry no flow once the demand has increased beyond a certain threshold. Regarding evolution of the Braess's paradox, a phenomenon which occurs when removal of a path from the network counterintuitively decreases the travel cost under equilibrium, we show that if a path displays Braess's paradox at a particular demand, then it is necessarily beneficial for the network at a lower demand. We also establish that the paradox happens only on a finite interval of demand.

Ashish Cherukuri

Universiteit Groningen
a.k.cherukuri@rug.nl

Jasper Verbree

University of Groningen
j.verbree@rug.nl

MS138

A General Blackbox Optimization Framework for

Hyperparameter Optimization in Deep Learning

The uprising of deep learning has brought new challenges in the blackbox optimization (BBO) community. Notably, tuning the hyperparameters of a deep model is a mixed-variable BBO problem with an unfixed structure. For instance, the number of hidden layers, which is itself a hyperparameter, affects the number of hyperparameters regarding the units in the architecture of the model. Moreover, the hyperparameter optimization problem (HPO) may simultaneously contain categorical, integer and continuous variables. In conjunction, the diversity of variable types and the unfixed structure of the HPO create a substantial challenge in a BBO context. To tackle the HPO, we meticulously developed a mathematical framework that properly models the class of problems of interest. Meta variables are introduced to model variables that influence which variables are included or excluded. In essence, meta variables model the unfixed structure of the problem, and they may affect the dimension of the problem. Many algorithmic subtleties and implications are outlined by the mathematical framework, which ease the development of optimization methods. A simple HPO problem is illustrated throughout the presentation and solution strategies to tackle meta and categorical variables are discussed with different approaches, such as direct search, heuristics and Bayesian optimization.

Edward Hallé-Hannan

Polytechnique Montréal

University of Montreal

edward.halle-hannan@polymtl.ca

MS139

On PETSc Amd GPU Performance Benchmarking Using the Hip Backend

In this talk, we present an overview of the HIP backend introduced into the PETSc library. This backend essentially completes the AMD GPU port of the Mat class hence allowing full GPU utilization of KSP and SNES solvers. Key math libraries like hipSPARSE and hipBLAS were leveraged to enable portable performance across all AMD hardware. Benchmarking results on MI250X GPUs demonstrate competitive performance. Insight into the general direction of AMD GPU hardware for scientific computing will also be discussed.

Justin Chang, Suyash Tandon

Advanced Micro Devices, Inc

justin.chang@amd.com, suyash.tandon@amd.com

MS139

A Multi-GPU Accelerated Linear Solver for Large-Scale Topology Optimization

Topology optimization is an increasingly popular tool for engineers to obtain lightweight yet stiff designs in an automated way. However, the design resolution is directly linked to the computational time required for the optimization. This is especially true for large-scale applications where a small design resolution is required. For instance, optimizing the structure of TU Delft's Flying V airplane (65 meter wingspan) [F Faggiano, *et al.*, Aerodynamic design of a flying-V aircraft, 2017] with a 10mm resolution requires multiple solutions to linear systems of equations with many billions of unknowns. This presents challenges both regarding memory usage and computation time. To address these challenges, we propose a matrix-free linear

solver, accelerated with GPUs. Throughout the domain, identical trilinear finite-elements (voxels) are used, thus only one unique element matrix needs to be calculated. Next, sectors of 888 voxels are used to model the entire geometry of the airplane, preventing redundant computations in regions that are not part of the airplane. Within each sector, regularity of the data is kept, thus ensuring coalesced memory access patterns for a high utilization of GPUs. This, together with multigrid preconditioning and multi-GPU strategies, results in the ability to perform extremely large-scale topology optimization at unprecedented speeds. The technique is showcased with optimization of the Flying V airplane.

Arnoud Delissen, Jun Wu

Delft University of Technology

a.a.t.m.delissen@tudelft.nl, j.wu-1@tudelft.nl

Matthias Moller

Delft Institute of Applied Mathematics (DIAM)

Delft University of Technology

M.Moller@tudelft.nl

Alejandro Aragón, Matthijs Langelaar

Delft University of Technology

a.m.aragon@tudelft.nl, m.langelaar@tudelft.nl

MS139

Distributed, GPU-Enabled Support Vector Machine Classification of Remote Sensing Data Using Permon and PETSc

Popular machine learning frameworks provide multi-threading parallelism on shared memory and are designed as off-the-shelf packages. They are commonly used as high-level building blocks for data processing pipelines; however, they typically lack extended settings of underlying optimization solvers or monitor convergence. Currently used machine learning applications perform parallelism on a task level or orchestration based on multiple container instances of an application running in Kubernetes or Microsoft Azure. In our contribution, we will present a software solution for distributed machine learning supporting computation on multiple GPUs running on the top of the PETSc framework (<https://petsc.org/>), which we will demonstrate in applications related to natural hazard localizations and detections employing supervised uncertainties modelling. It is called PERMON (<http://permon.vsb.cz>) and is designed for convex optimization using quadratic programming, and its extension PermonSVM implements maximal-margin classifier approaches associated with support vector machines (SVMs). Although deep learning (DL) is getting popular in recent years, SVMs are still applicable. Unlike DL, an underlying optimization problem associated with SVM is convex, and PermonSVM computes it in a deterministic way (no batch processing) through no intrinsic limitations imposed by a single node memory, convergence to a global minimum is guaranteed. Moreover, an attained model can be easily explainable.

Marek Pecha

Technical University of Ostrava

marek.pecha@vsb.cz

Zachary Langford

Oak Ridge National Laboratory

langfordzl@ornl.gov

David Horák

Technical University of Ostrava
david.horak@vsb.cz

Richard T. Mills
Argonne National Laboratory
rtmills@anl.gov

MS139

Leveraging Modern MPI+GPU Communication Strategies

Parallel applications are often dominated by MPI communication costs. MPI provides several abstractions to applications for describing how to move data across processes, such as the use of MPI Datatypes, MPI Neighborhood collectives, or both. However, the solution an application chooses for CPU based communication may experience different performance results when performing the same communication on GPUs. Furthermore, GPUs also introduce additional complexities by forcing the user to decide whether the CPU or the GPU is in charge of sending the data. This talk will look into some of the performance results modern applications and benchmarks are experiencing.

Derek Schafer
University of New Mexico
dschafer1@unm.edu

MS139

A GPU-Based DEM Solver for Large-Scale Granular Simulations

We present a new generation GPU-based solver for large scale granular dynamics simulations, using the Discrete Element Method (DEM). This simulation infrastructure is a step beyond our current DEM solver available in the open-source multi-physics package Chrono. The new solver uses two GPUs to accelerate the simulation, one dedicated to collision detection and related kinematic calculations, the other dedicated to advancing the dynamics of the granular material. We use just-in-time compiling to allow customization and to leverage the specific architecture of the execution hardware. Finally, the solver leverages CUB, an efficient CUDA utility package for parallel primitives. The resulting solver, distributed under a BSD license, provides an open-source efficient solution for DEM problems with millions of degrees of freedom. After a brief description of the underlying methods, we discuss the implementation of the new solver and illustrate its capabilities on several large-scale DEM problems. One of these is simulation of extraterrestrial rovers over deformable terrain; the terrain is a granular dynamics representation of the Glenn Research Center lunar soil simulant #1 (GRC-1), designed and manufactured to match geotechnical properties of lunar soil. The capabilities and performance of this new solver promise to enable, on relatively inexpensive hardware, simulations of a scale previously possible only on supercomputers.

Radu Serban
University of Wisconsin- Madison
serban@wisc.edu

Dan Negrut
University of Wisconsin - Madison
negrut@wisc.edu

Ruochun Zhang, Luning Fang

University of Wisconsin-Madison
rzhang294@wisc.edu, lfang9@wisc.edu

MS140

Proving Global Nonlinear Stability of Fluid Flows Past the ReynoldsOrr Energy Limit

A fundamental question in fluid stability is whether a laminar flow is nonlinearly stable to all perturbations. The typical way to verify this type of stability, called the Reynolds-Orr energy method, is to show that the energy of a perturbation must decay monotonically under a certain Reynolds number called the Reynolds-Orr energy stability limit. The Reynolds-Orr energy method is known to be overly conservative in many systems, such as in plane Couette flow. Here, we present a methodology to computationally construct Lyapunov functions more general than the energy, which is a quadratic function of the magnitude of the perturbation velocity. These new Lyapunov functions are not restricted to being quadratic, but are instead high-order functionals that depend explicitly on the spectrum of the velocity field in the eigenbasis of the energy stability operator. The methodology involves numerically computing energy eigenmodes and using them to solve a convex optimization problem through a semidefinite program (SDP) constrained by sums-of-squares polynomial ansatzes. We then apply this methodology to 2D plane Couette flow and under certain conditions we find a global stability limit higher than the Reynolds-Orr energy stability limit. For this specific flow, this is the first improvement in over 110 years.

Federico Fuentes
Institute for Mathematical and Computational
Engineering
Pontificia Universidad Católica de Chile, Santiago
federico.fuentes@uc.cl

Sergei I. Chernyshenko
Imperial College London
s.chernyshenko@imperial.ac.uk

David Goluskin
University of Victoria
goluskin@uvic.ca

MS141

Efficient Structure-Preserving Numerical Methods for Transport-Dominated Problems

Many transport-dominated problems can be modeled by partial differential equations with important secondary of interest, e.g., the total entropy of compressible flows. Recently, relaxation techniques have been proposed as modifications of standard time integration methods guaranteeing the conservation or dissipation of such secondary quantities of interest. Here, we investigate how to combine these relaxation techniques with efficient step size control mechanisms based on local error estimates. We demonstrate our results in several numerical experiments, focusing on nonlinear dispersive wave equations.

Sebastian Bleecke
Applied Mathematics
University of Hamburg, Germany
sebastian.bleecke@uni-hamburg.de

Hendrik Ranocha
Applied Mathematics, University of Münster

mail@ranocha.de

MS141

A Data Based Physics-Informed Neural Network Model for Simulation of Two-Dimensional Turbulence

Turbulence is a problem yet to be fully understood. It has many industrial applications in domains such as automotive, aerospace, to name a few. Lot of experimental and computational resources are used to characterize and study turbulent flow. With the advent of deep learning, a new method known as Physics Informed Neural Network (PINN) has been recently introduced which shows tremendous promise in not only reducing computational costs but could potentially save costs on experiments or measurements, which is a huge advantage. In this work, we will show solutions based on PINN methods to simulate two-dimensional turbulence. PINN is a powerful technique where the governing equation of the system is enforced in the neural network architecture. Along with traditional PINN, we also take a sparse set of training data inside the computational domain. We show two different PINN based neural network architectures and compare them with direct numerical simulation (DNS) results. For Reynolds number of order $\mathcal{O}(10^3)$, with the new PINN model which we call as PINN-FF (Fourier filter), we observe a close match in kinetic energy spectra even while using only 0.1% training data. We observe very good match of the large scale structures in comparison with DNS although the small scale structures are yet to be refined. Also, we observe similar results at higher Reynolds numbers of order $\mathcal{O}(10^4)$ which ensures robustness of the approach.

Venkatesh Gopinath, Vijay Kag
Bosch Research and Technology Centre, Bangalore
gopinath.venkatesh2@in.bosch.com,
fixed-term.vijay.kag@in.bosch.com

Kannabiran Seshasayanan
Indian Institute of Technology, Kharagpur, India
kannabiran.s@phy.iitkgp.ac.in

MS141

DoD Stabilization for Solving Hyperbolic Conservation Laws on Cut Cell Meshes

Cut cells methods have been developed in recent years for computing flow around bodies with complicated geometries. For mesh generation, the flow body is cut out of a regular Cartesian grid resulting in so called cut cells. Cut cells can have irregular shape and may be very small. For the solution of time-dependent hyperbolic conservation laws, this causes the small cell problem: explicit time stepping schemes are not stable on the arbitrarily small cut cells. In this talk we present the Domain of Dependence (DoD) stabilization for overcoming this issue for DG schemes. The stabilization is based on adding suitable penalty terms to the spatial discretization. In time, one can then use a standard explicit time stepping scheme. The terms are designed to restore proper domains of dependence. The terms are constructed such that the resulting stabilized scheme possesses desirable properties such as being monotone for scalar conservation laws for the first-order version or energy preserving for the semi-discrete scheme for all polynomial degrees. On cut cell meshes, this is a big challenge to guarantee these otherwise ‘standard’ properties. In this talk, we will present the idea behind the DoD stabilization and summarize the latest developments for non-linear

problems.

Sandra May
Uppsala University
sandra.may@it.uu.se

Gunnar Birke
University of Muenster, Germany
g_birk01@uni-muenster.de

Christian Engwer
INAM, University of Münster
christian.engwer@uni-muenster.de

Florian Streitbürger
TU Dortmund University, Germany
florian.streitbuerger@math.tu-dortmund.de

MS141

Invariant Domain Preserving Finite Element Methods for Surface-Quasi-Geostrophic Flows

We present a new finite element method for solving geostrophic flows. Our particular interest is the surface-quasi-geostrophic flow model (SQG), which is a scalar non-linear conservation law in two space dimensions. The velocity field and the viscosity term of the SQG equation involve calculating the fractional derivative of the solution, which is not trivial to compute numerically. For this purpose, we use Dunford–Taylor representations of fractional operators. Strong-stability-preserving Runge–Kutta methods are used in time and the continuous finite element method is used in space. The method retains the invariant domain property under the usual CFL condition and has a second-order accuracy in space. The method is tested by solving several benchmark problems. This work was done in collaboration with Andrea Bonito from Texas A&M University.

Murtazo Nazarov
Uppsala University, Department of Information
Technology
Sweden
murtazo.nazarov@it.uu.se

MS141

Bounds-Preserving Discontinuous Galerkin Methods for Multi-Component MHD Flows

We present a family of subcell limiting strategies for high-order discontinuous Galerkin spectral element methods (DGSEM) and apply them to the multi-component GLM-MHD equations: the magneto-hydrodynamics equations augmented with a generalized Lagrange multiplier technique to weakly enforce the divergence-free condition. In [Rueda-Ramrez et al., Subcell Limiting Strategies for DGSEM, 2022], the authors presented subcell limiting strategies for DGSEM discretizations of conservation laws based on the combination of finite volume (FV) schemes with DGSEM. However, the multi-component GLM-MHD equations impose two new challenges: (i) they require the use of non-conservative terms for entropy consistency, which adds complexity to the spatial discretization, and (ii) impose positivity constraints for several ion species. Our main contribution is the extension of our hybrid FV/DGSEM to non-conservative systems, which requires to write the DGSEM as a flux-differencing formula. To achieve it, we propose a generalization of the telescoping flux formula of Fisher and Carpenter [Fisher and Carpenter, High-order entropy stable finite difference schemes for

nonlinear conservation laws: Finite domains, 2013] to account for non-conservative terms. Furthermore, we show that the hybrid FV/DGSEM can be used to impose invariant domain preserving constraints (including positivity of all ion species). We test our new strategy to solve problems featuring shocks and turbulence.

Andres Rueda-Ramirez, Gregor Gassner
University of Cologne
aruedara@uni-koeln.de, ggassner@uni-koeln.de

MS142

A Systematic Particle Filter for Estimating Time-Varying Parameters in Nonstationary Inverse Problems

Estimating and quantifying uncertainty in unknown system parameters from limited data remains a challenging inverse problem in a variety of application areas. A subset of these problems includes estimating parameters that are known to vary with time but have unknown evolution models and often cannot be directly observed. Here we present a particle filtering algorithm that utilizes a hierarchical Bayesian approach to systematically estimate unknown time-varying parameters in nonstationary inverse problems, with several computed examples arising from deterministic dynamical systems.

Andrea Arnold
Worcester Polytechnic Institute
anarnold@wpi.edu

MS142

Correcting Convexity Bias in Function and Functional Estimate

A general framework with a series of different methods is proposed to improve the estimate of convex function (or functional) values when only noisy observations of the true input are available. Technically, our methods catch the bias introduced by the convexity and remove this bias from a baseline estimate. Theoretical analysis is conducted to show that the proposed methods can strictly reduce the expected estimate error under mild conditions. When applied, the methods require no specific knowledge about the problem except the convexity and the evaluation of the function. Therefore, they can serve as off-the-shelf tools to obtain good estimate for a wide range of problems, including optimization problems with random objective functions or constraints, and functionals of probability distributions such as the entropy and the Wasserstein distance. Numerical experiments on a wide variety of problems show that our methods can significantly improve the quality of the estimate compared with the baseline method.

Chao Ma
Stanford
chaoma@stanford.edu

Lexing Ying
Stanford University
Department of Mathematics
lexing@stanford.edu

MS142

Deep Learning for Bayesian Inverse Problems Governed by Nonlinear Differential Equations

erned by Nonlinear Differential Equations

We will discuss the use of deep neural networks for Bayesian inference and uncertainty quantification in optimization problems governed by nonlinear dynamical systems. In particular, we will consider non-linear control of a Hodgkin-Huxley model. We will use curvature information of the negative log posterior to approximate the covariance operator of the posterior. We will report results for the accuracy of the inference with respect to different architectures and loss functions.

Andreas Mang
University of Houston
Department of Mathematics
amang@Central.UH.EDU

Johann Rudi
Department of Mathematics
Virginia Tech
jrudi@vt.edu

German Villalobos
Department of Mathematics
University of Houston
gvillalobos@uh.edu

MS143

A Domain Decomposition Rayleigh-Ritz Algorithm for Symmetric Generalized Eigenvalue Problem

In this talk we propose a distributed-memory parallel algorithm for computing all the eigenvalues (and corresponding eigenvectors) of a large, sparse, real symmetric positive definite matrix pencil that lie within a target interval. The algorithm is based on Chebyshev interpolation of the eigenvalues of the Schur complement (over the interface variables) of a domain decomposition reordering of the pencil and accordingly exposes two dimensions of parallelism: one derived from the reordering and one from the independence of the interpolation nodes. The new method demonstrates excellent parallel scalability, comparing favorably with PARPACK, and does not require factorization of the mass matrix, which significantly reduces memory consumption, especially for 3D problems. Our implementation is publicly available on GitHub.

Vasileios Kalantzis
IBM Research
vkal@ibm.com

MS143

Computing Diffraction Anomalies As Nonlinear Eigenvalue Problems

When electromagnetic waves illuminate on a diffraction grating or other periodic structure, reflected and transmitted waves in different radiation channels are excited and propagate away from the structure. Such a diffraction problem can be studied by solving the Maxwell equations numerically. Diffraction anomalies such as zero reflection, zero transmission and perfect absorption may arise for incident waves with specific frequencies and/or wavevectors, and these phenomena can be used to manipulate electromagnetic waves and light. The frequencies and wavevectors of anomalies are discretely distributed in the spectrum, and some anomalies may appear only in structures with specific physical parameters. Existing methods for computing anomalies are usually computationally expen-

sive and not very effective, as they require repeatedly solving a boundary value problem for many parameter values. In this study, a number of diffraction anomalies are computed using an efficient numerical method based on nonlinear eigenvalue formulations and a contour-integral method (for solving the nonlinear eigenvalue problems). To demonstrate the new method, numerical examples for structures with periodic arrays of cylinders are provided.

Zitao Mai

Department of Mathematics, City University of Hong Kong
zmai2-c@my.cityu.edu.hk

Ya Yan Lu

City University of Hong Kong
Dept of Mathematics
mayylu@cityu.edu.hk

MS144

The Exascale Am Challenge Problem: Am Process Modeling at the Fidelity of the Microstructure

With the Exascale computers arriving at the DOE facilities, the ExaAM project will soon run its challenge problem. The ExaAM challenge problem is based on the NIST AMBench experimental builds and further characterization. These experiments both guide the model development and provide validation for the simulations. The project includes an integration of all the computational components of the AM process, where each component itself is an exascale simulation. What has emerged is that Exascale Computing will enable AM process modeling at the fidelity of the microstructure. This means tight coupling of Process-Structure-Property calculations. Macroscopic continuum codes (OpenFOAM) are used to simulate the metal melt-refreeze, within which mesoscopic codes (ExaCA, PFM) are used to simulate the development of material microstructure. This microstructure is then used by crystal plasticity codes (ExaConstit) to calculate local material properties. We present our coupled exascale simulation environment for additive manufacturing and its application to AM builds. Jim Belak (LLNL) and John Turner (ORNL)

James F. Belak

Lawrence Livermore Nat. Lab.
belak1@llnl.gov

MS144

A Parallel-in-Time Approach to Part-Scale Melt Pool Modeling of Metal Additive Manufacturing

One of the primary challenges in simulating an additive manufacturing process such as laser powder bed fusion (LPBF) for an entire component is the drastically varying length and time scales. As a consequence, the number of time steps required for a high fidelity LPBF process simulation of an entire component can easily reach into the billions, which is impractical with typical simulation codes. We present a computationally scalable approach to simulating the LPBF process for an entire component. This scalability is achieved by parallelizing the time domain integration using an iterative technique. The proposed technique offers a path to whole component simulation of the nonlinear heat transfer and phase change that occurs in the LPBF process. Matthew Bement, John Coleman, ORNL

Matt Bement

Oak Ridge National Laboratory
bementmt@ornl.gov

MS144

Super-Spectral Operator Recovery via the Fast Macroscopic Forcing Method

The macroscopic forcing method (MFM) was introduced by Mani and Park in 2021. It recovers lower-dimensional operators by successively forcing a high-dimensional direct numerical simulation. The MFM has already successfully recovered RANS-like turbulence models for fluid flows. Standard algorithms for MFM apply forcings to each coarse-scale degree of freedom and conduct a fine-scale simulation, which is expensive. We present an algorithm that is cheaper and more general. It applies sparse reconstruction to expose local features in the differential operator and reconstructs the coarse one in only a few matrix-vector products. For non-local operators, we prepend this approach by peeling long-range effects with dense matrix-vector products to expose a more local operator. We demonstrate the algorithm's performance on scalar transport and channel flow problems.

Spencer H. Bryngelson

University of Illinois
shb@gatech.edu

Florian Schaefer

Georgia Institute of Technology
florian.schaefer@cc.gatech.edu

Jessie Liu, Ali Mani

Stanford University
jeliu@stanford.edu, alimani@stanford.edu

MS144

Predictive Modeling And Uncertainty Quantification For Diblock Copolymer Self-Assembly

We present a Bayesian framework for the predictive modeling of diblock copolymer (Di-BCP) thin film self-assembly. It involves a procedure consisting of model calibration, selection, and validation of phase field models centered around Bayes rule and image data produced by microscopy or X-ray scattering characterizations. The aleatoric uncertainties of Di-BCP self-assembly represented by metastability and defectivity require introducing randomness into model predictions. These uncertainties, however, lead to integrated likelihoods that are generally intractable to evaluate. To tackle this challenge that hampers the execution of the predictive modeling procedure, likelihood-free inference approaches via pseudo-marginal methods and measure transport are considered in this study. To improve the efficiency of these inference methodologies, they are used in adjunct with carefully designed summary statistics that extract features of characterization images. Several Fourier- and energy-based summary statistics are proposed for top-down microscopy characterizations of Di-BCP thin films. Expected information gains are used to quantify the utilities of summary statistics choices, and we show that they can be computed via measure transport with no significant additional computational cost.

Lianghao Cao

The University of Texas at Austin
lianghao@oden.utexas.edu

Ricardo Baptista

MIT
rsb@caltech.edu

Joshua Chen
The University of Texas at Austin
joshuawchen@utexas.edu

Fengyi Li
Massachusetts Institute of Technology
fengyil@mit.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

John Tinsley Oden
Oden Institute for Computational Engineering and
Sciences
The University of Texas at Austin
oden@oden.utexas.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS144

Algorithms for Large Scale Simulations of Microstructural Evolution

I will describe new level set methods for simulating the motion of networks of interfaces (curves or surfaces) under geometric flows such as motion by mean curvature. These methods can be used in large scale simulations of microstructural evolution in polycrystalline materials, for instance grain growth and recrystallization that take place during many (including additive) manufacturing processes.

Selim Esedoglu
University of Michigan
esedoglu@umich.edu

MS145

Variationally Stabilized Isogeometric Analysis for Advection-Dominated Diffusion-Reaction Problems

We propose a new variationally stabilized isogeometric analysis for advection-dominated problems. Isogeometric analysis (IGA) is a widely-used analysis tool in modeling and simulation. Although IGA has advantages in comparison with the classical finite element method (FEM), one still requires stabilization techniques to approximate advection-dominated problems. We extend the edge-stabilization techniques by adding the jump of high-order derivatives to the standard bilinear forms of isogeometric analysis. Also, we show that the new bilinear form is consistent and coercive with respect to an induced norm. Next, using the proposed bilinear form and norm, we formulate a residual minimization as a saddle-point problem, which delivers a stable discrete solution. We use tensor-product B-splines of order p , with continuity of C^{p-1} to construct the trial space U_h . Then, we obtain a discrete approximation of the solution by minimizing the residual measured in a dual norm of a test space V_h constructed using B-splines of order p and C^{p-2} . The technique also provides an error representation deployed to guide the adaptivity. We provide numerical examples to show the performance of the

technique.

Pouria Behnoudfar
Commonwealth Scientific and Industrial Research
Organisation
(CSIRO)
pouria.behnoudfar@csiro.au

Victor M. Calo
Curtin University
School of Earth and Planetary Sciences
vmcalo@gmail.com

Quanling Deng
Australian National University
quanling.deng@anu.edu.au

MS145

C^1 Hierarchical Splines Spaces on Surface Multipatch Domains

Adaptive isogeometric methods for the solution of PDEs have been largely investigated in recent years (see [Bracco, Buffa, Giannelli, Vzquez, Discret. Contin. Dyn. S., 2019] for an overview). The combination of isogeometric analysis and local refinement makes them computationally attractive, and therefore it is relevant to study the extension of this approach to high-order problems on complex domains. In this work, we propose an adaptive isogeometric method for the numerical approximation of high-order PDEs based on C^1 hierarchical spline spaces defined on multipatch surfaces. We achieve this goal by combining the construction of C^1 multipatch spaces [Farahat, Jüttler, Kapl, Takacs, Isogeometric analysis with C^1 -smooth functions over multipatch surfaces, Comput. Methods Appl. Mech. Engrg., 2023] with the hierarchical framework. In particular, we define a refinement algorithm with linear complexity which guarantees the construction of suitably graded hierarchical meshes fulfilling the condition for linear independence of the set of hierarchical generating functions. A selection of numerical examples will confirm the potential of the adaptive scheme on different problems defined on multipatch surfaces.

Cesare Bracco
Department of Mathematics - University of Florence
cesare.bracco@unifi.it

MS145

Preconditioners for Adaptive Splines, Or Splines for Preconditioners?

The most common approach to achieve adaptive IGA is to break the tensor product structure of the tensor-product-spline space while keeping that of the basis functions: each basis function is a products of univariate functions. By choosing an appropriate set of basis functions local variation of the space resolution are possible. The examples of this approach are hierarchical-, T- and LR-splines. The change from tensor-product-splines to any of these adaptive-spline-space requires adapting the assembling, preconditioning and solving strategies in order to maintain efficiency. This is challenging and adds to the source code complexity. A space construction technique that is well suited to the subspace correction methods and preconditioned Krylov methods will be presented.

Andrea Bressan
IMATI CNR (Pavia)

andrea.bressan@imati.cnr.it

MS145

Refinement Strategies for Locally Linearly Independent LR B-Splines

Locally Refined (LR) B-splines are generalization of the tensor product B-splines to achieve adaptivity in the discretization process. Thereby, the approximation efficiency is dramatically improved as one avoids the wasting of degrees of freedom by increasing the number of basis functions only where rapid and large variations occur in the analyzed object. Nevertheless, the adoption of LR B-splines for simulation purposes, in the Isogeometric Analysis (IgA) framework, is hindered by the risk of linear dependence relations. Although a complete characterization of linear independence is still not available, the local linear independence of the basis functions is guaranteed when the underlying Locally Refined (LR) mesh has the so-called Non-Nested-Support property. The local linear independence not only avoids the hurdles of dealing with singular linear systems, but it also improves the sparsity of the matrices when assembling the numerical solution. Such a strong property of the basis functions is a rarity, or at least it is quite onerous to gain, among the technologies used for adaptive IgA. Only three refinement strategies have been proposed to build LR meshes with the Non-Nested-Support property so far: the Non-Nested-Support-Structured mesh refinement, the Effective Grading refinement and the Hierarchically Locally Refined mesh refinement. In this talk, we shall have an overview of them, highlighting and comparing their properties.

Francesco Patrizi
University of Florence
francesco.patrizi@ipp.mpg.de

MS145

Adaptive Isogeometric Analysis Based on Tchebysheffian Splines

Tchebycheffian splines are piecewise smooth functions with pieces in Tchebycheff spaces, the natural generalization of polynomial spaces. Under suitable assumptions, Tchebycheffian splines can be represented in terms of B-spline like basis functions, called Tchebycheffian B-splines, with structural similarities and all fundamental properties of classical polynomial B-splines (compact support, non-negativity, partition of unity, etc.). Tchebycheffian splines offer a huge flexibility compared to classical polynomial splines: they are equipped with parameters that can be selected according to a problem-oriented strategy, taking into account the geometrical and/or analytical issues of the specific addressed problem. While the tensor-product approach can easily build multivariate splines, they lack adequate adaptive local refinement, which could be important in both geometric modelling and numerical simulation. This triggered the interest in alternative (polynomial) spline structures supporting local refinement still retaining the local tensor-product structure. Since Tchebycheffian B-splines are plug-to-plug compatible with polynomial B-splines, they are naturally compatible with the local refinement structures known for polynomial B-splines (Hierarchical mesh, T-spline mesh and Locally Refined(LR-) mesh). In this talk, we discuss the use of Tchebycheffian splines equipped with a local tensor-product structure as a possible tool in IGA with adaptive refinement.

Krunal Raval

University of Rome Tor Vergata
raval@mat.uniroma2.it

MS146

Space-Time Shape Optimization of Rotating Electric Machines

Electric machines can often be modeled by the magneto-quasi-static approximation of Maxwell's equations in two space dimensions. We consider the simulation of a rotating electric machine by means of a space-time finite element method where the rotation is captured by the tetrahedral space-time mesh. We derive the shape derivative for a given cost function with respect to a perturbation of the (spatial) geometry and present a shape optimization algorithm for moving domains in space-time. Here, it is important to note that the optimized geometry is moving, but must not change its shape over time. Finally, for a realistic simulation, the initial condition of the evolution problem at hand is obtained as the solution to a static PDE on the given geometry and, thus, is shape-dependent as well. Accounting for this aspect yields a shape optimization problem that is constrained by a system of a static and a transient PDE. We show an extension of our sensitivity analysis and shape optimization algorithm to this setting and present numerical results for the optimization of a synchronous reluctance machine.

Peter Gangl
RICAM
Johann Radon Institute for Computational and Applied Mathema
peter.gangl@ricam.oeaw.ac.at

Alessio Cesarano
RICAM (Johann Radon Institute for Computational and Applied Mathematics)
alessio.cesarano@ricam.oeaw.ac.at

Charles Dapogny
Laboratoire Jean Kuntzmann
Université Grenoble Alpes
charles.dapogny@univ-grenoble-alpes.fr

MS146

Advanced Methods for Efficient Shape Optimization in Aeronautics

On the one hand, Sobolev gradient smoothing can considerably improve the performance of aerodynamic shape optimization and prevent issues with regularity. On the other hand, Sobolev smoothing can also be interpreted as an approximation for the shape Hessian. This paper demonstrates, how Sobolev smoothing, interpreted as a shape Hessian approximation, offers considerable benefits, although the parameterization is smooth in itself already. Such an approach is especially beneficial in the context of simultaneous analysis and design, where we deal with inexact flow and adjoint solutions, also called One Shot optimization. Furthermore, the incorporation of the parameterization allows for direct application to engineering test cases, where shapes are always described by a CAD model. The new methodology presented in this paper is used for reference test cases from aerodynamic shape optimization and performance improvements in comparison to a classical Quasi-Newton scheme are shown.

Nicolas R. Gauger
TU Kaiserslautern

nicolas.gauger@scicomp.uni-kl.de

Stephan Schmidt
Humboldt-Universität zu Berlin
s.schmidt@hu-berlin.de

Thomas Dick
TU Kaiserslautern
thomas.dick@scicomp.uni-kl.de

MS146

Quasi-Newton Methods for Fully Discrete Shape Optimization Problems

In this talk, we will consider fully discrete PDE-constrained shape optimization problems in 2D. These problems are defined on the complete manifold of planar triangular meshes initially proposed in [Herzog, Loayza-Romero, A Manifold of Planar Triangular Meshes with Complete Riemannian Metric, 2020, <https://doi.org/10.48550/arXiv.2012.05624>]. We propose an approximation for the Riemannian parallel transport and logarithm on the manifold of planar triangular meshes based on the discrete geodesic calculus proposed in [Rumpf, Wirth, Variational time discretization of geodesic calculus, 2015, doi:10.1093/imanum/dru027]. The approximation of the parallel transport will allow us to propose a Quasi-Newton Method for solving this problem. We will compare the results with the traditional Riemannian steepest descent method and show the acceleration in the convergence. Finally, since computing geodesics is too computationally expensive we propose a quality-preserving projection-like retraction obtained by exploiting the features of the manifold of planar triangular meshes as suggested in [Absil, Malik, Projection-like Retractions on Matrix Manifolds, 2012, <https://doi.org/10.1137/100802529>]. We will apply the proposed algorithm to solve compliance problems in structural mechanics.

Estefania Loayza-Romero
University of Münster
estefania.loayza-romero@uni-muenster.de

Benedikt Wirth
Universität Münster
benedikt.wirth@uni-muenster.de

MS146

Stabilization and Tracking Control Design for Stochastic Fuzzy Systems via Adaptive Event-Triggered Mechanism

An adaptive event-triggered mechanism-based output tracking control problem for stochastic T-S fuzzy systems with state delay and external disturbance is discussed in this work. An adaptive event-triggered mechanism is developed for the considered system to reduce the number of triggering and communication burdens. Adaptive feedback control is constructed to force the output trajectories to track the bounded reference input signal even in the presence of external disturbance. Further, the tracking control objective transformed into stochastic stabilization of the considered system. For the stabilization analysis, the stochastic fuzzy system does not share the same membership functions as the proposed control. The sufficient stabilization conditions are developed in the form of linear matrix inequalities together with the help of Lyapunov-Krasovskii functional and some integral inequalities. Finally, the proposed methodology is applied to some en-

gineering stochastic models to validate efficiency and practicability.

Vijayakumar Muthusamy
Anna university regional campus, Coimbatore
vijaymaths048@gmail.com

Marshal Anthoni Selvaraj
Anna University
smarshalanthoni@gmail.com

MS146

Dirichlet Control As a Topology Optimisation Problem

In this paper we study a special Dirichlet control problem where the control is supposed to be a piecewise constant function which is allowed to take $N \geq 2$ different values. As the piecewise functions do not belong to $H^{\frac{1}{2}}$ standard extension techniques to consider the weak solution of the Dirichlet problem do not apply. Therefore, as done in optimal control, we consider very weak solutions of the state equation. We then study the shape-to-control operator of this problem and derive the first order necessary optimality conditions using the topological derivative. Using the topological derivative we use a multimaterial levelset algorithm and present several numerical examples in dimension two and three for two and three different materials.

Kevin Sturm
Technische Universität Wien
kevin.sturm@asc.tuwien.ac.at

MS147

An RBF-FD Closest Point Method for Solving PDEs on Surfaces

In this work, we present a method for numerically solving PDEs on closed surfaces based on the combination of the closest point principles and radial basis function (RBF)-generated finite difference (FD) discretizations. The closest point principles allow the use of locally embedded stencils with an extension along the normal direction to approximate the surface derivatives in a RBF-FD meshfree setting. This approach only requires a scattered-node representation of the surface and the normal vectors on these scattered nodes. Several numerical tests are carried out to demonstrate the accuracy and effectiveness of our approach.

Victor Bayona
Universidad Carlos III de Madrid
vbayona@math.uc3m.es

Argyrios Petras
RICAM - Johann Radon Institute
argyrios.petras@ricam.oeaw.ac.at

MS147

Smooth Geometry Reconstruction of the Diaphragm from Noisy Data Using the RBF-PUM Method

We use the radial basis function partition of unity method (RBF-PUM) to make a smooth reconstruction of the diaphragm directly from noisy point cloud data obtained through the manual segmentation of the diaphragm from 3D medical images. The surface of the diaphragm is rep-

resented as the zero-level-set of a global function which is computed by covering the point cloud domain with overlapping patches and computing local solutions on each patch through RBF interpolation. The different solutions are then combined into a global function using weight functions and the data points are then moved to the zero-level-set surface through Newton iteration. The quality of the reconstruction is influenced by several different parameters, the most important of which are the number of stencil points in the local problems, the number of patches covering the domain and a scaling factor which determines the size of the separation between the data representing the lower and upper parts of the diaphragm.

Ulrika Sundin, Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
ulrika.sundin@it.uu.se, Elisabeth.Larsson@it.uu.se

Pierre-Frédéric Villard
Université de Lorraine, CNRS; Inria LORIA, France
pierrefrederic.villard@loria.fr

Igor Tominec
Uppsala University
Department of Information Technology
igor.tominec@it.uu.se

Nicola Cacciani
Karolinska Institutet
Department of Physiology and Pharmacology
nicola.cacciani@ki.se

MS147

RBF Modeling of Short Laser Pulses: Thermal Effects and Wave Collapse

Numerical simulation of the thermal blooming of laser pulses are modeled. The model includes terms representing the contributions of thermal blooming and the Kerr effect. In the context of the model, thermal blooming and the Kerr effect compete; with the Kerr effect driving and thermal blooming suppressing a phenomenon known as wave collapse. Numerically, wave collapse presents itself as a finite-time singularity. Such singularities present challenging resolution requirements for the split-step Fourier methods which have been commonly used in similar models. We implement a model including the use of radial basis functions to alleviate such challenges.

Tony Liu
Air Force Institute of Technology
Tony.Liu@afit.edu

MS147

Radial Basis Functions (RBFs) and Approximate Definite Integrals (Quadrature/Cubature)

A Radial Basis Function Generated Finite-Differences (RBF-FD) inspired technique for evaluating definite integrals over bounded domains in up to three dimensions is described. Such methods are necessary in many areas of Applied Mathematics, Mathematical Physics and myriad other application areas, where the task of evaluating definite integrals follows the solution of Partial Differential Equations or the collection of physical measurements. The RBF-FD methods outlined in this talk are efficient, easily

parallelizable and tunable to high orders of accuracy.

Jonah A. Reeger
Air Force Institute of Technology
jonah.reeger@afit.edu

MS148

Higher-Order Activation Functions in Neural PDE Theory

Neural networks have been used successfully in solving very high dimensional PDEs, overcoming the curse of dimensionality plaguing classical methods. A neural theory of PDE can explain why and when these neural networks can overcome the curse of dimensionality. The *Barron spaces* and *tree-like spaces* have been introduced as function spaces on which to build this neural theory of PDE. Their study has mostly been focused on the ReLU activation function. Although the ReLU is ubiquitous in deep learning, it is of limited use for solving PDEs due to its limited smoothness. In this work we unify previous work on different activation functions, and extend the list of relations between activation functions to create a taxonomy of the Barron spaces based on the activation function used. We show how the taxonomy can be used for operators.

Tjeerd Jan Heeringa, Len Spek, Felix Schwenninger,
Christoph Brune
Department of Applied Mathematics
University of Twente
t.j.heeringa@utwente.nl, l.spek@utwente.nl,
f.l.schwenninger@utwente.nl, c.brune@utwente.nl

MS148

Physics-Informed Two-Tier Neural Network for Non-Linear Model Order Reduction

In recent years, machine learning (ML) has had a great impact in the area of non-intrusive, non-linear model order reduction (MOR). However, the offline training phase still suffers from high computational costs since it requires numerous expensive full-order solutions as the training data. Furthermore, in state-of-the-art methods, neural networks trained by a small amount of the training data cannot be expected to generalize well enough, and the training phase generally ignores the underlying physical information. Moreover, state-of-the-art affine decomposition and hyper reduction techniques are intrusive or entail huge offline computational costs. To resolve these challenges, inspired by recent developments in physics-informed neural networks and the physics-reinforced neural networks, we propose a non-intrusive, physics-informed, two-tier deep network (TTDN) method. The proposed network, wherein the first tier achieves the regression of the unknown quantity of interest (QoI) and the second tier rebuilds the physical constitutive law between the unknown QoI's and derived quantities, is trained using pretraining and semi-supervised learning strategies. To showcase the efficiency of the proposed approach over the state-of-the-art methods, we perform numerical experiments on challenging non-linear and non-affine problems and study the behavior associated with the high or relatively low computational cost of the full-order data.

Yankun Hong, Harshit Bansal
Eindhoven University of Technology
y.hong@tue.nl, h.bansal@tue.nl

Karen Veroy-Grepl

TU/e Eindhoven University of Technology
k.p.veroy@tue.nl

MS148

IgaNets: Physics-Informed Machine Learning Embedded Into Isogeometric Analysis

In this talk we present a novel approach to embed physics-informed neural networks (PINN) into the framework of Isogeometric Analysis. IGA is an extension of the finite element method that integrates simulation-based analysis into the computer-aided design pipeline. In short, the same mathematical formalism, namely B-splines or NURBS, that is used to model the geometry is adopted to represent the approximate solution, which is computed following the same strategy as in classical finite elements. In contrast to classical PINNs, which predict point-wise solution values to (initial-)boundary-value problems directly, our IgaNets learn solutions in terms of their expansion coefficients relative to a given B-Spline or NURBS basis. This approach is also used to encode the geometry and other problem parameters such as boundary conditions and feed them into the network as inputs. Once trained, our IgaNets make it possible to explore various designs from a family of similar problem configurations efficiently without the need to perform a computationally expensive simulation for each new problem configuration. Next to discussing the method conceptually and presenting numerical results, we will shed some light on the technical details of our C++ reference implementation in Torch. In particular, we will discuss matrix-based implementation of B-splines that is particularly suited for efficient backpropagation.

Matthias Moller
Delft Institute of Applied Mathematics (DIAM)
Delft University of Technology
M.Moller@tudelft.nl

MS148

Data Driven Gradient Flows

We present a framework enabling variational data assimilation for gradient flows in general metric spaces, based on the minimizing movement (or Jordan-Kinderlehrer-Otto) approximation scheme. After discussing stability properties in the most general case, we specialise to the space of probability measures endowed with the Wasserstein distance. This setting covers many non-linear partial differential equations (PDEs), such as the porous medium equation or general drift-diffusion-aggregation equations, which can be treated by our methods independent of their respective properties (such as finite speed of propagation or blow-up). We then focus on the numerical implementation of our approach using an primal-dual algorithm. The strength of our approach lies in the fact that by simply changing the driving functional, a wide range of PDEs can be treated without the need to adopt the numerical scheme. We conclude by presenting detailed numerical examples.

Jan-Frederik Pietschmann
Faculty for Mathematics
TU Chemnitz
jfpietschmann@math.tu-chemnitz.de

Matthias Schlottbom
University of Twente

m.schlottbom@utwente.nl

MS148

Quantitative Phase-Field Modeling by Physics-Constrained Deep Learning

Phase-field models are recently known as standard approaches for the prediction of solidification microstructure. However, to perform a predictive phase-field simulation is remained still computationally challenge in practice, because numerical solutions exhibit mesh dependency issue. As a result, the quantitative phase field modeling is limited to very small spatial scales and low Peclet number regimes. Recent advances in deep neural networks revealed that the direct differentiation of neural networks makes it possible to solve nonlinear partial differential equations (PDEs) in a computational complexity and accuracy competitive to classic grid-based approaches like finite different and finite element methods. In physics-informed neural networks methods, instead of using existing data, the physics of problem is directly included into the training phase and the neural-network is trained without any required data. This work communicates results on the solution of PDEs corresponding to the quantitative phase-field modeling of pure elements solidification. Different network architectures and ways of training on the accuracy of results are studied. The results reveal that this method competes with existing mesh-based approaches in terms of accuracy and efficiency. In particular, we can train the network in a small space-time domain, and then can solve the corresponding PDEs with the already trained network in larger spatial domains and longer times.

Rouhollah Tavakoli
IMDEA MATERIALS INSTITUTE
IMDEA MATERIALS INSTITUTE
rouhollah.tavakoli@imdea.org

MS149

European Industrial Doctorate

European Industrial Doctorate is one of Doctoral networks financed by EU within Horizon Europe, Marie Skłodowska Curie Action, which are the flagship funding programme for doctoral education and postdoctoral training. It aims at implementing doctoral programs by partnership of universities and industries. Young researchers enrolled in this type of projects are jointly supervised by academic and industrial partners and work on research projects that are inspired by real industrial challenges. The program is an excellent opportunity for young people who are looking to step outside academic world and broaden their skills and it is open to students from all over the world. The details of the partnership between academia and industry will be presented in this talk as well as some examples of successful EID in mathematics.

Natasa Krejic
University of Novi Sad
Novi Sad, Serbia
natasak@uns.ac.rs

MS149

Industrial Mathematics at Fraunhofer ITWM: Career Opportunities at the Interface of Applied Research and Industry Projects

A brief look on the activities at ITWM will be given, fol-

lowed by a short overview of the activities of the department (Math for the Digital Factory), and how this relates to the activities of the ITWM spin-off company fleXstructures GmbH, the collaboration with FCC Gteborg and their spin-off IPS AB. Wr.t. this background, some reflections on whats important for math graduates to work successfully in industry, or research institutes like ours that collaborate intensively with industry will be shared, and how this differs from research work in academia.

Joachim Linn
Fraunhofer ITWM
Kaiserslautern, Germany
joachim.linn@itwm.fraunhofer.de

MS149

CSE Careers in the United States' DOE National Laboratories

The National Laboratories that are operated by the United States Department of Energy (DOE Labs) provide opportunities for excellent careers performing research and development in Computational Science and Engineering. The DOE Labs specialize in Big Science, including foundational and applied research that requires large, interdisciplinary teams, working on grand challenge applications such as fusion energy, climate modeling, battery technology, next-generation computing architectures, and a host of national security applications. Many of us have found a career at the DOE Labs to be the ideal middle ground between academia and industry, where we can work on challenging problems with excellent colleagues to advance technology to tackle some of the world's most pressing issues.

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

MS149

Lessons from 10 Years of Mathematical Consultancy

Sioux Technologies is a global technology partner that supports or acts as the RD department for high-tech companies. The company has a unique market position due to its dedicated applied mathematics department of about 60 engineers, branded the 'Mathware' department. This department applies mathematics in a very broad sense, from computational physics, through signal processing up to artificial intelligence. The Mathware department spun out of the mathematics department of the Technical University of Eindhoven a little over 10 years ago. Over these 10 years (and more), many lessons have been learned on how to approach industrial mathematics coming from a background of applied and numerical mathematics in an academic setting, as well as experiencing the growth from 10 people to 60 people. These experiences have led to principled approaches applied by Sioux regarding people management, technology management, software quality management and foremost project management through 'the Mathware method'.

Bas van der Linden
Sioux Technologies
Technical University of Eindhoven
bas.van.der.linden@sioux.eu

Robbert van Herpen

Sioux Technologies
robbert.van.herpen@sioux.eu

MS150

Semi-Implicit Time Integration for Multiscale Tokamak-Edge Plasma Dynamics

Tokamak edge plasmas are characterized by a large range of time scales, and the scales of interest are often much slower than the fastest scales. We present an implicit-explicit (IMEX) time integration framework in COGENT, a high-order code for tokamak edge simulations. The governing equations comprise collisional gyrokinetic equations for ion species and may include fluid equations for charged and neutral species as well as kinetic, fluid, or Boltzmann models for electrons. The large number of physics models implemented requires a flexible framework, where the implicit terms depend on the specific simulation. We implement additive Runge-Kutta (ARK) methods, where the implicit system is solved using the Jacobian-free Newton-Krylov method. We also consider a modified form of the ARK methods for ODEs with a nonlinear left-hand-side operator, $d[M(u)]/dt = L(u)$, that arise from a vorticity model of the electrostatic potential. The implicit term comprises multiple physics models; thus, we adopt a multiphysics preconditioning strategy that wraps individual preconditioners for each model with an operator-split approach to precondition the overall system. This allows tailored preconditioning strategies for each implicit physics. We demonstrate our algorithm on test cases representative of tokamak-edge plasma dynamics. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344.

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Livermore, CA
ghosh5@llnl.gov

Milo Dorr
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
dorr1@llnl.gov

Mikhail Dorf, Lee Ricketson
Lawrence Livermore National Laboratory
dorf1@llnl.gov, ricketson1@llnl.gov

MS150

A Structure Preserving, Conservative, Low-Rank Tensor Scheme for Solving the 1D2V Vlasov-Fokker-Planck Equation

We propose a hybrid low-rank tensor scheme for solving the 1D2V Vlasov-Fokker-Planck equation in Cartesian physical space and cylindrical velocity space. The solution is full rank in physical space and low rank in velocity space. By incorporating several robust methods into our proposed algorithm, we attain a scheme that is conservative, equilibrium preserving, relative entropy dissipative, and low-rank with low storage complexity. A kinetic ion - fluid electron model is assumed; the Leonard-Bernstein-Fokker-Planck operator is discretized using a structure preserving Chang-Cooper method; and the updated solution is truncated using a local macroscopic conservative low rank tensor method. Preliminary numerical results are presented

to demonstrate these properties.

Joseph Nakao
University of Delaware
nakaoj@udel.edu

MS150

A Time-Domain Preconditioner for the Helmholtz Equation: Relation with Other Methods

Several methods in the literature determine solutions to the Helmholtz equation by solving instances of a discrete time-domain wave equation. In this work we study a new method of this type, called a time-domain preconditioner. Given an indefinite linear system, a matrix recurrence relation is constructed, such that in the limit of infinitely many time steps the exact discrete solution is obtained. Using a large, finite number of time steps, an approximate solution is obtained. To improve the convergence, the process is used as a preconditioner for GMRES, and the time-harmonic forcing term is multiplied by a smooth window function. In this talk we will discuss in particular a comparison with other time-domain solvers, as well as some recent performance result using graphics processing units (GPUs).

Chris Stolk
University of Amsterdam
C.C.Stolk@uva.nl

MS150

Maximum-Entropy Quadrature for Moment Equations in Kinetic Gas Theory

Using the method of moments to approximate the solution of kinetic equations is a popular technique especially to model rarefied gases. One attractive example the maximum-entropy closure, which, however, in relevant scenarios is computationally hardly affordable due to its non-linearity and singular behavior. In this talk, we will review the maximum entropy approach and combine it with the quadrature method of moments. This allows to keep positive aspects from each approach while eliminating problematic features. Entropic quadrature calculates a sparse quadrature-based reconstruction of the unknown velocity distribution in a discrete way and the physical meaningful maximum-entropy principle is employed as a criterion for selecting a denser, possibly continuous quadrature among different quadratures fitting the given moments. We will discuss the construction of the closure and give several numerical examples of its performance in one and two dimensions.

Manuel Torrilhon
RWTH Aachen University
mt@mathcces.rwth-aachen.de

MS151

Approximation Theory of Deep Learning for Sequence Modelling

In this talk, we present some recent results on the approximation theory of deep learning architectures for sequence modelling. In particular, we formulate a basic mathematical framework, under which different popular architectures such as recurrent neural networks, dilated convolutional networks (e.g. WaveNet), encoder-decoder structures can be rigorously compared. These analyses reveal some interesting connections between approximation, memory, spar-

sity and low rank phenomena that may guide the practical selection and design of these network architectures.

Qianxiao Li
National University of Singapore
Singapore
qianxiao@nus.edu.sg

MS151

Stabilized Neural Ordinary Differential Equations for Long-Time Forecasting of Dynamical Systems

In data-driven modeling of spatiotemporal phenomena careful consideration often needs to be made in capturing the dynamics of the high wavenumbers. This problem becomes especially challenging when the system of interest exhibits shocks or chaotic dynamics. We present a data-driven modeling method that accurately captures shocks and chaotic dynamics by proposing a novel architecture, the stabilized neural ordinary differential equation (ODE). Here we learn the right hand side of an ODE by adding the outputs of two neural networks (NN) together where one learns a linear term and the other a nonlinear term. Specifically, we implement this by training a sparse linear convolutional NN to learn the linear term and a dense fully-connected nonlinear NN to learn the nonlinear term. This is in contrast with the standard neural ODE which involves training only a single NN for learning the RHS. Our method is applied for learning the viscous Burgers and Kuramoto-Sivashinsky equations where we outperform the standard variant not only in terms of metrics but also in terms of the capture of key physical characteristics such as discontinuities in the former and chaotic invariant manifolds in the latter.

Romit Maulik
Argonne National Laboratory
rmaulik@anl.gov

MS151

Transformer with Fourier Integral Attentions

Multi-head attention empowers the recent success of transformers, the state-of-the-art models that have achieved remarkable success in sequence modeling and beyond. These attention mechanisms compute the pairwise dot products between the queries and keys, which results from the use of unnormalized Gaussian kernels with the assumption that the queries follow a mixture of Gaussian distribution. There is no guarantee that this assumption is valid in practice. In response, we first interpret attention in transformers as a nonparametric kernel regression. We then propose the FourierFormer, a new class of transformers in which the dot-product kernels are replaced by the novel generalized Fourier integral kernels. Different from the dot-product kernels, where we need to choose a good covariance matrix to capture the dependency of the features of data, the generalized Fourier integral kernels can automatically capture such dependency and remove the need to tune the covariance matrix. We theoretically prove that our proposed Fourier integral kernels can efficiently approximate any key and query distributions. Compared to the conventional transformers with dot-product attention, FourierFormers attain better accuracy and reduce the redundancy between attention heads. We empirically corroborate the advantages of FourierFormers over the baseline transformers in a variety of practical applications: language modeling, machine translation, image classification, continuous control

tasks.

Tan M. Nguyen
University of California, Los Angeles
tanmnguyen89@ucla.edu

MS151

Physics-Inspired Machine Learning for Sequence Modeling

Combining physics with machine learning is a rapidly growing field of research. Thereby, most work focuses on leveraging machine learning methods to solve problems in physics. Here, however, we focus on the reverse direction of leveraging structure of physical systems (e.g. dynamical systems modeled by ODEs or PDEs) to construct novel machine learning algorithms, where the existence of highly desirable properties of the underlying method can be rigorously proved. In particular, we propose several physics-inspired deep learning architectures for sequence modelling based on coupled oscillators, Hamiltonian systems and multi-scale dynamical systems. The proposed architectures address central problems in the field of recurrent sequence modeling, namely the vanishing and exploding gradients problem as well as the issue of limited expressive power. Finally, we show that this leads to state-of-the-art performance on several widely used benchmark problems ranging from image recognition over speech recognition to NLP applications.

Konstantin Rusch
ETH Zurich
konstantin.rusch@sam.math.ethz.ch

MS151

Universal Approximation of Random Neural Networks

In this paper, we study single-layer feedforward neural networks with randomly initialized weight matrices and bias vectors, which is inspired by the works on extreme learning machines, random feature models, and reservoir computing. In this case, only the linear readout needs to be trained, which can be performed, e.g., by a linear regression. Despite the popularity of this approach in empirical tasks, only little is known about the approximation capabilities of such networks. By considering these so-called “random neural networks” as Banach space-valued random variables, we provide several universal approximation theorems within the underlying Bochner space. Moreover, we extend the results to more general function spaces such as L^p -spaces and also consider the simultaneous approximation including the derivatives, which can be used for the approximation in Sobolev spaces.

Ariel Neufeld
Nanyang Technological University
Singapore
ariel.neufeld@ntu.edu.sg

Philipp Schmocker
Nanyang Technological University
philippt001@e.ntu.edu.sg

MS152

Reduced Order Surrogate Modelling and Latent

Assimilation for Dynamical Systems

Reduced-order modelling and low-dimensional surrogate models generated using machine learning algorithms have been widely applied in high-dimensional dynamical systems to improve the algorithmic efficiency. In this work, we develop several systems which combines reduced-order surrogate models with data assimilation (DA) inside the latent space to incorporate real-time observations from different physical spaces. We make use of local smooth surrogate functions which link the space of encoded system variables and the one of current observations to perform variational DA with a low computational cost. The proposed model can benefit both the efficiency provided by the reduced-order modelling and the accuracy of data assimilation. A variety of different applications, including wildfire spread prediction, CFD modellings for multiphase flow and surrogate models for microfluidic drops will be presented in this talk.

Sibo Cheng
Imperial College London
sibo.cheng@imperial.ac.uk

Rossella Arcucci
Data Science Institute, Imperial College London
r.arcucci@imperial.ac.uk

MS152

Error Assessment for a Finite Elements - Neural Networks Approach Applied to Parametric PDEs

We consider a parametric PDE

$$\mathcal{F}(u(x; \mu); \mu) = 0 \quad x \in \Omega, \quad \mu \in \mathcal{P},$$

where Ω denotes the physical domain and \mathcal{P} the parameters domain. During an offline phase, time consuming approximations $u_h(\cdot; \mu_i)$ of $u(\cdot; \mu_i)$ are computed using a finite element mesh with size h , for given μ_i , $i = 1, \dots, N$. Then, the parameter-to-solution map $\mu \mapsto u_h(\cdot; \mu)$ is approximated using a deep neural network, producing an approximation $u_{h,\mathcal{N}}(\cdot; \mu)$ of $u_h(\cdot; \mu)$. The evaluation of the neural network, performed in an online phase, is instantaneous, or at least much faster than a single numerical simulation. We next decompose the L^2 error in $\Omega \times \mathcal{P}$ into

$$\|u - u_{h,\mathcal{N}}\|_{L^2(\Omega \times \mathcal{P})} \leq \|u - u_h\|_{L^2(\Omega \times \mathcal{P})} + \|u_h - u_{h,\mathcal{N}}\|_{L^2(\Omega \times \mathcal{P})}.$$

Both error terms can be estimated using Monte-Carlo type estimates over the parameters space and an a posteriori error estimator in the physical space for the first term. In the presentation, we discuss in more detail the estimation of these two terms and show numerical experiments both for a model problem and a more complex one.

Maude Girardin
EPFL
maude.girardin@epfl.ch

Alexandre Caboussat
University of Applied Sciences Western Switzerland
alexandre.caboussat@hesge.ch

Marco Picasso
EPFL, Switzerland
marco.picasso@epfl.ch

MS152

Efficient Synthesis of Antenna Arrays Using

Latent-Space Modelling

A popular antenna type is the phased-array antenna, an antenna type characterized by a large quantity of individually controllable antennas working in unison. Typically, the applications impose strict requirements on the radiation pattern from the array antenna, which require detailed optimisation of the antenna parameters. A major drawback of phased-array antenna design is the array synthesis problem, i.e., determining the phases of the array such that the requirements are satisfied. The synthesis problem poses several major challenges that are tied to, e.g., conflicting objectives of the underlying optimisation, many local optima, and the need to carry out computationally costly online simulations. To address these challenges, this work demonstrates how latent-space modelling by custom-tailored autoencoders can augment and improve state-of-the-art, model-based synthesis of phased arrays. We explore an encoder-decoder structure that combines highly accurate simulation-based synthetic data with application-tailored models that are embedded into the network architectures, thereby guiding the latent space to reflect the real-world antenna problem, allowing low-rank subspaces to be extracted from the data. Numerical results demonstrate that the autoencoder has a significant potential to accelerate array synthesis to near real-time by generating good starting guesses for specialized optimizers.

Lasse Hjuler Christiansen, Oscar Borries, Niels Skovgaard Jensen, Simon Aertssen
TICRA
lhch@ticra.com, ob@ticra.com, nsj@ticra.com, sa@ticra.com

MS152

Generative Network-Based Reduced Order Model for Data Assimilation and Uncertainty Quantification in the Latent Space

The production of numerous high-fidelity simulations has been a key aspect of research for many problems in computational physics. The computational resources and time required to generate these simulations can be so large and impractical. With several successes of generative models, we propose a new method in which generative neural networks within a reduced-order model (ROM) framework are used for prediction, data assimilation and uncertainty quantification. A method has been developed which enables a generative network to perform time series prediction and data assimilation by training it with unconditional simulations of a discretized partial differential equation (PDE) model. After training, the generative model can be used to predict the spatio-temporal evolution of the physical states and observed data can be assimilated. We also describe the process required in order to quantify uncertainty, during which no additional simulations of the high-fidelity numerical PDE model are required. These methods work in the latent space (reduced space), which improves efficiency, also they take advantage of the adjoint-like capabilities of neural networks and the ability to simulate forwards and backwards in time. The results show that the proposed Generative Network-Based ROM can efficiently quantify uncertainty and accurately match the observed data, using only few unconditional simulations of the high-fidelity numerical PDE model.

Vinicius L. Santos Silva, Claire Heaney
Imperial College London
v.santos-silva19@imperial.ac.uk, c.heaney@imperial.ac.uk

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

MS152

Structure Preserving Machine Learning: Energy-Conserving Neural Network for Turbulence Closure Modelling

In turbulence modelling, and more particularly in the LES framework, we are concerned with finding a suitable closure model to represent the effect of the unresolved subgrid-scales on the larger/resolved scales. In recent years, the scientific computing community has started to gravitate towards machine learning techniques to attempt to solve this issue. However, stability and abidance by physical structure laws of the resulting closure models is still an open problem. We apply a spatial averaging filter to a high-resolution reference simulation to reduce the degrees of freedom of the system and derive a new kinetic energy conservation condition. We then suggest a data-driven compression to represent the subgrid-scale content as latent variables, living on the coarse grid, in order to comply with our new conservation condition. Finally, a skew-symmetric energy-conserving convolutional neural network architecture is introduced that can be enhanced with dissipative terms to account for viscous flows. Combined with a structure-preserving discretization this framework is used to evolve both the filtered solution and the latent subgrid-scale representation in time in a structure-preserving fashion. This yields stability while still allowing for backscatter. We apply the methodology to both the viscous Burgers' equation and Korteweg-De Vries equation in 1-D and show increased accuracy and stability as compared to a standard convolutional neural network.

Wouter Edeling, Toby van Gastelen
CWI
wouter.edeling@cwi.nl, tvg@cwi.nl

MS153

Challenges and Requirements in Modeling Electrical Machines for Digital Twins

A Digital twin is a virtual representation of a system that uses physics based and data driven simulations to improve prediction, optimization, monitoring, controlling of the system. In this work the focus is on developing a digital twin for an electrical machine. Based on the functions of the digital twin, we present the requirements on the simulations that needs to be developed and the main challenges facing the construction of these models. These challenges include real-time simulations, multiphysics coupling, data assimilation and adaptive modeling. We present an initial investigation of those challenges by proposing solutions using hierarchical and port-Hamiltonian modeling. Port-Hamiltonian systems provide a framework that works across different physical domains and different scales and that can help in the development of the digital twin for electrical machines, in particular on the industrial scale. An outlook at potential solutions of the challenges and topics that are still open is also discussed.

Karim Cherifi
TU Berlin
cherifi@math.tu-berlin.de

Philipp Schulze
Chemnitz University of Technology

philipp.schulze@math.tu-chemnitz.de

Volker Mehrmann
Technische Universität Berlin
mehrmann@math.tu-berlin.de

MS153

Advanced Technologies Enabling Augmented Intelligence and Twins

The present paper revisits the state of the art and recent developments on the different components of a physics-informed digital twin, in particular model order reduction techniques enabling solving physics under the stringent real-time constraints; then, advanced machine learning techniques able to perform from quite reduced amount of data and operating online. The so-called hybrid-twin results from the hybridation of both components, making use of the existing knowledge and the collected data appropriately assimilated into the models (calibration) and on the machine learning training. The issue related to the optimal data collection (what data where and when) as well as the most informative one (active leaning) will be also addressed. The general framework will be illustrated with some major achievements.

Francisco Chinesta
ENSAM Institute of Technology
francisco.chinesta@ensam.eu

Elias Cueto
I3A, University of Zaragoza, Spain
ecueto@unizar.es

MS153

On a Digital Twin for Contact Mechanics Driven by Operator-Inference-Based Model Order Reduction Approach

Many industrial problems include mechanical contacts, which often are a pitfall in the context of real-time simulations. Model order reduction techniques are in general utilized for enabling fast simulations. For constrained problems such as contacts, an additional treatment is required. Novel physics-based reduction methods for contact problems were presented in [Manvelyan et al., *Comput Mech* 68, 12831295 (2021), Manvelyan et al., *Int J Numer Methods Eng*, 1-27, (2022)]. Within this work we introduce a new method, which is a generalization of the recent physics-based reduction frameworks. The generalized method is driven by operator inference method and adjoint methods. That way the unconstrained problem can be learned from data stemming from black-box models, whereas the contact behavior is reproduced within the adjoint system describing the Lagrange multipliers, i.e., the variables that enforce the contact conditions.

Diana Manvelyan
Siemens AG
diana.manvelyan@siemens.com

Utz Wever
None
utz.wever@t-online.de

MS153

Real-Time Probabilistic Learning and Virtualiza-

tion of Dynamic Systems

Among the large assortment of Structural Health Monitoring (SHM) techniques, vibration-based methods emerge as a promising means of automated assessment for structural and mechanical systems. These methods aim at providing a diagnostic suit, which is able to notify the operators on the occurrence of faults as well as on the necessity to undertake possible actions in real-time. Although such strategies may prove significantly more efficient than traditional time-based maintenance approaches, the challenge of implementation consists in the requirement for more sophisticated data processing techniques that will allow to distill the measured information and extract a set of robust and condition-sensitive indicators. This contribution is focused on the fusion on data-driven and physics-based modeling techniques, in order to devise a computationally efficient framework for online assimilation of output-only vibration measurements into the virtual instances of monitored systems, towards the materialization of real-time Digital Twins (DTs). The major challenges to be addressed in this respect are associated with the methodological tools and the modeling-related aspects, which comprise the uncertainty and complexity related to global model-based approaches, the nonlinear effects occurring on slender components, the system variability induced by the changing environmental and operational conditions and the existence of structural flaws.

Konstantinos Tatsis
ETH
tatsis@ibk.baug.ethz.ch

Eleni Chatzi
Institute of Structural ENgineering
ETH Zurich
chatzi@ibk.baug.ethz.ch

MS153

Adaptive Planning for Predictive Digital Twins

A digital twin is a virtual model dynamically updating which is able to mirror a physical asset of interest throughout its operational lifespan. Digital twins are rapidly spreading and are used in a diverse range of engineering applications. They allow predictive maintenance, optimization, and planning, among others. In this talk we present a digital representation of an unmanned aerial vehicle, focusing on its structural health, in the framework of autonomous aerial cargo missions. We show how to incorporate adaptive planning, assimilating information during the operational regime. We parametrize the state transition probability of the underlying Markov decision process with a beta distributed random variable. We exploit a Bernoulli process to update the posterior distribution without the need of numerical integration. This allows the selection of a new policy at every time step, adapting online the planning strategy to balance the structural self-preservation and time to arrive at client location.

Marco Tezzele
Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
marco.tezzele@austin.utexas.edu

Karen E. Willcox
University of Texas at Austin, U.S.

kwillcox@utexas.edu

MS154

The Eu-Maths-in OpenDesk: a European One Stop Shop for Researchers and Entrepreneurs

The EU-Maths-IN OpenDesk launched in 2022 will be presented as a European one-stop-shop to facilitate the connection between the more than 400 research centres participating in EU-Maths-IN and industry. The OpenDesk aims to be a driving force for European innovation for industries and companies from all sectors, small and medium-sized enterprises, start-ups, and administrations. In this talk, the portfolio of services offered to companies will be presented, as well as OpenDesk workflow for both customers - industry in the broadest sense - and service providers - the nearly 9000 researchers integrated around 400 research centres of excellence as nodes of the 20 national networks of EU-Maths-IN.

Manuel B. Cruz

LEMA - Engineering Mathematics Laboratory, School of Engineer
OpenDesk Committee of EU-Maths-IN
mbc@isep.ipp.pt

Veronique Maume-Deschamps
Université Claude Bernard Lyon 1 - Institut Camille Jordan
OpenDesk Committee of EU-Maths-IN
veronique.maume-deschamps@agence-maths-entreprises

Peregrina Quintela
CITMAga and University of Santiago de Compostela
OpenDesk Committee of EU-Maths-IN
peregrina.quintela@usc.es

Alexander Scherrer
Fraunhofer Institute for Industrial Mathematics ITWM
OpenDesk Committee of EU-Maths-IN
alexander.scherrer@itwm.fraunhofer.de

Antonino Sgalambro
Sheffield University Management School, UK
National Research Council of Italy and OpenDesk Committee
a.sgalambro@sheffield.ac.uk

Janusz Szwabinski
Wroclaw University of Science and Technology
OpenDesk Committee of EU-Maths-IN
janusz.szwabinski@pwr.edu.pl

MS154

Numerical Simulation of Innovative Processes for Silicon Production in the SisAl Project

SisAl Pilot is an Horizon 2020 funded project coordinated by Norwegian University of Science and Technology (NTNU) which comprises 22 partners from 9 countries. The main objective of this project is to demonstrate a patented novel industrial process to produce silicon. The actual carbothermic Submerged Arc Furnace (SAF) process is replaced by a far more environmentally and economically sustainable alternative: the aluminothermic reduction of quartz, which allows to use secondary raw materials such as aluminium (Al) EoL scrap and dross, instead of carbon reductant used today. To attain this goal, differ-

ent types of furnaces are being analyzed. In this contribution, we will focus on modelling and simulation of induction and rotary furnaces. Depending on the furnace, the simulations require to study several physical processes strongly coupled: heat transfer, multiphase fluid dynamics, electromagnetism, melting processes and chemical reactions. Thus, the challenge is to carry out numerical simulations based on these models that can support the experimental trials in plant of the industrial partners Elkem AS, Technology (Norway) and Fundiciones Rey SL (Spain).

Pilar Salgado
Universidade de Santiago de Compostela
Spain
mpilar.salgado@usc.es

scar Crego
Galician Centre for Mathematical Research and Technology
oscar.crego@usc.es

José Luis Ferrín
Universidade de Santiago de Compostela
Galician Centre for Mathematical Research and Technology
jose-luis.ferrin@usc.es

Branca García
Galician Centre for Mathematical Research and Technology
branca.garcia.correa@usc.es

Dolores Gómez
Universidade de Santiago de Compostela
Galician Centre for Mathematical Research and Technology
mdolores.gomez@usc.es

MS154

Real-Time Boundary Heat Flux Estimation in Continuous Casting Molds Using Data Assimilation

In continuous casting of steel, the most critical component is the mold. In the mold, the steel begins its solidification, and several complex physical phenomena happen. To ensure a proper control of the process, it is necessary to know how the steel is behaving inside the mold. However, it is not possible to make measurements inside the solidifying steel and the only available data are pointwise temperature measurements in the interior of the mold plates. To provide a tool for the proper control of the process, we developed a methodology for the real-time estimation of the heat flux at the steel-mold interface given the temperature measurements. With this tool, we allow the caster operator to quickly detect any malfunctioning in the casting increasing the safety and the productivity of continuous casters.

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

Giovanni Stabile
SISSA mathLab, Trieste, Italy
giovanni.stabile@uniurb.it

Umberto Emil Morelli
SISSA, USC, CITMAga

umbertoemil.morelli@usc.es

Peregrina Quintela
CITMAga and University of Santiago de Compostela
OpenDesk Committee of EU-Maths-IN
peregrina.quintela@usc.es

Patricia Barral
USC
patricia.barral@usc.es

MS154

A Digital Twin for Urban Bus Routes Comparison

In Barcelona, there is service called "Bus del Barri" with eminently social routes usually that have a meagre recovery ratio. Their itineraries typically respond more to the community's will than to recommendations found in specialized handbooks. There is a problem that consists on deciding if it is feasible the replacement of a group of district routes by a Bus on Demand model (DRT) by pursuing an alternative that offers a level of service equal to or higher than a group of local fixed routes at a comparable or lower cost. A DRT service is a new model of transportation based on the concept of a tailored ride with a starting and an ending point previously agreed by the user and the company. This new model aims to be a more economical and ecological alternative. The main objective of this talk is to present a real-time simulation approach by means of a digital twin (with a modular structure). These modules handle, for example, the generation of potential service requests by using the available data to feed forecasting models of event demand. The computation of all possible combinations of optimal routes is also considered and, also the digital twin is responsible of the choice of the route that best suits the service requirements (through a utility function) using a genetic algorithm. All together it allows to extract a series of indicators that help to decide if the change is useful or not.

Lluís Alsedà
Centre de Recerca Matemàtica
Universitat Autnoma de Barcelona
director@crm.cat

Miquel Barcelona
Universitat Autnoma de Barcelona
mbarcelona@mat.uab.cat

Josep Mensió
Transports Metropolitans de Barcelona
jmension@tmb.cat

Joan Pareras
Centre de Recerca Matemàtica
jpareras@crm.cat

David Romero
Universitat Autnoma de Barcelona
Spain
dromero@crm.cat

MS155

Quadratization of Differential Equations: the Connection with Analogic and Digital Computability of Their Solutions

In this talk a recent method, said 'exact quadratization',

of manifold immersion where the image vector field is always quadratic in the local, whatever chosen, coordinates is described, and the relevance of such kind of immersion is illustrated in problems of applied mathematics. In particular, it will be shown that exact quadratization allows a fast calculation of the Taylor coefficients in Taylor series methods for ODE numerical integration. The problem of classify quadratizable nonlinear functions is also sketched. It will be shown that this problem is closely related with computability, and in particular that the class of quadratizable functions consists of all the functions that can be generated though analog computer, which in turn consists (for ODEs) of all differentially algebraic functions. Connections with the algebraic theory of control is highlighted as well.

Francesco Carravetta

Istituto di Analisi dei Sistemi ed Informatica
Rome, Italy
francesco.carravetta@iasi.cnr.it

MS155

Quadratization for Periodic Solutions Continuation

The talk will be concerned with the quadratization process that we are using in the in-house software Manlab-4 devoted to the continuation of periodic solution branches of smooth ODE systems. Our numerical strategy relies on a successive use of two series expansion: a high order Fourier series for the periodic solution finding and a high order Taylor series for the continuation process (the so-called Asymptotic-numerical Method). In order to be able to easily and automatically derive the reduced equations (by collecting term with the same power), the ODE system is augmented with auxiliary variables and equations, so that each equation becomes either a quadratic one or involves only a single transcendental function (ex: $y=\exp(x)$). The differentiated form of the transcendental function is also written quadratically ($dy= x*dx$) for use in the series computation. The size of the augmented system may be much higher than the one of the original ODE system, but since we use a systematic condensation of the auxiliary variables each time we solve a linear system, this is not a drawback. The last implementation of this strategy makes an intensive use of sparse tensors of order one, two and three, and its computational performances allows to treat moderate size systems obtained by a finite element discretization of a continuous problem. Various examples ranging from solid vibrations, nonlinear acoustics and fluid-structure interactions will be presented.

Bruno Cochelin

CNRS, Marseille, France
bruno.cochelin@centrale-marseille.fr

Louis Guillot
LMA Marseille
guillot@lma.cnrs-mrs.fr

Christophe Vergez
LMA marseille
vergez@lma.cnrs-mrs.fr

MS155

On the Quadratization Problem Preserving the Set of Stable States Only

In previous work, we have shown the NP-hardness of the quadratization optimization problem for transforming a

polynomial ordinary differential equation (PODE) given in explicit form in non-succinct matrix representation, into a PODE of degree at most 2, preserving the solution of the original variables, and introducing a minimum number of monomial variables with bounded degrees. That problem is part of our compilation pipeline for transforming any computable real function presented by a PODE, into an elementary Chemical Reaction Network (CRN) which computes that function on a distinguished output species. In this talk, we are interested in a notion of online robust analog computation using stabilizing CRN/PODE, i.e. systems for which the input/output relation is solely defined by the set of stable fixed points. This leads us to study a variant of the quadratization problem which preserves the set of stable states of the original system instead of the complete trajectories. While this problem may seem simpler than the previous one, we show its NP-hardness in the non-succinct representation with a reduction of the cover vertex set problem as previously. We report however on satisfactory practical performance results obtained with our MAXSAT implementation on a benchmark of CRN synthesis problems.

Mathieu Hemery
INRIA Saclay, France
mathieu.hemery@polytechnique.org

Francois Fages
Inria Saclay
francois.fages@inria.fr

MS155 Towards Automatic Quadratization for PDEs

Quadratization problem is, given a system of ODEs with polynomial right-hand side, transform the system to a system with quadratic right-hand side by introducing new variables. Such transformations have been used, for example, as a preprocessing step by model order reduction methods and for transforming chemical reaction networks. We will present an extension of a recent algorithm for finding optimal quadratizations with new variables being monomials to the case of systems with inputs, and show its applications to models from the literature. We will use this to give an algorithm for quadratizing evolutionary polynomial PDEs. The talk is based on joint works with Andrey Bychkov and Boris Kramer.

Gleb Pogudin
Ecole Polytechnique
gleb.pogudin@polytechnique.edu

MS155 On Convergence Analysis of a First-Order Semi-Discrete Numerical Scheme for a Hydrodynamic Q-Tensor Model for Liquid Crystals

We present a convergence analysis of an unconditional energy-stable first-order semi-discrete numerical scheme designed for the hydrodynamic Q-tensor model based on the Invariant Quadratization Method (IEQ). This model couples a Navier-Stokes system for the flows and a parabolic type Q-tensor system governing the nematic crystal director fields. The invariant Quadratization Method enables us to construct a linear scheme for the systems, accelerating the computation speed. However, it introduces an auxiliary variable into the systems to replace the bulk potential energy. A typical problem is whether this new variable coincides with the original energy. In this work,

we prove the stability properties of the scheme and show its convergence to a weak solution of the coupled liquid crystal system. In the weak sense, we have also demonstrated the equivalence of the solution for the reformulated and original systems. Numerical experiments are also provided.

Yukun Yue
Carnegie Mellon University
yukuny@andrew.cmu.edu

MS156 Projection Scheme in a Discrete Exterior Calculus Framework for the Navier-Stokes Equations

Discrete exterior calculus (DEC) is a discretization within the family of geometric integrators. A fundamental property of DEC is that it naturally satisfies the relation $d^2 = 0$ to machine precision, where d is the discrete exterior differential operator. This relation is essential as it describes both identities $\text{curl grad} = 0$ and $\text{div curl} = 0$ of the vector calculus. Thus, ensuring these relations at the discrete level circumvents parasitic solutions. Many attempts have been made to predict flows by DEC, and the stream-function approach is often preferred due to its pressure-free formulation. For many other applications where the explicit computation of the pressure is required, we need to solve the governing equations in primary variables (velocity-pressure). In this case, Chorin's projection scheme based on the Helmholtz-Hodge decomposition can be used. It consists of predicting an intermediate velocity field, which is subsequently projected onto a solenoidal space via a Poisson equation for the pressure with Neumann boundary conditions. In this presentation, a projection scheme adapted to DEC is formulated. In particular, we show how Neumann boundary conditions can be incorporated into the DEC formulation, independent of the choice of the discrete Hodge operator. Validation experiments are performed in a benchmark simulation: lid-driven cavity flow. We show, for different Reynolds numbers, that our results very accurately match those in the literature.

Rama Ayoub
King Abdullah University of Science and Technology
Physical Science and Engineering Division
rama.ayoub@kaust.edu.sa

Aziz Hamdouni, Dina Razafindralandy
LaSIE, La Rochelle University
France
aziz.hamdouni@univ-lr.fr, dina.razafindralandy@univ-lr.fr

MS156 Generating DEC Simulations from Diagrammatic Equations

When assembling multiphysics simulations, the complexity of mathematics and software can overwhelm the simulation developer. By approaching the representation and assembly of multiphysics problems as a synthetic differential geometry problem and using de Rahm complexes to model partial differential equations, we developed Decapodes.jl, a software library for representing, composing, and solving partial differential equations. The diagrammatic representations provide an intuitive interface for specifying the relationships between variables in a system of equations, a method for composing systems equations into a multiphysics model using an operad of wiring diagrams, and an algorithm for deriving solvers using directed hypergraphs.

This approach yields a method of generating executable systems from these diagrams using the operators of discrete exterior calculus on a simplicial set. The generated solvers produce numerical solutions consistent with state of the art open source tools as demonstrated by benchmark comparisons with SU2. This software demonstrates the feasibility of synthetic approach to differential geometry as a foundation for numerical multiphysics simulation and identifies areas requiring further development, potentially via implementation of the finite-element exterior calculus.

Andrew Baas, Jesus Arias
Georgia Tech Research Institute
andrewbbaas@gmail.com, jesus.arias@gtri.gatech.edu

Evan Patterson
The Topos Institute
evan@topos.institute

James P. Fairbanks
Georgia Tech Research Institute
fairbanksj@ufl.edu

MS156

Naturality of Discrete Exterior Derivative and Discrete Wedge Product

In discrete exterior calculus, simplicial cochains play the role of discrete forms, the coboundary operator serves as the discrete exterior derivative and an antisymmetrized cup product like operator plays the role of wedge product. In smooth calculus on manifolds, an important property of the exterior derivative and wedge product is the naturality of these two operators with respect to smooth maps. That is, the exterior derivative commutes with pullbacks and so does the wedge product. Both these naturality properties are important in the smooth theory. For example, the chain rule in calculus is precisely the naturality of the exterior derivative. We show that in the discrete theory, abstract simplicial maps play the role of smooth maps in this context. That is, the discrete exterior derivative and wedge product commute with pullback via abstract simplicial maps.

Anil Hirani
Department of Mathematics
University of Illinois at Urbana-Champaign, IL, USA
hirani@illinois.edu

Mark Schubel
Apple, Inc.
mdschubel@gmail.com

Daniel Berwick-Evans
University of Illinois at Urbana-Champaign
danbe@illinois.edu

MS156

On the Role of the Weights in Polynomial Interpolation of Differential Forms

For the polynomial interpolation of differential k -forms on simplices, we present a framework that allows to generalize fundamentals concepts featuring in the classical scalar case. To this purpose, we rely on particular degrees of freedom that fit with the physical and geometrical nature of the field to interpolate: the weights on small k -simplices. A weight is the integral of a k -form on a k -simplex. As

nodal evaluations of a scalar field are done at the nodes of suitably selected sets in the mesh elements, similarly, weights of a field, intended as a k -form, are computed, in each mesh element, on suitably selected sets of simplices of dimension k , called small k -simplices (that are nodes for $k = 0$, edges if $k = 1$, faces for $k = 2$, volumes for $k = n$). Weights on small k -simplices play the role of interpolation coefficients when reconstructing scalar/vector fields in terms of a set of selected multivariate polynomial forms. As in the nodal case, a generalised Lebesgue constant pops up naturally to measure the stability of the interpolation and the Runge phenomenon may appear for particular distributions of small k -simplices.

Ana Alonso Rodriguez
Dipartimento di Matematica
Universita degli studi di Trento
ana.alonso@unitn.it

Francesca Rapetti
Université de Nice Cote d'Azur
francesca.rapetti@univ-cotedazur.fr

MS156

An Introduction to DEC and FEEC

Discrete Exterior Calculus (DEC) and Finite Element Exterior Calculus (FEEC) are two closely related frameworks for structure-preserving numerical modeling and computation. Both are based on the exterior calculus of differential forms, which generalizes the usual calculus of scalar and vector fields, making it possible to express a wide variety of PDEs in this language. As the opening talk of this minisymposium, this presentation will give a brief introduction to exterior calculus and a survey of the DEC and FEEC frameworks.

Ari Stern
Washington University in St. Louis
stern@wustl.edu

MS157

Escaping the Abstraction: a Foreign Function Interface for the Unified Form Language (ufl)

High level domain specific languages for the finite element method underpin high productivity programming environments for simulations based on partial differential equations (PDE) while employing automatic code generation to achieve high performance. However, a limitation of this approach is that it does not support operators that are not directly expressible in the vector calculus. This is critical in applications where PDEs are not enough to accurately describe the physical problem of interest. Examples include nonlinear implicit constitutive laws such as the Glen's flow law for glacier flow, the use of deep learning models to include features not represented in the differential equations, or closures for unresolved spatiotemporal scales. We introduce an interface within the Firedrake finite element system that enables the inclusion of arbitrary operators in PDE-based problems. This new feature composes with the automatic differentiation capabilities of Firedrake, enabling the automated solution of inverse problems.

Nacime Bouziani
Imperial College
n.bouziani18@imperial.ac.uk

David Ham

Department of Mathematics
Imperial College London
david.ham@imperial.ac.uk

MS157

Mixed-Domain Coupled Finite Elements in FEniCSx

Mixed-dimensional partial differential equations (PDEs) are equations coupling fields defined over distinct domains that may differ in topological dimension. Such PDEs naturally arise in a wide range of fields including geology, bio-medicine, and fracture mechanics. Mixed-dimensional models are also used to impose non-standard conditions through Lagrange multipliers. Finite element discretizations of such PDEs involve nested meshes of possibly heterogeneous topological dimension. The assembly of such systems is non-standard and non-trivial, and requires the design of both generic high level software abstractions and lower level algorithms. The FEniCS project aims at automating the numerical solution of PDE-based models using finite element methods. A core feature is a high-level domain-specific language for finite element spaces and variational forms, close to mathematical syntax. Lately, FEniCS gave way to its successor FEniCSx, including major improvements over the legacy library. An automated framework was developed in core FEniCS legacy libraries to address the challenges characterizing mixed-dimensional problems. These concepts were recently ported to FEniCSx, taking advantage of the underlying upgrades in the library features and design. This talk gives an overview of the abstractions and algorithms involved, and their implementation in the FEniCS project core libraries. The introduced features are illustrated by concrete applications in engineering and biomedicine.

Cécile Daversin-Catty
Simula
cecile@simula.no

Joseph P. Dean
Cambridge
jpd62@cam.ac.uk

Marie E. Rognes
Simula Research Laboratory
meg@simula.no

Garth N. Wells
University of Cambridge
gnw20@cam.ac.uk

MS157

xDSL: A Compiler Infrastructure for Python DSLs

Domain scientists simulate increasingly larger and more complex phenomena, and as a result, must reason about complex low-level parallelism to harness the full power of heterogeneous architecture. Domain Specific Languages (DSLs) abstract away lower-level decisions for multiple back-ends and architectures, allowing domain experts to reason only about high-level concepts. However, DSLs are often developed in isolation and require a lot of effort on the part of the DSL developer, because they cannot reuse pre-existing abstractions defined by other compilation pipelines. We present xDSL, an extensible Python compiler framework that allows compiler developers to reuse abstractions of existing DSLs. xDSL can also lever-

age the compiler back-end provided by MLIR, which is a new compiler framework from the LLVM community. Additionally, we will present our first integration of xDSL in the Devito and Psyclone compilers, to demonstrate the capabilities of xDSL towards representing different abstractions. The goal of this project is to help different communities share compiler abstractions, in order to build a rich ecosystem of shared intermediate representations and optimizations instead of a large set of isolated monolithic compilers.

Mathieu Fehr, Larisa Stoltzfus
University of Edinburgh
mathieu.fehr@ed.ac.uk, larisa.stoltzfus@ed.ac.uk

MS157

The Next Step: Interoperable Domain-Specific Programming

Even though domain-specific programming approaches allow for readable, scalable, and maintainable software without sacrificing performance, the new paradigm of learned physics-informed models calls for an interdisciplinary approach typically involving multiple domain-specific languages. Take for example the problem of inverting for the fluid-flow properties from time-lapse seismic data, which entails domain-specific programming on the intersection of wave simulators, matrix-free linear algebra, learned neural surrogates for two-phase flow, and prior and posterior distributions for the fluid-flow properties. While domain-specific solutions exist for each of these sub-disciplines, integrating these approaches which may involve different programming languages into a single coupled scalable inversion framework that supports algorithmic differentiation can be a challenge. However, we show that challenges like this can be met when working with proper abstractions. In our inversion example, this involves math-inspired symbolic abstractions for numerical solutions of the wave equation (Devito), matrix-free implementations for its Jacobians (JUDI.jl), abstractions for Automatic Differentiation (ChainRules.jl), and homegrown implementations for conditional Invertible Neural Networks (InvertibleNetworks.jl) and Fourier Neural Operators (ParametricOperators.jl).

Felix Herrmann, Mathias Louboutin, Thomas Grady, Ziyi Yin
Georgia Institute of Technology
felix.herrmann@gatech.edu, mlouboutin3@gatech.edu, tgrady@gatech.edu, ziyi.yin@gatech.edu

Rishi Kahn
Extreme Scale Solutions
rishi@extreme-scale.com

MS157

ProtoX: A First Look

Stencil operation is a key component in the numerical solution of partial differential equations. Developers tend to use different libraries that provide these operations for them. One such library is Proto. It is a C++ based domain specific library designed to provide an intuitive interface that optimizes the designing and scheduling of an algorithm aimed at solving various partial differential equations numerically. The high level of abstractions used in Proto can be fused together to improve its current performance. However, abstraction fusion cannot be performed easily by a compiler. In order to overcome this shortcom-

ing we present ProtoX, a code generation framework for stencil operation based on Proto and uses SPIRAL as its backend. SPIRAL is a GAP based code generation system that focuses on generating highly optimized target code in C/C++. We demonstrate the construction of ProtoX by considering two examples, the 2D Poisson problem and the Euler equations that appear in the study of gas dynamics. Some of the code generated for these two problem specifications is shown along with the initial speedup result.

Het Y. Mankad
Carnegie Mellon University
hmankad@andrew.cmu.edu

Sanil Rao
CMU
sanilr@andrew.cmu.edu

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

Brian Van Straalen
Lawrence Berkeley National Laboratory
Computational Research Division
bvstraalen@lbl.gov

Franz Franchetti
Carnegie Mellon University
franzf@andrew.cmu.edu

MS158

Learning-to-Prune: A Machine Learning Framework to Speed Up Combinatorial Optimisation Algorithms

In a large number of industrial applications, combinatorial optimization problems are repeatedly solved with datasets from similar distribution. In recent years, machine learning techniques have been shown to be quite effective in speeding up such computations. However, black-box end-to-end machine learning approaches suffer from poor interpretability and the requirement for a large amount of labelled data. In this talk, I will present a simple and highly effective way to incorporate the insights from the algorithmic and optimization literature on these problems into a machine learning framework to speed-up the solutions of these problems. We considered a range of optimization problems: Steiner trees, travelling salesperson problem, k-median and set cover. These problems are well studied and a number of algorithms have been designed for these problems. We look at the kind of quantities these algorithms employ and use these to derive useful features for training a classifier that helps quickly reduce the problem size by identifying the difficult core of the problem and pruning the remainder. The difficult core is then solved using an Integer Linear Programming (ILP) solver. A prime advantage of using such features is that we do not require much data to train the classifier. This results in algorithms returning a better trade-off between running time and solution quality compared to commercial ILP solvers.

Deepak Ajwani
University College Dublin
deepak.ajwani@gmail.com

Dena Tayebi
University College Dublin, Ireland

dena.tayebi@ucdconnect.ie

James Fitzpatrick, Paula Carroll
University College Dublin
james.fitzpatrick1@ucdconnect.ie, paula.carroll@ucd.ie

Saurabh Ray
NYU Abu Dhabi
sr194@nyu.edu

MS158

Structural Analysis of Unsatisfiable Cores for a Combinatorial Problem

Our work deals with the problem of searching for small subsets in the set $[n] = \{1 \dots n\}$ that do not contain any arithmetic triplets (i.e. elements a, b, c (sorted) such that $b - a = c - b$). Furthermore, the set has to be saturated in the sense that if any other number from $[n]$ is added to the subset, it forms an arithmetic triplet with (at least) two other numbers. We are interested in empirically computing the smallest such subset and finding interesting properties in the process. Instead of crafting our solver for this particular problem, we opted to construct an encoding to the SAT problem and use state-of-the-art solvers to solve it. We further explore Minimal Unsatisfiable Cores (MUS) to see what interesting properties could be extracted from them. We generated all the MUSes for small n and observe vast differences in the number of MUSes both for different sizes as well as different encodings. These expressions were then transformed into graphs and we explore their structural properties. We emphasize the main differences and connect them with the semantics of the original problem. The main goal was to compute a set of graph properties and see how they behave, and how they change in various contexts. By using common data-exploration techniques we identified interesting patterns and we connect these patterns through the increasing size of the problem. We present the results of this exploration and show the potential of such analysis for other similar problems.

Uros Cibej, Jurij Mihelic
University of Ljubljana
uros.cibej@fri.uni-lj.si, Jurij.Mihelic@fri.uni-lj.si

MS158

Graph Theory and Combinatorial Scientific Computing: Review and Perspectives

Problems on the discrete side of scientific computing are solved by various means. Sometimes we use well developed solvers like LP or SAT. And sometimes we develop specialized algorithms, as general solvers may have drawbacks like speed or robustness. In the course of developing such algorithms graphs and graph theoretical approach appear more and more often. In some cases graph emerge naturally as obvious models. In some chemical related problems the chemical structure represented by graph structure. Distances measured and clusters searched in network models for roads, social interactions or chemical interactions in complex systems as our body. In other cases graphs used as a modeling language alike numbers, and thus graph models could represent hidden relationships far from obvious. For analysis of the stock market the market graph is build from pairwise temporal correlation of stock price. Auxiliary graphs can represent scheduling problems, protein docking sites and many more. In this talk we will review the state of the art methodologies in the above sense with a special focus on algorithmic complexity. Concentrating on

hard problems we will highlight the modeling expressivity as well as the efficiency of different algorithmic engineering approaches. We will also outline the potential approaches for future research trends.

Miklos Kresz
Innorenw CoE
miklos.kresz@innorenw.eu

Bogdan Zavalnij
Alfred Renyi Institute of Mathematics
bogdan@renyi.hu

MS158

Graph-Based Matrix Completion Applied on Weather Data

Low-rank matrix completion is the task to recover unknown entries of a matrix from partial observations by assuming that the true matrix admits a good low-rank approximation. Sometimes additional information about the variables is known, and incorporating this information into a matrix completion model can lead to a better completion quality. This information between the column/row entities of the matrix is expressed as a graph. In our work, we are interested in the network of automatic stations of the Royal Meteorological Institute (RMI) in Belgium, which provides temperature measurements at 10-minute intervals. We evaluated the effectiveness and importance of good spatial and temporal graphs for improving the completion of missing data in these meteorological data.

Benot Loucheur
Université Catholique de Louvain
benoit.loucheur@uclouvain.be

Pierre-Antoine Absil
Université Catholique de Louvain
ICTEAM Institute
pa.absil@uclouvain.be

Michel Journée
Royal Meteorological Institute of Belgium
michelj@meteo.be

MS158

Faster Greedy Optimization of Resistance-Based Graph Robustness

The total effective resistance, also called the Kirchhoff index, provides a robustness measure for a graph G . We consider the optimization problem of adding k new edges to G such that the resulting graph has minimal total effective resistance (i. e., is most robust). The total effective resistance and effective resistances between nodes can be computed using the pseudoinverse of the graph Laplacian. The pseudoinverse may be computed explicitly, but this takes cubic time in practice and quadratic space. We instead exploit combinatorial and algebraic connections to speed up gain computations in established generic greedy heuristics. Moreover, we leverage existing randomized techniques to boost the performance of our approaches by introducing a sub-sampling step. Our different graph- and matrix-based approaches are indeed significantly faster than the state-of-the-art greedy algorithm, while their quality remains reasonably high and is often quite close. Our experiments show that we can now process large graphs for which the application of the state-of-the-art greedy approach was infeasible before. As far as we know, we are the first to be

able to process graphs with 100K+ nodes in the order of minutes.

Henning Meyerhenke
Humboldt-Universität zu Berlin
meyerhenke@hu-berlin.de

Maria Predari
Humboldt University of Berlin
predarimaria@gmail.com

Robert E. Kooij
Delft University of Technology
r.e.kooij@tudelft.nl

MS159

An Optimal Eighth-Order Three-Step Multiple Roots Solver For Nonlinear Equations

In this paper, we develop an optimal eighth-order multiple roots solver for nonlinear equations with multiplicity ' m '. The new scheme is established by using only four functional evaluations and three weight functions. The performance of presented scheme is tested through numerical examples of different nature. Numerical results of the proposed scheme have an edge over the other existing robust methods in the literature. Basins of attraction are also presented to compare the results of proposed methods with the existing ones.

Ranjita Guglani
Assistant Professor, Department of Mathematics,
D.A.V. College, Jalandhar
ranjitakapur@gmail.com

Rajni Sharma
Assistant Professor, Department of Applied Science,
D.A.V. Institute of Engineering and
rajni_daviet@yahoo.com

Ashu Bahl
Associate Professor, Department of Mathematics,
D.A.V. College, Jalandhar
bahl.ashu@rediffmail.com

MS159

High-Resolution High-Order Weighted Essentially Non-Oscillatory Schemes for Reshocked Single-Mode Richtmyer-Meshkov Instability

Turbulent mixing due to hydrodynamic instabilities occurs in a wide range of science and engineering applications. The experimental, theoretical, and numerical studies help us to understand the dynamics of hydrodynamically unstable turbulent flows such as Kelvin-Helmholtz, Rayleigh-Taylor, and Richtmyer-Meshkov Instabilities. In this talk, we present an increasingly accurate and robust front tracking method for the numerical simulations of single mode Richtmyer-Meshkov Instability (RMI) of an air/SF6 interface. The front tracking method is used to track the interface explicitly with high order accuracy. This technique stores and dynamically evolves a meshed front that partitions a simulation domain into two or more regions, each representing a different material or physics model. The simulations based on front tracking with the classical weighted essentially non-oscillatory (WENO) schemes in Jiang and Shu (1996) with Yang's artificial compression and Balsara and Shu (2000) are compared with Collins and

Jacobs (2002) shock tube experiments. We address verification and validation issues to achieve good agreement on the amplitude and the displacement of interface at all times including the pre-reshock and post-reshock.

Tulin Kaman
University of Arkansas
tkaman@uark.edu

Ryan Holley
University of Arkansas, U.S.
rh027@uark.edu

MS159

A Multivariate Spline Based Collocation Method for Numerical Solution of PDEs

We propose a collocation method based on multivariate polynomial splines over triangulation or tetrahedralization for the numerical solution of partial differential equations. We start with a detailed explanation of the method for the Poisson equation and then extend the study to the second-order elliptic PDE, Keller-Segel equations, Stokes, and Navier-Stokes equations. We shall show that the numerical solution can approximate the exact PDE solution very well. Then we present a large amount of numerical experimental results to demonstrate the performance of the method over the 2D and 3D settings. In addition, we present a comparison with the existing multivariate spline methods in [G. Awanou, M. -J. Lai, and P. Wenston, The multivariate spline method for scattered data fitting and numerical solution of partial differential equations. In *Wavelets and splines: Athens 2005*, pages 24–74. Nashboro Press, Brentwood, TN, 2006] and [M. -J. Lai and Wang, C. M., A bivariate spline method for 2nd order elliptic equations in non-divergence form, *Journal of Scientific Computing*, (2018) pp. 803–829] to show that the new method produces a similar and sometimes more accurate approximation in a more efficient fashion.

Jinsil Lee
University of Georgia
jl74942@uga.edu

MS159

High-Order Numerical Integration for Closed Surfaces

We propose a novel high-order integration method for surface integrals. The approach rests on curved triangulations, approximating a vast class of regular (smooth) surfaces. The novelty of our contribution is given by accurate interpolation of the closest point projection $\pi : \mathcal{M}_h \rightarrow \mathcal{M} \subset \mathbb{R}^3$, realising the curved triangles. To do so, we incorporate recent advances in multivariate interpolation, suppressing Runge’s phenomenon when interpolating π in proper chosen, *transformed Chebyshev-Lobatto nodes*. The specific transformation we propose here, yield the integration scheme to be largely independent of the mesh quality and maintains stability for high orders. Further, we combine our technique with the novel global polynomial level set parametrisation method (GPLS). GPLS only requires regular samples $P \subseteq \mathcal{M}$ of the surface \mathcal{M} in order to derive implicit global models of the surface $\mathcal{M} = Q_{\mathcal{M}}^{-1}(0)$. This allows us to apply our high-order integration scheme to a vast class of non-parameterised surfaces by deriving the closest point map from the global polynomial level set $Q_{\mathcal{M}}$. As our empirical demonstrations suggest, we expect the approach to be applicable for a broad class of compu-

tational tasks in numerical differential geometry arising for “real world” problems across disciplines, e.g., bio-physics, material sciences and computer graphics.

Gentian Zavalani
Center for Advanced Systems Understanding (CASUS)
Technische Universität Dresden, D-01062 Dresden,
Germany
g.zavalani@hzdr.de

Michael Hecht
Center for Advanced Systems Understanding (CASUS)
m.hecht@hzdr.de

MS160

Progressive Fusion for Multimodal Integration

Integration of multimodal information from various sources has been shown to boost the performance of machine learning models and thus has received increased attention in recent years. Often such models use deep modality-specific networks to obtain unimodal features which are combined to obtain ‘late-fusion’ representations. However, these designs run the risk of information loss in the respective unimodal pipelines. On the other hand, ‘early-fusion’ methodologies, which combine features early, suffer from the problems associated with feature heterogeneity and high sample complexity. We present an iterative representation refinement approach, called Progressive Fusion, which mitigates the issues with late fusion representations. Our model-agnostic technique introduces backward connections that make late stage fused representations available to early layers, improving the expressiveness of the representations at those stages, while retaining the advantages of late fusion designs. We test Progressive Fusion on tasks including affective sentiment detection, multimedia analysis, and time series fusion with different models, demonstrating its versatility. We show that our approach consistently improves performance, for instance attaining a 5% reduction in MSE and 40% improvement in robustness on multimodal time series prediction.

Madalina Fiterau
University of Massachusetts, Amherst
mfiterau@cs.umass.edu

MS160

Residual Multi-Fidelity Neural Networks

We present a residual multi-fidelity computational framework that effectively utilizes the approximation power of deep neural networks. Given a low-fidelity and a high-fidelity computational model, we formulate the correlation between the two models in terms of a residual function. Precisely, instead of searching for a direct correlation between the two models, we consider a possibly nonlinear relation between the low-fidelity model and the residual of the two models. The smaller magnitude of the residual function, compared to the size of the high-fidelity quantity, enables the construction of a residual network that learns the residual function using a small set of high-fidelity data. The trained residual network is then used to efficiently generate additional high-fidelity data. Finally, the set of all available and newly generated high-fidelity data are used to train a deep network that learns the high-fidelity quantity of interest. We present several numerical examples to demonstrate the power of the proposed framework. We show that dramatic savings in computational cost may be achieved when the output predictions are desired to be ac-

curate within small tolerances.

Mohammad Motamed, Owen Davis
University of New Mexico
motamed@unm.edu, daviso@unm.edu

Raul Tempone
RWTH Aachen University & KAUST
tempone@uq.rwth-aachen.de

MS160

Neural-Network Interpretability for Time Series Classification Task

Neural networks (NN) have been gaining significant traction for time series classification tasks over the past few years. Yet, they are frequently perceived as black-box tools, whose results may be difficult to interpret. To address this issue, several methods have been proposed to obtain maps of relevance scores highlighting the importance of different time steps for a given model. These methods were initially applied to images, and more recently to time-series data. Yet, interpretability of NN remains challenging. Trust in these interpretability methods will only be brought through a formal evaluation of their performance, which has been studied but with some drawbacks. These methods typically provide different evaluations of importance, sometimes even diametrically opposite results, and do not explain how neurons collaborate to represent specific patterns. Some studies, therefore, aim to introduce novel metrics to rank interpretability methods across the most commonly used machine learning (ML) models. Another aspect of the research focuses on developing novel interpretability methods to provide more meaningful insights for practitioners working with ML models. In this work, we extend interpretability methods to recognize patterns indicative of a given class across samples. We argue that this work is a critical step toward understanding NN-based decisions for a given classification task at a more holistic dataset scale, rather than the individual sample.

Jiawen Wei
National University of Singapore (NUS)
jiawenw@u.nus.edu

Hugues Turbé, Mina Bjelogrić
University of Geneva, Switzerland
hugues.turbe@unige.ch, mina.bjelogric@hcuge.ch

Adriano Gualandi
INGV-Istituto Nazionale di Geofisica e Vulcanologia, Italy
adriano.gualandi@ingv.it

Christian Lovis
University of Geneva, Switzerland
christian.lovis@hcuge.ch

Gianmarco Mengaldo
National University of Singapore
College of Design and Engineering
mpegim@nus.edu.sg

MS160

Measuring Uncertainty in Data Fusion Algorithms

In recent years, the fields of optimization, uncertainty quantification and machine learning have all been gravitating towards developing algorithms and models that can

enable the capacity to incorporate, or fuse, data from disparate sources or fidelities. In this talk, we give an overview of common themes in approaches for data fusion across applications and domains, setting the stage for subsequent discussions and talks in this minisymposia. In addition, we discuss current challenges in information fusion and discuss our recent work developing methods of uncertainty quantification for data fusion algorithms.

Tian Yu Yen
Sandia National Laboratories
tyen@sandia.gov

MS161

CAS4DL: Christoffel Adaptive Sampling for Deep Learning in Scientific Computing Applications

Many problems in computational science and engineering require the approximation of a high-dimensional function from data. In many such applications, data is costly to generate: for example, it each sample may require a costly PDE solve. Therefore, it is imperative to develop highly sample efficient algorithms. Recently, deep neural networks and deep learning have shown great promise to provide breakthrough performance in challenging function approximation tasks. In this work, we propose an adaptive sampling strategy, CAS4DL (Christoffel Adaptive Sampling for Deep Learning) to increase the sample efficiency of DL. Our novel approach is based on interpreting the second to last layer of a DNN as a dictionary of functions defined by the nodes on that layer. With this viewpoint, we then define an adaptive sampling strategy motivated by adaptive sampling schemes recently proposed for linear approximation schemes, wherein samples are drawn randomly with respect to the Christoffel function of the subspace spanned by this dictionary. We present numerical experiments comparing CAS4DL with standard Monte Carlo (MC) sampling. Our results demonstrate that CAS4DL often yields substantial savings in the number of samples required to achieve a given accuracy, particularly in the case of smooth activation functions. These results, therefore, are a promising step toward fully adapting DL to scientific computing applications.

Juan M. Cardenas
Simon Fraser University
jcardena@sfu.ca

Nick Dexter
Florida State University, U.S.
nick.dexter@fsu.edu

Ben Adcock
Simon Fraser University
ben_adcock@sfu.ca

MS161

A Neural Network Approach for Homogenization of Multiscale Problems

We propose a neural network-based approach to the homogenization of multiscale problems. The proposed method uses a derivative-free formulation of a training loss, which incorporates Brownian walkers to find the macroscopic description of a multiscale PDE solution. Compared with other network-based approaches for multiscale problems, the proposed method is free from the design of hand-crafted neural network architecture and the cell problem to calculate the homogenization coefficient. The exploration

neighborhood of the Brownian walkers affects the overall learning trajectory. We determine the bounds of micro- and macro-time steps that capture the local heterogeneous and global homogeneous solution behaviors, respectively, through a neural network. The bounds imply that the computational cost of the proposed method is independent of the microscale periodic structure for the standard periodic problems. We validate the efficiency and robustness of the proposed method through a suite of linear and non-linear multiscale problems with periodic and random field coefficients.

Jihun Han
Dartmouth College
jihun.han@dartmouth.edu

Yoonsang Lee
Dartmouth College
Department of Mathematics
yoonsang.lee@dartmouth.edu

MS161

Self-Supervised Deep Image Restoration via Adaptive Stochastic Gradient Langevin Dynamics

While supervised deep learning has been a prominent tool for solving many image restoration problems, there is an increasing interest on studying self-supervised or unsupervised methods to address the challenges and costs of collecting truth images. Based on the neuralization of a Bayesian estimator of the problem, this paper presents a self-supervised deep learning approach to general image restoration problems. The key ingredient of the neuralized estimator is an adaptive stochastic gradient Langevin dynamics algorithm for efficiently sampling the posterior distribution of network weights. The proposed method is applied on two image restoration problems: compressed sensing and phase retrieval. The experiments on these applications showed that the proposed method not only outperformed existing non-learning and unsupervised solutions in terms of image restoration quality, but also is more computationally efficient.

Hui Ji, Weixi Wang, Ji Li
National University of Singapore
matjh@nus.edu.sg, wangweixi@u.bus.edu,
matliji@nus.edu.sg

MS161

Operator Network Approximations for Elliptic PDEs

The application of neural networks (NNs) to the numerical solution of PDEs has seen growing popularity in the last five years: NNs have been used as an ansatz space for the solutions, with different training approaches (PINNs, deep Ritz methods, etc.); they have also been used to infer discretization parameters and strategies. In this talk, I will focus on the convergence of operator networks that approximate the solution operator of linear elliptic PDEs. I will, in particular, consider operator networks that, given a fixed right-hand side, map sets of diffusion-reaction coefficients into the space of solutions (coefficient-to-solution map). When the coefficients are smooth and with periodic boundary conditions, the size of the networks can be bounded with respect to the H^1 norm of the error, uniformly over the parameter set. Specifically, the number of non zero weights grows at most poly-logarithmically with respect to the error. The proofs of our approximation rates

combine elliptic regularity, classical and recent results in numerical analysis, and tools from NN approximation theory. Using the same techniques, we extend the analysis to linear elasticity and parametric problems.

Carlo Marcati
University of Pavia, Italy
carlo.marcati@unipv.it

Christoph Schwab
ETH-Zurich
schwab@sam.math.ethz.ch

MS162

Low-Rank Matrix and Tensor Factorization for Community Detection in Dynamic Networks

Many real-world systems and relational data can be modeled as networks or graphs, where the different entities of the system and the interactions between them are represented by the vertices and edges of a graph, respectively. With the availability of large amounts of network data, it is important to be able to reduce the networks dimensionality and extract useful information from it. A key approach to network data reduction is community detection. Early work in graph-based community detection methods has focused on static networks. This type of networks is usually considered as an oversimplification as many real-world complex systems exhibit variation in their community structure over time. Consequently, there is a growing need to develop algorithms that detect the community structure and track its evolution across time. In this talk, we will present low-rank matrix and tensor decomposition methods for community detection in dynamic networks. The first part of the talk will focus on a low-rank + sparse model for evolutionary spectral clustering to detect and track the community structure in temporal weighted and binary networks. We will then extend this framework by considering a tensor formulation of the time-varying networks and introduce a low-rank tensor model to identify and track the community structure. The performance of the proposed methods will be illustrated on various benchmark network models as well as real networks from a variety of applications.

Selin Aiyente
Michigan State University
aiyente@egr.msu.edu

Esraa Al-Sharoua
Jordan University of Science and Technology
alsharoua@msu.edu

MS162

Relational Learning on Temporal Knowledge Graphs

There has recently been increasing interest in learning representations of temporal knowledge graphs (tKGs), which record the dynamic relationships between entities over time. As tKGs dynamically evolves, the low-rank approximation embedding approaches should be empowered to capture temporal dynamics on tKGs. In this talk, we would like to introduce two approaches specified for tKG with two different perspectives. Firstly, the existing approaches model the tKGs in discrete state spaces while links on tKGs often vary continuously over time. To this end, we propose a novel continuum model by extending the idea of neural ordinary differential equations (ODEs) to modeling

tKGs. The proposed model preserves the continuous nature of dynamic multi-relational graph data and encodes both temporal and structural information into continuous-time dynamic embeddings. Secondly, tKGs often exhibit multiple simultaneous non-Euclidean structures. However, existing embedding approaches for temporal KGs typically learn entity representations and their dynamic evolution in Euclidean space, which might not capture such intrinsic structures very well. Thus, we propose DyERNIE, a non-Euclidean embedding approach that learns evolving entity representations in a product of Riemannian manifolds. We conduct extensive experiments on benchmark tKG datasets, demonstrating the proposed methods' superior performance.

Zhen Han

Amazon UK
hanzhen02111@163.com

Yunpu Ma, Zifeng Ding

LMU Munich
ma@dbs.ifi.lmu.de, zifeng.ding@campus.lmu.de

Volker Tresp

Ludwig Maximilian University of Munich
volker.tresp@siemens.com

MS162

Online Nonnegative CP-Dictionary Learning for Markovian Data

Online Tensor Factorization (OTF) is a fundamental tool in learning low-dimensional interpretable features from streaming multi-modal data. While various algorithmic and theoretical aspects of OTF have been investigated recently, a general convergence guarantee to stationary points of the objective function without any incoherence or sparsity assumptions is still lacking even for the i.i.d. case. In this work, we introduce a novel algorithm that learns a CANDECOMP/PARAFAC (CP) basis from a given stream of tensor-valued data under general constraints. We prove that our algorithm converges almost surely to the set of stationary points of the objective function under the hypothesis that the sequence of data tensors is generated by an underlying Markov chain. Our setting covers the classical i.i.d. case as well as a wide range of application contexts including data streams generated by independent or MCMC sampling. Our result closes a gap between OTF and Online Matrix Factorization in global convergence analysis for CP-decompositions. Experimentally, we show that our algorithm converges much faster than standard algorithms for nonnegative tensor factorization tasks on both synthetic and real-world data. Also, we demonstrate the utility of our algorithm on a diverse set of examples from image, video, and time-series data, illustrating how one may learn qualitatively different CP-dictionaries from the same tensor data by exploiting the tensor structure in multiple ways.

Hanbaek Lyu

University of Wisconsin Madison
hlyu@math.wisc.edu

MS162

Fusing Dynamic and Static Data Using Parafac2-Based Coupled Matrix and Tensor Factorizations

There is an emerging need to jointly analyze time-evolving data sets together with static data in many areas such as

social networks, omics data analysis, and temporal phenotyping. A promising approach in that direction is to jointly factorize data sets from multiple sources using coupled matrix and tensor factorizations. The PARAFAC2 model for tensor decompositions has been shown to be a promising alternative to the CANDECOMP/PARAFAC (CP) model for temporal data due to its capability of capturing time-evolving patterns. While there have been studies using a PARAFAC2 model coupled with a matrix decomposition for joint analysis of dynamic and static data, they rely on implicit estimation of the evolving mode or a *flexible* PARAFAC2 constraint, and are therefore limited in terms of possible regularizations on the evolving mode. In this talk, we present a flexible algorithmic framework for joint factorization of dynamic and static data based on the PARAFAC2 model. It allows the PARAFAC2 model to be coupled to a matrix or a CP decomposition in any static mode. The algorithmic framework utilizes Alternating Optimization with the Alternating Direction Method of Multipliers and facilitates the use of linear couplings and many different constraints on all modes, including the evolving mode in the PARAFAC2 model. In our framework, we allow for any proximable constraint, thus also enabling the use of time-aware constraints on the evolving mode.

Carla Schenker

Simula Metropolitan Center for Digital Engineering
Oslo Metropolitan University
carla@simula.no

Xiulin Wang

Affiliated Zhongshan Hospital of Dalian University
xiulin.wang@foxmail.com

Evrin Acar

Simula Research Laboratory
evrim@simula.no

MS163

Parametrized Model Order Reduction to a Heat Transfer Simulation of a Selective Laser Melting Process

Selective laser melting (SLM) is an additive manufacturing process in which a laser of high intensity sequentially melts layers of metal powder altogether to build an object with a specific geometry. The thermo-viscoplastic reaction occurring during the building process is dynamically intense and open to several issues. Many parameters come into play to control the quality of the designed material. Each one has its specific sensitivity to the final mechanical properties of the material and therefore needs to be adjusted. Using numerical simulation to analyze these models helps get a better understanding of the sensitivity of each parameter and its effect on the quality of the final material. However, the computational cost of these simulations is demanding in terms of time and resources. This work investigates the application of Model Order Reduction (MOR) to the thermal process of SLM. A hyper-reduction strategy, relying on the Energy Conserving Sampling and Weighting method, is presented before being implemented within an adaptive POD-Greedy algorithm to build a more general reduced order base composed of versatile thermal modes. To this end, an a posteriori error estimator is used alongside Ordinary Kriging to evaluate the degree of error committed during the parametrized-MOR process. The authors acknowledge SIM and VLAIO (Flanders, Belgium) for funding the "PROCSIMA" project (M3 research program).

Mohamed Amine Ben Yahmed

KU Leuven
mohamedamine.benyahmed@kuleuven.be

Frank Naets
Department of mechanical engineering, KU Leuven
frank.naets@kuleuven.be

MS163

Applications of a Space-Time First-Order System Least-Squares Formulation of Parabolic PDEs

While the common space-time variational formulation of a parabolic equation results in a bilinear form that is non-coercive, [Führer–Karkulik, Spacetime least-squares finite elements for parabolic equations, *Computers & Mathematics with Applications*, 2019] recently proved well-posedness of a space-time first-order system least-squares formulation of the heat equation. Least-squares formulations always correspond to a symmetric and coercive bilinear form. In particular, the Galerkin approximation from any conforming trial space exists and is a quasi-best approximation. Additionally, the least-squares functional automatically provides a reliable and efficient error estimator. In [Gantner–Stevenson, Further results on a space-time FOSLS formulation of parabolic PDEs, *ESAIM. Mathematical Modelling and Numerical Analysis*, 2021], we have generalized the least-squares method of Führer–Karkulik to general second-order parabolic PDEs with possibly inhomogeneous Dirichlet or Neumann boundary conditions. We employ the space-time least-squares method for parameter-dependent problems as well as optimal control problems [Gantner–Stevenson, Applications of a space-time FOSLS formulation for parabolic PDEs, *arXiv*, 2022]. In both cases, coercivity of the corresponding bilinear form plays a crucial role.

Gregor Gantner, Rob Stevenson
University of Amsterdam
gregor.gantner@inria.fr, rob.p.stevenson@gmail.com

MS163

A Certified Reduced Basis Method for Quasilinear Parabolic Equations

We propose a certified reduced basis (RB) method for quasilinear parabolic problems with strongly monotone spatial differential operator. We provide a residual based a posteriori error estimate for a space-time formulation and the corresponding efficiently computable bound for the certification of the method, where we use EIM to approximate the nonlinearity. We numerically validate our method at a parametrized magnetoquasistatic approximation of Maxwells equations (the eddy current model).

Michael Hinze
Universität Koblenz, Landau
hinze@uni-koblenz.de

Denis Korolev
Weierstraß Institut Berlin
korolev@wias-berlin.de

MS163

Residual Data-Driven Variational Multiscale Reduced Order Models for Parameter Dependent Problems

ROMs for fluid flows have been used to successfully min-

imize the computational cost of scientific and engineering applications dominated by a small number of recurring dominant spatial structures. For under-resolved simulations, using a relatively small number of ROM basis functions often yields an inaccurate approximation. Since numerical efficiency is one of the big advantages of the ROM, we want to increase the numerical accuracy while preserving the computational efficiency. Thus, we add a low-dimensional closure term $Closure(\mathbf{a})$ to the standard Galerkin ROM and solve the following closed ROM,

$$\dot{\mathbf{a}} = \mathbf{F}(\mathbf{a}) + Closure(\mathbf{a}), \quad (1)$$

where the $Closure(\mathbf{a})$ term models the interaction between the unresolved ROM basis functions $\{\varphi_{r+1}, \dots, \varphi_d\}$ and the resolved ROM basis functions $\{\varphi_1, \dots, \varphi_r\}$. In this talk, we propose a consistent data-driven (D2) variational multiscale (VMS) reduced order model (ROM) framework to increase the ROM accuracy at a modest computational cost for under-resolved regimes. To construct the new consistent D2-VMS-ROM, we need to model the closure term by using the residual term as:

$$Closure(\mathbf{a}) \approx \tilde{\mathbf{A}} Res(\mathbf{a}). \quad (2)$$

Birgul Koc
University of Seville
birgul.koc@inria.fr

Samuele Rubino
Universidad de Sevilla
samuele@us.es

Tomás Chacón Rebollo
Universidad de Sevilla
Spain
chacon@us.es

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MS163

Model Order Reduction for Parametric Optimal Control Problems in Space-Time Formulation

In this talk, we deal with optimal control problems governed by parabolic partial differential equations employing the certified reduced basis method. Specifically, we rely on a new easy-to-compute error estimator that guarantees a rigorous and efficient bound to the problem variables. We first introduce the space-time formulation for optimal control. To solve the space-time optimality system faster, we use a Greedy method. In the talk, we propose some insights on the derivation of such an error estimator. We test the performances of the proposed strategy on a physical parametrized distributed optimal control problem and a boundary optimal control problem with physical and geometrical parameters. The obtained results are compared to previously proposed bounds based expensive on the computation (or approximation) of the optimality system inf-sup constant. The comparison regards reliability and computational costs [M. Strazzullo, F. Ballarin, G. Rozza, ‘A Certified Reduced Basis Method for Linear Parametrized Parabolic Optimal Control Problems in Space-Time Formulation’, submitted, 2021]. Moreover, the findings hold

for steady problems, too: we will show some novel applications to varying boundary control, i.e. problems where the parameter changes the portion of the boundary where the control acts highlighting the pros and cons of the strategy in such a scenario.

Maria Strazzullo
mathLab, Mathematics Area, SISSA International School
for Advanced Studies
mstrazzu@sissa.it

Francesco Ballarin
Catholic University of the Sacred Heart, Italy
francesco.ballarin@unicatt.it

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS164

Diagrammatic Methods for Open Quantum Systems with Bold Lines

We consider the diagrammatic Monte Carlo methods for quantum systems coupled with a harmonic bath. The variance of the Monte Carlo methods grows exponentially with time, which is known as the dynamical sign problem. To mitigate the sign problem, we introduce the bold diagrams, which is a partial sum of the path integral that can be reused in future calculations. We will discuss various approaches to accelerating the algorithm, including different ways to use bold diagrams, fast implementation of the bath influence functional, and the sampling strategy to reduce calculations. The spin-boson model will be tested to validate our algorithms.

Zhenning Cai
National University of Singapore
matcz@nus.edu.sg

Jianfeng Lu
Mathematics Department
Duke University
jianfeng@math.duke.edu

Siyao Yang, Geshuo Wang
National University of Singapore
matsiya@nus.edu.sg, geshuowang@u.nus.edu

MS164

Quantum Perturbation Theory Using Tensor Cores and a Deep Neural Network

Time-independent quantum response calculations are performed using Tensor cores. This is achieved by mapping density matrix perturbation theory onto the computational structure of a deep neural network. The main computational cost of each deep layer is dominated by tensor contractions, i.e. dense matrix-matrix multiplications, in mixed precision arithmetics which achieves close to peak performance. Quantum linear response calculations using self-consistent charge density-functional tight-binding theory are presented where we demonstrate a peak performance of almost 200 Tflops using the Tensor cores of two Nvidia A100 GPUs. Additionally, a novel parameter-free convergence criterion is presented that is well-suited for

numerically noisy low precision floating point operations.

Joshua Finkelstein
Los Alamos National Laboratory
jdf@lanl.gov

Emanuel H. Rubensson
Uppsala University
emanuel.rubensson@it.uu.se

Susan Mniszewski, Christian Negre, Anders Niklasson
Los Alamos National Laboratory
smm@lanl.gov, cnegre@lanl.gov, amn@lanl.gov

MS164

Hierarchical Low-Rank Matrix Methods for the Numerical Solution of the Nonequilibrium Dyson Equation

We propose a method to improve the computational and memory efficiency of numerical solvers for the nonequilibrium Dyson equation (NDE), which appears in simulations of the effect of strong radiation fields on atoms and molecules, quantum materials, nuclear physics, and many other many-body quantum systems. The NDE is a coupled system of nonlinear Volterra integral equations, for which the history integrals lead to $\mathcal{O}(N^3)$ computational complexity and $\mathcal{O}(N^2)$ memory complexity for N time steps in traditional solvers. Our method is based on the empirical observation that for many problems of physical interest, the kernels of the integral operators in the NDE, as well as its solutions, can be represented as hierarchical off-diagonal low rank (HODLR) matrices. We present an algorithm which builds these HODLR representations on the fly during the course of time stepping, and uses them to reduce the cost of computing history integrals. For systems with the hierarchical low rank property, our method achieves $\mathcal{O}(N^2 \log N)$ computational complexity and $\mathcal{O}(N \log N)$ memory complexity. We present numerical examples demonstrating orders of magnitude speedup and memory reduction over previous methods, and reaching unprecedented propagation times.

Jason Kaye
Flatiron Institute
Simons Foundation
jkaye@flatironinstitute.org

Denis Gole
Joef Stefan Institute
Department of Theoretical Physics
denis.golez@ijs.si

MS164

Ternary Unitary Quantum Lattice Models and Circuits in $2 + 1$ Dimensions

We extend the concept of dual unitary quantum gates [Phys. Rev. Lett. 123, 210601 (2019)] to quantum lattice models in $2 + 1$ dimensions, by introducing and studying *ternary unitary* four-particle gates, which are unitary in time and both spatial dimensions. When used as building blocks of lattice models with periodic boundary conditions in time and space (corresponding to infinite temperature states), dynamical correlation functions exhibit a light-ray structure. We also generalize solvable MPS [Phys. Rev. B 101, 094304 (2020)] to two spatial dimensions with cylindrical boundary conditions, by showing that the analogous *solvable PEPS* can be identified with matrix product uni-

aries. In the resulting tensor network for evaluating equal-time correlation functions, the bulk ternary unitary gates cancel out. We delineate and implement a numerical algorithm for computing such correlations by contracting the remaining tensors. (Preprint: arXiv:2206.01499)

Richard Milbradt, Lisa Scheller, Christopher Assmus
TU Munich
Germany
r.milbradt@tum.de, lisa.scheller@tum.de,
chris.assmus@tum.de

Christian B. Mendl
TU Munich
mendl@in.tum.de

MS164

Numerical Methods for Simulating Non-Equilibrium Quantum Many-Body Dynamics

A practical way to compute time-dependent physical observables of a quantum many-body system is to focus on a two-time Green's function, which is a two-point correlator of the creation and annihilation field operators defined on the Keldysh contour. The equation of motion satisfied by the two-time non-equilibrium Green's function is a set of nonlinear integro-differential equations called the Kadanoff-Baym equations. We will describe numerical methods for solving this type of equations and techniques for reducing the computational complexity of this approach.

Chao Yang
Lawrence Berkeley National Lab
cyang@lbl.gov

Jia Yin, Yuanran Zhu
Lawrence Berkeley National Laboratory
jiayin@lbl.gov, yzhu4@lbl.gov

MS165

Fluctuating Hydrodynamics for Multiphase Mixtures

A wide range of mesoscale phenomena involve fluids that exhibit heterogeneous structure at the nanoscale. A key feature of these problems is that the dynamics occurs on nanometer to micrometer scales. At those scales, thermal fluctuations play an important role in the overall dynamics and are expected to have significant impact on overall behavior of the system, particularly for problems such as those discussed above that are far from equilibrium. Here, we introduce a model for multiphase, multicomponent mixtures based on an N-component form of the Flory-Huggins extension to regular solution theory. The thermodynamics of the system is described by a free energy that includes entropy and enthalpy of mixing as well as non-local terms representing interfacial tension. The multiphase model is incorporated into a fluctuating hydrodynamics (FHD) model for nonideal liquid mixtures. The FHD model represents thermal fluctuations by adding stochastic flux terms to both the species transport and momentum equations based on a model originally proposed Landau and Lifshitz. The presentation will discuss the basic formulation of the model and sketch the derivation of the equations of motion. Numerical results will be presented validating the model and illustrating the range of phenomena that it can represent. Finally, we will show how fluctuations shift the stability region Rayleigh-Plateau instability that describes

the breakup of fluid threads.

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Bryn Barker
University of North Carolina at Chapel Hill
brynbarker@unc.edu

Alejandro Garcia
San Jose State University
alejandro.garcia@sjsu.edu

MS165

Numerical Methods for Phase Separation of Surface Tension Dominated Immiscible Fluid Mixtures

Microscale particles—1 to 100 micrometer sized capsules—can enable the cheap and precise analyses of both single cells and individual molecules. Current research uses temperature-induced phase separation of aqueous polymer mixtures to efficiently fabricate microparticles of desired shapes and sizes. To better understand this process, we develop several mathematical models of microscale phase separation of ternary fluid mixtures. While the equilibrium configuration of the microparticles is described by a volume-constrained minimization of a Ginzburg-Landau free energy, a diffuse interface fluid mechanical model reveals a wide landscape of interesting parameter regimes; varying surface tensions, densities, viscosities, and concentrations can all influence the microparticle manufacture. We discuss some key features of a Cahn-Hilliard-Stokes model used to describe these effects, as well as its simulation with a pseudo-spectral method, and we highlight how fluid stresses can influence microparticle evolution.

Sean P. Carney
UCLA
Department of Mathematics
spcarney@math.ucla.edu

Eric W. Hester
UCLA
ehester@math.ucla.edu

Andrea L. Bertozzi
UCLA Department of Mathematics
bertozzi@math.ucla.edu

MS165

Modeling Microparticle Manufacture

Microparticles - 1 to 100 micrometer sized capsules - will enable cheap and precise single-cell analysis. Current research uses temperature-induced phase separation to efficiently fabricate polymer gel microparticles of desired shapes and sizes. To better understand this process, we develop several models of microscale phase separation of ternary fluids. A simple surface energy minimisation model describes possible equilibrium microparticle shapes. We then interrogate the dynamics of microparticle manufacture by combining a ternary Cahn-Hilliard model with surface tension- and buoyancy-driven Stokes flow in the Cahn-Hilliard-Stokes-Boussinesq (CHSB) model. By simulating the CHSB model in three dimensions using the efficient Dedalus spectral code, we show that all physical effects -

surface tension, fluid flow, and buoyancy - are necessary for microparticles to attain minimal energy crescent configurations. Without fluid flow or buoyancy, microparticles consistently settle on spherical shell equilibria, in contrast to experiments. We then outline how varying surface energies, densities, viscosities, and concentrations can each influence microparticle manufacture, before concluding with a discussion of how fluid stresses drive microparticle evolution.

Eric W. Hester
UCLA
ehester@math.ucla.edu

Sean P. Carney
UCLA
Department of Mathematics
spcarney@math.ucla.edu

MS165

Simulation of Micro-Scale Particulate and Droplet Motion in Gases

Low-speed gas flow around micro-scale particles (e.g. soot and other pollutants), and through suspensions of particles, are surprisingly rich in physics and challenging to simulate. The hydrodynamic reach of a single particle is extremely wide, preventing application of straightforward numerical-simulation approaches (e.g. finite-volume CFD). Furthermore, their scale renders the conventional Navier-Stokes equations, and associated boundary conditions, inaccurate, unable to capture rarefied effects such as thermophoresis and velocity slip. The problems are especially challenging in the transitional Knudsen regime for non-canonical geometries. In this paper, a framework is presented for simulating micro-scale particulate flows using the Method of Fundamental Solutions, applied to the linearised Grad's 13 moment equations, which provides a convenient and numerically efficient alternative to solving the full Boltzmann equation. A range of particulate geometry are considered, including agglomerates, such as soot. We present a methodology to couple heat transfer within the solid particle (i.e. for problems of finite particle conductivity) to that of the external rarefied gas flowfield. Results are compared to available experimental results and known analytical results, where they exist. The talk describes research funded in the UK by the EPSRC (EP/N016602/1).

Duncan Lockerby, Josiah Jordan
University of Warwick
duncan.lockerby@warwick.ac.uk,
j.jordan.2@warwick.ac.uk

MS165

A Realizable Scheme for Moment Models of Polydisperse Sprays of Non-Spherical Droplets

New models for polydisperse sprays with coupling capabilities to the separated phases zone have been proposed recently through a two-scale modeling of interfacial two-phase flows. In the present contribution, we design specific numerical methods for their resolution. The key ingredient in these models is a good choice of state variables, which describe both the polydisperse character of a spray as well as the geometrical dynamics of non spherical droplets. The resulting system of equations is hyperbolic and involves a complex structure: the continuous solutions evolves in a convex set called the realizability domain. This

induces constraints on the numerical methods. To achieve accuracy, robustness and realizability, we rely on kinetic finite volume schemes and Discontinuous Galerkin methods. We focus on a two-phase simulation of a polydisperse spray with oscillating droplets and assess the ability of the model and of the related numerical methods to capture the physics of such flows.

Teddy Pichard
École polytechnique
teddy.pichard@polytechnique.edu

Katia Ait-Ameur
Ecole Polytechnique
katia.ait-ameur@polytechnique.edu

Samuel Kokh
CEA - Maison de la Simulation
samuel.kokh@cea.fr

Marc Massot
CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau Cedex, France
marc.massot@polytechnique.edu

MS166

Novel Methods for Dynamic, Time-Dependent Simulation of Optical Propagation Through Atmospheric Turbulence

Applications of optics such as imaging, remote sensing, and adaptive optics can involve highly dynamic changes in propagation direction over extended time intervals. The standard method for modeling the turbulent random media involved in simulations of such scenarios is to construct a sequence of two-dimensional random phase screens generated based on the statistics of the medium. Current methods rely on Fourier series approximations and therefore suffer from numerical artefacts due to artificial periodicity and are limited to static scenarios. We develop new approaches using sinc bases, which are naturally suited to infinite domains, for both computing the diffraction integrals involved in propagation and for generating and extending the random phase screens. The efficient computational methods arising from this approach allow for arbitrarily long time-dependent simulations while capturing the correct evolution of the physical two-point statistics of the medium. We demonstrate the power of these methods in complex scenarios of practical interest to scientific applications.

Max Cubillos
Air Force Research Lab
Kirtland AFB
max.cubillos@gmail.com

Edwin Jimenez
Caltech
jimenez@caltech.edu

Kevin Luna
University of Arizona
kevinluna@arizona.edu

MS166

Multilevel Optimization for Inverse Problems in

Imaging

Image reconstruction suffers from the curse of dimensionality. One way of overcoming this challenge is through a multigrid type of method. In this work, we apply multigrid-based optimization framework (MG/OPT) to enhance the reconstruction performance. MG/OPT is designed to accelerate a traditional optimization algorithm applied to a high-fidelity problem by exploiting a hierarchy of coarser models. In the context of inverse problem, we design its hierarchical structure recursively in resolution of the image, as well as the measurement data. We provide several numerical results on the performance of MG/OPT in terms of reconstruction quality, as well as significant speedup and improvement of accuracy.

Zichao Wendy Di
Argonne National Laboratory
wendydi@anl.gov

MS166**Multilevel Proximal Methods for Image Restoration**

We propose a family of multilevel proximal gradient methods going from the simplest multilevel iterative soft-thresholding to inexact and inertial forward-backward extensions. Several algorithmic solutions are proposed to handle non-smooth functions. The convergence of the sequences generated in these different algorithmic variants is proved. We present promising numerical performances in the context of large image restoration highlighting the benefit of multilevel approaches for such a domain of application.

Guillaume Lauga
ENS Lyon
guillaume.lauga@ens-lyon.fr

MS167**Stream Triggered Communications for Mesh Applications on GPU Systems**

With the advancement of GPGPUs as powerful and reliable scientific computing tools, communication performance has become increasingly important. Due to both the relative increase in the amount of time being spent in communications as well as the often synchronous nature of these calls, improving the performance of communication kernels a high-value target for HPC simulation performance improvement. One new technology that may be able to improve performance of HPC communication kernels is stream triggered communication. This talk will give an overview of our investigation into stream triggered communication options and their application in AMReX. We will present the results from a toy halo exchange code used to identify the optimal algorithmic motifs and compare current stream triggered options, including MPI-X and MPI-ACX. We then present the results of adding stream-triggered communications to AMReXs FillBoundary halo exchange operation, using the optimal algorithm. Finally, higher level application improvements that are enabled by stream-triggered communications will be discussed, such as redesigning stream interactions between the CPU and GPU to eliminate synchronizations and maximize GPU uptime.

Kevin N. Gott
LBNL
kngott@lbl.gov

Yijian Hu
Georgia Institute of Technology
yhu438@gatech.edu

Brandon Cook
Lawrence Berkeley National Laboratory
bgcook@lbl.gov

MS167**Optimizing Wavefield-Modelling for Seismic Inversion Applications on HPC Architectures**

The seismic imaging process aims at creating an image of the subsurface layer interfaces by acoustic reflection measurements along the surface of the Earth. In many of such algorithms so-called one-way depth propagation of acoustic wavefields is the core mechanism, in order to reconstruct the acoustic waves as they would have occurred inside the Earth. Such process allows finding out where the reflections have occurred and with which strength. The computational efficiency of these applications resolves to the efficiency of the underlying wavefield propagator. An algorithm that has been proposed in the literature is the Phase-Shift Plus Interpolation, which handles wave propagation through inhomogeneous media by phase-shift propagation with a set of reference velocities - each assuming a local homogeneous medium - and a spatial interpolation mechanism. One objective of our work is to deliver a performance-portable wavefield propagation solver for various computer architectures. Our approach to this is to provide the algorithm through an architecture-agnostic interface, developed based on C++ templates, which at compile-time resolves to an architecture-specific back-end. We develop the OpenMP back-end for threaded parallelism on CPUs, CUDA for acceleration on Nvidia GPUs, and HIP for porting CUDA to AMD GPUs. To evaluate the efficiency of the implementations we use the roofline model. In this presentation we will describe our implementation and show some benchmark results.

Andreas Hadjigeorgiou
The Cyprus Institute
Cyprus
a.hadjigeorgiou@cyi.ac.cy

MS167**Performance Engineering in CSE: A Bird's Eye View**

Performance Engineering (PE) in Computational Science and Engineering (CSE) has two goals: First, to understand the performance bottlenecks of a code, and second, to use this insight to apply changes that will improve performance. There is a wide spectrum of workflows and tools in this field, but all have in common that the whole PE process is iterative, yielding improvements until some preset target is achieved or some other criterion is met (such as "I've done enough"). This talk gives an overview of methods, tools, and metrics that are in use for Performance Engineering in CSE. Some, but not all PE activities are well defined in the sense of a prescribed workflow, and we try to accommodate that by building mostly on examples from performance research, performance tools development, and performance-aware application development. Whichever strategy is applied, *performance patterns* can help to build a better understanding of the performance issues at hand. We will describe typical performance patterns and how they can serve to accelerate the PE process.

Georg Hager
Erlangen National High Performance Computing Center
georg.hager@fau.de

MS167

Performance of Low-Rank Tensor Algorithms

We discuss low-rank tensor algorithms and in particular algorithms for the tensor-train (TT) format (known as MPS in computational physics). We focus on the required building blocks and model their node-level performance on modern multi-core CPUs. More specifically, we consider the lossy compression of large dense data (TT-SVD), as well as linear solvers in TT format (TT-MALS, TT-GMRES). For the data compression, we derive the optimal roofline runtime for the complete algorithm based on the two main building blocks in an optimized implementation: Q-less TSQR and tall-skinny matrix-matrix multiplication. For the low-rank linear solvers, we categorize the different kinds of building blocks according to performance characteristics and show possible performance optimizations. While all required tensor operations can be mapped onto standard BLAS/LAPACK routines theoretically, faster implementations need specific performance optimizations: These include (1) avoiding costly singular-value decompositions (SVDs), and (2) employing special fused operations for sequences of memory-bound tensor-contractions and reshaping operations, as well as (3) tracking properties of tensors such as orthogonalities. We show the effect of the different optimizations and compare the runtime of our implementation with other tensor libraries.

Melven Roehrig-Zoellner
German Aerospace Center (DLR)
Institute for Software Technology
Melven.Roehrig-Zoellner@DLR.de

Achim Basermann
German Aerospace Center (DLR)
Simulation and Software Technology
Achim.Basermann@dlr.de

Jonas Thies
German Aerospace Center (DLR)
j.thies@tudelft.nl

MS167

Porting An Ocean Modeling Application to Fujitsu A64FX

In this presentation we report our preliminary efforts in porting ROMS (Regional Ocean Modeling System) to Fujitsu A64FX. ROMS is a free-surface, terrain-following, primitive equations ocean model widely used by the scientific community for a diverse range of applications. The arm-based Fujitsu A64FX processor developed by Fujitsu and RIKEN is used in Fugaku, which until June 2022 has been the fastest machine worldwide for two years. Its main features of SVE, HBM and being power efficient makes it unique in the world of HPC. ROMS is one of few Ookami user codes that demonstrate performance parity of A64FX compared to mainstream architectures. This implies a huge power advantage as the power usage on A64FX is around 2-3 times less than on Intel.

Eva Siegmann
Institute for Advanced Computational Science

Stony Brook, USA
eva.siegmann@stonybrook.edu

Katherine Gallagher
Institute for Advanced Computational Science
Stony Brook University
katherine.l.hudson@stonybrook.edu

David Carlson
Stony Brook University
david.carlson@stonybrook.edu

Robert J. Harrison
Brookhaven National Laboratory
Stony Brook University
robert.harrison@stonybrook.edu

MS168

Learning Green's Functions with Randomized Numerical Linear Algebra

Can one learn a differential operator from pairs of solutions and righthand sides? If so, how many pairs are required? These two questions have received significant research attention in differential equation learning. Given input-output pairs from an unknown partial differential equation, we will derive a theoretically rigorous scheme for learning the associated Green's function G . By exploiting the hierarchical low-rank structure of Greens functions and extending the randomized SVD algorithm to Hilbert-Schmidt operators, we will identify a learning rate associated with elliptic and parabolic partial differential operators and bound the number of input-output training pairs required to recover a Greens function approximately.

Nicolas Boule
Mathematical Institute
University of Oxford
nb690@cam.ac.uk

Alex J. Townsend
Cornell University
townsend@cornell.edu

MS168

An Overview of Randomized Linear Algebra in Scientific Computing

This talk gives an overview of the potential offered by randomization to speed up, improve the robustness, and increase the flexibility of numerical algorithms for solving large-scale linear algebra tasks in scientific computing. For more than a decade, randomization has proven its effectiveness in performing low-rank approximation and related tasks. Recently, the scope of randomization has been extended to essentially every part of numerical linear algebra, including linear systems, eigenvalue problems, model reduction, and matrix functions.

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

MS168

Randomized Compression Algorithms for Rank

Structured Matrices

The talk describes a set of recently developed randomized algorithms for computing a data sparse representation of a rank structured matrices (such as an H-matrix, or an HSS matrix). The algorithms are black box in the sense that they interact with the matrix to be compressed only through its action on vectors, making them ideal for tasks such as forming Schur complements or matrix matrix multiplication. In situations where the operator to be compressed (and its transpose) can be applied in $O(N)$ operations, the compression as a whole does in many environments have linear complexity as well.

Gunnar Martinsson

University of Texas at Austin, U.S.
pgm@oden.utexas.edu

MS168

Convergence of a Jacobi Iteration with Repeated Random Sparsification

The traditional methods of numerical linear algebra are prohibitively expensive for high-dimensional problems for which even a single matrix multiplication by a dense vector may be too costly. In this talk I will discuss a general framework for reducing the cost of classical iterative schemes like Jacobi iteration by randomly sparsifying the approximate solution at each iteration. I will provide a characterization of the error properties of Jacobi iteration with repeated random sparsification and show results of numerical tests applying the scheme to coupled cluster quantum chemistry calculations.

Robert Webber
California Institute of Technology
rwebber@caltech.edu

James Smith
Flatiron Institute
jsmith@flatironinstitute.org

Timothy Berkelbach
Columbia Univeristy
tcb2112@columbia.edu

Jonathan Weare
Courant Institute
weare@nyu.edu

MS169

Computational Tools for Koopman Spectral Analysis of Nonlinear Dynamical Systems

Dynamic Mode Decomposition (DMD) is a data driven spectral analysis technique for a time series. For a sequence of snapshot vectors $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_{m+1}$ in \mathbb{C}^n , assumed driven by a linear operator \mathbb{A} , $\mathbf{f}_{i+1} = \mathbb{A}\mathbf{f}_i$, the goal is to represent the snapshots in terms of the computed eigenvectors and eigenvalues of \mathbb{A} . We can think of \mathbb{A} as a discretization of the underlying physics that drives the measured \mathbf{f}_i 's. In a pure data driven setting we have no access to \mathbb{A} . Instead, the \mathbf{f}_i 's are the results of measurements, or e.g. \mathbb{A} represents PDE/ODE solver (software toolbox) that generates solution in high resolution, with given initial condition \mathbf{f}_1 . In this talk, we present recent development of numerically robust computational tools (numerical methods with analysis and robust software implementation) for dynamic mode decomposition and data driven Koopman spectral

analysis. This includes new implementation of the DMD and an improved methods for nonlinear system identification in data driven setting.

Zlatko Drmac

University of Zagreb
Department of Mathematics
drmac@math.hr

MS169

Compositionality-Based Computation of Control Lyapunov Functions Using Deep Neural Networks

Common methods for computing control Lyapunov functions are often confined to low dimensions due to an exponential growth of the computational effort in the state dimension. In this talk, we discuss the use of neural networks to avoid this curse of dimensionality. To this end, we build on the known fact that neural networks are capable of efficiently approximating so-called compositional functions and present conditions on the underlying control system that yield the existence or non-existence of compositional control Lyapunov functions. Therefore, we can identify systems where an efficient approximation of a control Lyapunov function via neural networks is possible. Moreover, suitable network architectures and training algorithms for the computation of control Lyapunov functions are presented. We also discuss aspects of the more general case of an optimal control problem and the approximation of a compositional optimal value function. Further, we demonstrate how these approaches perform in practice.

Mario Sperl
Chair of Applied Mathematics
University of Bayreuth
mario.sperl@uni-bayreuth.de

Lars Grüne

Mathematical Institute
Universitaet Bayreuth
lars.gruene@uni-bayreuth.de

MS169

Autoencoders and Pod for Very Low Dimensional Nonlinear Controller Models

The control of nonlinear large-dimensional systems is a challenging task. Linearization based approaches can cope well with the computational bottlenecks but, by design, are limited to regimes where the nonlinearities do not dominate. The idea of considering state dependent coefficients, i.e. factorizing nonlinear terms $f(x(t))$ as $f(x(t)) = A(x(t))x(t)$, where x is the system's state and where A is a matrix valued function of the state, seems capable to address control tasks in a nonlinear fashion while keeping a linear structure that is important for exploiting the many relevant tools from linear theory and numerical linear algebra. Still, the resulting formulation can be too complex for large-scale systems. Therefore, we propose an approximation by an affine-linear parametrization

$$f(x(t)) = A(x(t))x(t) \approx \sum_{i=1}^r \rho_i(x(t))A_i x(t) \quad (3)$$

with an encoding $\rho(x(t)) \in \mathbb{R}^r$ and r being much smaller than the actual state dimension and with constant matrices A_i , $i = 1, \dots, r$. In this talk we present a study how convolutional autoencoders can provide such parametrizations of very low dimensions and compare to more stan-

dard approaches like Proper Orthogonal Decomposition (POD). Furthermore, we show how the low-dimensional parametrization is applied for nonlinear controller design for incompressible Navier-Stokes equations and discuss the interplay of approximation quality and controller performance.

Jan Heiland

Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
heiland@mpi-magdeburg.mpg.de

Yongho Kim

MPI for Dynamics of Complex Technical Systems
ykim@mpi-magdeburg.mpg.de

MS169

A Data-Driven Approach to Stability Using the Koopman Operator

The Koopman operator offers a linear representation of nonlinear dynamical systems. This opens the door to systematic methods for global stability analysis, since global stability properties of nonlinear systems can be directly inferred from the linear stability analysis of the operator. Moreover, the Koopman operator framework is amenable to data-driven analysis, in the sense that a finite-dimensional approximation of the (infinite-dimensional) Koopman operator can be inferred from data (e.g. EDMD method). In this talk, we will leverage this framework to provide a novel method for data-driven stability analysis. In particular, we construct a candidate Lyapunov function for the data-driven approximation of the operator and estimate its validity region from the data. This allows to derive stability guarantees and compute an inner approximation of the attraction region in a purely data-driven fashion. The method will be illustrated with several examples and the use of the probabilistic scenario approach will be discussed in the case of high-dimensional systems.

Alexandre Mauroy

University of Namur
alexandre.mauroy@unamur.be

Aivar Sootla

Huawei Technologies Ltd
aivar.sootla@gmail.com

MS169

Learning of Euler-Lagrange Dynamics from Data

The principle of least action is one of the most fundamental physical principles. It says that among all possible motions connecting two points in a phase space, the system will exhibit those motions which extremise an action functional. Many qualitative features of dynamical systems, such as the presence of conservation laws and energy balance equations, are related to the existence of an action functional, i.e. the presence of variational structure. In this talk I will show how to learn dynamical systems from data while incorporating variational structure and why it is important to do so. Moreover, I will demonstrate that when trajectories of a learned system are computed numerically, then discretisation errors can be much bigger than expected rendering predictions unusable. As a remedy, I will introduce our new machine learning approach, coined *Lagrangian Shadow Integration*. It compensates for discretisation errors and enables highly accurate long-term

predictions of dynamics.

Christian Offen

Paderborn University
christian.offen@uni-paderborn.de

Sina Ober-Blöbaum

University of Paderborn
sinaob@uni-paderborn.de

MS170

Aaa-Eigs: Solving Nonlinear Eigenvalue Problems in Julia

The nonlinear eigenvalue problem (NEP) is a variant of the standard eigenvalue problem $Ax = \lambda x$ where the matrix can depend nonlinearly on the eigenvalue:

$$A(\lambda)x = 0$$

These NEPs arise in many engineering and physics applications such as acoustics and quantum mechanics. In Julia, a software package called NEP-PACK, has been developed by researchers at KTH university Stockholm to deal with and solve such NEPs. It contains an extensive library of state-of-the-art algorithms. One algorithm for solving NEPs that it does not yet provide, is AAA-EIGS. Its main application is to find eigenvalues in some specified subregion of the complex plane for large, sparse NEPs. AAA-EIGS uses AAA for setting up rational approximants of the nonlinearities in the NEP. Afterwards it linearizes the resulting rational eigenvalue problem and solves it using a compact version of the rational Krylov method (CORK). The eigenpairs computed by CORK are related to eigenpairs of the original NEP. When implementing AAA-EIGS in Julia, we numerically compare its performance to that of a similar method in NEP-PACK, named NLEIGS. They mainly differ in the way the rational approximation is constructed. Experiments show that the new AAA-EIGS implementation does significantly better than NLEIGS for most problems, which can be attributed to the use of AAA and a better implementation of CORK. Future work can be done to extend the AAA-EIGS implementation with more features to improve its performance.

Kobe Bruyninckx

KU Leuven
Department Computer Science, Division NUMA
kobe.bruyninckx@kuleuven.be

MS170

RandLAPACK: a Standard Library for Randomized Numerical Linear Algebra

Randomized numerical linear algebra, or RandNLA, concerns the use of randomization as a tool to solve large-scale linear algebra problems. RandNLA algorithms are unparalleled in their capacity for producing approximate solutions to problems such as low-rank approximation. For suitably structured problems, they can even offer substantial speedups over deterministic methods without sacrificing accuracy. This talk describes the philosophy of RandLAPACK – an aspiring standard library for RandNLA. Through careful identification of motifs occurring across RandNLA, it is possible to articulate many algorithms without making strong assumptions on the implementations of their core subroutines. This motivates RandLAPACK's object-oriented design: interfaces are used to specify the semantics of tasks, while *algorithms* manifest as interface-conformant classes. In order for an algorithm to

be executed it is necessary to instantiate it as an object in-memory. By allowing these objects to be stateful, we are able to achieve excellent performance without complicating the semantics of the task that an algorithm needs to carry out.

Riley J. Murray
University of California, Berkeley
rjmurray@berkeley.edu

MS170

An Overview on T-LAPACK

To this day, LAPACK remains a very popular tool for linear algebra. Because of the many wrappers and variants, a user can use the high performance library from almost any programming language. However, the package itself is still written in Fortran. Even Fortran enthusiasts must admit that it is less than ideal for writing a large software package. Instead of being another wrapper, T-LAPACK aims to be a complete rewrite of LAPACK in c++ using templates. C++ has major advantages over Fortran. Firstly, routines no longer need to be replicated to support different argument types. This eliminates a major source of bugs and inconsistencies. Secondly, there are better developer tools. Not only is it easier to configure debuggers, formatters and static code checkers for c++, but compiling Fortran code can also be difficult on some platforms like Apple silicon. Lastly, templates make it easier to support special storage formats like banded or even distributed storage. This talk will discuss some design decisions, show some major advantages of using templates and finally we will show an overview of what has already been implemented in T-LAPACK.

Thijs Steel
KU Leuven
thijs.steel@kuleuven.be

Wesley S. Pereira
Department of Mathematical and Statistical Sciences
University of Colorado Denver
wesley.pereira@ucdenver.edu

Julien Langou
University of Colorado Denver
julien.langou@ucdenver.edu

MS170

The Future of Linear Algebra in the C++ Standard

Linear Algebra support is coming to the ISO C++ standard, delivering a future where developers of scientific code do not need to explicitly deal with specific vendor APIs or Fortran/C++ interoperability just to call a matrix-vector multiply. In this talk I will provide an update on the status of this effort, and directions of further extensions of the proposal. I will also discuss some early experience with implementations of the proposal, including an early prototype preview which comes with NVIDIA's HPC-SDK. Attendees will also learn a bit about C++23 mdspar, as the foundational layer on which the linear algebra effort rests.

Christian Trott
Sandia National Laboratories

crtrott@sandia.gov

MS170

Portable Performance in the SLATE Linear Algebra Library

The SLATE project provides a dense linear algebra library designed to support distributed HPC machines with large numbers of cores and multiple hardware accelerators per node. We discuss the portability of the software design of SLATE, including the use of C++ templates and matrix abstractions, and the implementation of the BLAS++ and LAPACK++ packages for vendor and device abstraction. SLATE implements tile-based algorithms to execute effectively on HPC architectures using batch-BLAS, OpenMP and MPI to achieve performance and scalability. We will present some results to show how SLATE adapts to multiple architectures.

Asim YarKhan
Innovative Computing Laboratory
University of Tennessee
yarkhan@icl.utk.edu

Mark Gates
Computer Science Department
University of Tennessee
mgates3@icl.utk.edu

Dalal Sukkari
KAUST
sukkari@icl.utk.edu

Kadir Akbudak
Innovative Computing Laboratory
University of Tennessee
kadir@icl.utk.edu

Sebastien Cayrols
University of Tennessee, Knoxville
sebastien.cayrols@icl.utk.edu

Daniel R. Bielich
University of Colorado at Denver
drbieli@icl.utk.edu

MS171

Reduced Order Models for Inverse Scattering Problems

We introduce a novel, computationally inexpensive approach for imaging with an active array of sensors, which probe an unknown medium with a pulse and measure the resulting waves. The imaging function is based on the principle of time reversal in non-attenuating media and uses a data driven estimate of the internal wave originating from the vicinity of the imaging point and propagating to the sensors through the unknown medium. We explain how this estimate can be obtained using a reduced order model (ROM) for the wave propagation. We analyze the imaging function, connect it to the time reversal process and describe how its resolution depends on the aperture of the array, the bandwidth of the probing pulse and the medium through which the waves propagate.

Liliana Borcea
University of Michigan
borcea@umich.edu

Josselin Garnier
Ecole Polytechnique
josselin.garnier@polytechnique.edu

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

Jörn Zimmerling
University of Michigan
jtzimmerling@gmail.com

MS171

Waveform Inversion via Reduced Order Modeling

A novel approach to full waveform inversion (FWI), based on a data driven reduced order model (ROM) of the wave equation operator is introduced. The unknown medium is probed with pulses and the time domain pressure waveform data is recorded on an active array of sensors. The ROM, a projection of the wave equation operator is constructed from the data via a nonlinear process and is used for efficient velocity estimation. While the conventional FWI via nonlinear least-squares data fitting is challenging without low frequency information, and prone to getting stuck in local minima (cycle skipping), minimization of ROM misfit is behaved much better, even for a poor initial guess. For low-dimensional parametrizations of the unknown velocity the ROM misfit function is close to convex. The proposed approach consistently outperforms conventional FWI in standard synthetic tests.

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

Liliana Borcea
University of Michigan
borcea@umich.edu

Josselin Garnier
Ecole Polytechnique
josselin.garnier@polytechnique.edu

Jörn Zimmerling
University of Michigan
jtzimmerling@gmail.com

MS171

The Lippmann Schwinger Lanczos Algorithm for Inverse Scattering Problems

We combine data-driven reduced order models with the Lippmann-Schwinger integral equation to produce a direct nonlinear inversion method. The ROM is viewed as a Galerkin projection and is sparse due to Lanczos orthogonalization. Embedding into the continuous problem, a data-driven internal solution is produced. This internal solution is then used in the Lippmann-Schwinger equation, in a direct or iterative framework. The new approach also allows us to process non-square matrix-valued data-transfer functions, i.e., to remove the main limitation of the earlier versions of the ROM based inversion algorithms. We show numerical experiments for spectral domain data for which our inversion is far superior to the Born inversion.

Shari Moskow
Drexel University

Department of Mathematics
moskow@math.drexel.edu

Vladimir L. Druskin
Worcester Polytechnic Institute
vdruskin@gmail.com

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

MS171

Inverse Scattering Problems via Estimates of Internal Waves

In this talk, we consider inverse scattering problems in a reduced-order modeling framework. We show that a data-driven reduced-order model, constructed from observations of scattered waves, can be used to estimate the Greens function inside an unknown medium. Subsequently, we show how this estimation can be used in a Lippmann-Schwinger-based inversion approach.

Jörn Zimmerling, Liliana Borcea
University of Michigan
jtzimmerling@gmail.com, borcea@umich.edu

Josselin Garnier
Ecole Polytechnique
josselin.garnier@polytechnique.edu

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

MS172

Data-Driven Reconstruction of Ocean Wave Surfaces from Sparse Observations

Knowledge about water wave conditions in the near future is essential for a broad range of ocean engineering activities. The initialization of numerical wave prediction models first requires the reconstruction of water wave surfaces from sparse measurement data. However, common reconstruction methods rely on a computationally expensive optimization procedure, that often compromises the real-time capability of the entire wave prediction process. Since the inference of a trained deep learning model can usually be accomplished in a fraction of a second, we propose a novel data-driven method for ocean wave reconstruction. In the present case, two-dimensional radar snapshots are the input of a neural network, while the output is the reconstructed ocean wave surface at the same time instant. As these inputs and outputs are grid-structured data types, we employ an adapted U-Net-type convolutional neural network. The surface similarity parameter is used as the loss function for training, as it has been shown to be particularly useful for regression tasks concerning oscillatory data. The quality of the deep-learning-based reconstruction results is comparable to results achieved by classical ocean wave reconstruction methods. Nevertheless, collecting non-synthetic ground truth wave surfaces for this fully supervised approach is challenging. To address the problem of data availability, a physics-informed approach is promising for upcoming work.

Svenja Ehlers, Merten Stender, Marco Klein, Mathies Wedler

Hamburg University of Technology
 svenja.ehlers@tuhh.de, m.stender@tuhh.de,
 marco.klein@tuhh.de, mathies.wedler@tuhh.de

Norbert P. Hoffmann
 Hamburg University of Technology
 Institute for Structural Dynamics
 norbert.hoffmann@tuhh.de

MS172

Physics-Aware Convolutional Neural Networks for Computational Fluid Dynamics Simulation

CFD simulations are very costly to compute and have to be repeated if the geometry changes even slightly. Recently there have been a number of attempts to speed up this process using neural networks. Among these is the use of Convolutional Neural Networks (CNN) as surrogate models for CFD simulations with varying geometries; see, e.g., [1]. Here, the model is trained on images of high-fidelity simulation results. However, the generation of training data is expensive and this approach usually requires a large data set. Thus, it is of interest to be able to train a CNN in the absence of abundant training data with the help of physical constraints. First results have already been achieved for the heat equation on a fixed geometry and flow problems in parameterizable geometries; see [2, 3]. In this talk, we present a physics-aware approach to train CNNs as surrogate models for CFD simulations in varying geometries. The employed CNN takes an image of the geometry as input and returns images of the associated CFD simulation results, i.e., velocity and pressure, as output. Our CNN architecture is based on the structure of U-Net [?]. Since the model is trained on pixel images, it can be applied to a variety of different geometries. We show results for two-dimensional flows around obstacles of varying size and placement and in non-rectangular geometries, esp. arteries and aneurysms.

Viktor H. Grimm
 Department of Mathematics and Computer Science,
 University of
 viktor.grimm@uni-koeln.de

MS172

Incorporating Boundary Conditions into Physics-Informed Neural Networks for Solving Parametrized Time Dependent Partial Differential Equations

We explore new methods for efficiently solving parametrised time-dependent partial differential equations (PDEs) by modifying the architecture of physics informed neural networks (PINNs) and evolving their parameters in a time-stepping scheme. We use the eigenfunctions of the Laplace operator as an embedding layer, which allows us to encode geometrical information about domains with complicated shapes, e.g. with holes. We further demonstrate how to enforce Dirichlet and Neumann boundary conditions (BCs) exactly for the predicted PDE solution field. Our numerical experiments show that this can reduce training time and increase accuracy of the predicted solution compared to standard PINNs with Fourier Feature embedding. Evolutionary neural network (ENNs) use a time-stepping scheme based on the PDE residual to evolve the network parameters such that at each time step, the parameter state allows to predict the PDE solution. Our proposed modified PINNs eliminate the BC loss term and thus allow to extend this

paradigm to more general geometries. We discuss how to efficiently implement the proposed numerical integration schemes and the challenges in balancing accuracy and computational time. Computational experiments on a parametrised time-dependent heat equation are conducted and results are compared with a classical PDE solver.

Mariella Kast
 EPFL
 mariella.kast@epfl.ch

MS172

Deep Fourier Neural Operator for Solving a Class of Inverse Problems

In the present work, we develop a model, named Deep Fourier Neural Operator, to learn the map from operators to functions. The model is employed in the contexts of inverse problem for partial differential equation with unknown function coefficients. In particular, we aim at learning the underlying map from the Dirichlet-to-Neumann operator, defined for a wide collection of boundary data, to the PDE coefficient. The use of this model is supported by rigorously proven statements demonstrating the uniqueness and well-posedness of such map. Moreover, we also present extensive numerical experiments to support the robustness and accuracy of the model.

Roberto Molinaro
 Seminar for Applied Mathematics
 ETH Zurich
 roberto.molinaro@sam.math.ethz.ch

MS172

Advances in Model Reduction in Parametric CFD Problems with Machine Learning and Data-Driven Approaches

We deal with advances in reduced order modelling with a special focus on parametric CFD and applications, currently based on offline-online computing. Data-driven approaches are combined with online machine learning algorithms to improve computational performances, and to guarantee real-time capabilities in complex settings, incorporating turbulence, geometry parametrisation, as well as multi-physics.

Gianluigi Rozza
 SISSA, International School for Advanced Studies,
 Trieste, Italy
 grozza@sissa.it

MS173

Seeking Viable Paradigms for Hydrogen Peroxide Signalling in the Cytoplasm of Human Cells

Hydrogen peroxide (H₂O₂) regulates cell proliferation, inflammation, chemotaxis, apoptosis and other cellular processes with a key relevance for cancer progression and therapy, but the molecular mechanisms mediating this regulation remain unclear. It is known that the 2-Cys peroxidases, Prx1 and Prx2, not only control cytoplasmic H₂O₂ concentrations but can also mediate H₂O₂ signalling. Recent experiments indicate that H₂O₂ signalling at the cytosol of human cells is spatially localised. In this talk we will present and discuss reaction-diffusion model of the cytosolic human peroxidase (Prdx)-based system that accounts for the differential kinetic parameters. The results confirm that H₂O₂ signalling at the cytosol of

human cells is spatially localised and that site-specific scaffold proteins allow H₂O₂ released to the cytosol at distinct sites to independently regulate distinct genes and processes. Work financed by European Regional Development Fund, through COMPETE2020-Operational Program for Competitiveness and Internationalisation, and Portuguese funds via FCT-Fundao para a Ciéncia e a Tecnologia, under projects UIDB/04539/2020, UIDP/04539/2020, LA/P/0058/2020, UIDB/00313/2020, UIDP/00313/2020, UIDB/00324/2020, PTDC/MAT-APL/28118/2017, POCI-01-0145-FEDER-028118. MG funded by University of Bath grant NE/L002434/1.

Adérito Araújo
University of Coimbra, Portugal
alma@mat.uc.pt

Matthew Griffith
Department of Mathematics
University of Bath
M.J.Griffith@bath.ac.uk

Armindo J. Salvador
Centre for Neuroscience and Cell Biology
The University of Coimbra
a.salvador@mail.telepac.pt

Rui Travasso
Centro de Física Computacional, Universidade de Coimbra
Coimbra, Portugal
rui@teor.fis.uc.pt

MS173

Modelling the Coupling of Calcium Signalling and Mechanics in Embryogenesis

Calcium (Ca²⁺) signalling is one of the most important mechanisms of information propagation in the body. The coupling between Ca²⁺ signalling and mechanics plays a crucial role in fertilisation, embryogenesis, wound healing, and cancer. One of the most important stages of embryogenesis is neural tube closure (NTC), which is responsible for the formation of the neural tube, the structure that later develops into the central nervous system. NTC is driven, in part, by the complex interplay of Ca²⁺ signalling and mechanics through Apical Constriction (AC) and malformations like anencephaly and Spina Bifida can result when this coupling is disrupted. However, this mechanochemical coupling is poorly understood and few models are available. We have developed a novel cell-based mechanochemical model that integrates Ca²⁺ signalling into the vertex modelling framework. Using the vertex-based model (developed in Python), we investigate how the spatial and temporal patterning of Ca²⁺ can trigger changes in cell shape by regulating their mechanical properties, resulting in tissue deformation, and how this deformation in turn impacts the Ca²⁺ patterning. Crucially, the model reproduces key aspects of the NTC process: asynchronous, cell-autonomous Ca²⁺ flashes precede pulsed contractions at the cell level and a monotonic reduction of the tissue area.

Abhishek Chakraborty
School of Mathematics, Cardiff University, UK
chakrabortya5@cardiff.ac.uk

Katerina Kaouri, Timothy N. Phillips
School of Mathematics

Cardiff University, UK
kaourik@cardiff.ac.uk, phillipstn@cardiff.ac.uk

Philip K. Maini, Ruth E. Baker
Wolfson Centre for Mathematical Biology
Mathematical Institute, University of Oxford
philip.maini@maths.ox.ac.uk, baker@maths.ox.ac.uk

Neophytos Christodoulou, Paris Skourides
Department of Biological Sciences, University of Cyprus
christodoulou.neophytos@ucy.ac.cy, skourip@ucy.ac.cy

MS173

Scale Free Chaos in Insect Swarming

Swarms are examples of collective behavior of insects possessing strong correlations but no global order. They exhibit finite-size scaling and have dynamic critical exponent $z \approx 1$. Midge swarms have reproductive purposes, form near a marker placed on the floor, and comprise less than 900 flying male insects. An important model of flocking is Vicsks: N particles under discrete time dynamics align their velocities to the average velocity of all particles within their sphere of influence plus an alignment noise. To avoid aligned particles escaping to infinity, particles are usually confined within a box with periodic boundary conditions. Instead, we subject the particles to a confining harmonic potential in infinite space. For this model, appropriate noise and confinement strength, we have discovered a novel phase transition. The confinement versus noise critical line separates disperse single-cluster from multicluster chaos (distinguished by topological data analysis), swarms are scale-free, have minimal correlation time, correlation length proportional to swarm size. The largest Lyapunov exponent vanishes with confinement. Susceptibility, correlation length, dynamic correlation function, largest Lyapunov exponent obey power laws. Their critical exponents agree fairly with observations of natural midge swarms. These results contrast with those obtained from the order-disorder transition of the Vicsek model confined in a finite box under periodic boundary conditions.

Rafael González-Albaladejo
Universidad Complutense de Madrid, Spain
rafael09@ucm.es

Ana Carpio
Universidad Complutense de Madrid
Departamento de Matemática Aplicada
carpio@mat.ucm.es

Luis Bonilla
Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience & IndMath
bonilla@ing.uc3m.es

MS173

Predicting Biophysical Properties of Proteins

We report our recent work for the prediction of protein binding affinity using convolutional neural networks (CNNs). This work utilizes uniform features extraction from both topological data analysis and electrostatics from charged protein-ligand complexes. Our methods overcome the difficulties of involving electrostatics, which is expensive to compute due to its long range and pairwise natures. The simulation results show accurate prediction of binding

affinity and meanwhile demonstrate the significance of involving electrostatics in the machine learning framework.

Elyssa Sliheet
Southern Methodist University
esliheet@smu.edu

MS173

Active Vertex Model and Coarse Graining of Collective Cell Migration in a Monolayer

Due to its importance in wound healing, cancer progression and embryonic development, collective cell migration in a cellular monolayer has been extensively studied by mathematical and physical modeling, computation and experimentation. Confluent motion of epithelial cells can be described using simpler active vertex models, in which cells are polygons forming a Voronoi tessellation of the tissue. The dual Delaunay triangulation comprises cells centers that evolve according to dynamics including collective inertia and active forces. To model motion on supracellular length scales, we should develop continuum theories informed by models describing cell motion. Here we describe coarse graining strategies of the active vertex model that produce continuum equations for density, mean velocity, pressure and nematic order parameter and include active forces. We extract coarsened quantities directly from simulations of the active vertex model and compare with results from the continuum examples equations in simple cases of wound healing assays and formation of fingers.

Gloria Triguero
University of Heidelberg, Germany
gloria.triguero_platero@stud.uni-heidelberg.de

Falko Ziebert
Heidelberg University, Institute for Theoretical Physics
f.ziebert@thphys.uni-heidelberg.de

Luis Bonilla
Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience & IndMath
bonilla@ing.uc3m.es

MS174

Towards Integral Equations at Exascale in Rust

Over the last year we have started at University College London a project to develop a new set of core libraries to solve boundary integral equations at very large scale. Our developments are based on a suite of small core libraries developed in Rust for functionalities such as field translation of operators, octrees, kernel evaluation, etc. In this talk we motivate our approach, discuss software challenges and provide an overview of our current state of development and first benchmarking results.

Timo Betcke, Srinath Kailasa
University College London
t.betcke@ucl.ac.uk, srinathkailasa@gmail.com

Matthew Scroggs
University of Cambridge

matthew.scroggs.14@ucl.ac.uk

MS174

Bembel: The BEM-Based Engineering Library

Bembel, the BEM-based engineering library, is a flexible header-only C++ framework featuring higher order isogeometric Galerkin boundary element methods. To that end, functionality for Laplace, Helmholtz, and Maxwell-type problems is readily implemented and new problem types can be implemented in a few lines of code. Bembel is compatible with geometries from the Octave NURBS package, and provides an interface to the Eigen template library for linear algebra operations. For computational efficiency, it applies an embedded fast multipole method tailored to the isogeometric analysis framework and OpenMP parallelization.

Jürgen Dölz
University of Bonn
doelz@ins.uni-bonn.de

Helmut Harbrecht
University of Basel
Department of Mathematics and Computer Science
helmut.harbrecht@unibas.ch

Stefan Kurz
Technische Universität Darmstadt
Robert Bosch GmbH
stefan.kurz2@de.bosch.com

Michael Multerer
Università della Svizzera italiana
michael.multerer@usi.ch

Maximilian Nolte
TU Darmstadt
maximilian.nolte@tu-darmstadt.de

Sebastian Schöps
Computational Electromagnetics Group (CEM), TU Darmstadt
Centre for Computational Engineering (CCE), TU Darmstadt
schoeps@gsc.tu-darmstadt.de

Felix Wolf
Dassault Systèmes
wolf@temf.tu-darmstadt.de

MS174

Fast Direct Solvers for Helmholtz Problems Using Fmm3dbie

Recent years have seen the development of fast direct solvers which aim to reduce the inversion cost of the dense matrices that arise from boundary integral formulations to $O(N)$. In this talk we survey some recent approaches for the solution of oscillatory scattering problems and the software currently available, and present a recent fast direct solver for Helmholtz scattering problems developed with FMM3DBIE.

Srinath Kailasa
University College London

srinathkailasa@gmail.com

MS174

The Shifted Boundary Method for Embedded Domain Computations Using a High-Order Spectral Element Method for the 2D Poisson Problem

In recent years, the Shifted Boundary Method (SBM) have gained interest due to its ability to handle complex domains through embedded domain computations. The SBM address the problem of avoiding small cut cells and makes the meshing task close to trivial. The key feature - and the expense - of the SBM is how the boundary conditions (BCs) are applied on a surrogate/approximate boundary via the use of Taylor expansions to ensure that the convergence rates of the overall discrete formulation is preserved, see [Main and Scovazzi, 2017 and 2018]. This original work by Main & Scovazzi was exploiting the classical - second-order accurate - Finite Element Method (FEM), however, higher-order contributions have recently been made, see [Nouveau et al., 2019] and [Atallah et al., 2022]. A high-order numerical method is the Galerkin-based Spectral Element Method (SEM) [Xui et al., 2018] which can be viewed as a multi-domain version of the single-domain polynomial spectral method. We present a SEM-based model combined with the SBM for solving the Poisson equation in 2D on different domains/geometries with various BCs. Convergence studies are performed to establish the legitimacy of the work, including considerations of matrix conditioning when imposing Dirichlet BCs weakly via Nitsche's method. Also, the presented work is free of higher-order derivatives from the aforementioned Taylor expansions, as the formulation utilizes the polynomial behaving nature of the basis functions.

Jens Visbeck

Technical University of Denmark (DTU)
Department of Applied Mathematics and Computer Science
jvis@dtu.dk

Allan Peter Engsig-Karup
Department of Applied Mathematics and Computer Science
Technical University of Denmark
apek@dtu.dk

Mario Ricchiuto
INRIA
mario.ricchiuto@inria.fr

Guglielmo Scovazzi
Duke University
guglielmo.scovazzi@duke.edu

MS175

The Ubiquitous Sparse Matrix-Matrix Products

Multiplication of a sparse matrix with another matrix is a fundamental operation that captures the computational patterns of many data science applications, including but not limited to graph algorithms, sparsely connected neural networks, graph neural networks, clustering, and many-to-many comparisons of biological sequencing data. In the majority of these application scenarios, the matrix multiplication takes place on an arbitrary algebraic semiring where the scalar operations are overloaded with user-defined functions with certain properties or a more general

heterogeneous algebra where even the domains of the input matrices can be different. Here we provide unifying treatment of the sparse matrix-matrix operation and its rich application space.

Aydin Buluc

Lawrence Berkeley National Laboratory
abuluc@lbl.gov

MS175

Predicting ILU(0) Effectiveness for Sparse Matrix Systems via Explainable Machine Learning

Incomplete LU factorization of sparse matrices, which preserves the sparsity pattern of the original system, plays an important role in preconditioning iterative solvers. ILU(0) is a simple yet powerful preconditioner technique for many problems. However, generation of the ILU(0) preconditioner is costly, and it is thus useful to know whether it would be useful for a given problem beforehand. In this work we investigate the feasibility of machine learning techniques to aid in determining whether ILU(0) comprises a suitable preconditioner for a given matrix system. Our contributions are threefold: 1) We train a deep artificial neural network on examples from the SuiteSparse matrix collection to predict the effectiveness of ILU(0) based on meta-features describing the sparsity pattern. 2) Since deep learning approaches require huge amounts of training data, which cannot be provided by real-world matrix systems of appropriate sizes, we further evaluate the possibility of using partial matrices as a synthetic dataset for pretraining the model. 3) Additional insight into which characteristics are crucial for the successful ILU(0) is revealed through explainable machine learning via feature importance evaluation in tree-ensemble approaches. Our results show that ILU(0) effectiveness can be predicted with sufficient accuracy based on a few representative matrix characteristics, and that accuracy for very large matrices improves through learning on partial matrices.

Charlotte Debus

Karlsruhe Institute of Technology (KIT), SCC
charlotte.debus@kit.edu

Anna Ostrovskaya
Karlsruhe Institute of Technology, SCC
kontakt@annaos.dev

Markus Götz
Karlsruhe Institute of Technology, SCC
Helmholtz AI
markus.goetz@kit.edu

Hartwig Anzt
University of Tennessee
hanzt@tennessee.edu

MS175

2D S-Step Methods for Machine Learning

Stochastic gradient descent (SGD) is one of the most widely used optimization methods for solving various machine learning problems. SGD solves an optimization problem by iteratively sampling a few data points from the input data, computing gradients for the selected data points, and updating the solution. However, in a parallel setting, SGD requires interprocess communication at every iteration. We introduce an s-step SGD algorithm which adapts the s-step technique from Krylov methods to re-organize

the SGD computations into a form that communicates every s iterations instead of every iteration, where s is a tuning parameter. Furthermore, we develop 2D distributed-memory variants of the s -step SGD algorithm by leveraging existing work on 2D sparse matrix kernels and combining 1D s -step SGD with other existing SGD variants such as divide and conquer SGD and asynchronous SGD.

Aditya Devarakonda
Wake Forest University
devaraa@wfu.edu

MS175

Single-Factorization Interior Point Method for Constrained Optimization

This work proposes a new inexact interior point algorithm for convex quadratic programming problems with both equality and inequality constraints. Our method uses new preconditioned iterative methods to solve the linear systems that arise at every Newton step. These preconditioned conjugate gradient methods operate on an implicit Schur complement of the KKT system at each iteration. In contrast to standard approaches, the Schur complement we consider enables the reuse of the factorization of a fixed KKT subsystem across all interior point iterations. Further, the resulting reduced system admits preconditioners that directly alleviate the ill-conditioning associated with the strict complementarity condition in interior point methods. We propose two preconditioners that provably reduce the number of unique eigenvalues for the coefficient matrix (CG iteration count). One is efficient when the number of equality constraints is small, while the other is efficient when the number of remaining degrees of freedom is small. Numerical experiments with synthetic problems and problems from the Maros-Mszros QP collection show that our preconditioned inexact interior point solvers are effective at improving conditioning and reducing cost relative to the best alternative preconditioned method for each problem.

Samah Karim, Edgar Solomonik
University of Illinois at Urbana-Champaign
swkarim2@illinois.edu, solomon2@illinois.edu

MS175

Communication-Avoiding Algorithms for Sparse Triangular Matrices

Sparse triangular matrices occur extensively in scientific computing and data-intensive applications, most notably for solving systems of linear equations using Gaussian elimination. Such matrices have emerged as essential tools for artificial intelligence applications as they effectively model certain kinds of uni-directional information flow, such as the kind that occurs within feed-forward and graph neural networks or neuromorphic hardware. In this talk, I will discuss how we can exploit structured sparsity to reduce data transfer in parallel computation. Specifically, I will discuss new communication-avoiding algorithms for solving sparse triangular equations systems, which illustrate the more general idea of path decomposition and duplication two techniques that facilitate performing sparse triangular computation with reduced communication. The path-duplication technique, while generalizing traditional edge duplication-based methods like 3D matrix multiplication and ghost cells in stencil computation for reducing communication, also shows some differences that may be useful for broader sparse linear algebraic and graph com-

putations and beyond.

Piyush Sao
Oak Ridge National Laboratory
saopk@ornl.gov

Yang Liu
Lawrence Berkeley National Laboratory
liuyangzhuan@lbl.gov

Xiaoye S. Li
Computational Research Division
Lawrence Berkeley National Laboratory
xsli@lbl.gov

Ramakrishnan Kannan
Oak Ridge National Laboratories
kannanr@ornl.gov

Richard Vuduc
Georgia Institute of Technology
richie@cc.gatech.edu

MS176

Random Tree Besov Priors

We propose alternatives for Bayesian a priori distributions that are frequently used in the study of inverse problems. The aim is to construct priors that have same kind of good edge-preserving properties than total variation or Mumford-Shah but correspond to well defined infinite dimensional random variables and can be approximated with finite dimensional random variables. This is done by introducing a new random variable T that takes values in the space of trees, and which is chosen so that the realisations of the unknown have singularities only on a small set.

Hanne Kekkonen
Delft University of Technology
h.n.kekkonen@tudelft.nl

Matti Lassas
University of Helsinki
matti.lassas@helsinki.fi

Samuli Siltanen
University of Helsinki
Finland
samuli.siltanen@helsinki.fi

Eero Saksman
University of Helsinki
eero.saksman@helsinki.fi

MS176

Randomized MISMC Ratio Estimators

We consider the problem of estimating expectations with respect to a target distribution with an unknown normalizing constant, and where even the unnormalized target needs to be approximated at finite resolution. Under such an assumption, we extend a recently introduced multi-index Sequential Monte Carlo (SMC) ratio estimator, which provably enjoys the complexity improvements of multi-index Monte Carlo (MIMC) and the efficiency of SMC for inference. The present work leverages a randomization strategy to remove bias entirely, which simplifies estimation substantially, particularly in the MIMC context,

where the choice of index set is otherwise important. With theoretical results, the proposed method provably achieves the same canonical complexity of MSE^{-1} under appropriate assumptions as the original method, but without discretization bias. It is illustrated on examples of Bayesian inverse problems.

Xinzhu Liang, Kody Law
University of Manchester
xinzhu.liang@postgrad.manchester.ac.uk,
kodylaw@gmail.com

Simon Cotter
University of Manchester
School of Mathematics
Simon.cotter@manchester.ac.uk

Shangda Yang
University of Manchester
shangda.yang@manchester.ac.uk

MS176

Recent Advancements of Non-Gaussian Processes in Bayesian Inversion

This talk is focused on the recent work, which has considered the promotion of non-Gaussian priors for tasks in inverse problems. This is a challenging area as there are numerous applications where one requires learning rough features and edges, while not having a universal prior that works well for this. As a result this presentation will consider recently proposed prior and their developments, such as analysis and numerical evidence that they outperform Gaussian, and other non-Gaussian priors. This will be tested on a range of inverse problems such as tomography and PDE-inversion.

Lassi Roininen
Lappeenranta-Lahti University of Technology
lassi.roininen@lut.fi

MS176

Bayesian Inversion for Multiscale Materials with Limited Observations

Recovering material properties in the presence of multiscale features is particularly challenging for two reasons: the large amount of data needed to infer the small scale variations, and the very high computational cost of solving the forward problem. In this talk, we propose a procedure to overcome these difficulties for a stationary physical system with spatial multiscale features, that is, a system for which the measured quantities can be modeled as output functionals of the solution to an elliptic partial differential equation with multiscale diffusion coefficient. In order to estimate the uncertainties in the inversion procedure, we use a Bayesian approach. To handle the fact that the data are often too little informative to infer the multiscale coefficient, we aim at recovering effective properties of the system based on the local orthogonal decomposition (LOD), whose relationship with the microscale properties is approximated via a neural network surrogate constructed in a pre-processing procedure. In the inversion algorithm, the surrogate allows to solve the forward problem at the macroscale only, therefore reducing considerably the computational cost.

Fabian Kröppf
Universität Augsburg

fabian.kroepfl@math.uni-augsburg.de

Daniel Peterseim
Institute for Numerical Simulation
Rheinische Friedrich-Wilhelms-Universität Bonn,
Germany
daniel.peterseim@math.uni-augsburg.de

Laura Scarabosio
Department of Mathematics, IMAPP
Radboud University
scarabosio@science.ru.nl

MS177

Trade-Offs between Performance, Energy, and Accuracy of Non-Standard Computer Number Formats

Computer number formats like Posits or BFloat16 constitute alternative approaches to the IEEE floating-point standard that can potentially achieve reductions in terms of memory storage, bandwidth, and compute costs. Indeed, hardware vendors have incorporated half-precision data formats and have implemented mixed-precision (MP) instructions. This talk will discuss the advantages and limitations of these alternative formats with respect to the IEEE standard in terms of trade-offs between performance, energy, and accuracy.

Marc Casas
Barcelona Supercomputing Center
marc.casas@bsc.es

MS177

Verificarlo: Tuning and Debugging Floating-Point Computations Through Stochastic Rounding

Reducing the cost, both in time and energy, of computer simulations is critical. The numerical precision of simulations should be sufficient to provide scientific insights but as low as possible to save energy and computation time. We discuss the design of *verificarlo* (<https://github.com/verificarlo/verificarlo>), an open-source framework to verify and optimize numerical accuracy in complex programs. *Verificarlo* is a compiler tool built upon the LLVM infrastructure. It includes different floating-point backends simulating the effect of numerical errors and the impact of using lower precision. Before lowering precision, one must ensure that the simulation is numerically correct. In *Verificarlo*, we rely on alternative floating-point models, such as Stochastic Rounding, to pinpoint numerical bugs in simulation codes. A probabilistic definition of the number of significant digits allows us to estimate the accuracy of a computation. *Verificarlo* can then explore the compromise between precision and performance through its variable precision backend that simulates lower precisions on software. It can identify the parts of the code that benefit from smaller floating-point formats. *Verificarlo* has been applied in HPC codes such as neuroimaging pipelines, DFT quantum mechanical modeling, or structure simulations. In particular, we show how *Verificarlo* was used to optimize the speed and energy consumption of a deflated conjugate-gradient solver.

Pablo De Oliveira Castro
PRiSM & ECR - UVSQ
45 avenue des Etats-Unis, 78035 Versailles cedex

pablo.oliveira@uvsq.fr

MS177

Interflop: a Project for Interoperable Tools for Computing, Debugging, Validation and Optimization of Floating-Point Programs

The InterFLOP project aims at providing a modular and scalable platform to both analyze and control the costs of floating-point behavior of today's real programs facing those new paradigms (bigger problems, new architectures, new representation formats). The results of existing tools often generate new questions such as: Is this small error guaranteed? May this numerical error occur and how to reduce it? Some more costly analyses bring pieces of answers to these questions, but also require more expertise. In InterFLOP, we propose new analyses and combinations of existing ones to address the challenge of providing a quick and precise numerical diagnosis requiring little user expertise. For that, InterFLOP will collect and combine information on numerical instabilities, catastrophic cancellations, unstable tests, build various statistical analyses of program executions at minimal overhead. In this talk, we will describe advances made within the InterFLOP project which aims at 1) enlarging the class of possible applications by considering new front-ends and, therefore, new analyses; 2) providing finer numerical analyses based on formalized composite analyses; 3) verifying the accuracy in the context of precision auto-tuning to make applications more efficient and robust; 4) building statistical analyses tools and help the developer interpret the numerical behavior of program through graphical interpretation.

David Defour
DALI/LIRMM
University of Perpignan
david.defour@univ-perp.fr

MS177

Reliable and Sustainable Computations: A Brief Overview

A traditional view on computing and algorithms, offered in classic books on numerical linear algebra and analysis, often relies on sequential representation of algorithms as well as on rigorous error analysis of them. The former has been changing over the past 30 years, especially with increasing power of computing centers, while the later was largely untouched and still provides rather worse case scenarios (pessimistic estimates). In this talk, I present my work on accuracy and reproducibility assuring strategies for parallel iterative solvers that may not hold due to the non-associativity of floating-point operations; these strategies primarily rely on guarding every bit of resulting computation until the final rounding, hence they can be costly. Driven by the energy consumption constraint for computing centers and, hence, computations, scientists began a revision of applications, algorithms, as well as the underlying working / storage precision. The main aim is to make the computing cost sustainable and apply the lagom principle, especially when it comes to precision but also to the error estimates. Hence, I will also present preliminary results on probabilistic (aka optimistic) error analysis that closely captures the actual accumulated error.

Roman Iakymchuk
Umeå University

riakymch@cs.umu.se

MS178

Hotelling Deflation and sTRLED Algorithm for Computing Many Eigenpairs

Numerous applications in physical simulations and data analysis give rise to the need for computing many eigenpairs of large-scale Hermitian matrices. Subspace projection methods such as the Lanczos algorithm are expensive to apply directly due to the high costs of maintaining the orthogonality of a large number of projection basis vectors. To address the issue, we revisit Hotelling's explicit external deflation (EED) technique that sequentially shifts away converged eigenvalues while only maintaining a small number of basis vectors inside a subspace projection method. We first show that the EED with suitable shifts is numerically backward stable. We then introduce an eigensolver sTRLED, a combination of the EED with the s-step communication-avoiding thick-restarted Lanczos method (TRLan). sTRLED incrementally computes a batch of eigenpairs at a time and only uses a small projection subspace of TRLan. In addition, the sparse-plus-low-rank structure of the EED matrix allows for the design of specialized matrix power kernels that further reduces both communication and computation costs in sTRLED. This is joint work with Z. Bai, J. Dongarra, C. Lin, J. Wang, and I. Yamazaki.

Ding Lu
University of Kentucky
ding.lu@uky.edu

MS178

Contour Integration for a Class of Eigenvalue Problems with Eigenvector Nonlinearities

Solving polynomial eigenvalue problems with eigenvector nonlinearities (PEPv) is an interesting computational challenge, outside the reach of the well-developed methods for nonlinear eigenvalue problems. We present a natural generalization of these methods which leads to a contour integration approach for computing all eigenvalues of a PEPv in a compact region of the complex plane. Our methods can be used to solve any suitably generic system of polynomial or rational function equations.

Rob Claes, Karl Meerbergen
KU Leuven
rob.claes@kuleuven.be, karl.meerbergen@kuleuven.be

Simon Telen
MPI Leipzig
simon.telen@mis.mpg.de

MS178

TNL: Numerical Library for Modern Parallel Architectures

TNL (www.tnl-project.org) is a collection of building blocks that facilitate the development of efficient numerical solvers and HPC algorithms. It is implemented in C++ using modern programming paradigms in order to provide a flexible and user-friendly interface similar to, for example, the C++ Standard Template Library. TNL provides native support for modern hardware architectures such as multicore CPUs, GPUs, and distributed systems, which can be managed via a unified interface. In our presenta-

tion, we will demonstrate the main features of the library together with efficiency of the implemented algorithms and data structures.

Tomáš Oberhuber

Czech Technical University in Prague
tomas.oberhuber@fjfi.cvut.cz

Jakub Klinkovsk
FNSPE CTU in Prague
jakub.klinkovsky@fjfi.cvut.cz

Radek Fucík, Aleš Wodecki
Czech Technical University in Prague
radek.fucik@fjfi.cvut.cz, ales.wodecki@fjfi.cvut.cz

MS179

Uncovering the Dynamics of Liquid Metal Films via Efficient 3D GPU Simulations

This talk focuses on the influence that thermal effects have on the dewetting of liquid metal films deposited on thermally conductive substrates, and heated by an external heat source. Using asymptotic analysis we develop a mathematical model that simultaneously incorporates thermal effects in the metals, heat transfer in the substrate, and the evolution of the metallic films. In order to solve the underlying model we develop a 3D GPU code that implements an implicit-explicit finite difference scheme and allows us to simulate self-consistently free-surface evolution and heat conduction on large domains. In this talk, we focus in particular on the numerical algorithm and its parallelization. By utilizing this code, we find that the properties of the thermally conductive substrate – in particular its thickness and rate of heat loss – play a critical role in controlling the film temperature and dynamics. The authors acknowledge support from the USMA Dean's Faculty Research Fund and NSF DMS-1815613.

Ryan H. Allaire
United States Military Academy
ryan.allaire@westpoint.edu

Lou Kondic
Department of Mathematical Sciences, NJIT
University Heights, Newark, NJ 07102
kondic@njit.edu

Linda Cummings
New Jersey Institute of Technology
linda.cummings@njit.edu

MS179

Deep Reduced Order Models for Grain Microstructure Evolution

Predicting grain formation during alloy solidification is of great importance in additive manufacturing (AM). Numerical simulations require fine spatial and temporal discretizations that can be computationally expensive. In this talk we introduce GrainNN, an efficient and accurate reduced-order model for epitaxial grain growth in additive manufacturing conditions. GrainNN is a sequence-to-sequence long-short-term-memory (LSTM) deep neural network that evolves the dynamics of manually crafted features. We present results in which GrainNN can be orders of magnitude faster than phase field simulations, while delivering 5%15% pointwise error. This speedup includes the

cost of the phase field simulations for generating training data. This is joint work with Yigong Qin (UT Austin), Steve DeWitt (ORNL), and Balasubramanian Radhakrishnan (ORNL)

George Biros
The University of Texas at Austin
Oden Institute
biros@oden.utexas.edu

MS179

Assimilation of in-Situ Thermal Observations into Component-Scale Additive Manufacturing Simulations to Improve Thermal State and Model Parameter Estimates

Through controlled deposition of heat and mass, additive manufacturing offers the potential for complex component shapes and location-specific material properties. A downside of the flexibility in the geometry and the heat input into the system, however, is that thermal evolution is complex with rapid variation across the component. This thermal evolution is critical in determining factors such as the final microstructure in the component as well as the residual stress. Here we present a new approach for estimating the temperature throughout an additively manufactured component as it is being built. We combine a GPU-accelerated finite element thermal solver with in-situ infrared imaging and thermocouple data using an ensemble Kalman filtering approach with localization. We also explore augmenting the filtered state with model parameters (e.g. laser absorptivity, thermal diffusivity) to learn the values of those parameters that are most consistent with observations. Using synthetic observations where the ground truth is known, we demonstrate improved temperature estimates in the interior of the component using data assimilation, especially with the parameter-augmented approach. We also demonstrate assimilation of real infrared imaging and thermocouple data. Finally, we discuss future applications of this approach for qualification of components and for adaptive process control. Stephen DeWitt, Bruno Turcksin, and James Haley

Stephen DeWitt
Oak Ridge National Laboratory
dewittsj@ornl.gov

MS179

Application of Gaussian Process Auto-Regressive Models for Capturing Microstructure Evolution Paths in Ni Alloys

Phase field (PF) simulations using the MEUMAPPS code have been used to generate a set of microstructures representing the nickel-based superalloy Inconel 718 during a post additive manufacturing (AM) heat treatment process, covering a range of niobium concentration levels from 0.08 to 0.16 (atom fraction), characteristic of a solidification induced niobium segregation in the as-built AM microstructure. It has been demonstrated that varying Nb concentration influences the relative volume fractions of the delta and gamma precipitate phases in the alloy. The Materials Knowledge Systems (MKS) framework has been used to statistically quantify the microstructure arrays produced by the PF simulations. The resulting reduced-order representations capture all of the salient features of the microstructural morphology. Gaussian process autoregression (GPARG), a subset of Gaussian process regression, has been utilized to predict the time evolution of the reduced-order

microstructure representations. This result demonstrates that a combination of the MKS framework and machine learning modeling is a powerful tool for predicting the time evolution of microstructures in heat treatment processes. Work was supported by the Exascale Computing Project (17-SC-20-SC) at ORNL. This research used resources of the Oak Ridge Leadership Computing Facility, under the Office of Science of the United States Department of Energy under contract DE-AC05-00OR22725.

Balasubramaniam Radhakrishnan
Oak Ridge National Laboratory
radhakrishnb@ornl.gov

Surya Kalidindi
Georgia Institute of Technology
Materials Informatics for Engineering Design
surya.kalidindi@me.gatech.edu

Andrew Marshall
Georgia Institute of Technology
amarshall40@gatech.edu

James F. Belak
Lawrence Livermore Nat. Lab.
belak1@llnl.gov

MS180

Local h-, p- and k-adaptive Refinement through Non-Uniform Polynomial degree LR B-Splines with Application to Isogeometric Analysis

Locally Refined B-splines (LR B-splines) [Dokken et al., Polynomial splines over locally refined box-partitions, CAGD, 2013] have proven a very flexible and powerful framework in several applications, such as, interpolation and data reconstruction, Computer Aided Design (CAD) and isogeometric analysis (IGA). In this work we extend LR B-splines to allow for non-uniform polynomial degree. This capability, combined with local h-refinement, also allows for local k-refinement, that is, increasing the polynomial degree as well as the smoothness at newly introduced knot lines. The novel refinement schemes are based on degree elevation of individual basis functions by means of certain two-scale relations, and lead to nested spaces. We study the properties of the introduced spaces and apply them to several two-dimensional model problems, illustrating the efficacy of the proposed adaptive refinement methodology in the context of isogeometric analysis.

René Hiemstra
Technische Universität Darmstadt, Germany
hiemstra@mechanik.tu-darmstadt.de

MS180

Geometrically Smooth Surfaces and Basis for Isogeometric Analysis Simulations

Geometrically smooth spline functions are piecewise polynomial functions defined on a mesh, that satisfy properties of differentiability across shared edges; their unstructured nature provide them a wide range of application like, for example, Isogeometric Analysis approaches on surfaces of arbitrary topology. In this presentation, we consider G^1 splines on quadrangular meshes with given quadratic gluing data along shared edges. We describe briefly their properties, analyse their spaces, and provide dimension formula. Computing efficiently basis functions for these spaces is critical in the IGA approach; we investigate this problem

and show an explicit construction for such bases. A few experimentations illustrate these developments.

Michelangelo Marsala
INRIA Sophia Antipolis Méditerranée
France
michelangelo.marsala@inria.fr

MS180

Isogeometric Boundary Element Methods for Multipatch Geometries

The multi-patch IgA formulation of BEMs is considered and applied in particular to the discretization of conventional boundary integral equations coming from 3D Helmholtz problems. B-spline tailored quadrature rules developed for both singular and integrals allow us to perform numerical integration directly in the support of each B-spline generating the adopted IgA discretization space. The proposed quadratures combine higher order analytical singularity extraction process for the governing singular integrals and a B-spline quasi-interpolation numerical integration for the regular integrals. A special care is needed to correctly address integration across patch interfaces. On numerical examples we demonstrate that the expected orders of convergence for the approximate solutions are achieved with a small number of quadrature nodes. Joint work with Antonella Falini (University of Bari, Italy), Tadej Kanduč (University of Ljubljana, Slovenia), Alessandra Sestini (University of Florence)

Maria Lucia Sampoli
University of Siena
Department of Mathematics and Computer Science
sampoli@unisi.it

MS180

Smooth Isogeometric Discretizations for Fourth Order PDEs

In this talk we discuss the construction of smooth isogeometric discretization spaces and their applicability to solve fourth order PDEs. Here we focus mostly on C^0 -matching multi-patch constructions to represent complex geometries. In this case, the global continuity for the basis functions is usually only C^0 . To obtain exactly C^1 -smooth isogeometric functions, special constructions need to be employed. It could be shown in [Collin, Sangalli, Takacs. Analysis-suitable G^1 multi-patch parametrizations for C^1 isogeometric spaces. CAGD, 2016] that generic, C^0 -matching multi-patch parametrizations do not allow C^1 -smooth isogeometric discretizations of sufficient approximation power. Depending on the patch interior continuity, approximation orders are drastically reduced or convergence may be completely prevented. To circumvent this issue one can increase the polynomial degree locally or relax the smoothness conditions. Two such approaches are developed in [Weinmüller, Takacs. An approximate C^1 multi-patch space for isogeometric analysis with a comparison to Nitsche's method. CMAME, 2022] and [Takacs, Toshniwal. Almost- C^1 splines: Biquadratic splines on unstructured quadrilateral meshes and their application to fourth order problems. CMAME, 2023], respectively. We discuss these approaches and use them to solve fourth-order problems, such as the biharmonic equation or Kirchhoff-Love shell problems.

Thomas Takacs
Johann Radon Institut (RICAM)

thomas.takacs@ricam.oeaw.ac.at

MS181

Efficient Data-Driven Control Using Autonomous Surrogate Models

In recent years, significant progress has been made in the field of machine learning and data-based methods, including the modeling of dynamical systems with the help of data-based surrogate models. These methods have also found their way into the field of control engineering, especially in model predictive control. However, compared to the creation of surrogate models for autonomous models (without control input), the building of data-based models for controlled systems is typically much more difficult and data hungry. To address these issues, in this talk, we present a framework which is based on a discretization of the control inputs such that modeling techniques for autonomous systems can be used without adaptations. As a result, the modeling effort and also the amount of required data can be reduced. Relaxation and rounding techniques from the area of mixed-integer control are then used to solve the newly derived discrete control problem. We verify the presented approach with various examples where we use different techniques to build the data-based models.

Sebastian Peitz
Paderborn University
sebastian.peitz@upb.de

Katharina Bieker
Paderborn University
Paderborn, Germany
bieker@math.upb.de

MS181

Line-of-Sight Based Games in a Domain with Obstacles

We consider a pursuit-evasion game in which a moving pursuer seeks to maintain visual contact with an evader which is moving in an environment containing obstacles. We introduce a new visibility function to determine the region of the domain in which the line of sight between the players is not cut by any obstacle. Then, we use this approach to analyze the solution of the Hamilton-Jacobi-Isaacs (HJI) equation associated to the game, especially near the boundary, where the solution happens to be discontinuous. Finally, we show how the new visibility function can be used to design (sub)optimal feedback policies, which approximate the optimal one given by the (HJI) equation, and are computationally much more efficient. This is a joint work with Richard Tsai (UT Austin) and Carola-Bibiane Schnlieb (University of Cambridge).

Carlos Esteve-Yagüe
University of Cambridge
ce423@cam.ac.uk

MS181

Risk Neutral Quantum Optimal Control Using Quadratic Approximations

For the design of quantum logic gates – the fundamental building block of quantum algorithms – achieving high target fidelities is essential to the success of an algorithm. However, current Noisy Intermediate Scale Quantum (NISQ) era quantum devices are typically polluted

by noise. This work investigates performing optimal control of closed quantum systems for logic gate synthesis under time-dependent uncertainty. We consider an objective function that is the mean of the trace infidelity measure, leading to a risk-neutral optimal control problem. Rather than using a Monte Carlo estimate for the risk-neutral objective, we consider taking a quadratic approximation of the objective with respect to the uncertainty as a proxy. This leads to an explicit expression for the mean of the control objective in terms of the trace of the (preconditioned) Hessian with respect to the uncertainty, for which we leverage trace estimators to approximate with a modest number of Hessian-vector products. Finally, we reformulate the risk-neutral optimization problem as a constrained optimization problem in which we employ a quasi-Newton method to minimize the quadratic approximation with respect to the control. This approach yields a method for which the computation of the risk-neutral objective and its gradient (with respect to the control) has a cost that depends only on the number of trace estimators used which can be chosen independently of the dimension of the control and uncertainty.

Fortino Garcia
Courant Institute of Mathematical Sciences
New York City
fortino.garcia@cims.nyu.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Gabriele Ciaramella
Politecnico di Milano
gabriele.ciaramella@polimi.it

Dominique Sugny
Laboratoire Interdisciplinaire Carnot de Bourgogne
dominique.sugny@u-bourgogne.fr

MS181

Distributed Inexact Newton Method with Adaptive Step Size

In distributed optimization a set of computational agents has to cooperatively solve a given optimization problem. One of the main issues in this framework is the selection of the stepsize: classical globalization strategies such as line-search are not feasible in the multi-agent setting, while employing a fixed stepsize is known to often cause the method to be slow, and usually requires the knowledge of the regularity constants of the problem, which can be hard to estimate distributedly. We propose a distributed Inexact Newton method that combines the penalty formulation of the problem with Jacobi Overrelaxation method, and uses an adaptive stepsize that can be computed without a priori knowledge of the constants of the problem. Under suitable regularity assumptions over the objective function and on the connectivity of the underlying network, we prove both global and local fast convergence of the proposed method, with convergence order that depends on the choice of the parameters of the method. We also presents a set of numerical results that shows the comparison between the proposed algorithm and several methods from the literature.

Greta Malaspina, Dušan Jakovetic, Nataša Krejic
University of Novi Sad

Faculty of Sciences
 greta.malaspina@dmi.uns.ac.rs, du-
 san.jakovetic@dmi.uns.ac.rs, natasak@uns.ac.rs

MS181**A Multigrid Method for Challenging Optimal Control Problems Constrained by Random PDEs**

In this contribution we present a multigrid method to solve the full-space optimality systems that arise in Optimal Control Problems (OCPs) constrained by random Partial Differential Equations (PDEs). These optimality systems are of very large dimension, and are often characterized by N state PDEs, N adjoint PDEs, and a single optimality condition, where N is the total number of collocation points used to discretize the probability space. Our approach is based on a collective smoother that requires the solution (possibly in parallel) of several reduced linear systems, each of dimension $2(N + 1) \times 2(N + 1)$, which can be achieved with optimal linear complexity. The efficiency of the multigrid scheme is tested on several problems, including a simple unconstrained linear-quadratic OCP, an OCP with box constraints and a L^1 penalization on the control, and a risk-adverse OCP involving the smoothed CVaR risk measure. In addition, the multigrid solver presented is robust against several important parameters of the model problems considered.

Gabriele Ciaramella
 Politecnico di Milano
 gabriele.ciaramella@polimi.it

Fabio Nobile
 EPFL, Switzerland
 fabio.nobile@epfl.ch

Tommaso Vanzan
 École Polytechnique Fédérale de Lausanne
 tommaso.vanzan@epfl.ch

MS182**Positive Definite Kernels on General Regular Domains**

Positive definite functions and kernels are an important tool for interpolation and approximation of functions in several dimensions. They give rise to well-defined interpolation problems and provide a variety of further approximation methods including, for example, quasi-interpolation and radial basis function approximation. In this paper we construct and analyse positive definite kernels on special domains; we study the kernel functions on a particular list of regular domains. The list includes the unit ball, conic surfaces, hyperbolic surfaces, solid hyperboloids, and simplices.

Martin Buhmann
 Justus-Liebig-University
 Giessen
 buhmann@math.uni-giessen.de

Yuan Xu
 University of Oregon
 yuan@uoregon.edu

MS182**Global-Local-Integration-Based Radial Basis Func-****tion Methods for Numerical Simulations**

In this talk, the development in global, local, and integration-based meshless computational methods via the use of radial basis function (RBF) will be presented. The local radial basis function computational method (LRBFCM) is an extension to solve large scale problems which has hindered the practical application of the global RBF method for years due to the ill-conditioning of its resultant full coefficient matrix. The LRBFCM has recently been applied to solve 2D cavity flows problems with free surface and some nonlocal diffusion and phase field problems. The main idea of the integration-based radial basis function (FIM-RBF) method is to transform the original partial differential equation into an equivalent integral equation whose approximation can be sought by standard numerical integration techniques. Due to the advantages of unconditional stability of numerical integration and spectral convergence of RBF approximation, the FIM-RBF method can solve boundary value problem under irregular domain with various kinds of stiffness. Recently, the FIM method has been generalized to solve multi-dimensional options pricing problems. Numerical examples with sensitivity analysis will be given to verify the efficiency and effectiveness of these meshless computational methods.

Benny Hon
 City University of Hong Kong
 maychon@cityu.edu.hk

MS182**Solving Hyperelastic Problems in Solid Mechanics with RBF-FD Based Methods**

Even though the collocation version of RBF-FD has been applied in various fields, researchers still struggle to formulate a method that would be robust, especially in the presence of nonlinearities and Neumann boundary conditions. For elliptic problems, these deficiencies have recently been addressed by the introduction of oversampled RBF-FD methods that replace collocation with a solution procedure based on least squares. This has also facilitated the introduction of an unfitted version of the method that further reduced the constraints on node positioning and helped alleviate problems with skewed stencils near domain boundaries. In this talk an application of oversampled RBF-FD methods to problems in nonlinear solid mechanics will be presented, motivated by the desire to produce a meshless numerical model for the human diaphragm.

Boštjan Mavric
 Department of Information Technology
 Uppsala University
 bostjan.mavric@it.uu.se

MS182**A Non-Oscillatory Finite Volume Method Based on a Pointwise Smoothing RBF Approximation for Conservation Laws**

One of the properties of a conservation law problem is that even if the initial condition is smooth the solution may evolve into a discontinuity or a sharp front. These discontinuities and sharp gradients result in non-physical oscillations in the numerical solution. The purpose of numerical methods is to precisely capture and suppress these oscillations while achieving a solution with a high order of convergence. There are several approaches but one of the most significant and commonly used is a class of essen-

tially non-oscillatory (ENO) or weighted ENO (WENO) methods. The primary idea behind the ENO (or WENO) reconstruction is to determine a set of stencils surrounding a control volume, execute a reconstruction on each of these stencils, and finally pick up the smoothest (or a weighted) approximation as the desired reconstruction for that control volume. In this talk, we present an alternative approach which employs a single central stencil per any cell and suppresses oscillations by utilizing a new RBF approximate-interpolation method based on a pointwise smoothing approximation.

Davoud Mirzaei
Uppsala University
davoud.mirzaei@it.uu.se

Navid Soodbakhsh
University of Isfahan
ssh.navid@yahoo.com

MS182

A Parallel-in-Time Contour Integral Method for RBF-FD Discretizations

We study the numerical performance of the parallel-in-time contour integral method for one-dimensional differential equations of the reaction-diffusion-convection type. The method is based on the representation of a Dunford-Cauchy integral along a contour which encompasses the spectrum of an unbounded infinitesimal operator A corresponding to the equation. The discrete operator is constructed as a discretization of diffusion and advective, resp. convective terms by the Radial Basis Functions with generalized Finite Difference (RBF-FD) method. A contour is chosen to be of the elliptic shape with varying number of cubature points. Numerical performance of the contour integral method is compared with sequential-in-time realizations of the θ -scheme. The studied algorithm is $2N_{\text{time}}N_{\text{quad}}$ parallelizable, where N_{time} is the number of time points at which the numerical solution is computed and N_{quad} is the number of cubature points on the contour. The method is flexible with respect to the choice of a discretization method and can be extended to multiple dimensions. Corresponding numerical examples are demonstrated.

Andriy Sokolov
Technische Universitat Dortmund
andriy.sokolov@math.tu-dortmund.de

Pavel Barkhayev
B. Verkin Institute of Low Temperature
Academy of Science of Ukraine
pavlo.barkhayev@gmail.com

Stefan Turek
Institute for Applied Mathematics,
TU Dortmund
turek@featflow.de

MS183

Learning Filtered Discretisation Operators: Non-Intrusive vs Intrusive Approaches

Simulating multi-scale phenomena such as turbulent fluid flows is typically computationally very expensive. Filtering the smaller scales allows for using coarse discretizations, however, this requires closure models to account for the effects of the unresolved on the resolved scales. The common approach is to filter the continuous equations, but this gives

rise to several commutator errors due to nonlinear terms, non-uniform filters, or boundary conditions. We propose a new approach to filtering, where the equations are discretized first and then filtered. For a non-uniform filter applied to the linear convection equation, we show that the discretely filtered convection operator can be inferred using three methods: intrusive (explicit reconstruction) or non-intrusive operator inference, either via derivative fitting or trajectory fitting (embedded learning). We show that explicit reconstruction and derivative fitting identify a similar operator and produce small errors, but that trajectory fitting requires significant effort to train to achieve similar performance. However, the explicit reconstruction approach is more prone to instabilities.

Syver Agdenstein
Department of Scientific Computing
CWI Amsterdam
syver.agdenstein@cwi.nl

Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cwi.nl

MS183

Non-Intrusive Model Order Reduction for Nonlinear Coupled Problems

We are interested in thermo-electro-mechanical problems arising in applications like curing process and spark plasma sintering. These models consist of non-linear material models like viscoelasticity and viscoplasticity. High-fidelity simulations of such coupled multiphysics problems are prohibitively expensive for performing parametric studies, optimization, control or estimation. In order to solve such problems, accelerating the forward model solves by means of efficient reduced order models is necessary. However, developing ROMs addressing these non-linearities is challenging due to the history dependence and presence of internal variables. To address this, hyper-reduction techniques have been used in the literature to construct reduced order models for non-linear problems. In particular, Reduced Basis - Empirical Quadrature Procedure (RB-EQP) has been used for hyperelasticity problems. RB-EQP has been shown to provide offline efficiency by reducing the number of integration points and the required full-order solves which can be very expensive to compute. Furthermore, methods in scientific machine learning have demonstrated online efficiency by avoiding expensive assembly of the reduced nonlinear problem. Exploiting these advantages of hyper-reduction and ML methods, we propose a unified numerical framework for an offline and online efficient non-intrusive ROM technique. Finally, we demonstrate the performance of the proposed method on a macroscale sintering problem.

Rahul Dhopeswar, Harshit Bansal
Eindhoven University of Technology
r.dhopeswar@tue.nl, h.bansal@tue.nl

Karen Veroy
Department of Mathematics and Computer Science
Eindhoven University of Technology
k.p.veroy@tue.nl

MS183

Model Enhancement by Discovering Friction Terms Using Physics-Informed Neural Networks for Con-

trol

Often, technical systems cannot be fully modeled due to limited knowledge, limited resources or due to the complexity of the task to be solved. This leads to shortcomings in the resulting governing equations and to a lack of accuracy. In a control setting, this is often treated by a larger control input. However, modern physics-inspired machine learning does not only allow us to solve complex problems in a data-driven way, but also to gain new insights into the underlying physics. In this work, we use physics-informed neural networks (PINNs) as well as sparse regression learning to identify the missing physics of a complex system inside a control setting. In detail, we identify nonlinear friction terms inside the model of a 2-dof manipulator based on a limited number of training data, which is either obtained from simulation or experiment. Based on the original and enhanced model, we generate feed-forward control sequences for a trajectory tracking task. A simple PID controller is used to compensate for errors that occur in the feed-forward control. The results show that the controller has to make far fewer corrections to the control variable based on the improved model than was the case with the original model.

Joerg Fehr
Institute of Engineering and Computational Mechanics
University of Stuttgart
joerg.fehr@itm.uni-stuttgart.de

MS183

Physics Informed Neural Networks Assessment on Fluid Dynamics Problems

Machine learning and more specifically deep learning have revolutionised the way classification, pattern recognition, and regression tasks are performed in various application areas. Historically the solving of computational fluid dynamics problems relies on finite elements and finite volume methods, meaning complex meshes and an important number of degrees of freedom are used to estimate solutions to governing equations. Recently however the development of the Physics Informed Neural Network (PINN) has emerged as a promising tool to improve the solution methods to these equations. The purpose of this talk is to assess different variants of these methods on a few representative equations from fluid dynamics. We aim to show how the performance and the precision of different types of models are affected by altering networks hyper-parameters including altering number of hidden layers, number of neurons to a hidden layers and the use of different activation functions including locally adaptive functions which have been shown to improve the accuracy of estimators. We will propose then some alternatives to the state of the art.

Victorita Dolean Maini
University of Strathclyde, UK and Universite Côte dAzur
victorita.dolean@strath.ac.uk

Konstantinos Ritos
University of Strathclyde and University of Thessaly
konstantinos.ritos@strath.ac.uk

Robin Kamenicky
University of Strathclyde
robin.kamenicky@strath.ac.uk

Alexander Heinlein
Delft University of Technology
Delft Institute of Applied Mathematics

a.heinlein@tudelft.nl

MS183

A Bregman Learning Framework for Sparse Neural Networks

I will present a novel learning framework based on stochastic Bregman iterations. It allows to train sparse neural networks with an inverse scale space approach, starting from a very sparse network and gradually adding significant parameters. Apart from a baseline algorithm called LinBreg, I will also speak about an accelerated version using momentum, and AdaBreg, which is a Bregmanized generalization of the Adam algorithm. I will present a statistically profound sparse parameter initialization strategy, stochastic convergence analysis of the loss decay, and additional convergence proofs in the convex regime. The Bregman learning framework can also be applied to Neural Architecture Search and can, for instance, unveil autoencoder architectures for denoising or deblurring tasks.

Tim Roith
Friedrich-Alexander Universität Erlangen-Nürnberg
tim.roith@fau.de

Leon Bungert
University of Bonn
leon.bungert@hcm.uni-bonn.de

Martin Burger
Friedrich-Alexander University
Erlangen-Nürnberg
martin.burger@fau.de

Daniel Tenbrinck
Friedrich-Alexander Universität Erlangen-Nürnberg
daniel.tenbrinck@fau.de

MS184

Structure-Preserving Numerical Algorithms for Hyperbolic and Related Models with Uncertainty

In this talk, we will discuss design of structure-preserving numerical methods for hyperbolic and related models with uncertainty. As a primary example, shallow water systems with uncertainty will be considered, but the developed ideas can be extended to a wider class of models, including different models of conservation/balance laws. Shallow water equations are widely used in many scientific and engineering applications related to modeling of water flows in rivers, lakes and coastal areas. We will show that the developed structure-preserving numerical methods deliver high-resolution and satisfy necessary stability conditions. We will illustrate the performance of the designed algorithms on a number of challenging numerical tests. Current and future research will be discussed as well.

Yekaterina Epshteyn
Department of mathematics
University of Utah
epshteyn@math.utah.edu

MS184

Locally-Implicit Discontinuous Galerkin Methods for Kinetic Boltzmann-BGK that are Arbitrarily High-Order and Asymptotic-Preserving

The kinetic Boltzmann equation with the Bhatnagar-

Gross-Krook (BGK) collision operator allows for the simulation of gas dynamics over a wide range of Knudsen numbers with a simplified collision operator. Efficient numerical methods for Boltzmann-BGK should be asymptotic-preserving, which allows the numerical method to be stable at fixed mesh parameters for any value of the Knudsen number, including in the fluid (very small Knudsen numbers), slip flow (small Knudsen numbers), transition (moderate Knudsen numbers), and free molecular flow (large Knudsen numbers) regimes. In this work we develop a novel approach for solving the Boltzmann-BGK equation for achieving both arbitrary high-order and the asymptotic-preserving property. The proposed method is a locally-implicit discontinuous Galerkin (LIDG) scheme with careful modification in both the prediction and correction steps to achieve the asymptotic-preserving property. Some key advantages of the proposed schemes are: (1) no splitting between macroscale and microscale components of the distribution function is required; (2) only a single unified time-discretization is required; and (3) arbitrary high-order in both space and time can be achieved simply by increasing the spatial polynomial order in each element. Several numerical examples are shown to demonstrate the effectiveness of the proposed numerical scheme.

James A. Rossmannith
Iowa State University
Department of Mathematics
rossmani@iastate.edu

MS184

A Particle-in-Cell Method for Plasmas with a Generalized Momentum Formulation

In this talk, some of our recent work on new particle-in-cell-methods for the Vlasov-Maxwell system will be presented. We propose a formulation in terms of potentials, in which the first-order Maxwell system is cast as a collection of second-order wave equations with the aid of the Lorenz gauge condition. Then, the usual Lorentz force for the particles is expressed in terms of these wave potentials and their spatial derivatives through the use of generalized momentum. A new BDF wave solver is presented for the solution of these equations for the potentials that induces a discrete equivalence between the Lorenz gauge condition and the continuity equation. We also discuss the new methods used to construct the spatial derivatives required in the update for particles, which are inspired by the wave solver used to evolve the potentials. Numerical results will be presented to demonstrate the application of the new method. If time permits, some of our on-going work in extending these methods to high-order accuracy will be discussed.

William A. Sands
Michigan State University
william.sands@austin.utexas.edu

Andrew Chrislieb
Dept. of Mathematics
Michigan State University
christlieb@math.msu.edu

Stephen White
Michigan State University
whites73@msu.edu

MS184

Structure Preserving Numerical Methods for Gy-

rokinetic Models

For long time simulations, especially relevant in magnetic fusion devices like tokamaks and stellarators, the conservation of invariants is essential to get physically relevant results. As for the Vlasov-Maxwell equations, the drift-kinetic or gyrokinetic models which are generally used in magnetic fusion can be derived from an action principle or a non canonical hamiltonian formulation. We will show in this talk how geometric Particle-In-Cell methods that have been recently derived for the Vlasov-Maxwell equations can be extended to the more complex drift-kinetic and gyrokinetic models.

Eric Sonnendrücker
Max-Planck-Institut für Plasmaphysik, NMPP division,
and
Technische Universität München, Department of
Mathematics
eric.sonnendruecker@ipp.mpg.de

MS184

Well-Conditioned Mode Matching Method for Applications in Photonics

To analyze photonics devices such as optical waveguides, resonant cavities and diffraction gratings, it is necessary to solve the Maxwells equations numerically. Due to the complexity of the structure, conventional numerical methods, such as the finite difference and finite element methods, are not very efficient. Often, a photonic device consists of several simpler parts for which the dielectric function depends only on one spatial variable. The mode matching method and its many numerical variants have been widely used to analyze photonic devices with simple one-dimensional parts. Since the dielectric function is typically piecewise constant, numerical mode matching methods based on piecewise polynomials, such as the polynomial expansion mode matching method (PEMM) and the pseudo-spectral mode matching method (PSMM), are particularly efficient. However, the condition numbers associated with PEMM and PSMM are often very large, leading to limited accuracy and difficulty in convergence when iterative methods are used. In this paper, we develop a well-conditioned numerical mode matching method based on a spectral Galerkin scheme. We show that this method produces numerical modes with good orthogonality and significantly reduces the condition numbers of the final linear systems.

Nan Zhang
City University of Hong Kong
nzhang234-c@my.cityu.edu.hk

Yayan Lu
city university of Hong Kong
mayylu@cityu.edu.hk

MS185

Multifactor Sequential Disentanglement via Spectrally-Structured Koopman Autoencoders

Disentangling complex data to its latent factors of variation is a fundamental task in representation learning. Existing work on sequential disentanglement mostly provides two factor representations, i.e., it separates the data to time-varying and time-invariant factors. In contrast, we consider multifactor disentanglement in which multiple (more than two) semantic disentangled components are gener-

ated. Key to our approach is a strong inductive bias where we assume that the underlying dynamics can be represented linearly in the latent space. Under this assumption, it becomes natural to exploit the recently introduced Koopman autoencoder models. However, disentangled representations are not guaranteed in Koopman approaches, and thus we propose a novel spectral loss term which leads to structured Koopman matrices and disentanglement. Overall, we propose a simple and easy to code new deep model that is fully unsupervised and it supports multifactor disentanglement. We showcase new disentangling abilities such as swapping of individual static factors between characters, and an incremental swap of disentangled factors from the source to the target. Moreover, we evaluate our method extensively on two factor standard benchmark tasks where we significantly improve over competing unsupervised approaches, and we perform competitively in comparison to weakly- and self-supervised state-of-the-art approaches.

Omri Azencot

Ben Gurion University of the Negev
azencot@cs.bgu.ac.il

MS185

Long-Time Stability of Deep Learning-Based Forecast Models of Multi-Scale Chaos

There is growing interest in fully data-driven forecast models (aka digital twins) for multi-scale nonlinear dynamical systems such as Earths climate and turbulence. Recent studies have shown the promise of deep neural networks (DNNs) for this purpose, producing models (learned from data) with short-term forecast skills comparable to those of high-fidelity numerical solvers. However, these models are all long-time unstable, failing to reproduce the statistics of the system. The roots of these instabilities are not understood. Here, using observational weather data and simulation data of turbulent flows, we show that the root of the instability is spectral bias: the well-known inability of DNNs in learning small scales. We demonstrate how the DNNs fundamental inability to represent small scales instigates instability, which then propagates to the large scales. We further show that training a DNN using a Fourier spectral loss function that promotes the learning of small scales combined with a time-integration scheme that dampens error propagation leads to a long-time stable DNN for both testcases. The framework developed here is readily applicable to stabilizing digital twins of any other multi-scale chaotic or turbulent system. Such long-time stable digital twins have wide applications in science and engineering.

Pedram Hassanzadeh

Mechanical Engineering and Earth Science
Rice University
pedram@rice.edu

Ashesh K. Chattopadhyay
Rice University
akc6@rice.edu

MS185

Learning Physical Models that Can Respect Conservation Laws

Recent work in scientific machine learning (SciML) has focused on incorporating partial differential equation (PDE) information into the learning process. Most of this work has focused on relatively easy PDE operators (e.g., elliptic and parabolic), with less emphasis on relatively hard PDE

operators (e.g., hyperbolic). Within numerical PDEs, the latter problem class is difficult since one must control a type of volume element or conservation constraint, which is known to be challenging. Delivering on the promise of SciML, however, requires seamlessly incorporating both types of problems into the learning process. To address this issue, we propose PhysNP, a framework for incorporating constraints into a generic SciML architecture. We provide a detailed analysis of our approach on learning with the Generalized Porous Medium Equation (GPME), a widely-applicable parameterized family of equations that illustrates the qualitative properties of both easier and harder PDE operators. PhysNP is effective for easy variants of the GPME, performing well with state-of-the-art competitors with little additional cost; and for harder variants of the GPME it outperforms other approaches that do not guarantee volume conservation. PhysNP seamlessly conditions on known physical invariants, it maintains probabilistic uncertainty quantification, and it deals well with spatial/temporal shocks. In each case, it achieves superior predictive performance on downstream tasks.

Derek Hansen

University of Michigan
derekh@umich.edu

Danielle Maddix Robinson, Shima Alizadeh, Gaurav Gupta
AWS AI Labs
dmmaddix@amazon.com, alizshim@amazon.com,
gauravaz@amazon.com

Michael Mahoney
Amazon
zmahmich@amazon.com

MS186

Probabilistic Geometry Reconstruction for Additive Manufacturing Digital Twins

As more and more critical parts are being fabricated via additive manufacturing (AM), there is an increasing need to ensure that the geometric and material properties of the parts conform to certain requirements. Currently this need is addressed by expensive, time consuming, and environmentally wasteful destructive testing of physical replica components. However, with more and more connected sensors being installed in AM machines, new possibilities are emerging for creating digital replicas, or digital twins, of the physical parts, inline. Such digital twins can be used as a basis for qualifying parts digitally, removing the need for excess physical parts. In this talk we will discuss approaches for creating digital twins from inline sensor data, focussing on the problem of extracting the geometry from volumetric images. To do this we transfer techniques developed in the medical imaging domain, to the setting of AM. A digital twin is never an exact representation of the physical object or process, as there are unknowns inherent in both the dynamics of the process and in the capture of the sensor data. We therefore also investigate the use of techniques such as Kalman filters and Markov random fields to capture these uncertainties in the digital twin.

Oliver Barrowclough

SINTEF
Norway
oliver.barrowclough@sintef.no

Sverre Briseid

SINTEF ICT, Department of Applied Mathematics
sverre.briseid@sintef.no

Georg Muntingh
SINTEF
georg.muntingh@sintef.no

MS186

Real-Time Parametric Adaptation of Digital Twins

Digital twin frameworks more and more resort to projection-based reduced order models (ROM) to perform real time simulations of complex systems. Several reduced bases may be simultaneously manipulated, either to project the equations onto a lower dimensional subspace, or to approximate nonlinear terms resisting the real time constraint. These reduced bases are often computed with the Proper Orthogonal Decomposition (POD) which requires snapshots of the solution for a given set of parameters. A difficulty then arises when it comes to adapt both efficiently and accurately these multiple reduced bases for new set of parameters. In this contribution, we compare different state of the art approaches to build the parametric reduced order model of a gas bearings dynamical system. We revisit the Interpolation on the Tangent Space of the Grassmann Manifold (ITSGM) technique, and we propose several improvements. Numerical experiments are presented to demonstrate the potential of the developed model to perform fast and accurate simulations for future digital twin applications.

Dimitri Goutaudier
EPFL
dimitri.goutaudier@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Schiffmann Jürg
EPFL
Switzerland
jurg.schiffmann@epfl.ch

MS186

Creating Value with Models in the Physical World Through the Executable Digital Twin

The digital twin has become an intrinsic part of the product creation process. Basically, it is a virtual image of a real asset, integrating all data, models and other digital information over the lifecycle. Its true power lies in the relationship with its physical counterpart. Data acquired on the physical asset can validate and enrich the digital twin. Inversely, the digital representation can bring value to the physical asset. Dedicated digital twin encapsulations can be extracted to model a specific set of behaviors in a specific context, delivering a stand-alone executable representation referred to as the Executable Digital Twin. The process and basic methodologies of the Executable Digital Twin are reviewed. Key enabling technologies are fast simulation methods, Model Order Reduction and state estimation. Data reduction approaches include the use of machine learning to allow compact representations of complex non-linear systems. Hybrid co-simulation linking the executable models in an open runtime environment are key. The concept is illustrated by several examples in the automotive, aerospace and manufacturing industries. This in-

cludes embedded models for virtual sensing, model-based control, performance monitoring and X-in-the-loop hybrid testing. Also, the potential to use the Executable Digital Twin as a companion to the physical asset through its lifecycle for decision support is discussed. An outlook on open challenges will identify further research needs.

Herman Van der Auweraer
KU Leuven
herman.vanderauweraer@siemens.com

Dirk Hartmann
Siemens AG
Corporate Technology
hartmann.dirk@siemens.com

Wim Desmet
KU Leuven, Department of Mechanical Engineering
wim.desmet@kuleuven.be

MS186

Virtual Inspection of 3D Printing via Process-Scale Digital Twins and Computer Vision

In the age of digitalization, physical objects are increasingly being represented by digital twins. Our work at LLNL focuses on developing digital twins of additive manufacturing (AM) systems and processes with the goal of virtually designing and optimizing 3D printed components before needing to manufacture them. Currently, we manufacture mission-critical components from direct-ink-write AM systems and manually inspect and test parts. Using on-board sensor data from cameras and kinematic motion data, our work aims to automate the inspection and testing process by geometrically tolerancing parts with deep learning and computer vision. Furthermore, we aim to improve the accuracy of our process-scale digital twin of the 3D printing process by injecting real sensor noise into the simulator. We also compare our digital and physical twins in a virtual reality (VR) environment, allowing for intuitive interactions among multiple collaborators who need not be co-located. Using in situ process monitoring, data-driven prediction enhancement, and VR, we hope to yield insights into machine health, preventative maintenance, and corrective strategies with the goal of moving toward AM process certification. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, and funded by the Laboratory Directed Research and Development Program at LLNL under project tracking code 23-SI-003.

Brian Weston
University of California, Davis
Lawrence Livermore National Laboratory
weston8@llnl.gov

MS187

Lifting Transformations and One-Sided Model Reduction for Classes of Nonlinear Dynamical Systems

The problem of approximating and equivalently formulating generic nonlinear systems with analytic nonlinearities is studied. We are concerned with the study of lifting nonlinear systems to a polynomial (i.e., quadratic-bilinear, in short QB) structure, and with applying model order reduction (MOR) techniques for lowering the complexity (number of states) of such reformulated systems. Lifting tech-

niques have been successfully applied in recent years for MOR purposes, starting with the works of Gu, Benner, Breiten, Goyal (and collaborators), and more recently, with those of Willcox, Peherstorfer, Kramer, Qian (and collaborators). Systems with bilinear and quadratic dynamics can be used as surrogates for systems with more general nonlinearities. By means of lifting transformations, one is able to reformulate the original nonlinear dynamics into such a format. Then, the problem of MOR for large-scale lifted QB models arises. In particular, we propose a data-driven approach that requires samples of multivariate input-output mappings (in the frequency domain). The latter are (symmetric) generalized transfer functions (symGTFs), which are appropriately defined for QB systems. The MOR approach is data-driven, moment-matching-based, and it requires evaluations of the symGTFs. We show how to (explicitly) fit only a few sample values thereof, and also how to use the available degrees of freedom to derive QB ROMs, in a non-intrusive way. Finally, we present two numerical applications.

Ion Victor Gosea

Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

MS187

Nonlinear Dynamics in Gas Transport Networks - Are They Polynomial?

Gas flow in pipes is typically modeled via isothermal Euler equations. The nonlinear friction term is a product including an absolute value. Rewriting this into a polynomial setting is nontrivial. However the system can be interpreted as a switched quadratic system with a bit of a state dependent switching signal. We investigate quadratic model order reduction for switched systems in this context.

Sara Grundle

Max Planck Institute for Dynamics of Complex Technical Syst
grundel@mpi-magdeburg.mpg.de

MS187

Reachability of Weakly Nonlinear Systems Using Carleman Linearization

In this presentation we introduce a solution method for a special class of quadratic initial-value problems using set-propagation techniques. We employ a particular embedding (Carleman linearization) to leverage recent advances of high-dimensional reachability solvers for linear ordinary differential equations based on the support function. Using a global error bound for the Carleman linearization abstraction, we are able to describe the full set of behaviors of the system for sets of initial conditions and in dense time.

Christian Schilling

Aalborg University, Denmark
christianms@cs.aau.dk

MS187

Rewriting Almost Every Function As the Solution of Quadratic IVODEs

In this talk, we present an approach to recast non-polynomial IVODEs (e.g. elementary and piecewise differentiable) into polynomial differential systems that exploits auxiliary variables, and provide several examples. We also

discuss an *a priori* error bound for the solution to polynomial IVODEs of polynomial form. If there is time, we will also demonstrate an approach to recast a system of polynomial IVODEs into a single variable higher order IVODE.

Roger Thelwell

James Madison University, USA
thelwerj@jmu.edu

MS188

Optimal-Complexity Multigrid Solvers for High-Order FEM in the De Rham Complex

The Riesz maps of the L^2 de Rham complex frequently arise as subproblems in the construction of fast preconditioners for more complicated problems. In this work we present multigrid solvers for high-order finite element discretizations of these Riesz maps with optimal complexity in polynomial degree, i.e. with the same time and space complexity as sum-factorized operator application. The key idea of our approach is to build new finite elements for each space in the de Rham complex with orthogonality properties in both the L^2 and $H(d)$ inner products ($d \in \{\text{grad, curl, div}\}$) on the reference hexahedron. The resulting sparsity enables the fast solution of the patch problems arising in the generalised Arnold–Falk–Winther and Hiptmair space decompositions, in the separable case. In the non-separable case, the method can be applied to a spectrally-equivalent auxiliary operator. With exact Cholesky factorizations of the sparse patch problems, the application complexity is optimal but the setup costs and storage are not. We overcome this with the use of incomplete Cholesky factorizations with carefully specified sparsity patterns arising from static condensation. This yields multigrid relaxations with computational complexity and storage that are both optimal in the polynomial degree.

Patrick E. Farrell

Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

Pablo Brubeck

University of Oxford
brubeckmarti@maths.ox.ac.uk

MS188

Quasi-Interpolation for the Helmholtz-Hodge Decomposition

We propose a computationally efficient and stable quasi-interpolation based method for numerically computing the Helmholtz-Hodge decomposition of a vector field. To this end, we first explicitly construct a matrix kernel in a general form from polyharmonic splines such that it includes divergence-free/curl-free/harmonic matrix kernels as special cases. Then we apply the matrix kernel to vector decomposition via a convolution technique together with the Helmholtz-Hodge decomposition. More precisely, we show that if we convolve a vector field with a scaled divergence-free (curl-free) matrix kernel, then the resulting divergence-free (curl-free) convolution sequence converges to the corresponding divergence-free (curl-free) part of the Helmholtz-Hodge decomposition of the field as the scale parameter tends to zero. Finally, by discretizing the convolution sequence via certain quadrature rule, we construct a family of (divergence-free/curl-free) quasi-interpolants (defined both in the whole space and over a bounded domain) for approx-

inating divergence-free/curl-free part corresponding to the Helmholtz-Hodge decomposition of the field, respectively.

Nicholas Fisher

Minnesota State University, Mankato
nicholas.fisher@mnsu.edu

MS188

Discretization of the Incompressible Navier-Stokes Equations Based on the De Rham Complex

We are interested in the numerical resolution of the Navier-Stokes equations using the tools of exterior calculus and the De Rham complex. The use of the De Rham complex not only allows us to take advantage of the numerous results on its discretization [Arnold, Finite Element Exterior Calculus, 2018] but also of the naturalness of the construction to create structure preserving schemes. We present here an expression of the equations in the exterior calculus formalism and how the desired properties on the scheme guide the different choices made during the discretization. Finally we look more concretely at the resulting scheme by showing the well-posedness of the problem, by analyzing the error and the preservation of the quantities, and with a numerical simulation.

Marien-Lorenzo Hanot

Institut Montpelliérain Alexander Grothendieck
marien-lorenzo.hanot@umontpellier.fr

MS188

Background Independent Formulations of Conservation Laws in Structured Solids

The approximation of materials as continuous solids leads to the classical formulations of conservation of scalar (mass, energy, charge) and vector (linear and angular momentum) quantities in terms of their densities. Such formulations do not allow for describing localisations of properties and processes by underlying material structures. One approach to address the deficiency of the classical formulations is to consider explicitly the internal structures of engineering materials as assemblies of discrete, finite cells with different apparent dimensions. Calculus on such cell complexes was born at the beginning of this century. It offers great opportunities for revisiting and reformulating the descriptions and analyses of real materials. The talk will show the development of discrete analogues of the conservation laws relevant to the analysis of materials with complex internal structures. These are based on the notions of combinatorial vector fields and differential forms on cell complexes. Topological, i.e., metric-independent, as well as metric-dependent operation of such forms will be discussed, together with the strategy that combines them into intrinsic, i.e., background-independent, descriptions of the conservation laws.

Andrey Jivkov, Kiprian Berbatov

The University of Manchester
andrey.jivkov@manchester.ac.uk,
kiprian.berbatov@postgrad.manchester.ac.uk

Pieter Boom

King Fahd University of Petroleum & Minerals
pieter.boom@kfupm.edu.sa

Andrew Hazel

The University of Manchester

andrew.hazel@manchester.ac.uk

MS188

Towards Conforming Discretizations for Bending Shell Problems

We consider numerical simulations of large-deformation bending plates. These are models of thin mechanical objects (such as sheets of paper) that are constraint to deform isometrically. Previous research has focused on non-conforming approaches. As an alternative, we formulate the model as a mixed problem where one variable takes its values in a Stiefel manifold. A discretization by nonlinear Discrete Kirchhoff Triangle (DKT) finite elements then allows to interpret the discrete problem as an optimization problem on a manifold. For such problems, efficient solvers can be constructed and shown to converge. We explain the geometric construction and show numerical tests.

Oliver Sander

Institut fuer Numerische Mathematik
TU Dresden
oliver.sander@tu-dresden.de

MS189

High-Performance Matrix Computations: We Need More Than Fast Libraries

As matrix computations are the bottleneck in countless workflows and applications, many excellent libraries have been developed. These offer a broad range of computational building blocks, optimized for different aspects such as computing platform, problem size, precision, and matrix properties. However, in spite of the quality of such libraries, we observe that end users are increasingly less likely to use them, at least directly. In fact, realistic workflows are significantly more complex than the functionality of the kernels offered in such libraries, and the intelligent decomposition of a workflow in terms of a set of available kernels is in itself a challenging task. We refer to this task as "Linear Algebra Mapping Problem" (LAMP). In practice, the problem is often circumvented by adopting high-level languages such as Matlab, Python, and R, or C++ (in combination with libraries such as Armadillo or Eigen), which offer a convenient high-level syntax for matrix computations, thus boosting user productivity; by contrast, these languages are still immature with respect to the solution of the LAMP, and in terms of performance they suffer from vastly suboptimal choices. In a nutshell: High-performance libraries are a necessary, but not sufficient component for high-performance matrix computations.

Paolo Bientinesi

Umeå University
pauldj@cs.umu.se

MS189

High Performance Domain-Specific Language for Probabilistic Programming

In this talk, we will go through the design and implementation of the DynamicPPL library (the modelling language for Turing.jl). DynamicPPL is a high-performance, general-purpose domain-specific language for probabilistic programming. It allows applied scientists to specify probabilistic models using an intuitive syntax. The same model program can be transformed into many forms correspond-

ing to various statistical operations, such as sampling from the prior, evaluating the prior density, evaluating the log-likelihoods and evaluating the joint log density. These different model forms significantly reduce boiler-and-plate code for statistical programming, enhance communication, and accelerate the speed of iterating modelling workflows. DynamicPPL has been used in some critical settings such as Covid-19 modelling, and pharmaceutical modelling.

Hong Ge

University of Cambridge
hg344@cam.ac.uk

MS189

Data Layout in Domain-Specific Compilers and Code Generators

Domain-specific compiler technology bridges between the domain scientist's expression of their problem and architecture-specific code generation. In this talk, we will describe how domain-specific code generation is facilitated by a data layout description that is co-optimized with the computation. We will examine how to generate performance-portable code with data layouts designed to minimize data movement for the application domain.

Mary Hall

School of Computing
University of Utah
mhall@cs.utah.edu

MS189

Halide: Growing Dopsies from Research Projects to Industrial Adoption

While research into domain-specific programming systems ("dopsies") continues to expand, only a few of these systems have matured enough to find adoption outside of small research communities. One notable exception is Halide [Ragan-Kelley et al, SIGGRAPH 2012; Ragan-Kelley et al, PLDI 2013], a programming language and compiler for dense computations on images and tensors. Used in multiple shipping products, including Adobe Photoshop and Google Photos, Halide presents an interesting case study of a dopsy used for production while still forming the basis of continuing research projects. This talk outlines the history of Halide's adoption, the design decisions that enable this adoption, and the challenges that come with simultaneously supporting multiple production use cases across multiple (sometimes competing) companies.

Shoaib Kamil

Adobe Systems Inc
kamil@adobe.com

MS189

An Extension of C++ with a Memory-Centric Specification Language for HPC

The C programming language and its cousins such as C++ lean towards a memory-inefficient storage of structs: developers have to commit to one, invariant data model unless they rearrange the data structures manually, the natural data structure for object-based models is array of structs, and the compiler has to insert bits such that individual attributes align with byte boundaries. Furthermore, the language provides no native support for data exchange via MPI and precision formats beyond the IEEE standard.

We propose a language extension based upon C++ attributes which give developers more fine-grained control over the memory implementation of data structures: numerical types tailored to their value ranges with bit-level granularity, automatic bit packing of boolean and enumerable fields, customisable reduced-precision storage types for floating-point data, as well as local declarative conversions between Array-of-Structures to Structure-of-Arrays layouts for "for" and "for-each" loops. Finally, we propose attribute-driven MPI mappings generation for arbitrary data types, including the new data types that we offer. Our proposals are realised via a compiler-based approach by extending the Clang/LLVM compiler toolchain. In the presentation we demonstrate the capabilities of our extensions, as well as show the performance impacts on a set of large-scale smoothed-particle hydrodynamics simulation (SPH) codes.

Pawel Radtke

Computer Science, Durham University
pawel.k.radtke@durham.ac.uk

Tobias Weinzierl

Durham University
tobias.weinzierl@durham.ac.uk

MS190

Network-Based Framework for Estimating the Risk Factors of Epidemic Spreading

Network-based adaptations of traditional compartmental infection models such as SIR or SEIR can be used to model the spreading of diseases between cities, countries or other geographical regions. One of the common challenges arising in these applications is the lack of available transmission probabilities between these geographical units. Several methods have been proposed recently for solving this task. One of them is the Generalized Inverse Infection Model (GIIM) [A. Bta, L. M. Gardner: A generalized framework for the estimation of edge infection probabilities. arXiv:1706.07532]. GIIM offers a large amount of modelling flexibility and allows transmission probabilities to be defined as a function of known attributes, or risk factors in an epidemic context. In this presentation we will see how GIIM works in two specific real-life outbreaks. The first one considers the 2015-2016 Zika virus outbreak in the Americas [L. M. Gardner, A. Bta, N. D. Grubaugh, K. Gangavarapu, M. U. G. Kramer: Inferring the risk factors behind the geographical spread and transmission of Zika in the Americas. PLoS Neglected Tropical Diseases 12 (1), e0006194 (2018)], while the second application models the 2009 H1N1 outbreak between the municipalities of Sweden [A. Bta, M. Holmberg, L. M. Gardner, M. Rosvall: Socio-economic and environmental patterns behind H1N1 spreading in Sweden. Scientific Reports 11 (1), 1-14, 2021.]

Andras Bota

Lulea University of Technology
andras.bota@ltu.se

Lauren Gardner

Johns Hopkins University
l.gardner@jhu.edu

Martin Holmberg, Martin Rosvall
Umea University

mhmarhol@gmail.com, martin.rosvall@umu.se

MS190

Spatio-Temporal Dynamics of Disinformation with Committed Minorities

The spread of disinformation through social networks has led to many negative consequences, including manipulation of the 2016 presidential election, vaccine hesitancy during the COVID-19 pandemic, and the rise of QAnon. Disinformation campaigns with the presence of a committed minority are of particular concern because a minority with as low as 10% of the population can overtake the majority's opinion. This phenomenon is characterized by a strong tipping point (phase transition) separating two distinct outcomes. Previous studies employed the binary agreement model on Erdős-Rényi random, scale-free, and complete networks. Here, we examine the binary agreement model on hundreds of random graphs of several types (e.g., barbell, small world, etc.) to connect the tipping point to graph metrics (e.g., assortativity, betweenness, etc.). We also explore the dynamics on Facebook networks. The binary agreement model is run to steady state on all networks utilizing hundreds of core hours. Using the data generated in these simulations, we calculate correlation functions of individuals' opinions and calculate scaling exponents of the phase transitions. Gaining insights into the nature of the tipping points suggests strategies for mitigating the spread of disinformation.

David J. Butts, Sam Bollman, Michael Murrillo
Michigan State University
buttsdav@msu.edu, bollma13@msu.edu, murril-
lom@msu.edu

MS190

Optimal Link Scheduling via Subgraph Coloring

In datacenters, optical technologies allow the establishment of a limited set of additional direct connections of high bandwidth between pairs of top-of-rack switches. To maximize throughput, it is desirable to establish direct connections between those pairs that currently have a high communication demand and to reconfigure the links whenever the demand changes. We show how the problem of determining an optimal configuration can be cast as a subgraph coloring problem and study its complexity and various algorithmic solutions both in theory and practice.

Kathrin Hanauer
University of Vienna
kathrin.hanauer@univie.ac.at

MS190

Degree Sequence Problems in NMR Spectroscopy

The ¹H-NMR spectroscopy can measure how many hydrogen atoms a carbon atom can bind, furthermore, how many hydrogen atoms the neighbor carbon atoms altogether can bind. Each carbon atom can make four covalent bonds. Therefore, the number of neighbor carbon atoms to a particular carbon atom is four minus the number of its hydrogen atoms. Furthermore, the number of second neighbor carbon atoms is three times the number of neighbor carbon atoms minus the number of those hydrogen atoms that the neighbor carbon atoms bond. That is, from the ¹H-NMR spectroscopy, we can obtain the degree and neighbor degree sequences of hydrocarbons. When these hydrocar-

bons are saturated and acyclic (these properties are easy to check by simple, fast and cheap lab tests), the chemical structure prediction is equivalent to constructing bounded degree trees with prescribed degree and neighbor degree sequences. We will present a polynomial running time algorithm to the construction problem and we will also provide a polynomial delay enumeration algorithm to enumerate all possible solutions.

Istvan Miklos
Renyi Institute
miklosi@renyi.hu

MS190

A Clique Based Algorithm for Scheduling Coupled Tasks with Delay

The clique search algorithms has improved to such an extent that today we may compute clique numbers of highly non-trivial size graphs. These algorithms are even more efficient in the particular case when one wants to locate k -cliques in k -partite graphs. In this work we single out a scheduling problem in which tasks are coupled and the time delay between the first and second member of the couple is fixed by technological constraints. We will show that this scheduling problem can be reduced to the question to decide if a tactically constructed k -partite auxiliary graph contains a k -clique. The scheduling problems can be reformulated as an integer linear program and can be handled using powerful solvers. We argue that if the clique version of the scheduling problem is not overly large and the existing clique solvers can handle it, then the clique approach has certain advantages. These are the following. The clique solver does not use floating point arithmetic and so it is rounding error free. The clique reformulation technique is flexible enough to incorporate extra conflicts between tasks in case they would arise in practice. We will point out that before submitting the auxiliary graph to a clique solver it is useful to carry out various inspections in order to delete nodes and edges of the graph. In the lack of theoretical tools we will carry out numerical experiments to test the practicality of the clique approach.

Sandor Szabo
Institute of Mathematics and Informatics
University of Pecs
sszabo7@hotmail.com

MS191

Exposing AMR Application Interfaces for Time Integrators

We will discuss software design and lessons learned from our experience implementing SUNDIALS time integrators in the Chombo-based BISICLES ice sheet code. The C++-based software "adapter" was designed to be flexible enough for a variety of adaptive mesh refinement (AMR) applications, with abstractions that allow rapid prototyping and trying different time integration schemes. We will review some of the scientific software practices we used (from math and design, to testing and validation) in the process, then demonstrate the resulting flexibility and other design benefits on model problems and real ice sheet applications from BISICLES. We hope to provide this as a case study in mathematical software integration that others can benefit from.

Steven Roberts
Lawrence Livermore National Laboratory

roberts115@llnl.gov

Hans Johansen

Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

Daniel Martin

Lawrence Berkeley National Laboratory
dfmartin@lbl.gov

David J. Gardner, Carol S. Woodward
Lawrence Livermore National Laboratory
gardner48@llnl.gov, woodward6@llnl.gov

MS191

Overview of Fastmath Research for High-Performance Computing

In this presentation, I will give an introduction to the FASTMath Institute and an overview of our research activities in applied mathematics and high-performance numerical software for scientific applications. I will highlight some of the impacts that have been made on applications.

Todd Munson

Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

MS191

AMReX: An Overview of Recent Software Developments and Scientific Applications

In this talk I will give an overview of current features and scientific partnerships involving the AMReX software framework. AMReX is a key technology in the FASTMath structured mesh area, supporting block-structured adaptive mesh refinement algorithms on exascale systems, including native support for hyperbolic and parabolic/elliptic PDEs, subcycling approaches, particle-mesh algorithms, and embedded boundary representation of geometry. The framework is also interoperable with other FASTMath technologies including SUNDIALS and HYPRE. AMReX is the basis for numerous mature application codes across academia and industry; here I will detail some newly-funded SciDAC partnerships with Biological and Environmental Research (BER) and High Energy Physics (HEP) in applications such as atmospheric modeling, ocean modeling, particle accelerators, and cosmology.

Andrew Nonaka, Ann S. Almgren
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov, asalmgren@lbl.gov

John B. Bell

CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Weiqun Zhang

Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

Andrew Myers

Lawrence Berkeley National Lab

atmyers@lbl.gov

Jean M. Sexton

Lawrence Berkeley National Laboratory
jmsexton@lbl.gov

MS191

Adaptive Time Integrators and Their Use in Non-Equilibrium Quantum Simulations

Adaptive time integrators have long shown significant efficiency benefits over fixed step integrators in scientific applications. Despite these efficiency benefits, they are still not used in many application areas, including non-equilibrium quantum dynamics. In this work we explore the use of adaptive step integration methods for the Kadanoff-Baym equations (KBE) which give Greens function information for a quantum system. Because of the need for including the full time history in the solution of this system of two-time integral equations, constant step size methods have traditionally been used. This presentation will overview adaptive time integrator technologies and discuss their implementation within the FASTMath time integration software suite, SUNDIALS. We will then show the potential benefits of these adaptive methods for the KBE system and discuss targeted extensions within the SUNDIALS CVODE integrator to provide adaptive methods for two-time systems.

Thomas Blommel
University of Michigan
tblommel@umich.edu

David J. Gardner
Lawrence Livermore National Laboratory
gardner48@llnl.gov

Emanuel Gull
University of Michigan
egull@umich.edu

Carol S. Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

MS192

A Study of Prediction Performances of Econometric Arima and Soft Computing Ann Models in Forecasting Black Carbon Concentration Data

Fresh air is imminent for life and to thrive on this planet. However, this vital component of life is ill-effected by fast-paced industrialization, urbanization, automobiles, factories and coal-based thermal power generation as over the years these have jeopardized the air quality index. Due to hazardous impact of black carbon on the environment as well as human health, researchers have turned their attention towards its study. The comparative analysis of artificial neural network and econometric models used to predict the time series of monthly observations of black carbon emissions from three major coal mines located at Bokaro, Jharia and Raniganj in India. These coal mines have large emissions of particulate matter () that contribute significantly to pollution levels. A multilayer perceptron feedforward artificial neural network is used to predict the black carbon concentration data from these three coal mines. The neural network is trained using a Bayesian regularization backpropagation algorithm. The efficiency

of neural network models is evaluated by mean absolute deviation (MAD), root mean square (RMSE), and coefficient of determination (R^2) values. The results obtained are then compared with the well-known and widely used autoregressive integrated moving average (ARIMA) model. The results of the study reveal the effective performance of artificial neural networks over the ARIMA model in forecasting black carbon concentration data.

Kulwinder Singh, Jatinder Kaur
IITK Gujral Punjab Technical University
kulmaths@gmail.com, jatinderkaulkh@gmail.com

Sarbjit Singh
Guru Nanak Dev University
sarbasr@gmail.com

Kirti Soni
CSIR- Advanced Materials And Process Research
Institute (AM)
2006.kirti@gmail.com

MS193

Meta Net Ann: a New Fusion Algorithm

The fundamental characteristic of the Meta-Net consists of considering not only the positive credibility of its composing classifiers (i.e., this pattern is white), but also their negative credibility (i.e., this pattern is not white). So, the characterizing connection of the Meta-Net is to connect each output node of each composing classifier with each output class. Complete grid connections are planned between Meta-Net inputs and outputs, and each connection can be either excitatory (positive numbers), or inhibitory (negative numbers).

Massimo Buscema
Semeion
University of Colorado Denver
m.buscema@semeion.it

MS193

Deep Learning of Systems with Physical Structures

We propose a data-driven method to learn the dynamics of systems with physical structures using deep neural networks. We incorporate the physical structures into the design of the networks and show that such models can predict dynamics obeying desired physical laws of the system. To demonstrate our method, we focus on conservation laws in partial differential equations and derive conservative form networks. Numerical results show that models with physical structures outperform models without ones in both accuracy and physical property.

Zhen Chen
Dartmouth College
zhen.chen@dartmouth.edu

Yoonsang Lee, Anne Gelb
Dartmouth College
Department of Mathematics
yoonsang.lee@dartmouth.edu,
negelb@math.dartmouth.edu

MS193

Graph Construction in Multi-Fidelity Networked

Surrogates

This talk discusses graph-inference methodologies for MFNets – a paradigm for multifidelity information fusion via directed acyclic graphs. MFNets provide a flexible approach to modeling the relationships between unstructured ensembles of models and information sources by linking the outputs of each information source through a network of models. These structures can then be leveraged for forward and inverse problems. One of the challenges in this approach is determining good graphs to represent the information sources. In this talk we discuss and overview several graph-building strategies and apply them to MFNets. Specifically, we compare graph incremental construction, model selection, and model averaging. We take a probabilistic viewpoint for assessing and developing graph construction algorithms and investigate how the inner inference loops affect the resulting outcome. For example, we determine how posteriors based on variational inference affect the resulting graphs. Examples from both synthetic and physical models are provided.

Alex Gorodetsky
University of Michigan
goroda@umich.edu

MS193

Adaptive Random Fourier Features with Metropolis Sampling

The supervised learning problem to approximate a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ by a neural network approximation $\mathbb{R}^d \ni x \mapsto \sum_{k=1}^K \hat{\beta}_k e^{i\omega_k \cdot x}$ with one hidden layer is studied as a random Fourier features algorithm. Here the mean square loss problem can be solved easily, since it is convex in the amplitude parameters $\hat{\beta}_k$ given a density $p : \mathbb{R}^d \rightarrow [0, \infty)$ for independent frequencies ω_k . It is also well known that the corresponding generalization error is bounded by $K^{-1} \|\hat{f}\|^2 / ((2\pi)^d p) \|_{L^1(\mathbb{R}^d)}$, where \hat{f} is the Fourier transform of f . In my talk I will first show how the constant $\|\hat{f}\|^2 / ((2\pi)^d p) \|_{L^1(\mathbb{R}^d)}$ can be minimized by optimally choosing the density p and then how to approximately sample from this density, only using the data and certain adaptive Metropolis steps. I will also show results with other activation functions. [1] Kammonen, Aku and Kiessling, Jonas and Plechac, Petr and Sandberg, Mattias and Szepessy, Anders. Adaptive random Fourier features with Metropolis sampling. Foundations of Data Science, 2020. [2] Kammonen, Aku and Kiessling, Jonas and Plechac, Petr and Sandberg, Mattias and Szepessy, Anders and Tempone, Raul. Smaller generalization error derived for a deep residual neural network compared with shallow networks. IMA Journal of Numerical Analysis, 2022.

Aku Kammonen
KAUST
akujaakko.kammonen@kaust.edu.sa

Anders Szepessy
Department of Numerical Analysis and Computer Science
Royal Institute of Technology, Stockholm
szepessy@kth.se

Petr Plechac
University of Delaware
Department of Mathematical Sciences
plechac@math.udel.edu

Jonas Kiessling

KTH Royal Institute of Technology
jonas.kiessling@gmail.com

Mattias Sandberg
Department of Mathematics
Royal Institute of Technology, Stockholm
msandb@kth.se

MS193

Modeling Spatio-Temporal Processes in Climate Models via Functional Tensor Networks

We present a flexible framework for capturing high-dimensional non-linear interactions in computationally expensive models. In this talk we will use the land model component of E3SM to drive algorithm developments. Specifically we will rely on functional tensor networks to construct surrogate models for the land model dynamics at several observation sites. We will compare the performance of several network topologies to capture the interactions between model components and will use the resulting spatio-temporal surrogates to extract parameter sensitivity indices and for subsequent model calibration studies.

Cosmin Safta
Sandia National Laboratories
csafta@sandia.gov

John D. Jakeman
Sandia National Labs
jadjakem@sandia.gov

Alex Gorodetsky
University of Michigan
goroda@umich.edu

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

MS194

Deep Neural Network Approximations for Solutions of PDEs Based on Monte Carlo Algorithms

In the past few years deep artificial neural networks (DNNs) it has been proposed in the scientific literature to reformulate high-dimensional partial differential equations (PDEs) as stochastic learning problems and to employ DNNs together with stochastic gradient descent methods to approximate the solutions of such high-dimensional PDEs. One key argument in most of these results is, first, to employ a Monte Carlo approximation scheme which can approximate the solution of the PDE under consideration at a fixed space-time point without the curse of dimensionality and, thereafter, to prove then that DNNs are flexible enough to mimic the behaviour of the employed approximation scheme. Having this in mind, one could aim for a general abstract result which shows under suitable assumptions that if a certain function can be approximated by any kind of (Monte Carlo) approximation scheme without the curse of dimensionality, then the function can also be approximated with DNNs without the curse of dimensionality. The main result of this work, roughly speaking, shows that if a function can be approximated by means of some suitable discrete approximation scheme without the curse of dimensionality and if there exist DNNs which satisfy certain regularity properties and which approximate this discrete approximation scheme without the curse of

dimensionality, then the function itself can also be approximated with DNNs without the curse of dimensionality.

Phillip Grohs
University of Vienna
philipp.grohs@univie.ac.at;

MS194

The Kolmogorov Superposition Theorem can Break the Curse of Dimensionality When Approximating High Dimensional Functions

We explain how to use Kolmogorov's Superposition Theorem (KST) to overcome the curse of dimensionality in approximating multi-dimensional functions and learning multi-dimensional data sets by using neural networks of two layers. That is, there is a class of functions called K-Lipschitz continuous in the sense that the K-outer function of is Lipschitz continuous can be approximated by a ReLU network of two layers with widths to have an approximation order $\mathcal{O}(d^2/n)$. In addition, we show that polynomials of high degree can be expressed by using neural networks with activation function $\sigma_l(t) = (t_+)^l$ with $l \geq 2$ with multiple layers and appropriate widths. More layers of neural networks, the higher degree polynomials can be reproduced. Hence, the deep learning algorithm can well approximate multi-dimensional data when the number of layers increases with high degree activation function σ_l . Finally, we present a mathematical justification for image classification by using a deep learning algorithm.

Ming-Jun Lai
University of Georgia
Department of Mathematics
mjlai@math.uga.edu

Zhaiming Shen
University of Georgia
zhaiming.shen@uga.edu

MS194

pde-Learn: Using Deep Learning to Discover Partial Differential Equations from Noisy, Limited Data

Scientific progress is contingent upon finding predictive models for the physical world. In this paper, we introduce PDE-LEARN, a novel PDE-discovery algorithm that can identify PDEs directly from noisy, limited measurements of a physical system governed by a hidden PDE. PDE-LEARN uses a Rational Neural Network, U , to approximate the system response function and a sparse, trainable vector ξ to characterize the hidden PDE. Our approach couples the training of U and ξ using a specially designed loss function that (1) makes U approximate the system response function, (2) encapsulates the fact that U satisfies a hidden PDE that ξ characterizes, and (3) promotes sparsity in ξ using ideas from iteratively reweighted least-squares. This approach yields a robust algorithm that applies to many physical systems. We demonstrate the efficacy of PDE-LEARN by identifying several PDEs from noisy and limited measurements.

Robert R. Stephany
Cornell University

rrs254@cornell.edu

MS194

A Non-intrusive Domain-Decomposition Model Reduction Method for Linear Steady-State Partial Differential Equations with Random Coefficients

Domain decomposition methods have been proved to be an effective strategy to reduce the dimension of parametric partial differential equations (PDEs). However, existing domain decomposition methods for parametric PDEs are usually intrusive, which means domain decomposition based solvers need to be implemented from scratch for each target parametric PDE. To address this issue, we develop a new non-intrusive domain-decomposition model reduction method for linear steady-state PDEs with random-field coefficients. As a variant of our previous work by Mu and Zhang, the new method only needs access to the final linear system, that is, the global stiffness matrix and the right hand side, of a deterministic PDE solver, in order to build a domain-decomposition-based reduced model without intrusive implementation from scratch. The key idea is to remove the interface condition between sub-domains and rely on the correlation between columns of the linear system to couple the sub-domains. The non-intrusive feature enables the applicability of the proposed method to a broader class of uncertainty quantification problems, where many legacy codes/solvers can be fully reused by our method. Two numerical examples including diffusion equations with random diffusivity and convection-dominated transport with random velocity, are provided to demonstrate the effectiveness and efficiency of our method.

Guannan Zhang

Oak Ridge National Laboratory
zhangg@ornl.gov

Lin Mu

University of Georgia
linmu@uga.edu

MS195

Reduced-Order Modeling Using Hybrid Operator Inference for Nonlinear Thermo-Mechanics

Simulating large-scale nonlinear thermo-mechanics systems is crucial in many industrial problems. Hence, compact modeling is required. However, many challenges arise when dealing with nonlinearities. Typically, commercial FEM code does not allow the user to extract the operators describing the dynamics. Instead, only solution trajectories are available. As a result, nonlinear model reduction is often achieved using data-driven techniques. However, complex physical characteristics are usually not preserved with solely data-based approaches. Therefore, in this work a hybrid operator inference-based approach is proposed which combines both physics-based and data-based techniques. With this approach, it is demonstrated how nonlinear material properties of thermo-mechanical systems can be efficiently reproduced in a reduced-order model.

Pascal Den Boef

Eindhoven University of Technology
p.d.boef@tue.nl

Diana Manvelyan

Siemens AG
diana.manvelyan@siemens.com

Jos Maubach, Wil Schilders, Nathan van de Wouw
Eindhoven University of Technology

j.m.l.maubach@tue.nl, w.h.a.schilders@TUE.nl,
n.v.d.wouw@tue.nl

MS195

Symplectic Model Reduction of Hamiltonian Systems on Nonlinear Manifolds

Classical model reduction techniques project the governing equations onto linear subspaces of the high-dimensional state-space. However, for problems with slowly decaying Kolmogorov-n-widths such as certain transport-dominated problems, classical linear-subspace reduced order models (ROMs) of low dimension might yield inaccurate results. Thus, the reduced space needs to be extended to more general nonlinear manifolds. Moreover, as we are dealing with Hamiltonian systems, it is crucial that the underlying symplectic structure is preserved in the reduced model. To the best of our knowledge, existing literatures addresses either model reduction on manifolds or symplectic model reduction for Hamiltonian systems, but not their combination. In this talk, we bridge these two approaches by providing a novel projection technique called symplectic manifold Galerkin, which projects the Hamiltonian system onto a nonlinear symplectic trial manifold such that the ROM is again a Hamiltonian system. We provide numerical results which demonstrate the ability of the method to outperform linear-subspace ROMs. high-level commands. Do not include keywords, references or citations separately at the end of the abstract. All citations must be within the abstract text in general form "[Authorname, Title, etc]." Improper citations will be deleted.

Silke M. Glas

University of Twente
s.m.glas@utwente.nl

Patrick Buchfink

Universitaet Stuttgart
patrick.buchfink@ians.uni-stuttgart.de

Bernard Haasdonk

University of Stuttgart
haasdonk@mathematik.uni-stuttgart.de

MS195

Model Order Reduction for Large-Scale Seismic Wave Propagation

Several applications require repetitive numerical solutions of large-scale wave-type equations to resolve an inverse problem or produce a sizeable training dataset. Such numerical computations for multiple parameters are computationally expensive. A suitable approach to tackle such problems is model order reduction (MOR), where we solve the problem in an appropriate subspace of the high-dimensional discretization space. Standard MOR methods generally exhibit several issues for wave-type problems, including reducibility and potentially weak stability. We present a goal-oriented MOR approach for wave-type problems with band-limited output of interests (e.g., in seismology) or damping (e.g., in structural health monitoring). The proposed method combines numerical Laplace inversion within the Greedy algorithm to identify a parameter with the worst reduced basis (RB) approximation. A POD-based compression is then used in the frequency domain to iteratively enrich the RB space, where we exploit band-limited functions to limit high frequencies and

guarantee rapid convergence of the RB approximation. We also present a goal-oriented error estimator, targeting construction of the RB space to the output of interest. The proposed approach yields an amenable reduction and preserves the accuracy of the RB approximation in both the time and frequency domains. We also present supporting numerical results for the 2D parametric seismic wave equation with semi-realistic data functions.

Muhammad Hamza Khalid
University of Twente
m.h.khalid@utwente.nl

Kathrin Smetana
Stevens Institute of Technology, U.S.
ksmetana@stevens.edu

Matthias Schlottbom
University of Twente
m.schlottbom@utwente.nl

Rhys Hawkins
Utrecht University
r.p.hawkins@uu.nl

MS195

A Multi-Fidelity Ensemble Kalman Filter with Adaptive Reduced-Order Models

The use of model order reduction techniques in combination with ensemble-based methods for estimating the state of systems described by nonlinear partial differential equations has been of great interest in recent years in the data assimilation community. Methods such as the multi-fidelity ensemble Kalman filter (MFEnKF) and the multi-level ensemble Kalman filter (MLEnKF) have been developed and implemented in several papers and are recognized as state-of-the-art techniques. However, the construction of low-fidelity models in the offline stage, prior to solving the data assimilation problem, leads these methods into a trade-off between the accuracy and computational cost of the approximate models. In this work, we investigate the use of adaptive reduced-basis techniques in which the approximation space is modified (but not retrained) online based on the information extracted from the full-order solutions. This has the potential to simultaneously ensure good accuracy and low cost for the employed models and thus improve the performance of the multi-fidelity/multi-level methods.

Francesco Silva, Cecilia Pagliantini, Karen Veroy-Grepl
TU/e Eindhoven University of Technology
f.a.b.silva@tue.nl, c.pagliantini@tue.nl, k.p.veroy@tue.nl

MS195

Registration-Based Model Reduction of Advection-Dominated PDEs with Spatio-Parameter Adaptivity

We present a nonlinear registration-based model reduction procedure for rapid and reliable solution of parameterized two-dimensional steady conservation laws. This class of problems is challenging for model reduction techniques due to the presence of nonlinear terms in the equations and also due to the presence of parameter-dependent sharp gradient regions that cannot be adequately represented through linear approximation spaces. Our approach builds on the following ingredients: (i) a general (i.e., independent of the underlying equation) registration procedure for the compu-

tation of a parametric mapping that tracks moving features of the solution field; (ii) an hyper-reduced least-squares Petrov-Galerkin reduced-order model for the rapid and reliable estimation of the solution field; (iii) a greedy procedure driven by a residual-based error indicator for efficient exploration of the parameter domain; and (iv) an adaptive mesh refinement technique for the definition of an accurate discretization for all parameter values. We present results for a representative nonlinear problem in steady aerodynamics to demonstrate the effectiveness and the mathematical soundness of our proposal.

Tommaso Taddei
INRIA Bordeaux South-West
tommaso.taddei@inria.fr

MS196

Insights into Non-Equilibrium Spectroscopy from Theory

Time-resolved angle-resolved photoemission spectroscopy is one of the most powerful pump-probe measurements of materials driven far from equilibrium. Concomitantly, it is also difficult to interpret and model. In this talk, I will attempt to shed some light on the numerical simulation and the interpretation of time-resolved spectra. First, I will introduce the non-equilibrium Keldysh formalism and the methodology for solving it in an efficient manner. Following that, I will discuss some aspects of interpretation, primarily focusing on building a connection to the underlying equilibrium physics, as this is one of the goals of non-equilibrium experiments. Specifically, I will outline how a number of fundamental relations used in equilibrium break down when going out of equilibrium, where approximate relations hold, and how we can use these ideas to interpret non-equilibrium quantum measurements.

Alexander Kemper
North Carolina State University
akemper@ncsu.edu

MS196

Taming the Dynamical Sign Problem in Diagrammatic Methods for Quantum Dynamics

Numerical simulations for open quantum system dynamics is a profound challenge. In this talk, we will present some recent works on diagrammatic algorithms for open quantum systems. The focus will be an interplay between the dynamical sign problem and error amplification in numerical integration. In particular, our analysis demonstrates that the technique of partial resummation provides a tool to balance these two types of error, and the recently introduced inchworm Monte Carlo method is a successful case to suppress the numerical sign problem. This is joint work with Zhenning Cai and Siyao Yang.

Jianfeng Lu
Mathematics Department
Duke University
jianfeng@math.duke.edu

MS196

Non-Equilibrium Dynamics of Interacting Electrons, Phonons and Excitons from First Principles

Combining density functional theory with kinetic equations has advanced the modeling of ultrafast dynamics in materials out of equilibrium. This talk will focus on recent devel-

opments of this framework, including a numerical approach to evolve in time the coupled Boltzmann equations of electrons and phonons using ab initio interactions, extensions to model exciton interactions and non-equilibrium dynamics, simulations of ultrafast spectroscopies, and data-driven methods. Selected applications to bulk and 2D semiconductors will be presented. I will conclude with an overview of our Perturbo code, an open-source framework to study electron interactions and dynamics in materials, highlighting open problems and future directions.

Marco Bernardi
Caltech
bmarco@caltech.edu

Ivan Maliyov
California Institute of Technology
imaliyov@caltech.edu

MS196

Stochastic Real-Time Green's Function Theory for Neutral Excitations in Molecules

We present a real-time second-order Green's function method (TD-GF2) for computing neutral excitations in molecules and nanostructures. The framework is combined with the stochastic resolution of the identity to decouple the 4-index electron repulsion integrals (ERI) in the system Hamiltonian. This leads to the reduction of the computational cost to $O(N^3)$ with system size. The stochastic implementation recovers deterministic results for the electronic dynamics and excitation energies, and reproduces benchmark results from the analogous linear-response implementation in frequency. This approach is further combined with the Dynamic Mode Decomposition (DMD) technique to predict the nonlinear long-time dynamics of the density matrix. The statistical error due to the incorporation of the stochastic resolution of the identity and DMD extrapolation is analyzed in terms of the number of stochastic orbitals, system size, and propagation time. Overall, this approach offers an efficient route to investigate excited states in finite systems containing hundreds of electrons.

Leopoldo Mejia
Lawrence Berkeley National Laboratory
leopoldo.mejia@lbl.gov

MS197

Mathematical Framework for Basic Freeform Optical Systems

We present one mathematical framework for sixteen basic freeform optical systems. We consider configurations with a parallel or point source and a parallel, point, near-field or far-field target. This gives eight systems, if we use reflectors only and take the minimum number of freeform surfaces required. Similarly, we get eight basic lens systems if we only use freeform lens surfaces. Our goal is to find the shape and location of the freeform surfaces that convert a given source distribution into a desired target distribution. Many papers in freeform illumination optics cover only one particular optical system. Here we derive models based on Hamilton's characteristic functions and conservation of luminous flux for all sixteen systems. Some configurations lead to standard or generalized Monge-Ampere equations. The remaining systems are described by so-called generated Jacobian equations. Our numerical method for finding

the optical surface(s) is an iterative least-squares solver.

Martijn Anthonissen
Eindhoven University of Technology
m.j.h.anthonissen@tue.nl

MS197

The Second Boundary Value Problem for a Discrete Monge-Ampere Equation

In this work we propose a discretization of the second boundary condition for the Monge-Ampere equation arising in geometric optics and optimal transport. The discretization we propose is the natural generalization of the popular Oliker-Prussner method proposed in 1988. For the discretization of the differential operator, we use a discrete analogue of the subdifferential. Existence, unicity and stability of the solutions to the discrete problem are established. Convergence results to the continuous problem are given.

Gerard Awanou
University of Illinois at Chicago
awanou@uic.edu

MS197

Discretization of the Second Boundary Value Problem for the Monge-Ampere Equation Using Voronoi's First Reduction

The Monge-Ampere equation belongs to the class of fully nonlinear degenerate elliptic equations. A general argument often allows to prove the convergence of numerical schemes for such equations, provided that those schemes satisfy a property of monotonicity. I will show how Voronoi's first reduction, a tool originating from the theory of low-dimensional lattice geometry, allows to design a monotone and convergent finite difference scheme for the Monge-Ampere equation. I will also discuss the discretization of the optimal transport boundary condition, which is the relevant boundary condition for the Monge-Ampere equation in many applications, but which prevents the direct application of the theory of monotone numerical schemes to the resulting boundary value problem, due to the fact that the comparison principle does not hold in the usual sense for this problem. Finally, I will illustrate the talk by numerical results obtained in the context of an application to the far-field refractor problem in nonimaging optics.

Guillaume Bonnet
Scuola Internazionale Superiore di Studi Avanzati
gbonnet@sissa.it

Jean-Marie Mirebeau
CNRS and ENS Paris-Saclay
jean-marie.mirebeau@ens-paris-saclay.fr

MS197

Efficient Monotone Discretization of the Monge-Ampere Equation Through Quadrature

The Monge-Ampere equation is a fully nonlinear degenerate elliptic PDE used in a large number of applications including mesh generation, optical design, and medical image processing. The development of convergent numerical methods for this PDE has been guided by the Barles and Souganidis framework, which requires the use of monotone

schemes. However, existing monotone schemes are built using wide stencils, which are low accuracy and expensive to construct - in 3D just evaluating these schemes can be prohibitively expensive. In this talk, I will introduce an integral representation of the Monge-Ampère equation, which leads to a new monotone discretization through the use of numerical quadrature. By using higher order quadrature, we are able to improve on both the formal truncation error and the computational cost of evaluating the scheme, while preserving the monotonicity of the scheme. The result is a provably convergent scheme that allows for narrower stencils and better accuracy than existing methods. Computational results confirm the expected accuracy and efficiency.

Jake Brusca

Department of Mathematical Sciences
New Jersey Institute of Technology
jb327@njit.edu

Brittany Hamfeldt

New Jersey Institute of Technology
brittany.hamfeldt@njit.edu

MS197

Generated Jacobian Equations in Freeform Optical Design

Modern optical components are free-shape reflectors and lenses that transform the light from the LED source into the required light output of the lighting system. The inverse design of freeform optical surfaces involves solving a nonlinear PDE called a generalized Monge-Ampère equation. This PDE can be derived using the laws of geometrical optics and conservation of energy. The laws of geometrical optics define an optical map which connects coordinates on the source to coordinates on the target. Substituting the mapping into the relation for energy conservation leads to a nonlinear second-order PDE for the location of the optical surface. I developed a generic framework to derive the PDEs for a wide range of optical systems (e.g., involving parallel and point light sources with parallel, point and near- and far-field targets). This framework can be applied to optical systems that can be described using a cost function in optimal transport theory. However, many optical systems cannot be cast as optimal-transport problems, e.g., systems involving near-field targets or multiple freeform surfaces. For these systems I generalized the concept of a cost function to a generating function. In these cases, the PDE for the location of the optical surface is a so-called generated Jacobian equation. The PDEs can be solved using a least-squares numerical algorithm. I will present 16 optical systems that all fit into this framework.

Lotte Romijn

Signify
lotte.romijn@signify.com

MS198

High-Order Adaptive Multi-Domain Time Integration Scheme for Microscale Lithium-Ion Batteries Simulations

The modeling of ionic transport and charge conservation in lithium-ion batteries (LIBs) at microscale is a multi-physics problem which involves a wide range of time scales. The associated computational challenges motivate the investigation of numerical techniques that can decouple the

time integration of the governing equations in the liquid electrolyte and the solid phase (active materials and current collectors). First, it is shown that semi-discretisation in space of the governing equations leads to a system of *index-1* differential-algebraic equations (DAEs). Then, a new generation of strategies for multi-domain integration is presented, enabling high-order adaptive coupling in time of both domains. A simple 1D LIB code is implemented as a demonstrator of the new strategy for the simulation of charge, discharge, and relaxation. The integration of the decoupled subsystems is performed with high-order accurate implicit nonlinear solvers. A fully coupled monolithic time integration scheme is also presented as reference. The accuracy of the space discretisation is assessed by comparing numerical results to analytical solutions. Then, temporal convergence studies demonstrate the accuracy of the new multi-domain coupling approach. Finally, considerations on its computational efficiency are also discussed. It is believed that this new approach will constitute a key ingredient for full-scale 3D simulations based on actual electrode microstructures.

Ali Asad

Ecole Polytechnique
TotalEnergies, OneTech R&D
ali.asad@polytechnique.edu

Romain JWBRL de Loubens

TotalEnergies OneTech,
One R&D, 91120 Palaiseau Cedex, France
romain.de-loubens@totalenergies.com

Laurent Séries

CMAP, CNRS, École polytechnique, Institut
Polytechnique
de Paris, 91120 Palaiseau Cedex, France
laurent.series@polytechnique.edu

Laurent François

ONERA, DMPE, 6 Chemin de la Vauve aux Granges,
91120 Palaiseau Cedex, France
laurent.francois@onera.fr

Marc Massot

CMAP, CNRS, École polytechnique, Institut
Polytechnique
de Paris, 91120 Palaiseau Cedex, France
marc.massot@polytechnique.edu

MS198

Fluctuating Diffuse Interface Method: a Stochastic Mesoscale Model to Numerically Address the Bubble Nucleation Process

Vapour bubble nucleation is a multiscale and still puzzling process: its origin is found at the molecular scale but it develops up to millimetric scales, where bubbles evolve. Thermal fluctuations of the liquid molecules play a crucial role in the process. Most often these fluctuations lead to the formation of small vapour embryos, not large enough to survive and grow. However, rare events occur where the vapour cluster is big enough to trigger the transition to a stable phase change, leading to the appearance of a macroscopic bubble. The intrinsic stochasticity of this process, together with the wide range of length scales involved, makes the investigation of nucleation events a major challenge both from an experimental and a numerical viewpoint. Focusing on the numerical methodologies, the common toolkit dealing with the interface tracking/capturing

and the surface tension modelling of pre-existing bubbles are not enough to address nucleation. Capturing the spontaneous formation of a new vapour region with the proper statistical is critical to succeed. The Fluctuating Diffuse Interface (FDI) model, we recently developed, represented an important breakthrough in the numerical analysis of bubble nucleation. The FDI method, in fact, enabled the investigation from the very first inception of the vapour embryos up to the bubble growth. In this talk the application of this methodology to cavitation and boiling conditions will be presented.

Francesco Magaletti, Mirko Gallo, Anastasios Georgoulas, Marco Marengo
University of Brighton
f.magaletti2@brighton.ac.uk, m.gallo2@brighton.ac.uk,
a.georgoulas@brighton.ac.uk, m.marengo@brighton.ac.uk

Carlo Massimo Casciola
Sapienza Università di Roma
carlomassimo.casciola@uniroma1.it

MS198

Structure and Dynamics of Magnetically Responsive Particle-Stabilized Emulsion Gels

Complex multicomponent fluids continue to attract interest for their nonlinear rheological properties that enable mixture formulations with applications in materials processing. The interplay of phase behavior, fluid dynamics, and interfacial thermodynamics in complex fluids gives rise to interesting structure-property relations such as interfacial assembly and kinetically arrested phase morphologies. One particular example of kinetically arrested liquid mixtures are interfacially jammed emulsion gels (bijels) that emerge due to colloidal jamming during spinodal decomposition, resulting in an out-of-equilibrium liquid morphology with gel-like properties. Computational modeling and simulations can help gain a fundamental understanding of the structure and dynamics of particle-stabilized emulsions. In this talk, I will present results from lattice Boltzmann simulations of binary fluid mixtures with anisotropic suspended particles. Magnetic particles can be used for controlled deformation and coalescence of particle-stabilized droplets by applied fields. I will further present simulations of bicontinuous interfacially jammed emulsion gels. The simulations shed light on the effect of magnetic fields on the domain size and morphology of liquid mesophases. Potential applications include the development of control mechanisms that enable design of fluid templates with tailored structure, which can be leveraged to fabricate porous membranes and fibers for filtration and separation.

Nikhil Karthikeyan, Erin O'Neill, Fang Wang,
Ulf D. Schiller
Clemson University
nkarthi@g.clemson.edu, eoneil@g.clemson.edu,
fwang8@clemson.edu, uschill@clemson.edu

MS198

Rogue Nanowaves: A New Route to Thin Film Rupture

Understanding the stability of thin liquid films and their potential to spontaneously rupture is crucial for a wide range of fluid-based applications. Over the last thirty years, there has been significant progress in understanding the stability of ultra thin films, where disjoining pressure creates a linear instability that leads to a hole forming

(the so-called 'spinodal regime'). Notably, it is known that thermal capillary nanowaves, driven by Brownian motion in the bulk, play a key role in speeding up this instability. Here, we consider an alternative and complementary possibility - that the formation of rogue nanowaves can lead to the rupture of linearly stable films. Experimentally, such a 'thermal regime' has been observed, but a theoretical framework and predictive capability for this regime is absent. In this talk, it will be shown that the film's dynamics can be described by a stochastic thin film equation, which naturally captures the generation of nanowaves and, once solved computationally, predicts the rupture of linearly stable nanofilms. Then, using large deviation theory, which has recently been applied to model rogue ocean waves, we are able to formulate a theory which describes the most likely route to rupture and the timescale associated with this process, which compare well to our direct simulations. Finally, comparisons to experimental analyses and molecular simulations will be presented and directions of future research discussed.

James E. Sprittles, Jingbang Liu, Duncan Lockerby
University of Warwick
j.e.sprittles@warwick.ac.uk, jingbang.liu@warwick.ac.uk,
duncan.lockerby@warwick.ac.uk

Tobias Grafke
Warwick Mathematics Institute
University of Warwick
T.Grafke@warwick.ac.uk

MS198

Multi-Component and Multi-Phase Diffuse Interface Model: Objectivity and Thermodynamic Consistency

Multi-component and multi-phase fluids are omnipresent in our daily lives as well as in industrial applications. Two typical examples are coffee-ring effect and inkjet printing for electronics, where the fluid often contains more than two components. The fluid behavior is governed sometimes by diffusion, sometimes by convection, and sometimes by a mixture of them. Moreover, the fluid usually contacts a solid wall, where the fluid-solid interaction has to be considered. In this presentation, we will derive the evolution equations of multi-component and multi-phase fluids coupling with wetting phenomenon based on two basic facts: (i) Thermodynamic consistency and (ii) objectivity. Specifically speaking, we will derive the evolution equation starting from the free energy formulation to fulfill the second law of thermodynamics as well as the Gibbs-Duhem equation. We will show that the derived evolution equations are consistent with Fick's diffusion equation, Allen-Cahn (AC) model, Cahn-Hilliard (CH) model, and the coupled AC-CH model. The associated boundary condition coincides with Young's law. Next, we will show a generalized Navier-Stokes equation for systems with a large density ratio, where the origin of the Kortweg stress is discussed. At last, we will demonstrate that all the derived equations are objective, i.e., if an Euclidean transformation is performed, the evolution equations remain the respective form.

Fei Wang
Karlsruhe Institute of Technology
fei.wang@kit.edu

Britta Nestler
Karlsruhe Institute of Technology (KIT)
Institute of Applied Materials (IAM-ZBS)

britta.nestler@kit.edu

MS199

Accelerating Sparse Iterative Solvers and Preconditioners Using Race

Sparse linear iterative solvers are indispensable for large-scale simulations. In this talk, we present methods to accelerate some of the existing solvers and preconditioners by using the concept of levels as developed in the context of our RACE library framework. Levels are constructed using breadth-first search on the graph related to the underlying sparse matrix. These levels are then used to implement cache blocking of the matrix elements for high spatial and temporal reuse. The approach finds its use in kernels like sparse-matrix-power vector multiplication, which perform back-to-back sparse-matrix vector multiplication (SpMV)-type iterations without global synchronizations in between. The method is highly effective and achieves performance levels of 50-100 GF/s on a single modern Intel or AMD multicore chip, providing speedups of typically 2x - 4x compared to a highly optimized classical SpMV implementation. After briefly introducing the optimization strategy, we shed light on the application of these optimized kernels in iterative solvers. To this end, we discuss the coupling of the RACE library with the Trilinos framework and address the application to communication-avoiding s-step Krylov solvers, polynomial preconditioners, and algebraic multi-grid preconditioners. We then dive into the performance benefits and challenges of the RACE integration and show that our optimization produces numerically identical results and improves the total solver time by 1.3x - 2x.

Christie Louis Alappat
Friedrich-Alexander-Universität Erlangen-Nürnberg
Erlangen, Germany
christie.alappat@fau.de

Georg Hager
Erlangen National High Performance Computing Center
georg.hager@fau.de

Gerhard Wellein
Friedrich-Alexander-Universität Erlangen-Nürnberg
gerhard.wellein@fau.de

MS199

Performance Engineering of Madness

MADNESS (multiresolution numerical environment for scientific simulation) is a long-lived community software project that employs multiresolution analysis (MRA) and an advanced parallel runtime to solve differential and integral equations in broad areas of physics and chemistry with guarantees of both accuracy and speed. Its irregular and fine grain computational patterns have proven challenging to port to modern hybrid computer architectures. In this presentation, we discuss experience porting and tuning the software to the Fujitsu A64FX processor and also the complete re-design and new implementation using the Template Task Graph programming model in order to exploit GPUs and the massive size of modern exascale supercomputers.

Robert J. Harrison
Brookhaven National Laboratory
Stony Brook University
robert.harrison@stonybrook.edu

Poornima Nookala
Stony Brook University
poornimavinaya.nookala@stonybrook.edu

Mahdi Javanmard
Meta
javanmard.mahdi@gmail.com

MS199

Exploiting Tensor Product Structure to Accelerate Eigenvalue Problem Solvers in Few-Body Physics

In this presentation a computationally and numerically efficient method to compute both binding energies and their corresponding wave function in quantum mechanical few-body physics. The linear eigenvalue solver takes exploits the tensor product structure of the multidimensional stationary Schrödinger equation. The application of the Hamiltonian Operator is represented by matrix-matrix methods and then combined with a newly-designed preconditioner for the Jacobi-Davidson QR. This tensor method allows for significantly faster, highly-accurate computation of the three-body energies. For higher dimensions, we introduced a hybrid distributed/shared memory parallel approach. A GPU accelerated computing implementation was also considered. We will analyse the advantages of this implementation by looking at throughput, latency, memory usage as well as the numerical performance considerations

Moritz T. Hof
Institute of Software Technology
German Aerospace Center (DLR)
moritz.hof@dlr.de

MS199

An Optimal GPU Implementation for Computational Fluid Dynamics with the Finite Element Code Alya

Alya is a high-performance computational mechanics code for complex coupled multiphysics engineering problems. It is one of the two CFD codes of the Unified European Applications Benchmark Suite (UEBAS). On its path to Exascale, it has needed significant changes to adapt to new architectures, such as GPUs. This work describes recent improvements focused on large-scale incompressible flow problems. The momentum equation is treated explicitly while the pressure is solved implicitly. A fractional step scheme is used to enable elements that do not satisfy the inf-sup condition. The Laplacian matrix for the pressure remains fixed during the simulation for typical cases with a fixed mesh. Therefore, the two main kernels are calculating the right-hand side term for the momentum equation and the solution of a linear system for the pressure. This work focuses on the first step, but some minor comments on the linear solver will also be presented. Compared to our previous GPU implementation, more than an order of magnitude reduction in computational time for the assembly has been obtained. OpenACC has allowed us to be close to 50% of the maximum floating-point performance on an A100 Nvidia GPU. While optimizing the GPU version, improvements to the CPU implementation have also been obtained. Comparing the current CPU and GPU energy consumption shows good agreement with the ratios one can expect from the Top500 or Green500 lists.

Herbert Owen
BSC
herbert.owen@bsc.es

Georg Hager
Erlangen National High Performance Computing Center
georg.hager@fau.de

Gerhard Wellein
Erlangen Regional Computing Center (RRZE), Germany
Gerhard.Wellein@rrze.uni-erlangen.de

Dominik Ernst
FAU university
dominik.ernst@fau.de

MS199

Effective Simd Vectorisation for Finite Element Kernels

We examine the performance of finite element kernels for a wide range of problems on different CPU architectures. When the code is structured appropriately, modern compilers can auto-vectorize very effectively, and almost ideal vectorization can be achieved. As an added benefit, the maintenance and portability burden of special intrinsics and special programming instructions can be avoided. The structure of the code may also need to vary depending on the machine architecture. In general, memory bandwidth is a more important factor than raw floating-point performance. A performance model is presented, based on empirical data such as cache bandwidths and flops, to guide the optimization process and algorithm selection. Finally, automatic code generation is used to write suitable code for different kernels, whilst incorporating the insights from the performance model and compiler-friendly code design. We present the node-level performance of our generated code for three different architectures: Intel Icelake, Fujitsu A64fx, and AMD Milan processors.

Chris N. Richardson
University of Cambridge
chris@bpi.cam.ac.uk

Igor Baratta, Garth Wells
Department of Engineering
University of Cambridge
ia397@cam.ac.uk, gnw20@cam.ac.uk

MS200

Norm and Trace Estimation Using Rank-One Random Vectors

A few matrix-vector multiplications with random vectors are often sufficient to obtain reasonably good estimates for the norm of a general matrix or the trace of a symmetric positive semi-definite matrix. Several such probabilistic estimators have been proposed and analyzed for standard Gaussian and Rademacher random vectors. In this talk, we discuss the use of rank-one random vectors, that is, Kronecker products of (smaller) Gaussian or Rademacher vectors. It is not only cheaper to sample such vectors but it can sometimes also be much cheaper to multiply a matrix with a rank-one vector instead of a general vector. We provide theoretical and numerical evidence that the use of rank-one instead of unstructured random vectors still leads to good estimates. In particular, we show that our rank-one estimators multiplied with a modest constant constitute, with high probability, both upper and lower bounds of the quantity of interest. We illustrate the application of our techniques to condition number estimation for matrix

functions.

Zvonimir Bujanovic
University of Zagreb
Faculty of Science, Department of Mathematics
zbujanov@math.hr

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

MS200

Speeding Up the Computation of Matrix Functions via Randomized Methods

In this talk we discuss randomized algorithms for computing the action of a matrix function $f(A)$, such as the matrix exponential or the matrix square root, on a vector b . For a general (non symmetric) matrix A , this can be done by computing the projection of A onto a suitable Krylov subspace. Such projection is usually computed by forming an orthonormal basis of the Krylov subspace using the Arnoldi method. In this talk, we propose to use non-orthonormal bases of the Krylov subspace, which are faster to construct. Then we use a fast randomized algorithm for least squares problems to compute the projection of A onto the Krylov subspace. We present some numerical examples which show that our proposed algorithms can be faster than the standard Arnoldi method while achieving comparable accuracy.

Alice Cortinovis
EPFL
alicecor@stanford.edu

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

Yuji Nakatsukasa
University of Oxford
nakatsukasa@maths.ox.ac.uk

MS200

Randomized Contour Integral Methods for Eigenvalue Problems

Randomized NLA methods have recently gained popularity because of their easy implementation, computational efficiency, and numerical robustness. In this talk, we present the analysis of a randomized version of a well-established FEAST algorithm that enables computing the eigenvalues of the Hermitian matrix pencil (\mathbf{A}, \mathbf{B}) located in the given real interval $\mathcal{I} \subset [\lambda_{min}, \lambda_{max}]$. First, we establish new structural as well as probabilistic error analysis of the accuracy of approximate eigenpairs and subspaces obtained using the randomized FEAST algorithm, i.e., bounds for the canonical angles between the exact and the approximate eigenspaces, and for the accuracy of the eigenvalues and the corresponding eigenvectors. Since this part of the analysis is independent of the particular distribution of an initial subspace, we denote it as *structural*. In the case of the starting guess being a Gaussian random matrix, we provide more informative, probabilistic error bounds. Our modified algorithm allows to improve the accuracy of orthogonalization when \mathbf{B} is ill-conditioned, efficiently apply the rational filter by using MPMRES-Sh [Bakhos, Ki-

tanidis, Ladenheim, Saibaba and Szyld, 2016] method to accelerate solving shifted linear systems and estimate the eigenvalue counts in a given interval. Finally, we illustrate numerically the effectiveness of presented error bounds and proposed algorithmic modifications.

Agnieszka Miedlar
Virginia Tech
amiedlar@vt.edu

Arvind Saibaba
North Carolina State University
saibab@ncsu.edu

Eric De Sturler
Virginia Tech
sturler@vt.edu

MS200

Randomization Techniques Applied to Krylov Methods for Solving Sparse Symmetric Systems of Equations

We present our study on randomized orthogonal projection methods (ROPs) applied to symmetric systems of linear equations. We discuss both their absolute performance and their performance compared to that of deterministic orthogonal projection methods (OPs), and their possible implementations, such as full-orthogonalization algorithms (FOM, randomized FOM).

Edouard Timsit
INRIA
edouard.timsit@inria.fr

Laura Grigori
INRIA
France
laura.grigori@inria.fr

Oleg Balabanov
INRIA
oleg.balabanov@inria.fr

MS201

A Quantile Conserving Ensemble Filtering Framework: Regressing Quantile Increments to Update Unobserved Variables

Ensemble Kalman filters are commonly used for data assimilation in geoscience applications. A novel efficient algorithm that allows the use of arbitrary continuous priors and likelihoods for an observed variable was presented at several recent SIAM meetings. The key innovation was to select posterior ensemble members with the same quantiles with respect to the continuous posterior distribution as the prior ensemble had with respect to the prior continuous distribution. While this method led to significant improvements in analysis estimates for observed variables, those improvements can be lost when using standard linear regression of observation increments to update other state variables. However, doing the regression of observation quantile increments in a bivariate quantile space guarantees that the posterior ensembles for state variables also have all the advantages of the observation space quantile conserving posteriors. For example, if state variables are bounded then posterior ensembles will respect those bounds and eliminate most bias near the boundary. The

posterior ensembles also respect other aspects of the continuous prior distributions. Examples are shown for a variety of bivariate prior ensembles including bounded quantities and multimodal distributions. The method has potential to significantly improve data assimilation for distinctly non-Gaussian quantities in Earth system models like tracers.

Jeffrey Anderson
National Center for Atmospheric Research
Institute for Math Applied to Geophysics
jla@ucar.edu

MS201

Time-Limited Balanced Truncation for Data Assimilation

Balanced truncation is a well-established model order reduction concept in system theory that has been applied to a variety of problems. Recently, a connection between linear Gaussian Bayesian inference problems and the system theoretic concept of balanced truncation was drawn for the first time. Although this connection is new, the application of balanced truncation to data assimilation is not a novel concept: It has already been used in four-dimensional variational data assimilation (4D-Var) in its discrete formulation. In our work, the link between system theory and data assimilation is further strengthened by discussing the application of balanced truncation to standard linear Gaussian Bayesian inference, and, in particular, the 4D-Var method. Similarities between both data assimilation problems allow a discussion of established methods as well as a generalisation of the state-of-the-art approach to arbitrary prior covariances as reachability Gramians. Furthermore, we propose an enhanced approach that allows to balance Bayesian inference for unstable systems and improves the numerical results for short observation periods.

Josie König, Melina Freitag
University of Potsdam
josie.koenig@uni-potsdam.de, melina.freitag@uni-potsdam.de

MS201

Efficient, Scalable and Parallel Uncertainty Quantification for Electromagnetic Data via the RTO-TKO

The goal when inverting electromagnetic (EM) data is to understand the composition of the solid Earth down to several kilometers below the ocean floor. The main issue in EM inversions is that the data do not constrain all aspects of the (resistivity) model. For this reason, regularization has a very large effect on the inversion and associated uncertainty quantification (UQ). In this talk, I will explain how to adapt ideas from "randomize-then-optimize" (RTO) to the inversion and UQ of EM data. Specifically, I will argue that one can easily extend RTO so that the inversion incorporates a data-driven regularization and that one can do so at an acceptable computational cost, leveraging parallelism and high-performance computers. We call the overall algorithm the RTO-TKO, because it essentially amounts to performing two optimizations per sample. The TKO stands for "technical knock out" because the RTO-TKO algorithm samples an approximate Bayesian posterior distribution, but a "knock out" would be to sample without bias.

Matthias Morzfeld

University of California, San Diego
Scripps Institution of Oceanography
mmorzfeld@ucsd.edu

MS201

An Adaptive Ensemble Gaussian Mixture Filter

The ensemble Gaussian mixture filter (EnGMF) attempts to bridge the gap between Kalman filter-based methods and particle filters. Its major downside is that it relies on a largely heuristic bandwidth parameter to describe the prior and posterior Gaussian mixtures. We provide a methodology by which this parameter can be determined in an online fashion, and generalize this methodology to a larger class of parameterized covariances in the EnGMF.

Andrey A. Popov
University of Texas at Austin
apopov@vt.edu

Renato Zanetti
University of Texas at Austin
Department of Aerospace Engineering and Engineering
Mechanic
renato@utexas.edu

MS201

Efficient Randomization Techniques for Solving Bayesian Inverse Problems

Large-scale inverse problems involve fusing incomplete and noisy information from multiple sources, such as model simulations, measurements from sensors, and physical experiments, to obtain a consistent description of the state of the underlying physical system. Solving the Bayesian formulation of these problems enables quantifying the uncertainties associated with the solution. However, solving Bayesian problems presents a major challenge: Solving Bayesian inverse problems is computationally demanding, often requiring hundreds to thousands of expensive simulations to accurately estimate the parameters and their uncertainties. Randomized algorithms provide an attractive means to reduce the computational cost. In this work, we will explore efficient randomization techniques as a means to develop scalable solvers and pre-conditioners to mitigate the computational costs associated with solving Bayesian inverse problems.

Amit Subrahmanya
Virginia Tech
amitns@vt.edu

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Adrian Sandu
Virginia Polytechnic Institute and State University
Computational Science Laboratory
sandu@cs.vt.edu

Arvind Saibaba
North Carolina State University

asaibab@ncsu.edu

MS202

Data-Enabled Predictive Control

We consider the problem of optimal and constrained control for unknown systems. A novel data-enabled predictive control (DeePC) algorithm is presented that computes optimal and safe control policies using real-time feedback driving the unknown system along a desired trajectory while satisfying system constraints. Using a finite number of data samples from the unknown system, our proposed algorithm uses a behavioral systems theory approach to learn a non-parametric system model used to predict future trajectories. We show that, in the case of deterministic linear time-invariant systems, the DeePC algorithm is equivalent to the widely adopted Model Predictive Control (MPC), but it generally outperforms subsequent system identification and model-based control. To cope with nonlinear and stochastic systems, we propose salient regularizations to the DeePC algorithm. Using techniques from (distributionally) robust stochastic optimization, we prove that these regularization indeed robustify DeePC against corrupted data. We illustrate our results with nonlinear and noisy simulations and experiments from power electronics and power systems.

Florian Dörfler
ETH Zürich
D - ITET Automatic Control Laboratory
dorfler@ethz.ch

MS202

Analysis and Validation of Parameter Varying Genetic Toggle Switches Using Koopman Operators

The genetic toggle switch is a well-known model in synthetic biology that represents the dynamic interactions between two genes that repress each other. The mathematical models for the genetic toggle switch that currently exist have been useful in describing circuit dynamics in rapidly dividing cells, assuming fixed or time-invariant kinetic rates. As cells transition from one growth phase to another, kinetic rates can be modeled with time-varying parameters. We propose a novel class of parameter varying nonlinear models to capture the dynamics of genetic circuits, including the toggle switch, as they transition from different phases of growth. We show that there exists unique solutions for this class of systems. We show that the domain of these systems, which is the positive orthant, is positively invariant. We showcase a theoretical control strategy for these systems that would grant asymptotic monostability of a desired fixed point. We then take the general form of these systems and analyze their stability properties through the framework of time-varying Koopman operator theory. A necessary condition for asymptotic stability is also provided as well as a sufficient condition for instability. Furthermore, we use real data from a genetic toggle switch which contains a hybrid promoter to validate the models using an algorithm which constructs time varying parameters to model the toggle switch's behavior.

Jamiree M. Harrison
University of California, Santa Barbara
jamiree@ucsb.edu

Enoch Yeung
University of California Santa Barbara

eyeung@ucsb.edu

MS202

A Data-Based Approach for Optimal Planning of a Surgical Cannula

For future medical interventions, concentric tube continuum robots are considered promising approaches. This requires mathematical research in both, modeling and control. In this presentation, we focus on a path planning task for a multi-tube cannula as it could be used in deep-brain neurosurgery. A system model derived by first principles has to be refined by data-based approaches. Improving the accuracy of the tool tip position, a data-based iterative control scheme is applied.

Kathrin Flaßkamp
Systems Modeling and Simulation
Saarland University, German
kathrin.flasskamp@uni-saarland.de

Matthias K. Hoffmann

Systems Modelling and Simulation
Saarland University
matthias.hoffmann@uni-saarland.de

MS202

Data-Driven Predictive Control with Stochastic Uncertainties

Commonly, Model Predictive Control (MPC)—which is also called receding-horizon control—is based on the repeated solution of an optimal control problem. As such MPC usually relies on *model* of underlying dynamics. However, there is rapidly growing line of research which replaces the model with data-driven system descriptions. Of particular interest are methods which can work with output-data only. So far most of the research on data-driven predictive control has focused on noisy measurement data but not on the consideration of stochastic disturbances. In this talk, we discuss data-driven predictive control with stochastic uncertainties. We show how a tailored variant of the ubiquitous fundamental lemma in combination with polynomial chaos expansions enables data-driven output-feedback stochastic predictive control. We present sufficient stability conditions and we draw upon numerical examples to illustrate our findings.

Timm Faulwasser, Ruchuan Ou, Guanru Pan
TU Dortmund University
timm.faulwasser@tu-dortmund.de, ruchuan.ou@tu-dortmund.de, guanaru.pan@tu-dortmund.de

MS202

Distributed Control of Partial Differential Equations Using Convolutional Reinforcement Learning

We present a convolutional framework which significantly reduces the complexity and thus, the computational effort for distributed reinforcement learning control of partial differential equations (PDEs). Exploiting translational invariances, the high-dimensional distributed control problem can be transformed into a multi-agent control problem with many identical agents. Furthermore, using the fact that information is transported with finite velocity in many cases, the dimension of the agents' environment can be drastically reduced using a convolution operation over the state space of the PDE. In this setting, the complexity can

be flexibly adjusted via the kernel width or using a stride greater than one. A central question in this framework is the definition of the reward function, which may consist of both local and global contributions. We demonstrate the performance of the proposed framework using several standard PDE examples with increasing complexity, where stabilization is achieved by training a low-dimensional DDPG agent with small training effort.

Sebastian Peitz, Jan Stenner, Vikas Chidananda
Paderborn University
sebastian.peitz@upb.de, jan.stenner@upb.de,
cvikas@mail.upb.de

Steve Brunton
University of Washington
sbrunton@uw.edu

Kunihiko Taira
University of California, Los Angeles
Mechanical and Aerospace Engineering
ktaira@seas.ucla.edu

MS203

Ginkgo's In-Register Compression for Accelerating Memory-Bound Linear Algebra

Ginkgo is an open source linear algebra library for high performance on modern HPC architectures. As much of sparse linear algebra is memory bound, accelerating algorithms means reducing the data access volume. We present the idea of decoupling the memory format from the arithmetic format and use in-register compression to improve the performance of numerical linear algebra while preserving the final output accuracy.

Hartwig Anzt
University of Tennessee, U.S.
hantz@icl.utk.edu

Terry Cojean, Pratik Nayak, Thomas Grützmacher,
Tobias Ribizel, Marcel Koch, Yu-Hsiang Mike Tsai
Karlsruhe Institute of Technology
terry.cojean@kit.edu, pratik.nayak@kit.edu,
thomas.gruetzmacher@kit.edu, tobias.ribizel@kit.edu,
marcel.koch@kit.edu, yu-hsiang.tsai@kit.edu

MS203

Trilinos: Portable and Scalable Linear Algebra using MPI+Kokkos

The Trilinos collection of software packages provide all the functionality required to build high-performance scientific applications. It includes a wide variety of scalable linear algebra data structures, solvers, and preconditioners. Trilinos builds upon MPI for efficient inter-process communication and Kokkos for portable, shared-memory parallel computational kernels. With the MPI+Kokkos model, the same source code performs well on all common HPC hardware, including the upcoming DOE exascale machines: Intel, AMD and ARM CPUs and NVIDIA, AMD and Intel GPUs. We present performance results from low-level vector migration to high-level FEM graph assembly and preconditioned sparse linear solves using algebraic multigrid. We also share some of the software engineering techniques that have helped us develop portable code without sacrificing performance or specializing for each type of hardware.

Brian Kelley, Christopher Siefert
Sandia National Laboratories
bmkelle@sandia.gov, csiefer@sandia.gov

MS203

Portable C++ for Modern, Heterogenous Mixed-Precision Methods

Mixed precision methods are increasingly popular in numerical linear algebra. However, developing and implementing these methods often requires repeatedly implementing the same logic for different combinations of data types. Furthermore, adding support for hardware accelerators often compounds this issue due to differences in API and hardware behaviors. Fortunately, modern C++ features, such as templating, allow for generic implementations that provide both performance and readability.

Neil Lindquist
University of Tennessee
nlindqu1@icl.utk.edu

MS203

Eigen: System Overview, Maintenance Challenges, and Lessons Learned

Eigen is an open-source C++ template library for linear algebra, supporting matrix and tensor operations, geometric transforms, numerical solvers, and related algorithms. It is the principal C++ math library used internally across Google, underpinning TensorFlow, Waymo, and Android ARCore. As such, performance is critical across a wide range of devices. To accomplish this, Eigen uses expression templates to build compile-time computation graphs, and a generic SIMD abstraction layer to allow explicit vectorization for multiple platforms. In this talk, we will provide an overview of the system architecture, its usage in TensorFlow and Google at large, and some of the maintenance challenges due to both the current design and use at scale.

Antonio Sanchez, Rasmus Larsen
Google
cantonios@google.com, rmlarsen@google.com

MS204

A Data-Driven ROM Algorithm for Synthetic Aperture Radar Imaging in Multi-Scattering Environment

The data-driven reduced order models (ROMs) have recently emerged as an efficient tool for the solution of the inverse scattering problems with applications to seismic and sonar imaging. One specific of this approach is that it requires the full square multiple-output/multiple-input (MIMO) matrix valued transfer function as the data for the multidimensional problems. The synthetic aperture radar (SAR) however is limited to the single input/single output (SISO) measurements corresponding to the diagonal of the transfer function. We present a ROM based Lippmann-Schwinger approach overcoming this drawback. It also simplifies the data-driven computations of the internal field by formulating it as a composition of Gram-Schmidt transformations on the known background field. Our algorithm allows us to obtain the solution of the non-linear inverse scattering problem at the cost of the linearized Born method however suppressing multi-scattering

artifacts such as ghost echoes. Efficiency of the proposed approach is demonstrated on 2D and 2.5D numerical examples.

Vladimir L. Druskin
Worcester Polytechnic Institute
vdruskin@gmail.com

Shari Moskow
Drexel University
Department of Mathematics
moskow@math.drexel.edu

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

MS204

Finding the Conductivity on a Graph from Internal Power Measurements

We consider the problem of finding the resistors in a resistor network from knowing the power dissipated by each of the resistors under different imposed voltages at certain nodes. We give conditions under which the linearized problem is injective, including in cases where some of the imposed voltages may leave some resistors with zero current traversing them. The approach is inspired by the Bal uniqueness proof for certain linearized continuum inverse problems with internal functionals. The same approach can be applied to a discrete version of the Schroedinger problem and is illustrated with numerical experiments.

Fernando Guevara Vasquez
University of Utah
fguevara@math.utah.edu

MS204

Machine Learning and Computational Methods for Blood Flow in Organs

Mapping the vascular system in the liver organ helps practitioners with diagnosis, treatment, and surgical planning. There is growing interest to use deep learning techniques to more efficiently and effectively model the vascular system and to use computational methods to provide insight to the flow of blood in the organ. Using PocketNet paradigm, we create a 3D neural network for vessel segmentation on CT images. From these segmentations, the blood vessels are skeletonized and their centerlines are computational domains for reduced 1D models while the liver is converted to a 3D mesh to be the simulation domain. Using the 3D liver and the 1D vasculature we model the flow of solute from the vessels into the liver using finite element method. Such automation of vessel segmentation and simulation of flow modeling can be translated into the clinical setting to improve patient outcome and care. This is particularly beneficial for the treatment of hepatocellular carcinoma where the delivery of chemotherapy to hepatic tumors can be modeled.

Bilyana Tzolova
Rice University
bmt3@rice.edu

Beatrice Riviere
Department of Computational and Applied Mathematics
Rice University
riviere@rice.edu

David Fuentes
University of Texas MD Anderson Cancer Center
dtfuentes@mdanderson.org

MS204

Distance Preserving Model Order Reduction of Graph-Laplacians for Cluster Analysis

Graph-Laplacians play an important role in multiple areas of machine learning directly related to imaging, such as image segmentation and, in general, discovering the structure of large data sets that may not have explicit geometric interpretation. Spectral imbedding of graph-Laplacians provides a low-dimensional parametrization of the data manifold which makes the subsequent task (for example, clustering, with k-means or any of its approximations) much easier. However, despite reducing data dimensionality, the overall computational cost may still be prohibitive for accurate clustering of large data sets. In this talk we present a novel multi-level divide-and-conquer approach for clustering of big data sets. At each level, independent clustering tasks are performed for small data subsets. Unlike conventional greedy graph coarsening algorithms, the crucial part of our approach is efficient clustering of each of these subsets (target subsets) that ensures the consistency of subset clustering with the clustering of the full data set (if one would perform such). To achieve this goal we construct a low-dimensional representation of the original graph that respects important distances between the vertices of the target subset. Our main tool is the so-called reduced-order graph Laplacian that approximates these distances with spectral accuracy. Numerical results will be provided to illustrate the performance of the approach.

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

Vladimir L. Druskin
Worcester Polytechnic Institute
vdruskin@gmail.com

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

MS205

Algebraic Multilevel Domain Decomposition Preconditioners

Multilevel Domain decomposition (MDD) methods are among the most efficient iterative methods for solving sparse linear systems. One of the main technical difficulties in using efficient MDD methods (and most other efficient preconditioners) is that they require information from the underlying problem which prohibits them from being used as a black-box. I will present a framework that simplifies the analysis of the algebraic two-level overlapping Schwarz preconditioner. Based on this framework, I will present a series of robust and fully algebraic MDD methods, i.e., that can be constructed given only the coefficient matrix and guarantee a priori prescribed convergence rate. The series consists of preconditioners for least-squares problems, sparse SPD matrices, general sparse matrices, and certain saddle-point systems. Numerical experiments on challenging discretized PDEs and systems from the SuiteSparse Collection illustrate the effectiveness, wide applicability, scalability of the proposed preconditioners. A comparison of each one against state-of-the-art preconditioners is also

presented.

Hussam Al Daas
STFC, Rutherford Appleton Laboratory
Computational Mathematics Group
hussam.al-daas@stfc.ac.uk

MS205

How to Recover from Node Failures in s-step Conjugate Gradient Methods

With the massive growth of concurrency in large-scale computer systems, global communication becomes a major bottleneck and the likelihood of component failures increases. The Preconditioned Conjugate Gradient (PCG) method is an iterative solver for large sparse linear systems facing these challenges. To reduce the number of global synchronizations by a factor of $O(s)$, communication-avoiding s-step PCG methods compute the iterations of PCG in blocks of s . Our focus is on resilience against node failures, for which the commonly used approach is periodically saving the state of the solver (Checkpoint-Restart). Other approaches exploit algorithm-specific properties to achieve scalable resilience on large-scale parallel computers. During the failure-free phase of resilient PCG, a small part of the states data can be stored redundantly by exploiting the inherent data redundancy in each iteration. In the event of node failures, the lost parts of the state are recovered using an exact state reconstruction (ESR) strategy after retrieving this redundantly stored data. We incorporate resilience against multiple node failures in two s-step PCG methods using ESR ideas and illustrate very low overhead compared to the non-resilient methods. By providing theoretical and experimental evaluations, we confirm that the communication-avoiding properties and therefore scalability of the s-step PCG methods are preserved.

Viktoria Mayer
University of Vienna
viktoria.mayer@univie.ac.at

Wilfried N. Gansterer
Department of Computer Science
University of Vienna
wilfried.gansterer@univie.ac.at

MS205

Block Preconditioning for Magnetic Confinement Fusion Relevant Resistive MHD Simulations

A base-level mathematical basis for the continuum fluid modeling of dissipative plasma system is the resistive magnetohydrodynamic model. This model requires the solution of the governing partial differential equations (PDEs) describing conservation of mass, momentum, and thermal energy, along with various reduced forms of Maxwells equations for the electromagnetic fields. The resulting systems are characterized by strong nonlinear and nonsymmetric coupling of fluid and electromagnetic phenomena, as well as the significant range of time- and length-scales that these interactions produce. These characteristics make scalable and efficient iterative solution, of the resulting poorly-conditioned discrete systems, extremely difficult. In this talk we consider the use of block preconditioners for solving the coupled physics block systems. The block preconditioner considered here is based on an approximate operator splitting approach which can isolate certain coupled systems, allowing them to be handled independently. Here we use the splitting to create two independent 2×2 block sys-

tems, a magnetics-flow system and a magnetics-constraint system. We demonstrate this approach for various resistive MHD problems that are relevant to magnetic confinement fusion applications and compare the performance with alternative preconditioning methods.

Peter Ohm

Computational Mathematics Dept.
Sandia National Laboratory
peter.ohm@riken.jp

John N. Shadid
Sandia National Laboratories
jnshadi@sandia.gov

Jesus Bonilla
Los Alamos National Laboratory
jbonilla@lanl.gov

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

Edward G. Phillips
Sandia National Laboratories
egphill@sandia.gov

Jonathan J. Hu
Sandia National Laboratories
Livermore, CA 94551
jhu@sandia.gov

MS205

Efficient Multigrid Reduction Strategies for Sub-surface Flow Applications

Simulation of subsurface fluid flow involves solving a multi-physics problem in which multiphase flow and transport are tightly coupled. To capture this dynamic interplay, fully implicit methods, also known as monolithic approaches, are usually preferred. However, this requires solution of large linear systems that result from the discretization and linearization of the governing balance equations. Such systems are non-symmetric, indefinite, and highly ill-conditioned, which make them difficult to solve. These problems become even more challenging when the flow is coupled to mechanical processes such as rock fracture and deformation. Thus, designing preconditioners that can handle the couplings between different physics is critical for fast convergence. Furthermore, the emergence of new HPC hardware with accelerators, such as GPUs, provide opportunities to further improve the performance of the solution process, when adequately exploited. This work will present our efforts to develop scalable linear solver strategies for subsurface flow applications, designed to be efficient on modern HPC systems. We present a framework based on multigrid reduction (MGR) that is suited for tightly coupled systems of PDEs and demonstrate its applicability to handle challenging problems from multiphase flow applications. We show that the framework is flexible to accommodate a wide range of scenarios, as well as efficient and scalable for large problems on state-of-the-art HPC architectures.

Daniel Osei-Kuffuor, Victor Magri
Lawrence Livermore National Laboratory
oseikuffuor1@llnl.gov, paludettomag1@llnl.gov

Francois P. Hamon
Total Exploration & Production
francois.hamon@total.com

Nicola Castelletto
Lawrence Livermore National Laboratory
Livermore, CA, USA
castelletto1@llnl.gov

Matteo Cusini, Joshua A. White
Lawrence Livermore National Laboratory
cusini1@llnl.gov, white230@llnl.gov

Randolph R. Settigast
Lawrence Livermore National Laboratory
settigast1@llnl.gov

MS206

Iterative Coupling of PDEs and Neural Networks

Coupled processes govern many real-world systems. In the computational sciences, the robust and efficient numerical solution of coupled systems has become therefore of increased interest in the recent past. Consequently, the field has seen many theoretical and practical advances, mainly for problems explicitly stated as PDEs and solved using high-fidelity simulators. New challenges appear when coupling data-driven models with high-fidelity simulators. In this talk, we investigate to which extent widely studied concepts from poroelasticity (the coupling of flow and deformation in porous media) can be applied also in this context. Particular focus lies on iterative coupling using stabilization and extrapolation methods. PINNs are utilized as data-driven models, while a finite volume simulator is used as physics-based model.

Jakub Both

University of Bergen, Norway
jakub.both@uib.no

MS206

Deep Neural Network Approximation in Shape Uncertainty Quantification for Acoustic and Electromagnetic Scattering

We consider the acoustic and electromagnetic scattering problems by random domains and surfaces. Our goal is to construct an efficient approximation of the corresponding domain-to-solution map. After considering a domain or surface parametrization depending on countably many parameters, one obtains a high-dimensional parametric map describing the problem's solution manifold. To tackle the efficient approximation of these maps, we adopt the approach proposed by Hesthaven and Ubbiali (Journal of Computational Physics 363 (2018): 55-78.). Firstly, by using a collection of so-called high-fidelity solutions or *snapshots*, a reduced basis is constructed to approximate the solution manifold. Then, we use deep feed-forward neural networks to compute the coefficient of the reduced basis solution. This procedure renders the evaluation of the approximate parameter-to-solution in the online phase independent of the offline part. Provided that the parameter-to-solution map satisfies certain smoothness properties, which have been verified for a range of models used in acoustic and electromagnetic scattering, we show that one can use quasi-Monte Carlo quadrature rules in the computation of the snapshots used in the reduced basis and also in the training of the involved deep neural networks.

Finally, we provide numerical results for different models arising in acoustic and electromagnetic scattering.

Fernando Henriquez

École polytechnique fédérale de Lausanne
fernando.henriquez@epfl.ch

MS206

Enhancing Training of Scientific Machine Learning Applications

Scientific machine learning (SCL) approaches employ supervised and unsupervised deep learning techniques in order to model the dynamics of complex multiscale and multiphysics problems, and are designed to provide exceptionally cheap surrogates. Despite their performance during model evaluation, SCL approaches suffer from a computationally exhaustive training phase. In this talk, we will leverage ideas from multilevel minimization methods to efficiently train deep neural networks (DNNs) employed in SCL applications. The main idea behind multilevel methods is to employ a hierarchy of auxiliary problems, which are minimized internally to accelerate the solution process of the original problem. We will discuss how to construct a multilevel hierarchy and transfer operators by exploring the properties of the loss function and the structure of the DNN architecture. The convergence properties of our novel training method will be analyzed using a series of numerical experiments, including physics-informed neural networks, and deep operator learning approaches. A comparison with the stochastic gradient and Adam optimizers will be also presented, showing a significant reduction in terms of the computational cost for the training phase.

Alena Kopanicakova
Università della Svizzera italiana
Switzerland
alena.kopanicakova@brown.edu

Hardik Kothari
Università della Svizzera italiana
hardik.kothari@usi.ch

Rolf Krause
Università della Svizzera italiana, Switzerland
rolf.krause@usi.ch

MS206

A Hybrid Finite Element / Neural Network Method for the Navier-Stokes Equations

In this talk we discuss the use of deep neural networks for augmenting classical finite element simulations in fluid-dynamics. Classical simulation methods often reach their limits. Even if the finite element method is highly efficient and established for the discretization of the Navier-Stokes equations, fundamental problems, such as the resolution of fine structures or a correct information transport between scales, are still not sufficiently solved. We discuss approaches to connect the finite element method with neural networks to overcome these obstacles. The paradigm is to use classical simulation techniques where their strengths are eminent, such as in the efficient representation of a coarse, large-scale flow field. Neural networks are used where a full resolution of the effects does not seem possible or efficient. The Deep Neural Network Multigrid Solver takes up these ideas by combining a geometric multigrid solver and a deep neural network. We show the efficiency,

generalizability and scalability by 2D and 3D simulations. We introduce an approach to retrain the network adaptively based on the uncertainty of the predictions. Based on extensions of the Deep Neural Network Multigrid Solver we offer perspectives on how finite element simulations can benefit from neural networks in the future. In addition to the desired increase in efficiency, the focus is particularly on issues of stability, generalizability and error accuracy.

Nils Margenberg
Helmut Schmidt University Hamburg
margenbn@hsu-hh.de

Christian Lessig
Otto-von-Guericke Universität Magdeburg
christian.lessig@ovgu.de

Thomas Richter
University of Magdeburg
Germany
thomas.richter@ovgu.de

Robert Jendersie
University of Magdeburg
robert.jendersie@ovgu.de

MS206

Scientific Machine Learning Models in Cardiac Electrophysiology

Cutting-edge machine- and deep-learning technologies for processing and analyzing clinical data enable new approaches to personalized diagnosis and treatment based on individual patient characteristics. These approaches require large volumes of high-resolution data for training the models, which are not always available; measurements may be incomplete, affected by noise, or give a partial reconstruction of the phenomenon. In this talk, we present physics-aware and physics-informed learning approaches that leverage clinical data with the physical knowledge of underlying phenomena. Specifically, we adopt dimensionality reduction techniques built from simulation snapshots in physics-aware artificial neural networks to improve prediction accuracy and training efficiency. We also consider additional physics-informed constraints to data-driven loss functions. These terms directly penalize the residual of partial differential equations or boundary/initial conditions. In this way, we simultaneously counterbalance incomplete data with physical knowledge and improve the performance of low-complexity architectures based on a limited number of hyperparameters. We finally show some numerical results of physics-aware and physics-informed neural networks with applications to cardiac electrophysiology.

Stefano Pagani, Francesco Regazzoni
Politecnico di Milano
stefano.pagani@polimi.it, francesco.regazzoni@polimi.it

Riccardo Tenderini, Simone Deparis
Ecole Polytechnique Federale de Lausanne
riccardo.tenderini@epfl.ch, simone.deparis@epfl.ch

Alfio Quarteroni
Politecnico di Milano

alfio.quarteroni@polimi.it

MS207

Self-Supervised, Probabilistic Surrogates of PDEs Without Labeled Data

The cost of repeatedly solving a PDE under different parametric inputs is the main computational bottleneck in many-query applications such as uncertainty quantification and inverse problems. While several data-driven strategies for the construction of inexpensive surrogates have been proposed, their majority is hampered by the following difficulties: - they rely on a large, training dataset of input-output pairs which in turn implies a significant computational cost. - the inputs for which training data is generated are decided a priori without regard to their informational content in constructing the surrogate. - deterministic surrogates are mostly considered which cannot quantify predictive uncertainty, especially in out-of-distribution settings. We propose a fully probabilistic approach in which the governing PDE is treated as a source of virtual data which is inquired iteratively by considering appropriate weighted residuals. We formulate this as a Variational Inference optimization objective, the output of which is a fine-tuned, probabilistic surrogate. We consider several variations of the main strategy and demonstrate its efficacy in a variety of problems.

Phaedon S. Koutsourelakis
Technical University Munich
p.s.koutsourelakis@tum.de

Matthaios Chatzopoulos
Technical University of Munich
matthaios.chatzopoulos@tum.de

MS207

Sparse Sampling for Accelerating On-The-Fly Reduced-Order Modeling with Time-Dependent Subspaces

We present a model reduction framework for the computation of finite-time nonlinear sensitivities in dynamical systems. Unlike solving a linearized system that assumes infinitesimal perturbations around a base state, this framework considers nonlinear sensitivity equations (NLSE) for finite perturbations. To solve the NLSE, we propose a low-rank approximation that leverages a time-dependent basis by extracting correlations between sensitivities on-the-fly. To this end, we derive forward low-rank evolution equations for an orthonormal state basis, correlation matrix, and orthonormal parametric basis. The resulting equations are Jacobian-free and leverage the same nonlinear solver that is used to compute the evolution of the base state. To enable efficient computation of the low-rank equations, we employ a rank-adaptive sparse sampling strategy. This strategy constructs a low-rank approximation for the right hand side of the NLSE via the discrete empirical interpolation method (DEIM). Not only does this approach reduce the computational cost, it is significantly less intrusive as it does not require implementing the reduced operators term by term. For nonlinear sensitivities with arbitrarily time dependent base state, we demonstrate that low-rank structure often exists, and can be extracted in real time by solving evolution equations. We consider two case studies: (i) 2D compressible flow and (ii) species transport in a turbulent reacting flow.

Michael Donello, Mohammad Hossein Naderi, Hessam

Babae
University of Pittsburgh
mid47@pitt.edu,
h.babae@pitt.edu

m.h.naderi@pitt.edu,

MS207

Machine Learning Based Structural Design and Analysis with Uncertainty Quantification

Machine learning for structural design and optimization problems have been receiving significant interest in recent years. While much progress has been done, it is becoming challenging to discern where classical versus ML-based approaches should be used. In this talk, I will share our past and ongoing studies trying to answer this question. Specifically, I will introduce our work in data-driven as well as on-line neural network based structural topology optimization approaches and compare how they perform relative to the classical approaches. Also, I will introduce a new ML-driven scalar field prediction approach that aims to perform fast predictions of engineering quantities of interest on geometrically and topologically different structures. I will discuss our approach to uncertainty quantification in these predictions, which becomes critical for an industrial adoption of such techniques.

Levent Burak Kara
Carnegie Mellon University
lkara@cmu.edu

MS207

Constrained Optimization of Neural Networks for Physics-Based Learning of Partial Differential Equations

Physics-informed neural networks (PINNs) are typically trained using a composite objective function, which is a weighted sum of the residuals of a governing partial differential equation (PDE) and its boundary conditions. A major drawback of this approach is that boundary conditions are not properly used to constrain the solution and the weighting factors in the objective function are problem specific and not known a priori. Additional complications arise when multi-fidelity observations are introduced into the objective function as well. To address these limitations in a principled fashion, we pursue equality constraint optimization to determine the parameters of neural networks. In our approach, we adopt the augmented Lagrangian method (ALM) to constrain PDE loss with its boundary conditions and any high-fidelity observations that may be available for the PDE at hand. Furthermore, we present a strategy to learn non-smooth solutions, where we introduce auxiliary parameters to reduce the order of differential operators and further constrain the solution. We demonstrate the efficacy and versatility of our framework by applying it to several forward and inverse problems involving multi-dimensional PDEs. We achieve orders of magnitude improvements in accuracy levels in comparison with state-of-the-art physics-informed neural networks.

Inanc Senocak
Mechanical & Biomedical Engineering
Boise State University
senocak@pitt.edu

Shamsulhaq Basir
University of Pittsburgh

shb105@pitt.edu

MS207**Data and Operator-Adapted Basis for Predictive Reduced Order Modeling**

This work presents advances towards the development of effective projection-based reduced order models (ROMs) for complex multi-scale multi-physics problems. In particular, we emphasize the process of accomplishing true predictivity in a scalable setting. Dimension reduction approaches based on static manifolds - linear or non-linear - are not effective in predictive modeling of multi-scale problems with significant transport effects. To address this issue, we present an adaptive formulation in which the basis vectors and sampling points are adapted online using a novel non-local procedure. We compare and contrast adapting the basis to the solution manifold, and to the underlying PDE operator. Results are presented in a combustion dynamics problem involving compressible turbulence, reactions, and acoustics.

Christopher R. Wentland
University of Michigan, Ann Arbor
chriswen@umich.edu

Cheng Huang
University of Kansas
Department of Aerospace Engineering
chenghuang@ku.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS208**Analysis of the Shallow Water Model with Two Velocities**

The Shallow water equations (also called Saint-Venant's equations) are the usual model governing fluid flow in the rivers, channels or the oceans. They are used, for example, for the protection of the environment, the prediction of tides and storm surges, the transport of the sediment or the study of floods. Some references in the literature propose an improvement of the Shallow water equations to take into account the vertical profile of the horizontal velocity. The objective of this work is to develop a scheme of the model with two velocities in the vertical profil based on a recent analysis of the Riemann problem. We look for a scheme able to exactly recover any steady solution in 1D over arbitrary topography. To do so, first we analyse the moving steady solutions following the Bernoulli's principle for C^1 regular solutions and the Rankine-Hugoniot conditions as well as the dissipation of entropy at a point of discontinuity. We then propose several well-balanced and positivity preserving Approximate Riemann solvers. Finally, we validate our results with numerical simulations.

Nelly Boulos Al Makary
INRIA Paris
nelly.boulos-al-makary@inria.fr

MS208**A New Unstaggered Locally Divergence-Free Finite Volume Scheme for Ideal and Shallow Water****Magnetohydrodynamics**

Many mathematical models of fundamental interest are constrained by the infamous divergence-free condition of magnetic fields or velocity profiles. Seen within a wide range of astrophysical, geophysical, and engineering applications, such divergence-free constraints are physically and analytically exact. On a discrete level, however, an improper treatment of an identically-zero divergence may lead to large instabilities in solutions even when applying methods that already successfully simulate fluids without this constraint. Thus, a careful algorithmic construction is required to ensure these constraints are exactly preserved within numerical approximations. In this talk, we restrict our attention to an important sub-class of divergence-free systems the magnetohydrodynamic equations, keeping in mind that the developed method can be extended to other models constrained by zero-divergence. We present a new method that exactly preserves the divergence-free condition of the magnetohydrodynamic system on an unstaggered mesh. The designed method has been successfully tested on several examples.

Michael Redle
North Carolina State University, U.S.
mtredle@ncsu.edu

Alina Chertock
North Carolina State University
Department of Mathematics
chertock@math.ncsu.edu

Alexander Kurganov
Southern University of Science and Technology of China
alexander@sustech.edu.cn

Vladimir Zeitlin
ENS/University of Paris 6, France
zeitlin@lmd.ens.fr

MS208**Moment Model Cascades for Non-Hydrostatic Shallow Flow**

We will discuss some recent results regarding a hierarchy of nonhydrostatic moment equations for shallow flow. These include energy balances and numerical solutions for stationary flow. Shallow flow models are of high practical relevance, however, in certain situations they reach their limits. Due to an averaging process, information on the vertical profile of the flow variables is generally lost. This paper presents a framework for the systematic derivation of dimensionally reduced dispersive equation systems that hold information on the vertical profiles of the flow variables. The derivation from a set of balance laws is based on a splitting of the pressure followed by a same-degree polynomial expansion of the velocity and pressure fields in vertical direction. Dimensional reduction is done via Galerkin projections with weak enforcement of the boundary conditions at the bottom and at the free surface. The resulting equation systems of order zero and one are presented in linear and nonlinear form for Legendre basis functions and an analysis of dispersive properties is given. A numerical experiment shows convergence towards the resolved reference model in the linear stationary case and demonstrates the reconstruction of vertical profiles.

Ullika Scholz
RWTH Aachen University

scholz@mathcces.rwth-aachen.de

MS208

Transient Velocity Profiles in Complex Shallow Flow

Shallow water models are successfully used for simulating geophysical flows like river floods, tsunamis, sediment transport, or debris flows. One important characteristic in all these applications is the underlying material closure, e.g. a Newtonian fluid for water. In order to model more complex fluids like snow or debris as granular material, different closure relations are required. The $\mu(I)$ -rheology is one such closure for granular flow that has already been applied successfully in depth-averaged models. An extension of classical shallow water models, which typically require a constant or predefined vertical velocity profile, is introduced in the *shallow moment method*. This method retains transient information about the vertical flow profile using a finite Legendre expansion to resolve the vertical velocity with time-dependent coefficients. The *shallow moment* approach allows to include more information systematically and generates a hierarchy of models that in the limit, recover the reference equations before depth-averaging and are therefore vertically fully resolved. However, even a low number of basis functions significantly increases the predictive power of the model compared to classical shallow water systems. The *shallow moment method* has mostly been studied for Newtonian fluids. Therefore, this talk will focus on current developments and challenges in modeling complex fluids within the *shallow moment* framework.

Ingo Steldermann

RWTH Aachen University
steldermann@mbd.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

Julia Kowalski
RWTH Aachen University
kowalski@mbd.rwth-aachen.de

MS208

Two-Dimensional Hyperbolic Shallow Water Moment Equations

The recently derived Shallow Water Moment Equations (SWME) are an extension of the Shallow Water Equations and allow for vertical changes in the horizontal velocity. This results in a system that is more accurate in situations where the horizontal velocity varies considerably over the height of the fluid. Unfortunately, these models lack global hyperbolicity. This has previously been observed for the 1D SWME and we show that the loss of hyperbolicity also occurs in the 2D systems. We derive 2D Hyperbolic Shallow Water Moment Equations (HSWME) by modifying the system matrix. After the derivation of the 2D equations, an equilibrium stability analysis of the new system is performed. Equilibrium manifolds are derived and structural stability conditions are checked for these manifolds. The analysis is a direct extension of the work done on the one-dimensional equations. We show that the 2D HSWME contain stable and unstable equilibrium states. Next, a formulation of the SWME in cylindrical coordinates is presented. The model is useful in situations where cylindrical coordinates are more appropriate, e.g., when the application exhibits circular symmetry, as is the case in the

modelling of a tsunami, for example. Finally, a hyperbolic system for axisymmetric flow is derived.

Rik Verbiest

University of Groningen
r.verbiest@rug.nl

MS209

Detecting Inclusions in Soil Profiles

We propose a computational framework to evaluate the presence of inclusions of different kinds of materials in specific soil profiles by full waveform inversion. We construct a cost functional using the data recorded at a grid of receivers in response to signals emitted by a set of sources. Next, we propose possible shapes and locations for the inclusions by topological energy techniques, which we parametrize in terms of a moderate number of parameters. Then, we implement constrained optimization schemes which decrease the cost starting from the selected initial guess and subject to wave equation constraints. Numerical simulations illustrate the performance of the method.

Carolina Abugattas

Universidad Complutense de Madrid, Spain
cabugatt@ucm.es

Ana Carpio

Universidad Complutense de Madrid
Departamento de Matematica Aplicada
carpio@mat.ucm.es

Elena Cebrián

Universidad de Burgos
elenac@ubu.es

Gerardo E. Oleaga
Universidad Complutense de Madrid, Spain
oleaga@mat.ucm.es

MS209

Reconstruction of the Mechanical Properties in Optical Coherence Elastography

In this talk we present a mathematical model to reconstruct the mechanical properties of an elastic medium, in the optical coherence elastography imaging technique. We start by addressing the numerical simulation of the mechanical wave propagation and induced displacements. This direct problem is the computational basis to solve the inverse problem which consists of determining the parameters that characterize the mechanical properties of the medium, knowing the displacement field for a given excitation. We formulate the inverse model problem as a PDE-constrained optimization problem, where the objective function measures the discrepancy between observations and predictions. We will discuss different strategies for learning the space varying elasticity coefficients, which include the use of neural networks. We report several computational results that illustrate the behavior of the proposed methods in terms of accuracy and efficiency. This work was supported by FEDER Funds through the Operational Program for Competitiveness Factors - COMPETE and by Portuguese National Funds through FCT - Foundation for Science and Technology under the PTDC / EMD-EMD / 32162/2017 project "Optical Coherence Elastography, for imaging of the mechanical properties of the retina and by the Centre for Mathematics of the University of Coimbra - UIDB/00324/2020, funded by the Portuguese

Government through FCT/MCTES.

Silvia Barbeiro

CMUC, Department of Mathematics, University of
Coimbra
silvia@mat.uc.pt

MS209

The Same, But Different Discovering Differences Between Neutrophil Elastase and Cathepsin G

Although the general mechanism for serine protease catalysis is well established, some questions still remain. For instance, the two enzymes, neutrophil elastase and cathepsin G, have a lot of structural resemblances; however, elastase degrades virulence factors, while cathepsin G does not. But their crystal structures are remarkably similar. Molecular dynamics simulations have been employed to study the two enzymes in solvent and probe their conformational differences. The two protein structures do not have the same number of residues, which renders the measurement of the similarity over the duration of simulations problematic. In order to tackle the problem, we established and implemented a three-step-protocol: The Frchet distance is used to superimpose the two protein structures, such that only internal motions will be probed in subsequent steps. The dynamic time warping algorithm and its penalty matrix is used to generate suitable pairs of amino acids between the two proteins. Lastly, a symmetrized version of the so-called Kullback-Leibler divergence is employed to quantify the difference and similarities between the two protein dynamics trajectories. The approach showed a subtle difference in a specific residue in the two protein trajectories, which might explain the difference in specificity.

Fabian Schuhmann

Carl von Ossietzky University Oldenburg
fabian.schuhmann@uol.de

Xiangyin Tan

University of Copenhagen
jcv814@alumni.ku.dk

Luca Gerhards

Carl von Ossietzky University Oldenburg
luca.gerhards@uol.de

Heloisa N. Bordallo

University of Copenhagen
bordallo@nbi.ku.dk

Iliia A. Solov'yov

Carl von Ossietzky University Oldenburg
ilia.solovyov@uol.de

MS209

Data Processing and Reduced Order Modeling in Semiconductor Electron Transport

Self-sustained oscillations associated with electron transport in a single miniband semiconductor superlattice are analyzed by means of efficient data processing techniques. The system is modeled by a drift-diffusion equation in one spatial dimension, which describes the temporal evolution of the scalar current density and the distributed electric field in the voltage-biased superlattice. Numerically computed data show that, for some parameter regimes, the dynamics become temporally periodic after a transient:

a traveling solitary wave for the electric field cyclically moves from cathode to anode. The higher order dynamic mode decomposition is applied to uncover the broad frequency spectrum of the spatio-temporal dynamics driving the transport mechanisms. The purely decaying transient can be isolated from the attractor, which is approximated by considering snapshots in short time intervals. On the other hand, the pattern associated with the electric field traveling wave is further investigated by the spatio-temporal Koopman decomposition. The related dispersion diagram and an estimation of the overall propagation velocity are accurately recovered. Finally, the construction of a data-driven reduced order model of the physical system is illustrated. The outcome allows to fast simulate the dynamical response of the superlattice over a range of different configurations.

Filippo Terragni

Universidad Carlos III de Madrid, Spain
fterragn@ing.uc3m.es

Luis Bonilla

Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience
&IndMath
bonilla@ing.uc3m.es

Jose Manuel Vega

Universidad Politecnica de Madrid
josemanuel.vega@upm.es

MS209

Band-Based Binary Indices for the Analysis of Gene Expression Data

The statistical concept of depth induces an ordering from centre outwards in multivariate data. Most depth definitions are unfeasible for dimensions larger than three or four, but the Modified Band Depth (MBD) is a notable exception. It relates the centrality of each individual to its (partial) inclusion in all possible 'j-bands' formed by j distinct elements of the data set. We assess (dis)similarity between pairs of observations by accounting for such bands and constructing binary matrices associated to each pair. From these, contingency tables are calculated and used to derive standard similarity indices. Our approach is computationally efficient and can be applied to bands formed by any number of observations from the data set. We have evaluated the performance of several band-based similarity indices with respect to that of other classical distances in standard tasks in a variety of high dimension, simulated and real gene expression data sets. Our experiments show the benefits of our technique, with some of the selected indices outperforming classical distances.

Aurora Torrente

G. Millan Institute of Fluid Dynamics, Nanoscience
etorrent@est-econ.uc3m.es

MS210

A Boundary-Integral Computational Framework for Molecular Electrostatics Modeling

Even though in molecular solvation small atomic-level scales are important, continuum models can take us a long way. In the case of electrostatics, they give rise to a system of partial differential equations where the Poisson and Poisson-Boltzmann equations are coupled on a molecular surface, which interfaces the solute and solvent regions. A

boundary-integral approach is a natural choice to simulate such system, as the interface and Dirac delta-like charge distributions are accurately represented. There are a number of Poisson-Boltzmann solvers available that are widely used, however, they fall short in terms of usability in a larger computational workflow, and extensibility to incorporate modifications to the model. In this talk, we will introduce the Poisson-Boltzmann & Jupyter (PB&J) code, a boundary integral solver based on Bempp-cl, designed to be easy to use and extend without sacrificing accuracy or speed. We will review the main design principles and present application examples showing its efficiency in terms of user and computer time. Moreover, this Python-based approach is aligned with current developments in computational chemistry and biophysics, making it possible to integrate into existing computational workflows for multi-scale modeling or analysis software, straight from a Jupyter notebook.

Christopher D. Cooper
Department of Mechanical Engineering
Universidad Tecnica Federico Santa Maria
christopher.cooper@usm.cl

MS210

Automatic Synthesis of Fast Algorithms for the Solution of Elliptic PDEs

When using the Fast Multipole Method to evaluate a potential satisfying a Elliptic PDE, we can exploit the knowledge about the potential and the PDE to reduce the cost of the FMM. In this talk, we present two methods that make use of the symbolic form of the potential and the symbolic form of the PDE given by the user to synthesize efficient FMM code. First we describe a method that automatically synthesizes low complexity translation operators for arbitrary potentials satisfying an elliptic PDE using a compression scheme. These operators have the same complexity as the state of the art spherical/plane wave based translation operators. We observe the asymptotic time complexities for these operators and also give an upper bound on the error in the compression scheme. Next we outline a method that automatically synthesizes execution plans for an integral equation representation involving multiple input Green's function and multiple outputs. We introduce tools to reduce the number of FMM calls needed by doing algebraic manipulations on the symbolic integral equation representation. Finally, we describe applications where our methods are able to synthesize efficient execution plans and generate efficient code for each FMM in the execution plan.

Isuru D. Fernando
University of Illinois at Urbana-Champaign
idf2@illinois.edu

MS210

Scalable Parallelization of the Phonon Boltzmann Transport Equation

The Boltzmann Transport Equation (BTE) for phonons is often used to predict thermal transport in semiconductors. The BTE is a seven-dimensional nonlinear integro-differential equation, resulting in difficulty in its solution even after linearization under the single relaxation time approximation. Furthermore, parallelization and load-balancing are challenging given the high dimensionality and variability of the linear systems' conditioning. This work presents a "synthetic" scalable method for solving

the BTE on large-scale clusters. The method includes cell-based parallelization and combined cell-based/band-based techniques. The cell-based parallelization consists of a sparse matrix-vector product (SpMV) that can be integrated with an existing linear algebra library like PETSc. Additionally, we developed a batched SpMV enabling multiple linear systems to be solved simultaneously. The batched method reduces the communication overhead as the grain size becomes smaller. We present numerical experiments to demonstrate our method's excellent scalability and speedups compared with previously existing schemes.

Han D. Tran
School of Computing, University of Utah
hantran@cs.utah.edu

Siddharth Saurav
Ohio State University
siddharthsaurav.1@buckeyemail.osu.edu

Ponnuswamy Sadayappan
University of Utah
saday@cs.utah.edu

Sandip Mazumder
Ohio State University
mazumder.2@osu.edu

Hari Sundar
School of Computing
University of Utah
hari@cs.utah.edu

MS211

Towards a Modified Nested Dissection Ordering to Enhance Low-Rank Compressibility in Sparse Direct Solvers

The advent of rank-structured compression techniques for the solution of large-scale sparse linear systems has been demonstrated to significantly reduce both the computational cost and the memory footprint of sparse direct solvers. The block low-rank (BLR) compression format exploits the blockwise low-rank property of sparse matrices that arise in many scientific applications. In this talk, we investigate the potential of new combinatorial algorithms that enhance BLR solvers through the computation of a smarter ordering of the unknowns prior to the factorization phase. Our purely algebraic approach aims to bridge the gap between traditional dense and sparse linear algebra methods, and improve compression rates at the cost of additional fill-in. Instead of building low-rank clusters on top of nested dissection, we adopt an alternative approach by clustering the unknowns before exhibiting a nested dissection on top of it. We demonstrate the efficiency of this new approach by means of experiments that have been conducted for the PaStiX solver.

Grégoire Pichon
Université Claude Bernard Lyon 1
gregoire.pichon@univ-lyon1.fr

Brian Bantsoukissa
Inria Lyon
brian.bantsoukissa@ens-lyon.fr

Bora Uçar
CNRS

LIP ENS de Lyon
bora.ucar@ens-lyon.fr

MS211

Sampling Algorithms for Distributed-Memory Sparse CP Decomposition

Low rank Candecomp / Parafac (CP) Tensor Decomposition is a powerful computational tool to extract patterns from sparse data, but its computational cost may become intractable for massive sparse tensors. Recently, a number of randomized sketching algorithms have been proposed to drive down the cost of Alternating Least Squares, a popular heuristic to compute the CP decomposition. We extend two algorithms based on statistical leverage-score sampling to the distributed-memory setting, where processor-to-processor communication overhead is a major obstacle to achieving the speedup enjoyed in a shared-memory system. We investigate multiple strategies to combat this problem by adapting techniques from distributed sparse matrix algorithms and produce high-performance implementations of these algorithms. Using NERSC Perlmutter, our communication-avoiding algorithms produce low-rank CP decompositions of billion-scale sparse tensors in seconds.

Vivek Bharadwaj
University of California, Berkeley
vivek_bharadwaj@berkeley.edu

Osman Malik
Lawrence Berkeley National Laboratory
oamalik@lbl.gov

Riley J. Murray
University of California, Berkeley
rjmurray@berkeley.edu

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

James W. Demmel
UC Berkeley
Division of Computer Science
demmel@berkeley.edu

MS211

Incremental Graph Clustering in Parallel

We develop a distributed memory graph clustering algorithm to find clusters in a graph where new nodes and edges are being added incrementally. At each stage of the algorithm, we maintain a summary of the clustered graph computed from all incremental batches received thus far. As we receive a new batch of nodes and edges, we cluster the new graph and merge new clusters with the previous summary clusters. We use sparse linear algebra to perform these operations. Our algorithm would make it possible to find clusters in very large graphs for which regular graph clustering algorithms could not run due to computation/communication bottlenecks. We use this algorithm to cluster billions of metagenomic proteins that are being collected incrementally in various databases.

Md Taufique Hussain, Ariful Azad
Indiana University

mth@indiana.edu, azad@iu.edu

MS211

Decoupling Multivariate Functions

While linear functions are well-understood, for nonlinear (multivariate vector) functions it is unclear how to i) define their complexity, ii) reduce the complexity and iii) increase their interpretability. To answer these questions, we propose a decomposition of nonlinear functions [P. Dreesen, M. Ishteva, and J. Schoukens. Decoupling multivariate polynomials using first-order information and tensor decompositions. SIMAX, 36:864–879, 2015], which can be viewed as a generalization of the singular value decomposition. In this decomposition, univariate nonlinear mappings replace the simpler scaling performed by the singular values. We discuss the computation of the decomposition, which is based on tensor techniques. We also mention an application in nonlinear system identification.

Mariya Ishteva
Vrije Universiteit Brussel
mariya.ishteva@kuleuven.be

Philippe Dreesen
KU Leuven
philippe.dreesen@kuleuven.be

MS211

Optimizing Locality in Dynamic Graph Data Structures

This talk discusses dynamic graph data structures and how to choose data structures that optimize for locality, which is key to performance in graph computations. The focus is on how different use cases, which can cause different data patterns and access patterns, can influence the design of these structures and how to exploit these differences to maximize performance.

Helen Xu
Lawrence Berkeley National Laboratory
hjxu@lbl.gov

Brian Wheatman
JHU
bwheatman@gmail.com

Prashant Pandey
University of Utah
prashant.prashn@gmail.com

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

MS212

Using Machine Learning in Geophysical Data Assimilation

In recent years, data assimilation, and more generally the climate science modelling enterprise have been influenced by the rapid advent of artificial intelligence, in particular machine learning (ML), opening the path to various form of ML-based methodology. In this talk we will schematically show how ML can be included in the prediction and DA workflow in three different ways. First, in a so-called non-intrusive ML, we will show the use of supervised learning

to estimate the local Lyapunov exponents (LLEs) based exclusively on the systems state. In this approach, ML is used as a supplementary tool, added to the given physical model. Our results prove ML is successful in retrieving the correct LLEs, although the skill is itself dependent on the degree of local homogeneity of the LLEs on the systems attractor. In the second and third approach, ML is used to substitute fully or partly a physical model with a surrogate one reconstructed from data.

Alberto Carrassi

University of Reading
n.a.carrassi@reading.ac.uk

MS212

Gradient Flows and Randomised Thresholding: Sparse Inversion and Classification

Sparse inversion and classification problems are ubiquitous in modern data science and imaging. They are often formulated as non-smooth minimisation problems. In sparse inversion, we minimise, e.g., the sum of a data fidelity term and an L1/LASSO regulariser. In classification, we consider, e.g., the sum of a data fidelity term and a non-smooth Ginzburg–Landau energy. In this work, we study splitting from a stochastic continuous-time perspective. Indeed, we define a differential inclusion that follows one of the two subtarget function’s negative subgradient at each point in time. The choice of the subtarget function is controlled by a binary continuous-time Markov process. The resulting dynamical system is a stochastic approximation of the underlying subgradient flow. We investigate this stochastic approximation for an L1-regularised sparse inversion flow and for a discrete Allen-Cahn equation minimising a Ginzburg–Landau energy. In both cases, we study the longtime behaviour of the stochastic dynamical system and its ability to approximate the underlying subgradient flow at any accuracy. We illustrate our theoretical findings in a simple sparse estimation problem and also in a low-dimensional classification problem.

Jonas Latz

Heriot-Watt University
j.latz@hw.ac.uk

MS212

A Bayesian Approach to Modelling Biological Pattern Formation with Limited Data

Pattern formation plays an important role in the development of living organisms. Since the classical work of Alan Turing, a pre-eminent way of modelling has been through reaction-diffusion mechanisms. Alternative models have been proposed, that link dynamics of diffusing molecular signals with tissue mechanics. Model validation is complicated as in many experimental situations only the limiting, stationary regime of the pattern formation process is observable, without any knowledge of the transient behaviour or the initial state. To overcome this problem, the initial state of the model can be randomised. But then fixed values of the model parameters correspond to a family of patterns rather than a fixed stationary solution, and standard estimation approaches, such as the least squares method, are not suitable. Instead, statistical characteristics of the patterns should be compared, which is difficult given the typically limited amount of available data in practical applications. To deal with this problem, we extend the recently developed statistical approach (the Correlation Integral Likelihood method) for parameter identi-

fication by pattern data. We introduce modifications that allow increasing the accuracy of the identification process without resizing the data set. The proposed approach is tested using different classes of pattern formation models and severely limited data sets.

Robert Scheichl

Universitaet Heidelberg
r.scheichl@uni-heidelberg.de

MS212

Convergence Guarantees for MCMC in Statistical Inverse Problems: Upper and Lower Bounds

In this presentation we exhibit examples of high-dimensional unimodal posterior distributions arising in non-linear regression models with Gaussian process priors for which worst-case (cold start) initialised MCMC methods typically take an exponential run-time to enter the regions where the bulk of the posterior measure concentrates. The counter-examples hold for general MCMC schemes based on gradient or random walk steps, and the theory is illustrated for Metropolis-Hastings adjusted methods such as pCN and MALA.

Sven Wang

Massachusetts Institute of
Technology
svenwang@mit.edu

MS213

An Overview of Adaptive Mesh Refinement: Algorithms and Applications

Adaptive mesh refinement (AMR) is one of several techniques for adapting the spatial resolution of a simulation in particular regions of the spatial domain. Block-structured AMR specifically refines the mesh by defining locally structured regions with finer spatial, and possibly temporal, resolution. This combination of locally structured meshes within an irregular global hierarchy is in some sense the best of both worlds in that it enables regular local data access while enabling greater flexibility in the overall computation. Algorithms that evolve the solution on a hierarchical mesh are more complex than algorithms operating on a single uniform domain. The nature of the complexity typically depends on the mathematical character of the governing equations and the additional physics being represented across levels. In this talk I will give a very brief overview of the different types of block-structured AMR, different strategies for creating the mesh hierarchy, different approaches to time advancement, and some of the ways that algorithms must be adapted for use on an AMR hierarchy.

Ann S. Almgren

Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

MS213

A Fast, Adaptive, Matrix-Free Direct Elliptic Solver on Quadtree Meshes

Elliptic partial differential equations (PDEs) are difficult to solve on adaptive meshes due to their global nature and the dynamic mesh. Direct solvers, like the Hierarchical Poincar-Steklov (HPS) method, work by forming a factorization of the system matrix that can be applied to

as many right-hand sides or time steps as needed. However, when the mesh changes, the factorization must be completely reformed. The HPS method attempts to solve this by forming a set of solution operators that are built on an adaptive, quadtree structure. The HPS method can achieve linear $O(N)$ factorization performance in 2D and is matrix-free as it does not need to form the system matrix. It also allows for a dynamic update to the solution operators when the mesh changes, solving the so-called near-by problem. We present our implementation of the HPS method on quadtree meshes with the p4est software library. We will show results solving Poissons equation for eventual coupling with a hyperbolic solver through the ForestClaw software library. In addition, we will present our progress on a parallel implementation on distributed memory architectures.

Damyn M. Chipman, Donna Calhoun
Boise State University
DamynChipman@u.boisestate.edu, donnacal-
houn@boisestate.edu

MS213

Insights into Compressible Reacting Flows in Complex Geometry Using Adaptive Mesh Refinement Simulations

With the increased availability of exascale computing hardware, detailed simulations of realistic devices can be performed at practically relevant time and length scales. Insights into the multiscale driving mechanisms in compressible reacting flow systems with complex geometry, such as combustors, can be used for design optimization and technology improvements. However, to effectively perform these simulations, advanced numerical algorithms must be used to maintain solution accuracy without incurring undue computational costs. PeleC, part of the Pele suite of codes, leverages block-structured adaptive mesh refinement (AMR) through the AMReX library to capture fine-scale flow features in compressible reacting flows. In this talk, we discuss recent improvements to the numerical algorithms, particularly in regard to describing flows at complex boundary structures, and PeleC's performance on exascale computing hardware. We will demonstrate that PeleC is well-suited for modern, extreme-scale, heterogeneous compute platforms.

Marc Henry de Frahan, Jon Rood, Nicholas Wimer,
Bruce A. Perry, Hariswaran Sitaraman, Shashank
Yellapantula
National Renewable Energy Laboratory
marc.henrydefrahan@nrel.gov, jon.rood@nrel.gov,
nicholas.wimer@nrel.gov, bruce.perry@nrel.gov,
hariswaran.sitaraman@nrel.gov,
shashank.yellapantula@nrel.gov

Lucas Esclapez
National Renewable Energy Lab
lucas.esclapez@nrel.gov

Sreejith Appukuttan, Marc Day
National Renewable Energy Laboratory, Colorado
sreejith.nadakkalappukuttan@nrel.gov, mar-
cus.day@nrel.gov

MS213

Adaptive Grids for Algebraic Stabilizations of

Convection-Diffusion-Reaction Equations

Non-linear discretizations are necessary for convection-diffusion-reaction equations for obtaining accurate solutions that satisfy the discrete maximum principle (DMP). Algebraic stabilizations belong to the very few finite element discretizations that satisfy this property. This talk considers three algebraically stabilized finite element schemes for discretizing convection-diffusion-reaction equations on adaptively refined grids. These schemes are the algebraic flux correction (AFC) scheme with Kuzmin limiter, the AFC scheme with BJK limiter, and the recently proposed Monotone Upwind-type Algebraically Stabilized (MUAS) method. Both conforming closure of the refined grids and grids with hanging vertices are considered based on a residual-based a posteriori error estimator. A non-standard algorithmic step becomes necessary before these schemes can be applied on grids with hanging vertices. The assessment of the schemes is performed with respect to the satisfaction of the global discrete maximum principle (DMP), the accuracy, e.g., smearing of layers, and the efficiency in solving the corresponding nonlinear problems.

Abhinav Jha
University of Stuttgart
jha.abhinav0207@gmail.com

Volker John
Weierstrass Institute for Applied Analysis and Stochastics
Research Group Numerical Mathematics and Sci.
Computing
john@wias-berlin.de

Petr Knobloch
Charles University
knobloch@karlin.mff.cuni.cz

MS213

Multiscale Fluid Dynamics Enabled by Adaptive Mesh Refinement in Aerospace and Wind Energy Applications

Many fluid dynamics problems found in aerospace and wind energy applications require multi-scale simulation capabilities to track turbulent vortical fluid structures such as the wakes generated by wind turbines or helicopter rotors. To resolve these complex multiscale phenomena, parallel adaptive mesh refinement paired with overset grid technologies provides a computationally efficient and scalable solution. We highlight the octree-based AMR approach in tandem with overset grid assembly for simulating these relative-body motion problems. Further, we leverage characteristics of octree-based AMR algorithms to accelerate the overset grid assembly process for implicit hole cutting. Lastly, we'll discuss future capabilities needed for AMR and overset to leverage next-generation heterogeneous computing systems.

Andrew Kirby
University of Wyoming
akirby@uwyo.edu

Dimitri Mavriplis
Department of Mechanical Engineering
University of Wyoming
mavripl@uwyo.edu

MS214

Exploration Strategies for Control of Chaotic Dy-

namical Systems Using Reinforcement Learning

Predicting and controlling the future state of a complex physical system is a cornerstone in a wide range of situations. This aim however often remains difficult to achieve. In particular, for a high-dimensional multiscale nonlinear system, a model is not necessarily available or usable due to real-time control constraints which drastically limit the affordable computational complexity. Data-driven control circumvents the need for an a priori model and constitutes an attractive approach. In this talk, we will focus on a Reinforcement Learning strategy fully leveraging its theoretical foundation in optimal control theory while accounting for challenges of practical situations such as weak observability. We identify the multiple factors and their influence on the policy identification without relying on physics constraints or prior knowledge. In particular, we develop consistent exploration strategies allowing to carefully sample new observations in order to improve knowledge of the system. As a result, better control policies could be obtained with fewer interactions with the system, resulting in a faster learning. To achieve this, we rely on tools from information theory and optimal design of experiments to quantify the knowledge gathers during the process. Our framework will be illustrated on simplified models of turbulent flows such as the Kuramoto-Sivashinsky equations.

Remy Hosseinkhan Boucher

LISN-CNRS, Univ. Paris-Saclay, France
remy.hosseinkhan@universite-paris-saclay.fr

Michele-Alessandro Bucci

Inria Tau, LISN, CNRS, Université Paris-Saclay
michele-alessandro.bucci@safrangroup.com

Onofrio Semeraro

LISN, CNRS, Université Paris-Saclay
semeraro@lisn.upsaclay.fr;

Lionel Mathelin

LISN, CNRS, Univ. Paris-Saclay
lionel.mathelin@lisn.upsaclay.fr

MS214

Machine-Learning-Based Spectral Methods for Partial Differential Equations

Spectral methods are an important part of scientific computing's arsenal for solving partial differential equations (PDEs). However, their applicability and effectiveness depend crucially on the choice of basis functions used to expand the solution of a PDE. The last decade has seen the emergence of deep learning as a strong contender in providing efficient representations of complex functions. In the current work, we present an approach for combining deep neural networks with spectral methods to solve PDEs. In particular, we use a deep learning technique known as the Deep Operator Network (DeepONet) to identify candidate functions on which to expand the solution of PDEs. We have devised an approach that uses the candidate functions provided by the DeepONet as a starting point to construct a set of functions that have the following properties: (1) they constitute a basis, (2) they are orthonormal, and (3) they are hierarchical, i.e., akin to Fourier series or orthogonal polynomials. We have exploited the favorable properties of our custom-made basis functions to both study their approximation capability and use them to expand the solution of linear and nonlinear time-dependent PDEs. The

proposed approach advances the state of the art and versatility of spectral methods and, more generally, promotes the synergy between traditional scientific computing and machine learning.

Saad Qadeer

University of North Carolina Chapel Hill
saad.qadeer@pnml.gov

Panos Stinis

Pacific Northwest National Laboratory
panos.stinis@pnml.gov

MS214

Learning Homogenized Constitutive Models in Viscoelasticity and Viscoplasticity

The macroscopic behavior of materials is governed in part by small-scale rapidly-varying material properties. Fully resolving these features within the balance laws thus involves expensive fine-scale computations which need to be conducted on macroscopic scales. The theory of homogenization provides an approach to derive effective macroscopic equations which eliminates the small scales by exploiting scale separation. An accurate homogenized model avoids the computationally-expensive task of numerically solving the underlying balance laws at a fine scale. In simple settings the homogenization produces an explicit formula for a macroscopic constitutive model, but in more complex settings it may only define the constitutive model implicitly. In these complex settings machine learning can be used to learn the constitutive model from localized fine-scale simulations. In the case of one-dimensional viscoelasticity, the linearity of the model allows for a complete analysis. For this case, we derive a homogenized constitutive model and develop a theory to prove that the model may be approximated by a recurrent neural network (RNN) model that captures the memory; this may be thought of as discovering appropriate internal variables. Simulations are presented which validate the theory, and additional numerical experiments demonstrate extension of the methodology to higher dimensions and to nonlinear viscoplasticity.

Margaret K. Trautner

California Institute of Technology
mtrautne@caltech.edu

MS214

Data in Deep Operator Learning

In this talk, I will show the importance of data in deep operator learning. First, I will show the accuracy of a MOD-Net (Model-operator-data) with physics-informed loss to learn the operator of a PDE could be greatly improved by using few data points computed by traditional numerical schemes on coarse-grained grids. Second, I will use data-driven DeepONet to solve an inverse problem of fractional problem. Finally, I will show a good sampling scheme is extremely important for learning the operator of a stiff ODE system, such as combustion chemical system, where our neural network model is accurate and efficient in many different working conditions for large reaction systems.

Zhi-Qin Xu

Shanghai Jiao Tong University

xuzhiqin@sjtu.edu.cn

MS215

From Perspective Maps to Epigraphical Projections

The projection onto the epigraph or a level set of a closed proper convex function can be achieved by finding a root of a scalar equation that involves the proximal operator as a function of the proximal parameter. We develop the variational analysis of this scalar equation. The approach is based on a study of the variational-analytic properties of general convex optimization problems that are (partial) infimal projections of the the sum of the function in question and the perspective map of a convex kernel. When the kernel is the Euclidean norm squared, the solution map corresponds to the proximal map, and thus the variational properties derived for the general case apply to the proximal case. Properties of the value function and the corresponding solution map including local Lipschitz continuity, directional differentiability, and semismoothness are derived. An SC1 optimization framework for computing epigraphical and level-set projections is thus established. Numerical experiments on 1-norm projection illustrate the effectiveness of the approach as compared with specialized algorithms.

Tim Hoheisel

McGill University
Montreal, QC, Canada
tim.hoheisel@mcgill.ca

Michael P. Friedlander
Dept of Computer Science
University of British Columbia
mpf@cs.ubc.ca

Ariel Goodwin
McGill University
ariel.goodwin@mail.mcgill.ca

MS215

An Inexact Trust-Region Algorithm for Nonsmooth Nonconvex Optimization

In this talk, we develop a new trust-region method to minimize the sum of a smooth nonconvex function and a nonsmooth convex function. Our method is unique in that it permits and systematically controls the use of inexact objective function and derivative evaluations. We prove global convergence of our method in Hilbert space and analyze the worst-case complexity to reach a prescribed tolerance. Our method employs the proximal mapping of the nonsmooth objective function and is simple to implement. Moreover, when using a quadratic Taylor model, our algorithm represents a matrix-free proximal Newton-type method that permits indefinite Hessians. We demonstrate the efficacy of our algorithm on various examples from data-science and PDE-constrained optimization.

Drew P. Kouri

Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS215

Surrogation Methods for Nonconvex Nondifferen-

tial Problems in Finite Dimensions

Optimization problems with coupled nonconvexity and nondifferentiability are particularly challenging to solve. Foremost is the question of what kind of ‘stationary solutions’ can be computed in practice. In this talk, we present several classes of these problems that arise in various topical areas of contemporary interests. We give an overview of the family of surrogation methods and illustrate their applications to a few applied problems. If time permits, we will also briefly discuss some challenging nonconvex stochastic optimization problems that so far have been minimally studied at best. Much of the materials to be presented is documented in the recent research monograph on Modern Nonconvex and Nondifferentiable Optimization published by SIAM and co-authored by the presenter and Dr. Ying Cui of the University of Minnesota.

Jong-Shi Pang

University of Southern California
jongship@usc.edu

Ying Cui

University of Minnesota
yingcui@umn.edu

MS215

Inexact Proximal Newton Methods in Hilbert Spaces

We consider proximal Newton methods in Hilbert spaces for the solution of nonsmooth variational problems. In contrast to classical proximal gradient methods on \mathbb{R}^n these methods enjoy fast local convergence and thus need only few iterations until convergence. However, the subproblems are typically harder, and we apply an inner nonsmooth multigrid solver (TNNMG) for their efficient inexact solution. We discuss inexactness criteria and their relation to global and fast local convergence of our algorithm. Finally we present some applications and numerical examples.

Anton Schiela, Bastian Pötzl

University of Bayreuth
anton.schiela@uni-bayreuth.de,
bastian.poetzl@uni-bayreuth.de

Patrick Jaap

TU Dresden
patrick.jaap@tu-dresden.de

MS215

Regularized Smoothing for Solution Mappings of Convex Problems, with Applications to Two-Stage Stochastic Programming and Some Hierarchical Problems

Many modern optimization problems involve in the objective function solution mappings or optimal-value functions of other optimization problems. In most/many cases, those solution mappings and optimal-value functions are nonsmooth, and the optimal-value function is also possibly nonconvex (even if the defining data is smooth and convex). Moreover, stemming from solving optimization problems, those solution mappings and value-functions are usually not known explicitly, via any closed formulas. We present an approach to regularize and approximate solution mappings of fully parametrized convex optimization problems that combines interior penalty (log-barrier) with Tikhonov

regularization. Because the regularized solution mappings are single-valued and smooth under reasonable conditions, they can also be used to build a computationally practical smoothing for the associated optimal-value function and/or solution mapping. Applications are presented to two-stage (possibly nonconvex) stochastic programming, and to a certain class of hierarchical decision problems that can be viewed as single-leader multi-follower games.

Mikhail V. Solodov

IMPA-Instituto de Matematica Pura e Aplicada, Brazil
solodov@impa.br

Pedro Borges

BigO Tech

pedro.borges.melo@gmail.com

Claudia Sagastizabal

IMECC-Unicamp and CEMEAI

sagastiz@unicamp.br

MS216

Quadratures for Edge and Corner Discretization Schemes in Three Dimensions

In this talk, we present an RCIP-type scheme for discretizing corner and edge singularities in three dimensions. The resulting scheme attempts to resolve the geometric singularities in a weak sense. We demonstrate the performance of the algorithm in the context of FMM-based field solvers.

Zydrunas Gimbutas

National Institute of Standards and Technology

zydrunas.gimbutas@nist.gov

MS216

Fast and Accurate Boundary Element Methods for Large-Scale Computational Acoustics

The Helmholtz equation can be numerically solved using the boundary element method (BEM) by reformulating the volumetric partial differential equation as a boundary integral equation on the surface of the scatterers. The explicit use of Greens functions avoids domain truncation of unbounded regions and accurately models wave propagation through homogeneous domains. Triangular surface grids with few linear elements per wavelength suffice to discretize the boundary integral operators with a Galerkin method. However, the dense matrices require a large amount of memory, and the convergence of iterative linear solvers deteriorates at high frequencies and large material contrasts relative to the exterior medium. This talk presents novel preconditioning techniques for coupled BEM formulations modelling multiple scattering and transmission through on-surface radiation conditions that are especially effective at high frequencies. Also, we designed novel boundary integral formulations that are well-conditioned when different domains have high contrasts in density or speed of sound. We used our fast BEM implementation to simulate focused ultrasound propagation in the human body, which can be translated to important biomedical applications such as the non-invasive treatment of liver cancer and neuromodulation of the brain. We validated the methodology within an international benchmarking exercise and implemented all functionality in our open-source Python library, OptimUS.

Elwin Van 't Wout

Pontifica Universidad Catolica de Chile

e.wout@uc.cl

Pierre G elat, Reza Haqshenas

University College London

p.gelat@ucl.ac.uk, r.haqshenas@ucl.ac.uk

MS216

High-Order Boundary Integral Methods for Slender Bodies

The dynamics of active and passive filaments in viscous fluids is frequently used as a model for many complex fluids in biological systems such as: microtubules which are involved in intracellular transport and cell division; flagella and cilia which aid in locomotion. The numerical simulation of such systems is generally based on slender-body theory which give asymptotic approximations of the solution. However, these methods are low-order and cannot enforce no-slip boundary conditions to high-accuracy, uniformly over the boundary. Boundary-integral equation methods which completely resolve the fiber surface have so far been impractical due to the prohibitive cost of current layer-potential quadratures for such high aspect-ratio geometries. In this talk, we will present new quadrature schemes which make such computations possible and new integral equation formulations which lead to well-conditioned linear systems upon discretization. We will present numerical results to show the efficiency of our methods.

Dhairya Malhotra

Flatiron Institute

dmalhotra@flatironinstitute.org

Alex H. Barnett

Flatiron Institute

Simons Foundation

abarnett@flatironinstitute.org

MS216

Kinematic Boundary Element Approaches to Imaging Earthquake Cycle Activity

Large earthquakes occur infrequently but have substantial impacts on both society and the evolution of seismogenic fault systems. Forecasting these events is difficult because of small sample sizes, uncertain mechanics, and the questionable identification of precursor signals. Here, we describe a physics-based approach to imaging fault systems during the slow interseismic phase of the earthquake cycle, when elastic strain is temporarily stored before it is released in great earthquakes. We do this through the development of large-scale kinematic fault system models that integrate the motions of tectonic plates with a quasi-static representation of earthquake cycle kinematics. To accurately represent elastic deformation in geometrically complex fault systems, we exploit closed form boundary integral solutions to relate observations of surface deformation to fault system activity at depth.

Brendan J. Meade

Harvard University

meade@fas.harvard.edu

MS216

A Boundary Integral Method for the Klein-Gordon Equation with Interface Waves

In this talk, we analyze the behavior of solutions to the

time-harmonic Klein-Gordon equation in the presence of an interface. For certain jump conditions, solutions are localized near, and propagate along, the interface. We present a well-posed boundary integral formulation for this PDE and describe a fast iterative algorithm for solving it. We illustrate our theoretical results with several numerical examples of solutions and scattering effects.

Solomon Quinn
University of Chicago
solomonquinn@uchicago.edu

Guillaume Bal
University of Chicago
Departments of Statistics and Mathematics
guillaumebal@uchicago.edu

Jeremy Hoskins
University of Chicago
jeremyhoskins@uchicago.edu

Manas N. Rachh
Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS217

Outreach for Better Scientific Software: The IDEAS Productivity Project

Since 2017, the IDEAS Productivity project has been supported by the DOE Exascale Computing Project (ECP) to help improve developer productivity and software sustainability and reliability. With roughly 1,000 participants involved in the ECP, it was immediately clear that the small IDEAS project team could not provide direct hands-on support for all of the software development efforts in the ECP. A major portion of IDEAS effort, therefore, goes into outreach activities, aimed at (1) raising the level of awareness and discussion of software-related issues in the community, (2) providing seminars and training to help developers improve their skills, (3) an online portal for the exchange of information and resources on scientific software development, and (4) a fellowship program to help bring new voices and new ideas to the scientific software community. In this talk, I will provide more detail about why and how IDEAS approaches these areas and discuss some of the impacts we've seen from them.

David E. Bernholdt
Oak Ridge National Laboratory
Computer Science and Mathematics Division
bernholdtde@ornl.gov

MS217

Continuous Deployment on HPC Clusters Using Spack

Installing highly optimized software stacks on specific HPC architectures is a tenuous task. Spack solves the dependency hell problem, configuring and installing libraries. The remaining challenge is to exploit Spack automatic installations with already installed vendor software. Moreover, we want to keep track of all installed software. We use git as a deployment notebook and database of machine configurations and rely on Gitlab-CI to drive installation processes. We will present this infrastructure named Spack Organizer and will show how this tool can boost software

quality and development productivity.

Cedric Chevalier
CEA, France
cedric.chevalier@cea.fr

Jordan Galby, Julien Racine, Franck Ledoux
CEA
jordan.galby@cea.fr, julien.racine@cea.fr,
franck.ledoux@cea.fr

MS217

E4S Extreme-Scale Scientific Stack

The Extreme-scale Scientific Software Stack (E4S) is a collection of open source software packages for high performance computing. The E4S stack comes with up to 100+ HPC applications, libraries and tools, MPI, development tools such as TAU, HPCToolkit, PAPI, math libraries including PETSc and Trilinos, and node-level programming libraries including Kokkos and RAJA. E4S is available for use via containers, buildcache, AWS EC2 image, and facility tuned spack based deployments. E4S provides a new model for providing a standard set of software stack to HPC centers with dedicated support to help bridge the gap between HPC facilities and application developers of E4S products.

Sameer Shende
Dept. of Computer and Information Science
University of Oregon
sameer@cs.uoregon.edu

MS217

Achieving and Maintaining Performance and Performance Portability Within the Albany Multi-Physics Code: Perspectives and Tools

The focus of this talk is the Albany multi-physics finite element code, a software framework created using the Agile Components code development strategy. This strategy enables the rapid development of parallel efficient multi-physics tools by requiring that: (1) application codes be built primarily from modular pieces (independently developed libraries, e.g., Trilinos), abstract class hierarchies, and template-based generic classes, and (2) projects both leverage and contribute to a comprehensive set of software components. Specifically, I will discuss our experience with and perspectives on evaluating, improving and maintaining performance and performance-portability within Albany. For performance portability, Albany relies primarily on the Kokkos library and programming model, which enables the same code to run correctly and efficiently, with reasonable scalability, across a variety of computer architectures including the most current generation of GPUs. I will present some key performance-portability developments for the land ice modeling application implemented within Albany and known as Albany Land Ice (ALI), covering topics including both finite element assembly and linear solvers. I will also discuss some recently-developed tools that have helped us maintain performance of Albany across a variety of platforms, namely automatic parameter tuning and automated performance testing/analysis.

Irina K. Tezaur, Max Carlson, Jerry Watkins
Sandia National Laboratories
ikalash@sandia.gov, maxcarl@sandia.gov,

jwatkin@sandia.gov

MS218

High-Performance Computing in Tensor Networks: the Current Status and the Challenges Ahead

Numerical techniques based on tensor networks has been used in the past twenty years to tackle problems in quantum many-body physics. In recent years, tensor network methods have witnessed applications beyond quantum physics to include quantum circuits, artificial intelligence and lattice gauge theory, just to name a few. This expansion in application domain is accompanied by an increase in the scale of the simulations performed and their computational complexity. In order to address complexity and scale, there have been few initiative towards the introduction of high-performance and parallel computing methods in tensor networks computations. The aim of this contribution is to highlight the current status of HPC in tensor networks and give a measure of some of the challenges ahead, setting the stage for all the contributions to this minisymposium.

Edoardo Di Napoli
Juelich Supercomputing Centre, Germany
e.di.napoli@fz-juelich.de

MS218

Low-Rank Tensor Approximation for the Numerical Solution of Elliptic and Parabolic Problems

Low-rank tensor decompositions are an attractive tool for the numerical solution of differential equations, reducing the complexity of the problem by the adaptive low-parametric approximation using techniques of numerical linear algebra and optimization. We focus on a multilevel approximation scheme based on the tensor decomposition known as matrix product states (MPS) in computational quantum physics and as tensor train (TT) in computational mathematics. This scheme has been shown, both theoretically and experimentally, to efficiently approximate functions with algebraic singularities and highly-oscillatory solutions to multiscale diffusion problems, achieving exponential convergence with respect to the total number of representation parameters. We present recent results on the use of the multilevel MPS-TT representation for the numerical solution of elliptic and parabolic PDE problems, which involves the multilevel preconditioning and the iterative solution of resulting linear systems in low-rank form. The approach is based on generic finite-element spaces and amounts to the adaptive computation of a collection of quasi-optimal low-dimensional subspaces. The use of low-rank tensor approximation in combination with robust multilevel preconditioning allows to work with extravagantly large finite-element spaces and leads to data-driven computations with effective discretizations adapted to the data, which contrasts with the use of approximation spaces designed analytically.

Vladimir Kazeev
University of Vienna
vladimir.kazeev@univie.ac.at

MS218

Introduction to the ITensor Software Library for Tensor Network Calculations

ITensor is a software library for developing tensor network

algorithms. Modeled on tensor diagram notation, ITensor has a unique memory-independent tensor interface which allows users to focus on the connectivity of a tensor network without manually bookkeeping tensor indices. ITensor has primarily been used in applications to quantum many-body physics, but has found applications in a variety of fields such as quantum computing, machine learning, chemistry, and differential equations. Here we will introduce the ITensor interface and discuss recent advancements in the library such as multithreaded block sparse operations, general tensor network optimization, GPU support, and more as time allows.

Karl Pierce
Flatiron Institute
kpierce@flatironinstitute.org

Matthew Fishman
Flatiron Institute, New York
mfishman@flatironinstitute.org

Miles Stoudenmire
Flatiron Institute
mstoudenmire@flatironinstitute.org

MS218

2D Tensor Network Optimization via Automatic Differentiation

Two-dimensional quantum many-body systems exhibit some of the most elusive problems in condensed matter physics, whose solution is foreseen to impact material science and information processing. In recent years, practical progress has been made towards finding an approximation to their ground-state(s) by means of tensor-network ansatzes, in particular (infinite) projected entangled pair states (iPEPS). Based on the variational principle, such ground-state search can be formulated as a highly non-linear optimization problem, for which well established gradient-based non-linear optimization strategies like conjugate gradient or quasi-Newton methods could be employed. Computing explicitly the gradient via a plethora of tensor contractions to be programmed case-by-case has traditionally obstructed this approach. Nowadays, algorithmic differentiation has considerably eased the path, and the variational procedure has been demonstrated to yield lower energies and less biased many-body configurations than alternative methods like imaginary-time evolution. Here, we report on our implementation of a flexible and efficient library for this state-of-the-art approach in the programming language Julia, based on the TensorKit package, and we showcase some results on frustrated 2d quantum spin models.

Matteo Rizzi
Forschungszentrum Juelich
Germany
m.rizzi@fz-juelich.de

Erik Weerda
University of Cologne
weerda@thp.uni-koeln.de

MS218

The Importance of Mode Ordering of Intermediate Tensors

The contraction of tensor networks dominates the runtime of many scientific fields (e.g., quantum circuit simulations,

condensed matter physics), requiring thousands of compute hours on sophisticated hardware systems. While the inputs and the output of a tensor network contraction are fixed, the intermediate tensors are not; this provides both additional complexity but also ample opportunity for optimizations. This talk outlines a constructive way of choosing different mode orders (i.e., data layouts) for the intermediate tensors and highlights its impact on end-to-end performance on the simulation of Google's 53-bit Sycamore quantum circuit.

Paul Springer
Nvidia
pspringer@nvidia.com

Ali Charara, Markus Hohnerbach
NVIDIA
acharara@nvidia.com, mhohnerbach@nvidia.com

MS219

Brownian HydroDynamics in Doubly-Periodic Confined Geometries

We develop a linearly-scaling variant of the Force Coupling Method for computing hydrodynamic interactions among particles confined to a doubly-periodic geometry with either a single bottom wall or two walls (slit channel) in the aperiodic direction. Our spectrally-accurate Stokes solver uses the Fast Fourier Transform (FFT) in the periodic xy plane and Chebyshev polynomials in the aperiodic z direction normal to the wall(s). We decompose the problem into a doubly-periodic subproblem in the presence of particles (source terms) with free-space boundary conditions in the z direction, which we solve by borrowing ideas from a recent method for rapid evaluation of electrostatic interactions in doubly-periodic geometries [Maxian et al., J. Chem. Phys., 154, 204107, 2021], and a correction subproblem to impose the boundary conditions on the wall(s). Instead of the traditional Gaussian kernel, we use the exponential of a semicircle kernel to model the source terms (body force) due to the presence of particles, and provide optimum values for the kernel parameters that ensure a given hydrodynamic radius with at least two digits of accuracy and rotational and translational invariance. The computation time of our solver, which is implemented in graphical processing units, scales linearly with the number of particles, and allows computations with about a million of particles in less than a second for a sedimented layer of colloidal microrollers.

Aref Hashemi
Courant Institute of Mathematical Sciences
Mathematics Department
aref@cims.nyu.edu

Raul P. Pelaez
Universidad Autónoma de Madrid
raul.perez@uam.es

Sachin Natesh
University of Colorado Boulder
Applied Mathematics
sachin.natesh@gmail.com

Brennan Sprinkle
Colorado School of Mines
bsprinkl@mines.edu

Ondrej Maxian

Courant Institute
New York University
om759@nyu.edu

Zecheng Gan, Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
zg17@nyu.edu, donev@courant.nyu.edu

MS219

A Novel Polydisperse Gaussian-Moment Model for the Air Dispersion of Virus-Containing Aerosols

The COVID-19 pandemic has highlighted the need for improved models to describe the transport of virus-containing droplets and aerosols. Although Lagrangian-based models have been widely adopted for the study of such bio-aerosols, they remain prohibitively expensive to generate statistically relevant predictions often requiring an excessive number of particles and numerical experiments. This limitation is further exacerbated as the physics describing the droplet transport become more complex, specifically as it pertains to inertial interactions, evaporation, and the naturally encountered virus decay. The proposed new model of the polydisperse Gaussian-moment model (PGM) family provides an Eulerian treatment for virus-laden aerosols. The model notably provides a direct treatment for local higher-order statistics such as covariances between particle distinguishable properties (e.g., diameter, temperature, viral load, and non-volatile solute) and particle velocity. Derived within an entropy-maximization moment-closure formulation, the new model yields a set of first-order robustly-hyperbolic balance laws which are efficiently solved with a high-order, massively parallel, discontinuous-Galerkin-Hancock method. This talk provides a summary of the model derivation and its mathematical properties, and illustrates the predictive capabilities for several scenarios of bio-aerosol dispersion.

Francois Forgues
Canadian Nuclear Laboratories
francois.forgues@cnl.ca

Benoit Allard
University of Ottawa
balla069@uottawa.ca

Lucian Ivan
Canadian Nuclear Laboratories
lucian.ivan@cnl.ca

James McDonald
Department of Mechanical Engineering
University of Ottawa
james.mcdonald@uottawa.ca

MS219

Modeling Electrokinetic Flows with the Discrete Ion Stochastic Continuum

In this talk we present a methodology for the mesoscale simulation of fluid/particle systems such as strong electrolytes. This is an extension of the Fluctuating Immersed Boundary (FIB) approach that treats a solute as discrete Lagrangian particles that interact with Eulerian hydrodynamic and electrostatic fields. In both cases the Immersed Boundary (IB) method of Peskin is used for particle-field coupling. Hydrodynamic interactions are taken to be over-

damped, with thermal noise incorporated using the fluctuating Stokes equation, including a "dry diffusion" Brownian motion to account for scales not resolved by the coarse-grained model of the solvent. Long range electrostatic interactions are computed by solving the Poisson equation, with short range corrections included using a novel immersed-boundary variant of the classical Particle-Particle Particle-Mesh (P3M) technique. This approach is designed to enable scaling to large problems which are difficult to tackle using many existing mesoscale methods. It has been implemented using the AMReX framework for use on large scale HPC systems, including heterogeneous CPU+GPU architectures.

Daniel R. Ladiges, Jialun G. Wang
Center for Computational Science and Engineering
Lawrence Berkeley National Laboratory
drladiges@lbl.gov, jgalenwang@lbl.gov

Ishan Srivastava
Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory
isriva@lbl.gov

Sean P. Carney
UCLA
Department of Mathematics
spcarney@math.ucla.edu

Andrew Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Alejandro Garcia
San Jose State University
alejandrogarcia@sjsu.edu

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

MS219

A Spectral Discretization for Fluctuating Fiber Hydrodynamics

Semiflexible filaments are ubiquitous in nature and cell biology, including in the cytoskeleton, where reorganization of actin filaments allows the cell to move and divide. Most methods for simulating semiflexible bending fluctuations are based on worm-like chain models, which become prohibitively expensive in the slender limit when computing hydrodynamics. In this talk, we develop a novel approach for simulating fluctuating slender filaments. Our numerical method is based on using a non-uniform spectral discretization of the filament centerline combined with a spectrally-accurate "slender-body" quadrature for hydrodynamics. We apply our method to simulate equilibrium and non-equilibrium dynamics of semiflexible fibers, and conclude by analyzing the role of semiflexible fluctuations in actin filament bundling.

Ondrej Maxian
Courant Institute

New York University
om759@nyu.edu

Brennan Sprinkle
Colorado School of Mines
bsprinkl@mines.edu

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

MS219

Efficient Brownian Dynamics for Suspensions of Rigid Bodies or Semiflexible Fibers

I will describe a new "one-solve" midpoint temporal integrator for efficiently solving the equations of Brownian dynamics for rigid bodies or semiflexible fibers. The technique is similar, but different from, a drift-corrected scheme of Westwood, Delmotte and Keaveny. The scheme I will present has particular advantages when calculating Brownian contributions to suspension rheology, and I'll present examples that demonstrate this.

Brennan Sprinkle
Colorado School of Mines
bsprinkl@mines.edu

Ondrej Maxian
Courant Institute
New York University
om759@nyu.edu

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

MS220

Exascale Computing for Thermal Fluid Applications

Simulating thermal fluid for full-core nuclear reactors has been a great interest for nuclear engineers. In particular, performing wall-resolved large eddy simulations (LES) can predict the steady state performance and help improve the reactor design. High order method is known to be able to capture the turbulence with minimal numerical dissipation. Our spectral element methods code, Nek5000/NekRS, uses curved hexahedral elements that takes $O(n)$ storage where $n = EN^3$ is the number of grid points for E elements and polynomial order N , and with an efficient preconditioner, the computational cost is only $O(nN)$. However, elevating N requires 2-3 orders of magnitude fewer elements than standard finite elements, which brings challenges to meshing complex geometries at large scales. We present an all-hex meshing strategy for randomly packed spheres in pebble bed reactors. With Voronoi decomposition, spheres are bounded by cells and the facets are tessellated into quads that are swept to the sphere to generate a base mesh. The singularities at contact points between spheres are covered by the overlapping region to achieve the target void fraction that realizes the correct pressure drop. We generate high quality meshes with $10^2 10^5$ spheres using 300 elements per sphere. We present mesh quality metrics and its impact to the solver performance under different preconditioners which realizes fluid simulations for full-core pebble

bed reactors at exascale platforms.

Yu-Hsiang Lan
Argonne National Laboratory
lan14@illinois.edu

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Paul Fischer
University of Illinois, Urbana-Champaign
Argonne National Laboratory
pffischer@gmail.com

MS220

Advanced Simulation Methods for the Prediction of Aerodynamically Generated Noise

The mitigation of aerodynamically generated sound is an important technical field. It usually requires an accurate prediction of the turbulent flow and a computational aeroacoustics method to predict the noise generation and the propagating sound field. In this paper, a direct-hybrid simulation method is presented, in which a solver for the flow field is directly coupled to a solution of the acoustic perturbation equations. It uses an interleaved execution pattern such that an efficient parallelization is achieved also on high-performance computing (HPC) systems, see [Michael Schlottke-Lakemper, Ansgar Niemller, Matthias Meinke, Wolfgang Schröder, Efficient parallelization for volume-coupled multiphysics simulations on hierarchical Cartesian grids, *Computer Methods in Applied Mechanics and Engineering*, Volume 352, 2019, Pages 461-487, <https://doi.org/10.1016/j.cma.2019.04.032>] [Ansgar Niemller, Michael Schlottke-Lakemper, Matthias Meinke, Wolfgang Schröder, Dynamic load balancing for direct-coupled multiphysics simulations, *Computers & Fluids*, Volume 199, 2020, 104437, <https://doi.org/10.1016/j.compfluid.2020.104437>]. Several applications are presented for which the HPC performance is analyzed. For a landing gear noise mitigation is studied by applying fairings made of porous material. Other large scale applications will be discussed, such as the noise generated by jets from chevron nozzles or from propellers interacting with a downstream located wing.

Matthias Meinke, Miro Gondrum, Sutharsan Satcunanathan, Zhe Yang, Wolfgang Schröder
Institute of Aerodynamics
RWTH Aachen University
m.meinke@aia.rwth-aachen.de, m.gondrum@aia.rwth-aachen.de,
s.satcunanathan@aia.rwth-aachen.de,
z.tang@aia.rwth-aachen.de, office@aia.rwth-aachen.de

MS220

Analysis and Compression of Temporal CFD Data Sets Using Streaming and Parallel SVD

Computational Fluid Dynamics is characterized by its production of high dimensional data, which must often be treated to extract useful and interpretable information. For this purpose, Singular Value Decomposition (SVD), a powerful tool used in various fields under different names, is often utilized. The main characteristic of the method, as used in fluid dynamics, is to decompose an original data set into a set of basis functions dependent on space and coefficients dependent on time. Given that SVD provides an

optimal representation of the original data set, it is often used for data analysis and compression, since low energetic modes can be discarded while keeping a reasonable reconstruction of the original fields. The present work takes advantage of the temporally dynamic and spatially partitioned nature of the domains in turbulence simulations by calculating local SVDs in different regions of the flow field. This approach exploits the locality of the flow features which allows SVD to find a sparser basis to describe the flow and, in turn, increase compression. Apart from the inherently parallel nature of the approach, we calculate the local SVDs in a streaming manner such that, in general, the original data does not need to be stored but can be analyzed at simulation runtime. Thus, our approach provides a way to compress data with local SVDs, obtain global POD modes, and a simple way to calculate DMD, another modal decomposition used for understanding flow dynamics.

Adalberto Perez
KTH Royal Institute of Technology
adperez@kth.se

Saleh Rezaeiravesh
University of Manchester
saleh.rezaeiravesh@manchester.ac.uk

Stefano Markidis
KTH, Sweden
markidis@kth.se

Philipp Schlatter
KTH Royal Institute of Technology
pschlatt@mech.kth.se

MS220

Hierarchical Bayesian Multifidelity Modelling Applied to Turbulent Flows

Conducting high-fidelity studies in fluid mechanics can be prohibitively expensive, particularly at high Reynolds numbers. Thus, it is necessary to develop accurate yet cost-effective models for outer-loop problems involving turbulent flows. One way is to use multifidelity models (MFMs) which aim at accurately predicting quantities of interest (QoIs) and their stochastic moments by combining the data obtained from different fidelities. When constructing MFMs, a balance is sought between a few expensive (but accurate) simulations and many more inexpensive (but potentially less accurate) simulations. Two main characteristics have to be considered. 1) there is a distinguishable hierarchy in the fidelity of the turbulence simulation approaches such as RANS, hybrid RANS-LES, LES, and DNS. 2) the outcome of any of these approaches can be potentially uncertain due to various inherent uncertainties. We further develop and apply a MFM that is consistent with these characteristics and is based on the Gaussian processes (GP). The fidelity-specific calibration parameters as well as the hyperparameters appearing in the GPs are simultaneously estimated within a Bayesian framework. The Bayesian inference is done using a Markov Chain Monte Carlo (MCMC) approach, and the predictions of the MFM are accompanied by estimates of the confidence intervals. The efficiency and performance of this MFM are illustrated for various problems involving turbulent flows.

Saleh Rezaeiravesh
University of Manchester
saleh.rezaeiravesh@manchester.ac.uk

Timofey Mukha, Philipp Schlatter
KTH Royal Institute of Technology
tmu@kth.se, pschlatt@mech.kth.se

MS220

Neko: A New Spectral Element Code Applied to the Simulation of a Flettner Rotor

Ways to reduce the fuel consumption of ships have been gaining increased interest. One potentially relevant approach is a Flettner rotor, i.e. a rotating cylinder that uses the Magnus effect to generate lift. While Flettner rotors have been considered in experiments and RANS simulations, no direct numerical simulation (DNS) of a Flettner rotor, focusing on the interaction with the surrounding turbulent boundary layer has been carried out. Our simulation demonstrates our new CFD code Neko, a solver that owes its homage to Nek5000. Neko is based on similar numerical methods, but differs in that it also accommodates modern computer architectures such as GPUs and has an object-oriented codebase written in modern Fortran. Special emphasis has been put on portability, scalability and the possibility to extend to code easily with new features; thereby keeping the complexity for the domain scientist to a minimum. The flow case under consideration is a Flettner rotor submerged in a turbulent boundary layer, consisting of 1M spectral elements, which turns into 0.5B unique grid points. We discuss the strong scaling efficiency with comparing several architectures, including latest GPUs. We observe excellent parallel scaling, up to hundreds of nodes. Our initial findings for the lift are in excellent agreement with experimental data. We further observe that there is a strong interaction between the rotor and the turbulent boundary layer, in terms of modified coherent structures.

Philipp Schlatter, Martin Karp
KTH Royal Institute of Technology
pschlatt@mech.kth.se, makarp@kth.se

Daniele Massaro
KTH Engineering Mechanics
dmassaro@kth.se

Niclas Jansson
KTH Royal Institute of Technology
njansson@kth.se

Stefano Markidis
KTH, Sweden
markidis@kth.se

MS221

Scalable Computation of Mixed Dimensional Problems in FEniCSx

Problems involving fields defined over domains of different topological dimensions arise in many areas of mathematics, science, and engineering. Examples include hybrid discontinuous Galerkin methods, where function spaces defined only over the facets of the mesh are introduced, and problems with boundary or interface conditions enforced by Lagrange multipliers. Solving such problems for real-world applications typically requires large scale simulations to provide sufficient accuracy. We present recent developments in the FEniCSx core libraries, inspired by earlier work in FEniCS [C. Daversin-Catty et al., Abstractions and Automated Algorithms for Mixed Domain Finite Element Methods, ACM Trans. Math. Softw. 47, 4, 31],

that allows mixed dimensional problems to be solved in an efficient manner. Subdomains, of possibly differing topological dimensions, can be created by specifying a subset of mesh entities, yielding a submesh. Forms can then be assembled involving functions defined over different meshes. The relationships between the meshes can be nested, and high-order geometry is supported. Parallelism and scalability are at the heart of the design philosophy, and code prototyped on a laptop can be deployed on an HPC system without modification. We present scaling results for several applications, including hybridised discontinuous Galerkin methods and Lagrange multiplier problems.

Cécile Daversin-Catty
Simula
cecile@simula.no

Joseph P. Dean
Cambridge
jpd62@cam.ac.uk

Jørgen Dokken, Marie E. Rognes
Simula Research Laboratory
dokken@simula.no, meg@simula.no

Garth N. Wells
University of Cambridge
gnw20@cam.ac.uk

MS221

BDDC Preconditioners for Hybrid Discontinuous Galerkin Discretizations in Cardiac Electrophysiology

The drive to better understand complex phenomena like cardiac functions has led to mathematical models that are able to describe the multiscale structure and dynamics of the heart. The numerical simulation of these complex cardiac models is very challenging since it requires high-resolution space and time discretizations, as well as numerical tools that can handle natural discontinuities (for instance, disconnected cells). Here we consider a cardiac microscopic model, which is able to describe a cell-by-cell framework. We show how physiological discontinuities can be naturally included through a standard space discretization. Then, we build an ad-hoc non-overlapping domain decomposition preconditioner for the solution of the associated algebraic system. This preconditioner relies on the one introduced by [Dryja et al. (2007)], but differs since it handles a global discontinuous solution across the domain. We prove theoretical convergence of the proposed solver and we provide several numerical tests that show scalability and robustness, as well as optimality.

Ngoc Mai Monica Huynh
Università degli Studi di Pavia
ngocmaimonica.huynh@unipv.it

MS221

Monolithic Multiphysics Preconditioning

Numerous engineering, biology and medicine systems consist of components governed by their respective physical laws interacting across a common interface. The interface is typically a manifold of co-dimension 1 or 2 with the latter being the case in reduced-order models. The mixed-dimensional nature of the models, the coupling between the subsystems and the large ranges of physical parameters found in the applications of interest present a challenge

for designing efficient numerical schemes. In this talk, we develop monolithic preconditioners leading to order optimal and parameter robust solvers for coupled multiphysics models where the coupling is realized through a dedicated Lagrange multiplier on the interface as well as for the formulations where it is enforced weakly. A crucial component in the analysis is correctly posing the interface conditions between the subsystems. In particular, we show that robust preconditioning requires weighted fractional order Sobolev spaces and their intersections. The performance of the preconditioners will be illustrated in biomedical applications in cardiac and brain mechanics.

Miroslav Kuchta
Simula Research Laboratory
Oslo, Norway
miroslav@simula.no

MS221

Multiresolution and Iterative Schemes for the Discrete Element Method Using Triangulated Shapes

Discrete Element Methods (DEM) simulate the interaction of large numbers of rigid, incompressible objects with each other. We present a family of algorithmic ideas where computationally cheap iterative approximations are used to solve most instances of aspects of the code. Where convergence isn't reached we fall back to a more expensive but robust method. This pattern allows us to take advantage of modern parallel architectures. Geometry distance checks are calculated using a distance minimisation algorithm. The precision of the result from this minimisation can be tuned to produce a result only as precise as needed. We support triangulated particles with efficient implicit DEM code using a multiresolution and multiprecision hierarchy. Early iterations of the implicit solve yield an initial guess and thus reduce the number of later, more expensive iterations that are required to converge to the correct solution. Finally, we show a local time stepping scheme where the time step size is selected optimistically. Where this results in invalid physical states the state of each particle is rolled back and a more simple method (half the time step size until a valid state is found) is used. The optimistic selection of time step size for local clusters of particles can end up advancing some particles too far and missing interactions with new particles entering the local cluster. In this case we roll back to the time of the new interaction and resimulate from there.

Peter J. Noble
University of Durham
peter.j.noble@durham.ac.uk

Tobias Weinzierl
Durham University
tobias.weinzierl@durham.ac.uk

MS221

Immersed Mesh Methods for the Solution of Coupled Finite Element Problems

The ever-increasing computational power of modern supercomputers enables the numerical simulation of complex and coupled large-scale multiphysics problems, e.g., arising from fluid-structure interaction, or geometrically complex interactions, e.g., in contact mechanics. Common to these coupled and complex problems is the need for the transfer of data or information between the different models, meshes, or approximation spaces. The transfer of discrete

fields, such as stresses, pressure, displacements, or velocities, might be required both along surfaces or within volumes. Variational transfer techniques might also play an important role at the level of the discretization, e.g., within non-conforming domain decomposition or mortar methods for the transfer along surfaces, or on the level of the solution method, e.g., within multigrid or multi-level methods for the transfer between different volume meshes. In this talk, we introduce our parallel variational transfer technique and its application to a wide variety of problems.

Patrick Zulian
Euler institute, Università della Svizzera italiana
Switzerland
patrick.zulian@usi.ch

MS222

Modeling Multi-Physics Coastal Hydrodynamics Using Operator Learning

High-fidelity numerical models are commonly used to simulate physical phenomena in coastal and near-shore regions to obtain an improved understanding of hydrodynamic processes. Despite the tremendous progress in computational resources and methodologies over the last two decades, these high-fidelity models are still computationally expensive and require a significant amount of expertise. Operator learning frameworks like Deep Operator Networks (DeepONets) have recently become popular due to their demonstrated capability to approximate nonlinear operators between infinite dimensional Banach spaces, thus allowing them to learn solution operators to the partial differential equations that describe various hydrodynamic processes. We explore the use of such operator inference methods to construct reduced order models (ROMs) that can approximate coupled coastal wave-circulation systems governed by shallow water hydrodynamics and high-fidelity wave approximation models. The efficiency and accuracy of these models are studied using both benchmark problems with analytical results as well as physically realistic 2D problems. We also investigate the optimal combination of physics-based regularization losses and required high-fidelity training data, as well as perform automated hyperparameter searches to formulate design guidelines for future efforts in this area.

Sourav Dutta
University of Texas at Austin
sourav.dutta@austin.utexas.edu

Peter Rivera-Casillas
US Army Corps of Engineers ERDC
Vicksburg, MS
peter.g.rivera-casillas@erdc.dren.mil

Matthew W. Farthing
US Army Engineer Research and Development Center
matthew.w.farthing@usace.army.mil

Clint Dawson
University of Texas at Austin
clint.dawson@austin.utexas.edu

MS222

Physics-Informed Machine Learning for Heterogeneous Underground Reservoir Pressure Manage-

ment

Avoiding over-pressurization in subsurface reservoirs is critical for applications like CO₂ sequestration and wastewater injection. Managing the pressures by controlling injection/extraction is challenging because of complex heterogeneity in the subsurface. The heterogeneity typically requires high-fidelity physics-based models to make predictions on CO₂ fate. Furthermore, characterizing the heterogeneity accurately is fraught with parametric uncertainty. Accounting for heterogeneity and uncertainty makes this a computationally-intensive problem challenging for current reservoir simulators. To tackle this, we use differentiable programming with a full-physics model and machine learning to determine the fluid extraction rates that prevent over-pressurization at critical reservoir locations. We use DPFEHM framework, which has trustworthy physics based on the standard two-point flux finite volume discretization and is also automatically differentiable like machine learning models. Our physics-informed machine learning framework uses convolutional neural networks to learn an appropriate extraction rate based on the permeability field. We also perform a hyperparameter search to improve the models accuracy. We constructed and tested a sufficiently accurate simulator that is 400 000 times faster than the underlying physics-based simulator, allowing for near real-time analysis and robust uncertainty quantification.

Aleksandra Pachalieva, Daniel O'Malley
Los Alamos National Laboratory
apachalieva@lanl.gov, omalled@gmail.com

Dylan Harp
The Freshwater Trust
dharp73@gmail.com

Hari Viswanathan
Los Alamos National Laboratory
Earth and Environmental Sciences Division
viswana@lanl.gov

MS222

Probabilistic Sampling for High Dimensional Kernel Based Approximation

Kernel based algorithm can be formulated in a Reproducing Kernel Hilbert Spaces, where functional analytical results can be used to provide guarantees on the performances of the algorithm. For data-based algorithms, the problem of sampling is of utmost importance. The P-greedy algorithm allows one to obtain in this setting a quasi-optimal sampling of the space, that can be used in its own or as a discretization to reduce infinite dimensional problem to finite ones. In this talk, we prove that the optimization step required by this algorithm may be reduced to a finite search over a sequence of randomly drawn points, thus significantly accelerating its computation especially in high dimension, while maintaining essentially the same error guarantees as the abstract algorithm, even if only in high probability. We furthermore discuss connections to the coresets for the computation of kernel maximum mean discrepancy (MMD).

Gabriele Santin
Center for Information and Communication Technology
Bruno Kessler Foundation, Trento, Italy

gsantin@fbk.eu

MS222

Two Layered Kernel Machines for Sparse Surrogate Modeling

Kernel methods are popular, especially using easily implementable radial basis function (RBF) kernels. They usually make use of a length scale hyperparameter, which is usually tuned via costly cross validation. Here we leverage approaches and tools from the machine learning community to introduce structured deep kernel networks and two-layered kernel machines, whereby the latter ones which can be seen as using hyperparameter optimized RBF kernels. These two-layered kernel machines allow to obtain smaller expansion sizes, thus yielding faster and cheaper surrogate models. The use of the proposed approach is highlighted on both synthesized as well as real data sets.

Tizian Wenzel
University of Stuttgart
tizian.wenzel@mathematik.uni-stuttgart.de

Emma Perracchione
Polytechnic University of Torino
emma.perracchione@polito.it

Francesco Marchetti
University of Padova (Italy)
francesco.marchetti@math.unipd.it

Gabriele Santin
Center for Information and Communication Technology
Bruno Kessler Foundation, Trento, Italy
gsantin@fbk.eu

Bernard Haasdonk
University of Stuttgart
bernard.haasdonk@ians.uni-stuttgart.de

MS222

Efficient Solution of Coupled Problems Through Deim-Based Data Projection Across Non-Conforming Interfaces

Expensive multi-scale and multi-physics problems are nowadays part of cornerstone problems of applied sciences and engineering. One of the major challenges of coupled problems is to manage non-conforming meshes at the interface between two models and/or domains, due to different numerical schemes or domains discretizations employed. Moreover, very often complex submodels depend on (e.g., physical or geometrical) parameters. Knowing how outputs of interest are affected by parameter variations is key to gain useful insights of the described physics. In this talk we propose a parametric reduced order modeling (ROM) technique based on Reduced Basis methods, to fully reduce multi-scale and multi-physics problems featuring interface non-conformity. We combine Proper Orthogonal Decomposition and discretized empirical interpolation method (DEIM) to represent each sub-problem solution, whereas Dirichlet and Neumann parameter dependent interface conditions are reconstructed and interpolated across non-conforming interfaces by means of DEIM. The proposed method can be regarded as an efficient alternative to numerical schemes such as domain-decomposition methods, reducing the costs of multi-query simulations. Numerical results on the coupled heart-torso model, a

multi-physics model arising in cardiac electrophysiology to simulate relevant clinical outputs (e.g. the ECG), show the feasibility of the proposed strategy in terms of both accuracy and efficiency.

Elena Zappon, Andrea Manzoni
Politecnico di Milano
elena.zappon@polimi.it, andrea1.manzoni@polimi.it

Alfio Quarteroni
Ecole Polytechnique Fédérale de Lausanne, Switzerland
alfio.quarteroni@epfl

MS223

Using the FFTX Library in Applications

FFTX is a performance-portable, open-source FFT software system for CPUs and GPUs analogous to FFTW for CPU systems. It supports application-specific optimizations corresponding to integrating more of the algorithms into the analysis / code-generation process. FFTX is based on the use of Spiral, an open-source analysis and code-generation tool chain for FFTs and tensor algebra algorithms, developed at Carnegie-Mellon University and SpiralGen, Inc.; an FFTX user API implemented in standard C++; and a factored design that allows FFTX / Spiral to be more easily ported across multiple architectures. In FFTX, we can represent larger integrated algorithms that include FFTs composed with algorithmic operations such as multiplication by a (possibly matrix-valued) symbol and batching. By combining substeps in an integrated algorithm, the amount of data traffic can be reduced by significant amounts. FFTX applies the size-specific analysis and automatic code generation techniques in Spiral to generate code, leading to implementations with far higher performance than obtainable from approaches based on black-box FFT implementations. In this talk, we will demonstrate how to install the FFTX library and use it in applications written in C++ and Python, and present results obtained on supercomputers with NVIDIA and AMD GPU systems.

Peter McCorquodale
Lawrence Berkeley National Laboratory
PWWMcCorquodale@lbl.gov

MS223

Enabling End-to-End Accelerated Simulations in the Exascale Era Using PETSc

The Portable Extensible Toolkit for Scientific computation (PETSc) library provides scalable solvers for nonlinear time-dependent differential and algebraic equations and for numerical optimization. Throughout its history, a major goal of PETSc development has been achieving scalability required to fully utilize leadership-class supercomputers. In the past, this meant a great deal of focus on inter-node scalability, but, as we enter the exascale era, the primary challenge to enabling extreme-scale computation has become harnessing abundant fine-scale parallelism within compute nodes – primarily in the form of GPU-based accelerators. We will discuss how the PETSc design for performance portability addresses these challenges while stressing flexibility and extensibility by separating the programming model used by the application from that used by the library, and we will present some examples of how this facilitates end-to-end utilization of GPUs for PETSc-based simulations on cutting-edge high-performance computing architectures. Additionally, we will discuss recent develop-

ments in PETSc's communication module, PetscSF, that enable flexibility and scalable performance across large GPU-based systems while overcoming some of the difficulties posed by working directly with the Message Passing Interface (MPI) on these systems.

Richard T. Mills
Argonne National Laboratory
rtmills@anl.gov

Mark Adams
Lawrence Berkeley National Laboratory
mfadams@lbl.gov

Satish Balay
Argonne National Laboratory
balay@anl.gov

Jed Brown
University of Colorado Boulder
jed@jedbrown.org

Alp Dener
Argonne National Laboratory
Mathematics and Computer Science Division
adener@anl.gov

Jacob Faibussowitsch
University of Illinois Urbana-Champaign
jacob.fai@gmail.com

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Scott Kruger
Tech-X Corporation
kruger@txcorp.com

Hannah Morgan
University of Chicago
hmmorgan@uchicago.edu

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

Karl Rupp
TU Wien
me@karlrupp.net

Barry Smith
Center for Computational Mathematics
Flatiron Institute
bsmith@petsc.dev

Stefano Zampini
King Abdullah University of Science and Technology
stefano.zampini@kaust.edu.sa

Hong Zhang, Junchao Zhang
Argonne National Laboratory

hongzhang@anl.gov, jczhang@mcs.anl.gov

MS223

Can Planewave DFT Codes Scale on Massively Parallel GPU-Based Supercomputers?

In the age of exascale computing, scientific simulations must make use of massively parallel GPU-based supercomputers. For example, two of the top supercomputers, Summit and Frontier, offer huge amounts of computing capabilities meant to enable scientists to execute bigger scientific problems in shorter amounts of time. However, these new systems require existing scientific software to be re-designed and re-implemented to fully utilize the hardware features. Previous assumptions that made software execute efficiently on the supercomputers from 10 years ago, may not comply with recent hardware capabilities. In this work, we will focus on re-designing some algorithms that are heavily used in Planewave Density Functional Theory (DFT) simulations. More specifically, we will focus on four algorithms meant to solve a non-linear eigenvalue problem, namely Conjugate Gradient, DIIS, Jacobi Davidson and Unconstrained. For each method we identify the bottlenecks, and devise strategies to overcome the issues and scale the codes to thousands of GPUs. We will highlight the similarities and differences between the four algorithms using our approach. In addition, we will compare the codes against the theoretical machine peak and discuss each implementation in detail.

Doru Thom A. Popovici, Mauro Del ben, Andrew M. Canning
Lawrence Berkeley National Laboratory
dtpopovici@lbl.gov, mdelben@lbl.gov, acanning@lbl.gov

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

MS223

Recent Advances in Hypra to Support Exascale Computers

With the increasing inclusion of accelerators into current and future high-performance computers, it has become important to enable parallel mathematical libraries to take advantage of the increased performance potential of GPUs. The hypra software library provides high performance preconditioners and solvers for the solution of large sparse linear systems on massively parallel computers with focus on algebraic multigrid (AMG) methods. One of its attractive features is the provision of conceptual interfaces, which include a structured, a semi-structured interface, and a traditional linear-algebra-based interface. These interfaces give application users various means for describing their linear systems and provide access to methods such as structured multigrid solvers, which can take advantage of the additional information beyond just the matrix, as well as unstructured multigrid solvers. This talk will discuss new developments in hypra.

Ulrike Meier Yang, Rui Peng Li, Victor Magri, Wayne Mitchell, Robert D. Falgout
Lawrence Livermore National Laboratory
yangl1@llnl.gov, li50@llnl.gov, paludettomag1@llnl.gov,

mitchell82@llnl.gov, falgout2@llnl.gov

MS224

Hybrid Stochastic/Deterministic Methods for Linear Algebra on Advanced Heterogeneous Architectures

In this talk, we are investigating the applicability of hybrid Monte Carlo methods for solving systems of Linear Algebraic Equations to variety of problems. We present the results of our exploration of using the Markov Chain Monte Carlo Matrix Inversion (MCMCMI) method as preconditioner to a variety of linear systems arising from the domains of quantum chromodynamics, plasma physics and engineering. The latter two are represented by matrices extracted from BOUT++ and Nektar++ codes with specific problem statements. We then study the performance of pure Python implementation against C++/CUDA one and present the corresponding scaling behaviour of the corresponding implementations on GPUs and CPUs.

Vassil Alexandrov
Hartree Centre, STFC
UK
vassil.alexandrov@stfc.ac.uk

Emre Sahin
Hartree Centre, STFC UK
emre.sahin@stfc.ac.uk

Anton Lebedev
Hartree Centre STFC, UK
anton.lebedev@stfc.ac.uk

MS224

Cost-Efficient Integrated Analysis of Distributed Data in Secure Environments

Amidst the growing demand for integrating vast data silos in research and business to achieve a competitive advantage, various methods are proposed for distributed data analysis. However, most of them come at high computational costs requiring iterative communication rounds among the clients. We analyze the cost-effectiveness along with data security of several methods for distributed data analysis. We further highlight the advantages of Data Collaboration, which offers a one-pass solution by securely integrating latent representations of distributed data.

Tetsuya Sakurai, Akira Imakura
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp, imakura@cs.tsukuba.ac.jp

Xiukai Ye, Anna Bogdanova, Yasunori Futamura,
Yukihiko Okada
University of Tsukuba
yexiukai@cs.tsukuba.ac.jp,
bogdanova.anna.fw@u.tsukuba.ac.jp,
futamura@cs.tsukuba.ac.jp, okayu@sk.tsukuba.ac.jp

Takeya Shigezumi, Hiroyuki Eguchi, Hiroyuki Yamanaka
NTT Data Corporation
takeya.shigezumi@nttdata.com, hiroyuki11.eguchi@nttdata.com,

hi-
hi-

royuki.yamanaka@nttdata.com

MS224

Pmex: A Post-Modern Approach to the Matrix Exponential for Exponential Integrators

Higher order exponential integrators require computing linear combination of exponential like $\varphi(A)$ -functions of large matrices A . Recent advancements in computational linear algebra allows efficient approximation of the φ -functions by exploiting the equivalence with a linear ordinary differential equation. The resulting adaptive Krylov methods can be very efficient in serial when little or no prior knowledge of the spectrum of A is available. However, in parallel, communication is of paramount importance in the formulation of methods for distributed architectures. Arnoldi iteration in Krylov subspace methods is a computational bottleneck because global communication for normalization and inner products is needed at each Arnoldi step. We aim to improve the scalability of exponential integrators by using a low-synchronization Post-Modern Arnoldi method in conjunction with higher-order adaptive sub-stepping techniques via control theory. The resulting orthogonalization algorithm has an accuracy comparable to modified Gram-Schmidt yet better suited for distributed architecture as there is only 1 global communication per iteration. We present geophysical based numerical experiments which validate that reducing global communication and having a smoother sub-step sequence in the adaptivity lead to better parallel scalability and reduce the computational complexity of the adaptive Krylov methods, which in turn improve the performance of exponential integrators.

Tanya V. Tafolla
University of California, Merced
ttafolla@ucmerced.edu

Stéphane Gaudreault
Environment Canada
stephane.gaudreault@ec.gc.ca

Mayya Tokman
University of California, Merced
School of Natural Sciences
mtokman@ucmerced.edu

MS224

Collaborative Future Selection for Distributed Data

Feature selection is an efficient dimensionality reduction technique, which aims to extract the important features and eliminate the noisy ones. On the other hand, the data is sometimes distributed to multiple parties, and it is difficult to share the original data due to privacy requirements. As a result, the data in one party may lack useful information to learn the most discriminative features. We introduce a collaborative feature selection method for multiple parties without revealing their original data. Each party constructs the intermediate representations of the original data and shares the intermediate representations for collaborative feature selection. Based on the shared intermediate representations, the original data from multiple parties are transformed into a common latent space. The feature ranking of the original data is learned by imposing row sparsity on the transformation matrix simultaneously. Experimental results on real-world datasets demonstrate

the effectiveness of the proposed method.

Xiucui Ye
University of Tsukuba
yexiucui@cs.tsukuba.ac.jp

Akira Imakura, Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
imakura@cs.tsukuba.ac.jp, sakurai@cs.tsukuba.ac.jp

MS224

High-Performance Linear Algebra and Anomaly Detection

In this talk, we explore anomaly detection methods and focus on those based on spectral computation. An original approach for anomaly detection and the performance of the corresponding method are presented.

Nahid Emad
LI-PaRAD laboratory and Maison de la Simulation
University of Paris Saclay / Versailles
Nahid.Emad@uvsq.fr

Zineb Ziani
University of Paris-Saclay & Numeryx
z.ziani@numeryx.fr

Bouaziz Ahmed Abdelkader
NUMERYX
aa.bouaziz@numeryx.fr

MS225

Boosting Your Creativity

We describe our personal experiences of using creativity techniques in mathematics research, which have been honed through participating in creativity workshops. Some simple tips will be described that can help you to boost your creativity.

Stefan Güttel
The University of Manchester
stefan.guettel@manchester.ac.uk

Nicholas J. Higham
School of Mathematics
The University of Manchester
nick.higham@manchester.ac.uk

MS225

Developing Mentoring Relationships

Mentoring has become a ubiquitous feature in career and professional development, but what type of mentoring is the most effective and what are the keys to developing productive mentoring relationships? How many mentors or mentees should you have? How do you find a good mentor or mentee and what are keys to being a good mentor or mentee? This session will explore the benefits for the mentee and the mentor and how to establish, navigate, and maintain mentoring relationships over time. We will briefly explore some of the research on mentoring and share lessons learned from participating, organizing, and facilitating mentoring programs over the years.

Mary Ann E. Leung

Sustainable Horizons Institute
mleung@shinstitute.org

MS225

Planning and Managing Research Projects

Planning and managing research projects are key aspects of work in computational science and engineering (CSE), necessary for students as well as early-career, mid-career, and senior professionals who may work in academia, national labs, and industry. This presentation will discuss strategies for planning and managing multidisciplinary CSE research, from the perspectives of individuals, research teams, and teams of teams (that is, aggregate teams who provide complementary expertise in CSE research collaborations). We will consider issues related to planning and managing work overall, as well as issues specific to research software, which encapsulates expertise across applied mathematics, computer science, and domains of science and engineering to enable sustained CSE collaboration and scientific progress.

Lois Curfman McInnes

Mathematics and Computer Science Division
Argonne National Laboratory
curfman@anl.gov

MS225

Building Your Professional Network

Networking is one of the most important ingredients of a successful career. Both in academia and in industry, knowing people and people knowing you is extremely important. An extensive professional network will stimulate collaborations, it will lead to joint projects, jointly organized conferences, invited presentations, increased number of citations, and much more. Especially in this day and age, social media are part of building a professional network, with LinkedIn probably the most suitable. Social media can boost your career if used in the right manner. In this presentation, I will show how I experienced building a professional network, what the advantages were for me. There will also be some general advice on how to build a professional network and how to use it for growing your career.

Wil Schilders

Eindhoven University of Technology
w.h.a.schilders@TUE.nl

MS226

Finite Element Discretizations of Three-Dimensional Time-Harmonic Maxwell's Equations with Two Sign-Changing Coefficients

I will consider time-harmonic Maxwell's equations in three-dimensional domains containing metamaterials. Mathematically, the originality of this setting is that the two physical coefficients entering the PDE may simultaneously change sign through an interface. I will first analyze the zero-frequency case by presenting an extension of the T-coercivity approach that includes curl-preserving Piola mappings. Then, I will introduce new duality arguments which, in conjunction with the proposed T-coercivity technique, allow to show the stability and convergence of finite element approximations in the time-harmonic case.

Théophile Chaumont-Frelet
Inria Sophia-Antipolis

theophile.chaumont@inria.fr

Patrick Ciarlet

Laboratoire POEMS (CNRS-ENSTA-INRIA)
patrick.ciarlet@ensta-paris.fr

MS226

Capturing Plasmonic Behaviors On Spheres Using Enriched Finite Element Techniques

We present a Finite Element Method (FEM) based approach to capture plasmonic behaviors in light scattering by metallic spheres. Surface plasmons are highly oscillatory waves localized to the interface between a dielectric (air, vacuum) and a metal (gold, silver). As surface plasmons lead to large field enhancements, they are useful for high-resolution imaging and other applications. It is challenging to capture them numerically, and standard methods do not succeed. In the context of spherical scatterers, we identify where plasmonic excitations can arise, propose a strategy to enrich the solution space with the plasmonic behavior, and construct a mesh to allow us to treat both the highly oscillatory and regular behaviors. In the discrete context, this leads to an enriched finite element method.

Benjamin Latham

University of California Merced
blatham@ucmerced.edu

MS226

High-Order Polyhedral Methods for Electromagnetics

This talk will review the recent developments in the HHO approximation of (three-dimensional) electromagnetics problems. Introduced in 2014, the Hybrid High-Order (HHO for short) method provides a novel framework for the discretization of models based on PDEs. Its construction hinges on (hybrid) polynomial unknowns attached to both the cells and the faces of the mesh. HHO methods have several assets: (i) they are applicable on general polyhedral meshes, (ii) they allow for an arbitrarily large approximation order, and (iii) they are amenable to static condensation, thus allowing for a reduced computational burden. This talk will focus in particular on magnetostatics models. In that case, the well-posedness of HHO schemes hinges on (discrete) hybrid Weber inequalities, whose statement strongly depends on the topology of the problem at hand. This is a joint work with F. Chave and D. A. Di Pietro.

Simon Lemaire

Inria, Univ. Lille
simon.lemaire@inria.fr

MS226

T-conforming Meshes for Transmission Problems with Sign-Changing Coefficients and Polygonal Interfaces

In electromagnetism, we can have materials characterized by unusual properties such as a negative permittivity. When surrounded by classical materials (with a positive permittivity), particular behaviors arise and the interface between the materials, especially in presence of corners. This motivates the study of transmission problems with sign-changing coefficients and polygonal interface. Even when the problem is well-posed (which can be established via the T-coercivity theory), the numerical computation

based on standard finite element method (FEM) still can fail. FEM convergence can be guaranteed as long as T-coercivity arguments are satisfied at the discrete level, which can boil down to use locally T-conforming meshes. T-conforming meshes rely on specific symmetry requirements on the mesh, and in practice, it is relatively cumbersome to construct for polygonal interfaces. We develop and provide a T-conforming mesh generator for (almost) arbitrary polygonal interfaces in two dimensions, and provide a battery of tests to illustrate the need and efficiency of the designed meshes.

Zois Moitier
ENSTA Paris, France
zois.moitier@ensta-paris.fr

MS227

Model Order Reduction for Wave Energy Converter Farms

Wave energy can play an important role in the actual energetic scenario, but to become competitive as other, more classical, renewable energy sources, wave energy converters (WECs) must be deployed in farms, thus reducing the cost of power production. The complexity of the design of WEC farms is mainly related to the large areas that multiple devices cover and the multi-scale character of the hydrodynamics involved. Numerical simulations completely based on high-fidelity models, like the CFD, are accurate but prohibitive in terms of computational cost. To overcome this limitation, a Galerkin Proper Orthogonal Decomposition (POD) Reduced Order Model (ROM) is presented for a bi-fluid application. Furthermore, for a proper, yet computationally efficient, treatment of the different dynamical features present in a WEC farm, a methodology is proposed coupling the CFD for the fluid-structure interaction around the device, and a Galerkin-free POD ROM for the wave propagation in the far-field. The significant drop in the computational cost of the numerical simulation of wave energy converter farms permitted by this versatile coupling methodology facilitates the design process.

Beatrice Battisti
Inria Bordeaux South-West
Politecnico di Torino
beatrice.battisti@inria.fr

MS227

Energy-Stable ROM-Based Closure Modelling

We propose a ROM-based closure model for turbulent flow that conserves the total kinetic energy. For this purpose, we start by considering an energy-conserving DNS discretization which is too expensive to be used in multi-query applications. To reduce the computational cost, we introduce a discrete filter matrix. Applying this discrete filter to the DNS velocity vector, we get a discretely filtered velocity vector with significantly smaller dimension. The information that is lost in this filtering process we call subgrid velocity. When simulating the dynamics of the filtered velocity, we cannot neglect the effects of the subgrid velocity. To take these subgrid effects into account, we approximate the subgrid velocity within a reduced basis. We obtain this reduced basis via Proper Orthogonal Decomposition of subgrid velocity snapshots. The resulting model consists of the filtered velocity and the approximated subgrid velocity and has a significantly smaller dimension than the DNS model. In contrast to classical model order reduction, we do not aim at reproducing the complete dynam-

ics of the DNS model but only the filtered velocity. By preserving the structure of the underlying DNS discretization, we guarantee the proposed reduced order model to be energy-conserving and hence Lyapunov-stable.

Henrik Rosenberger
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
hker@cw.nl

Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cw.nl

MS227

Nonlinear Model Reduction for Parametrized Transonic Aerodynamic Flows

We present a projection-based model reduction formulation that uses nonlinear approximation spaces to provide an efficient solution of transonic flows with parameter-dependent shocks. The formulation is designed to overcome the slow decay of the Kolmogorov n -width in the presence of parameter-dependent discontinuities. The key ingredients are as follows. The first ingredient is the nonlinear approximation spaces; to construct the spaces, we transform, for each solution snapshot, the underlying mesh to transport the snapshot such that the shocks are aligned and then apply proper orthogonal decomposition. The second ingredient is the Galerkin projection accelerated by point-wise reduced-quadrature rule constructed using the empirical quadrature procedure, which enables efficient reduced-residual evaluation. The third ingredient is the goal-oriented error estimate based on the dual-weighted residual method, which provides an online-efficient and quantitative estimate of the reduced model error for any parameter value. We demonstrate the formulation using transonic aerodynamic flows with parameter-dependent shocks.

Masayuki Yano
University of Toronto
myano@utias.utoronto.ca

MS227

Adaptive Model Reduction for High-Speed Flows

To avoid the fundamental linear reducibility limitation associated with high-speed fluid flows (or other convection-dominated problems), we construct a nonlinear approximation by composing a low-dimensional linear space with a parametrized domain mapping. The linear space is constructed using the method of snapshots and POD; prior to compression, each snapshot is composed with a mapping that causes its local features to align (same spatial location) with the corresponding features in all other snapshots. The parametrized domain mapping is chosen such that the local features present in the linear space deform to the corresponding features in the solution being approximated, effectively removing the convection-dominated nature of the problem. The domain mapping is determined implicitly through the solution of a residual minimization problem, rather than relying on explicit sensing/detection. We provide several numerical experiments to demonstrate the effectivity of the proposed method for high-speed aerodynamic flows.

Matthew J. Zahr

University of Notre Dame
mzahr@nd.edu

MS228

The Art to Repeatedly Project Your Problems The Variable Projected Augmented Lagrangian Method

Inference by means of mathematical modeling from a collection of observations remains a crucial tool for scientific discovery and is ubiquitous in application areas such as signal compression, imaging restoration, and supervised machine learning. With ever-increasing model complexities and growing data size, new specially designed methods are urgently needed to recover meaningful quantities of interest. We consider the broad spectrum of linear inverse problems where the aim is to reconstruct quantities with a sparse representation on some vector space; often solved using the (generalized) least absolute shrinkage and selection operator (lasso). The associated optimization problems have received significant attention, in particular in the early 2000s, because of their connection to compressed sensing and the reconstruction of solutions with favorable sparsity properties using augmented Lagrangians, alternating directions and splitting methods. We provide a new perspective on the underlying ℓ_1 regularized inverse problem by exploring the generalized lasso problem through variable projection methods. We arrive at our proposed variable projected augmented Lagrangian (vpal) method. We provide numerical examples demonstrating the computational efficiency for various imaging problems.

Matthias Chung
Emory University, U.S.
matthias.chung@emory.edu

Rosemary A. Renaut
Arizona State University
School of Mathematical and Statistical Sciences
renaut@asu.edu

MS228

Approximative Bayesian Optimal Experimental Design in Inverse Problems

We study the stability properties of the expected utility function in Bayesian optimal experimental design. We provide a framework for this problem in the case of expected information gain criterion in an infinite-dimensional setting, where we obtain the convergence of the expected utility with respect to perturbations. To make the problem more concrete we demonstrate that non-linear Bayesian inverse problems with Gaussian likelihood satisfy necessary assumptions in our theory. Some numerical simulations with different examples are explored.

Duc-Lam Duong, Tapio Helin, Rodrigo Rojo Garcia
LUT University
duc-lam.duong@lut.fi, Tapio.Helin@lut.fi, ro-
drigo.rojo.garcia@lut.fi

MS228

Transport Map-Based Bayesian Optimal Experimental Design

The Bayesian optimal experimental design is essential in many fields of science and engineering, especially when each experiment is expensive and resources are limited.

Given a prior and a design-dependent likelihood function, we would like to choose the design that maximizes the expected information gain (EIG) in the posterior. Efficient and accurate estimation of EIG therefore becomes crucial. We introduce a flexible transport-map based framework that enables fast estimation of EIG by solving only convex optimization problems. This framework is also compatible with implicit models, where one can simulate from the likelihood but the conditional probability density function of the data is unknown. Several estimators naturally appear within our framework—in particular, positively and negatively biased estimators that provide bounds for the true EIG. We explore the bias and variance of our estimators and study the optimal allocation between the training and the evaluation samples given a fixed number of samples. We then demonstrate the performance of our approach using both a linear and a nonlinear example.

Fengyi Li
Massachusetts Institute of Technology
fengyil@MIT.EDU

Ricardo Baptista
MIT
rsb@caltech.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS228

Adaptive Batching for Gaussian Process Surrogates

We develop adaptive replicated designs for Gaussian process surrogates of stochastic experiments. Replication offers efficiency in learning tasks that require large scale (thousands) stochastic simulations by reducing the size of the effective covariance matrix. Adaptive batching is a natural extension of sequential design heuristics with the benefit of replication growing as response features are learned, inputs concentrate, and the metamodeling overhead rises. Our algorithms propose novel acquisition functions that simultaneously determine the sequential design inputs x^n and respective number of replicates r^n . We consider both schemes that gradually ratchet up the replication quantities r^n as the number of inputs n grows, and those that allocate additional replicates to existing inputs $x^k, k \leq n$. Illustrations with synthetic experiments and an application in discrete-action stochastic control (learning noisily observed level sets that determine optimal actions) show that adaptive batching brings significant computational speed-ups with minimal loss of modeling fidelity.

Michael Ludkovski
UC Santa Barbara
Dept of Statistics & Applied Probability
ludkovski@pstat.ucsb.edu

MS228

Deep Reinforcement Learning for Experimental Design of Material Model Calibration

This work presents a deep reinforcement learning (RL) algorithm for design of experiments that uses the Kullback-Leibler (KL) divergence, measured by a Kalman filter (KF), to maximize information gain. We apply the RL algorithm to optimally calibrate continuum plastic flow material models that capture the strain history dependence

of materials exhibiting elastoplastic response. Although a subject matter expert can select experiments based on prior knowledge, as the complexity (anisotropy) of the model grows, so does the parameter space, and therefore it becomes increasingly difficult for a human expert to design a minimal set of experiments that simultaneously satisfy budget constraints and accuracy requirements. We cast the design of experiments as a game of traversing a decision tree to reach an optimally calibrated model at the end of the tree. Each decision represents an experimental control action. The material response is passed to an extended switching Kalman filter, which performs a Bayesian update of the model parameters and relays the information gain reward to the agent. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Ruben Villarreal
Sandia National Laboratories
rubvill@sandia.gov

Nikolaos Vlassis, Nhon Phan
Columbia University
nrv2102@columbia.edu, nnp2116@columbia.edu

Thomas A. Catanach, Reese Jones, Nathaniel Trask,
Sharlotte Kramer
Sandia National Laboratories
tacatan@sandia.gov, rjones@sandia.gov,
natrask@sandia.gov, slkrame@sandia.gov

WaiChing Sun
Columbia University, New York
wsun@columbia.edu

MS229

Implicitly Constituted Materials: Fluids, Boundary Conditions, Solids, and Mixtures

In order to understand nonlinear responses of materials to external stimuli of different sort, be they of mechanical, thermal, electrical, magnetic, or of optical nature, it is useful to have at ones disposal a broad spectrum of models that have the capacity to describe in mathematical terms a wide range of material behavior. It is advantageous if such a framework stems from a simple and elegant general idea. Implicit constitutive theory of materials provides such a framework: while being built upon simple ideas, it is able to capture experimental observations with the minimum number of physical quantities involved. It also provides theoretical justification in the full three-dimensional setting for various models that were previously proposed in an ad hoc manner. From the perspective of the theory of nonlinear partial differential equations and corresponding numerical methods, implicit constitutive theory leads to new classes of challenging mathematical problems. After introducing the basic concepts of implicit constitutive theory, established by K.R. Rajagopal twenty years ago, we provide an overview of its main achievements in modeling and the analysis of corresponding problems.

Josef Malek
Charles University, Prague
Mathematical Institute
malek@karlin.mff.cuni.cz

MS229

On Mathematical Foundations for Implicitly Constituted Fluids with Implicitly Constituted Bound-

ary Conditions

We study systems of nonlinear partial differential equations of parabolic type, in which the elliptic operator is replaced by the first order divergence operator acting on a flux function, which is related to the spatial gradient of the unknown through an additional implicit equation. Formulating four conditions concerning the form of the implicit equation, we first show that these conditions describe a maximal monotone p -coercive graph. We then establish the global-in-time and large-data existence of a (weak) solution and its uniqueness. The theory is tractable from the point of view of numerical approximations. For details, we refer to [1]. [1] M. Bulicek, E. Maringova, J. Malek: On nonlinear problems of parabolic type with implicit constitutive equations involving flux, *Math. Models Methods Appl. Sci.*, 31, (2021), no.10, 2039–2090.

Erika Maringova
Institute of Science and Technology, Austria
erika.maringova@ist.ac.at

Miroslav Bulicek
Charles University
Faculty of Mathematics and Physics
mbul8060@karlin.mff.cuni.cz

Josef Malek
Charles University, Prague
Mathematical Institute
malek@karlin.mff.cuni.cz

MS229

A Thermodynamic Framework for Non-Isothermal Phenomenological Models of Mullins Effect

The Mullins effect is a common name for a family of intriguing inelastic responses of various solid materials, in particular filled rubbers. Given the importance of the Mullins effect, there have been many attempts to develop mathematical models describing the effect. However, most of available models focus exclusively on the mechanical response, and are restricted to the idealised isothermal setting. We lift the restriction to isothermal processes, and we propose a full thermodynamic framework for a class of phenomenological models of Mullins effect. In particular, we identify energy storage mechanisms (Helmholtz free energy) and entropy production mechanisms that on the level of stress-strain relation lead to the idealised Mullins effect or to the Mullins effect with permanent strain. The models constructed within the proposed framework can be used in modelling of fully coupled thermo-mechanical processes, and the models are guaranteed to be consistent with the laws of thermodynamics.

David Cichra
Charles University, Prague
Czech Republic
cichra.david@gmail.com

Pablo Alexei Gazca Orozco
Albert-Ludwigs-Universität Freiburg
alexei.gazca@mathematik.uni-freiburg.de

Vit Prusa

Charles University, Faculty of Mathematics and Physics
Sokolovská 83, 18675 Prague 8, Czech Republic
prusv@karlin.mff.cuni.cz

Karel Tuma
Charles University, Prague, Czech Republic
ktuma@karlin.mff.cuni.cz

MS229

The Oof Finite-Element Solver for Materials Science: Plasticity and MI

The Object-Oriented Finite Element code (OOF) is a long-standing project developed at NIST to deliver high-quality mathematics and computational capabilities to an audience of materials scientists. The code features tools to easily mesh complex 3D microstructures, and a scheme for the addition of custom constitutive rules, allowing users to conduct sophisticated structure-property explorations. The development team has recently added a crystal-plastic modeling capability to this tool, expanding its scope beyond divergence equations, but retaining the modularity and extensibility of the tool's architecture. In addition, in partnership with machine-learning experts, we are now actively exploring means by which machine learning techniques might be used to construct high fidelity crystal-plastic surrogate models, compatible with the underlying finite-element technique. If successful, the result will be a very high performance modeling capability for a wide range of yielding systems, with high-fidelity meshes accurately capturing complex microstructures.

Andrew Reid

Center for Theoretical and Computational Materials Science
National Institute of Standards and Technology
andrew.reid@nist.gov

MS230

Static and Kinetic Depinning of Domain Walls in a Notched Ferromagnetic Nanostrip

In this talk, we delineate the investigation of the static and kinetic depinning field of a domain wall in a notched magnetic nanostrip under the framework of the Landau-Lifshitz-Gilbert equation combined with inertial and nonlinear dissipative effects. We focus our attention on the impact of viscous, dry-friction, and inertial damping on the domain wall propagation. We execute the investigation using the Walker trial function approach and observe that the dry friction and inertial damping significantly impact the depinning field. Also, the domain wall velocity and displacement can be manipulated by suitably tuning the key parameters. The obtained results qualitatively agree with the recent observations.

Sarabindu Dolui

NATIONAL INSTITUTE OF TECHNOLOGY
ANDHRA PRADESH
doluisarabindu@gmail.com

MS230

Artificial Intelligence for Quantum Processes: from Quantum Technologies to Taming Complexity

From face recognition to the development of autonomous vehicles, machine learning (ML) is permeating technology and science in an infectious manner, opening new perspectives for the management, mining and manipulation of Big Data. In this talk, I will review the applications of ML to quantum processes, highlighting the opportunities for advances in both fundamental and applied quantum physics

that the ML-enhanced processing of quantum information is offering. I will address specific instances linked to quantum computing, quantum simulation, and quantum dynamics of complex adaptive systems, discussing both theory and recent experimental endeavours.

Mauro Paternostro

Queens University Belfast
m.paternostro@qub.ac.uk

MS230

Numerical Modeling of Memristive Devices for Neuromorphic Computing

In recent years, memristive devices have shown considerable potential for realizing synaptic functionalities in neuromorphic computing systems. Mobile defects and electronic processes in memristive elements cause a hysteretic current-voltage characteristic and enable resistive switching of the device. Details of the mechanisms underlying the switching process depend on a complex interplay of different phenomena and can decisively affect the resulting device properties. Identifying and understanding these mechanisms is key for adapting the switching characteristics to the requirements of specific computational schemes and improving performance and reproducibility. Here, we present numerical models to analyze the memristive behavior of memristive devices. A finite volume model self-consistently solves the semiconductor equations for the electrostatic potential and the quasi-Fermi levels of electrons and holes. Further, we include the dynamics of mobile dopants and briefly discuss the boundary conditions implemented for the calculation of the switching dynamics and the current-voltage characteristics. A simplified physics-based device model is introduced to reduce the computational effort and enable large-scale simulations of arrays at the expense of physical accuracy. Simulations are compared with measurements on various types of memristive devices, including vertical TiO_x/AlO_x/HfO_x thin-film stacks and quasi-two-dimensional flake devices.

Benjamin Spetzler

Ilmenau University of Technology
benjamin.spetzler@tu-ilmenau.de

MS230

Neuromorphic Computing with Memristive Devices

Memristive devices (memristors) have attracted considerable attention in the electronic device community due to their inherent memory effect, which allows them to mimic the function of biological synapses. As such, they are a key building block for the realization of artificial neural computing schemes in hardware, so-called neuromorphic systems. However, there are a variety of different neuromorphic architectures with different requirements for memristive devices, requiring optimization of electrical properties and materials as well as the technological framework for each specific application. In this talk, I discuss the challenges and prospects of memristive devices for neuromorphic computing in general and in some selected examples. I introduce redox-based memristive elements and show how their properties can be tailored for applications in neuromorphic computing architectures through systematic design variations. In this context, the influence of different oxide layer systems and electrode materials on the device properties is analyzed to assess their suitability for neuro-

morphic computing.

Martin Ziegler
Ilmenau University of Technology
martin.ziegler@tu-ilmenau.de

MS231

Numerical Methods for the Hyperbolic Monge-Ampere Equation

Applying the method of characteristics (MOC) to a hyperbolic second order partial differential equation (PDE) in two variables yields two systems of ordinary differential equations (ODEs) describing the evolution of the solution along two families of characteristics. In general the ODE systems are mutually coupled, implying a characteristic depends on both the other characteristic and on the solution of the PDE. We applied the MOC to the 2D hyperbolic Monge-Ampere equation using the x-coordinate as parametrization. The solution evolves along the characteristics and henceforth the number, type and location of the necessary boundary conditions (implicitly) depends on the characteristics over the entire domain. A criterion for the boundary conditions is formulated. Based on the coupled systems of ODEs we formulate a numerical solution method. We employ explicit one-step methods (Euler, Runge-Kutta) to integrate the ODEs systems. Subsequently, we apply B-spline interpolation along vertical grid lines to approximate the solution in grid points. Using an adaptive integration step we bound the errors of the interpolation. Subsequently, the order of accuracy of the numerical scheme is known and stability of the scheme is obtained. The scheme attains the theoretical order of convergence and converges to computer precision for multiple examples. Moreover, the numerical method is able to handle complicated cases with shocks in- and varying numbers of boundary conditions.

Maikel Bertens
Eindhoven University of Technology
m.w.m.c.bertens@tue.nl

MS231

Adaptive Methods for the Numerical Solution of Some Fully Nonlinear Equations

We focus on adaptive mesh refinement methods for first- or second order fully nonlinear equations in two space dimensions. As prototypical examples, the canonical Monge-Ampere equation (second order) and the nonlinear maps problem (first order) are addressed. Splitting methods are developed via variational and least-squares approaches. They allow to solve a sequence of linear variational problems and of algebraic eigenvalue problems. Mixed, low order, finite element approximations are used for the discretization. The linear variational problems are solved with a conjugate gradient algorithm, while the algebraic eigenvalue problems are solved locally with dedicated mathematical programming techniques. The goal is to extend the methodology by coupling the splitting methods with mesh adaptation. Low order, piecewise linear, finite elements are used for space discretization. Residual-based, a posteriori error estimates are derived based on the linear variational problems only. A mesh refinement strategy is implemented within the iterative loop. When dealing with line singularities, anisotropic adaptive techniques allow to obtain refined triangulations near these edges. When dealing with point singularities, isotropic adaptive techniques are used instead. Numerical experiments show that the

method exhibits appropriate convergence orders and a robust behavior. Adaptive mesh refinement proves to be efficient and accurate to tackle test cases with singularities.

Alexandre Caboussat
University of Applied Sciences Western Switzerland
alexandre.caboussat@hesge.ch

Dimitrios Gourzoulidis
University of Applied Sciences Western Switzerland (HES-SO)
dimigourz@gmail.com

Marco Picasso
EPFL, Switzerland
marco.picasso@epfl.ch

MS231

Least Squares Galerkin Methods for Monge-Ampere Type Equations

Least squares recovery methods provide a simple and practical way to approximate linear elliptic PDEs in nondivergence form where standard variational approach either fails or requires technically complex modifications. This idea allows the creation of relatively efficient solvers for fully nonlinear elliptic equations, the linearization of which leaves us with an equation in nondivergence form. An important class of fully nonlinear elliptic PDEs is that of Hamilton–Jacobi–Bellman form. Suitable functional spaces and penalties in the cost functional must be carefully crafted in order to ensure stability and convergence of the scheme with a good approximation of the gradient and Hessian which is useful, for example, for Newton–Raphson, semismooth Newton, or a policy iteration (Howard) approximation of a Hamilton–Jacobi–Bellman equation. We prove convergence and provide convergence rates under a Cordes condition.

Omar Lakkis
University of Sussex
o.lakkis@sussex.ac.uk

MS232

Modelling High Mach Number Low-Collisional Flows Using the Maximum-Entropy Moment Method

High Mach number and low-collisional flows are frequently encountered during the high-velocity atmospheric entry of blunt bodies, such as meteoroids, satellites and spacecrafts. In the first part of the entry trajectory, the flow is highly rarefied, and convection and surface interaction dominate over particle collisions. Macroscopically, one observes a diffuse shock in front of the body and a very-low density shadow behind it. In both these regions, the gas is out of thermodynamic (translational) equilibrium, and the particle velocity distribution function deviates from the Maxwellian distribution. Traditional fluid methods such as the Euler and the Navier-Stokes-Fourier (NSF) equations are unable to reproduce these situations accurately. In this work, we investigate the application of maximum-entropy moment methods to this problem, considering the 10 and 14-moment systems. These methods consist in a system of hyperbolic partial differential equations, that extend the validity of the NSF equations to strongly non-equilibrium conditions. We analyze the accuracy of the maximum-entropy solutions and the associated computational cost,

that depends on the trajectory of the solution in moment space. A BGK collision model is employed for simplicity. Maximum-entropy results are compared to a particle-based solution of the kinetic equation.

Stefano Boccelli
University of Ottawa
stefano.boccelli@polimi.it

Pietro Parodi
VKI and KU Leuven
pietro.parodi@vki.ac.be

Thierry Magin
von Karman Institute, Belgium
magin@vki.ac.be

James McDonald
Department of Mechanical Engineering
University of Ottawa
james.mcdonald@uottawa.ca

MS232

14 and 17 Moments Systems for Polyatomic Gases - Comparison Regarding Transport Coefficients

In this talk, we will consider the kinetic model of continuous type describing a polyatomic gas [Bourgat et al., 1994; Desvillettes et al., 2005]. Such a model introduces a single continuous variable supposed to capture all the phenomena related to the more complex structure of a polyatomic molecule. We then focus on the kinetic model for which the rigorous existence and uniqueness result in the space homogeneous case is recently proven [Gamba and Pavic-Colic, 2020]. Using extended version of the cross-section proposed in that analysis we establish macroscopic models of fourteen and seventeen fields. In addition, the model introduces frozen collisions when a polyatomic molecule behaves as a monatomic one [Djordjic et al., 2021]. For the fourteen and seventeen moments model, using the publicly available computer algebra code [Djordjic et al., 2021], we compute relaxation times and transport coefficients such as shear viscosity, bulk viscosity and thermal conductivity as well as its ratio - the Prandtl number. Next, we examine differences in transport coefficients from these two moments systems and find equivalence for the specific choice of parameters from the cross-section model. However, we expect that with an increasing number of moments the transport coefficients will converge for all possible parameter choices in the cross-section.

Vladimir Djordjic
Applied and Computational Mathematics, RWTH Aachen University
djordjic@acom.rwth-aachen.de

Milana Pavic-Colic
University of Novi Sad, Serbia and
ACoM, RWTH Aachen University, Germany
milana.pavic@dmi.uns.ac.rs

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

MS232

Entropy-Stable Methods for Rarefied Gas Dynam-

ics

In this talk, we develop an entropy-stable finite-element-moment method for the Boltzmann equation with full collision operator. To that end, we employ a discontinuous-Galerkin method in space and time, and a moment method in velocity. We base our method on a converging sequence of approximations to the collision operator. We associate with each member of this sequence a normalisation map for the distribution function and an entropy. We show that each approximate collision operator inherits salient properties from Boltzmann's operator, such as the preservation of the collision operator, the H-theorem and that the linearisations near equilibrium agree. We prove that our method is entropy-stable for each member of the sequence of approximations to the collision operator. Finally, we apply our method to the Boltzmann equation with full collision operator and demonstrate the corresponding approximation properties, using benchmark test cases, in comparison to Direct Simulation Monte Carlo.

Michael Abdelmalik
Eindhoven University of Technology
m.abdel.malik@tue.nl

Irene M. Gamba
Department of Mathematics and ICES
University of Texas
gamba@math.utexas.edu

Torsten Kessler, Sergej Rjasanow
Saarland University
t.kessler@tue.nl, rjasanow@num.uni-sb.de

MS232

Analysis of Implicit PIC Schemes and Their Application to Electric Propulsion

Simulation of plasmas for space electric propulsion is a complex task because of two reasons: first, the size of interest of the system is usually comparable to the collisional length scale. As a result, velocity distributions are often in nonequilibrium. Using the hybrid Particle-in-Cell - Direct Simulation Monte Carlo (PIC-DSMC) method allows to overcome this problem, since phase space can be efficiently represented using computational particles. Secondly, the length- and time scales of interest are often much larger than the smallest natural scales of the plasma, which are dictated by electrons. Traditional explicit PIC schemes need to resolve these small scales, otherwise disrupting numerical instabilities appear. These requirements can make simulation of an entire physical device extremely expensive, and practically unfeasible above a certain plasma density. Recently, a new class of implicit and semi-implicit methods have been devised, mostly for applications in astrophysics. These feature exact total energy conservation and are stable for cell sizes and time steps well in excess of the Debye length and plasma frequency, allowing the simulation of higher density plasmas in larger domains. We present our finite element implementation of the electrostatic version of these schemes, and analyze their properties through simple numerical experiments. Finally, we investigate how these methods can be applied to the simulation of electric propulsion and laboratory plasmas.

Pietro Parodi
VKI and KU Leuven
pietro.parodi@vki.ac.be

Thierry Magin

von Karman Institute, Belgium
magin@vki.ac.be

Giovanni Lapenta
Center for Plasma Astrophysics
K.U.Leuven
giovanni.lapenta@wis.kuleuven.be

MS232

Asymptotic Preserving Discontinuous Galerkin Methods for a Linear Boltzmann Semiconductor Model

A key property of the linear Boltzmann semiconductor model is that as the collision frequency tends to infinity, the phase space density $f = f(x, v, t)$ converges to an isotropic function $M(v)\rho(x, t)$, called the drift-diffusion limit, where M is a Maxwellian and the physical density ρ satisfies a second-order parabolic PDE known as the drift-diffusion equation. Numerical approximations that mirror this property are said to be asymptotic preserving. In this paper we build two discontinuous Galerkin methods to the semiconductor model: one with the standard upwinding flux and the other with a ε -scaled Lax-Friedrichs flux, where $1/\varepsilon$ is the scale of the collision frequency. We show that these schemes are uniformly stable in ε and are asymptotic preserving. In particular, we discuss what properties the discrete Maxwellian must satisfy in order for the schemes to converge in ε to an accurate h -approximation of the drift diffusion limit. Discrete versions of the drift-diffusion equation and error estimates in several norms with respect to ε and the spatial resolution are also included.

Victor DeCaria
Oak Ridge National Laboratory
vpdecaria@gmail.com

Cory Hauck
Oak Ridge National Laboratory and University of Tennessee
hauckc@ornl.gov

Stefan Schnake
Oak Ridge National Laboratory, U.S.
schnakesr@ornl.gov

MS233

Total Generalized Variation for Geometric Inverse Problems

The total variation of the normal is discussed in the context of triangulated meshes embedded in 3D. This non-smooth regularizer requires an advanced algorithm to be used in (inverse) shape optimization problems. A split Bregman/ADMM method is used for this purpose. There, the non-smooth objective is spitted into a smooth shape optimization problem and a simple non-smooth problem. The smooth shape problem is solved by a globalized Newton method. Due to the nature of the regularizer, the first and second-order shape derivatives can not be computed by algorithmic differentiation. Therefore, their analytic form is derived and some of their properties are discussed. Numerical results are presented for mesh denoising problems as well as a simple (PDE constraint) electrical impedance tomography problem. An extension, the total generalized variation of the normal, to counteract the so-called stair-casing effect is presented as an extension.

Lukas Baumgärtner

Humboldt University, Berlin
lukas.baumgaertner@hu-berlin.de

Stephan Schmidt
Humboldt-Universität zu Berlin
s.schmidt@hu-berlin.de

Roland Herzog
Heidelberg University
roland.herzog@iwr.uni-heidelberg.de

Ronny Bergmann
Norwegian University of Science and Technology
ronny.bergmann@ntnu.no

Manuel Weiß
Heidelberg University
manuel.weiss@iwr.uni-heidelberg.de

Jose Vidal-Nunez
University of Alcalá de Henares
j.vidal@uah.es

MS233

$W^{1,\infty}$ -Approach for Shape Optimizations Using ADMM

In this contribution we present a first order approach to shape optimization problems for Lipschitz domains featuring mesh quality preserving deformations. The approach utilizes an Alternating Direction Method of Multipliers (ADMM) to determine the deformation field from the directional shape derivative of the objective function. The considered shape deformation field corresponds to the steepest descent direction in $W^{1,\infty}$ -topology. We use the surface formulation of the shape derivative of the objective function to obtain the descent direction and a standard second-order finite volume method for the discretization. Applications refer to parameter-free shape optimizations to minimize the drag of a wetted obstacle in fluid dynamics Navier-Stokes flow problems. To avoid trivial or undesired solutions, the deformation field inherently fulfills a set of geometric constraints, i.e. conserving displacement and barycenter/center of buoyancy of the obstacle. The method is first applied to a 3D configuration in laminar, single-phase flow at low Reynolds number ($Re = 10$), where the characteristic (optimal) shape of the solution is known. Subsequently, we minimize the drag of a free floating vessel exposed to turbulent (Reynolds-averaged) two-phase flows in realistic operating conditions at $Re = 1.43 \cdot 10^7$ and $Fn = 0.26$.

Peter M. Müller
Hamburg University of Technology
peter.marvin.mueller@tuhh.de

Philip J. Herbert
Heriot Watt University
p.herbert@hw.ac.uk

Michael Hinze
Universität Koblenz, Landau
hinze@uni-koblenz.de

NIKLAS Kühn, Thomas Rung
Hamburg University of Technology
niklas.kuehl@tuhh.de, thomas.rung@tuhh.de

Martin Siebenborn
University Hamburg
Mathematics Department
martin.siebenborn@uni-hamburg.de

MS233

Approximating the Inverse Hessian in 4D-Var Data Assimilation

Large-scale variational data assimilation problems are commonly found in applications like numerical weather prediction and oceanographic modelling. The 4D-Var method is frequently used to calculate a forecast model trajectory that best fits the available observations to within the observational error over a period of time. One key challenge is that the state vectors used in realistic applications could contain a very large number of unknowns so, due to memory limitations, in practice it is often impossible to assemble, store or manipulate the matrices involved explicitly. In this talk we present a limited memory approximation to the inverse Hessian based on multilevel eigenvalues and eigenvectors computed using the Lanczos method. We illustrate one use of this approximation by showing its potential effectiveness as a preconditioner within a Gauss-Newton iteration and investigate the use of randomised eigenvalue algorithms to reduce computational costs.

Alison Ramage
University of Strathclyde
a.ramage@strath.ac.uk

MS233

Shape Optimization Using Material Derivative Shape Hessians

A variational approach for higher order shape optimization based on the distributed formulation of the shape derivative is considered. Within this approach, the adjoint is used to eliminate material derivatives instead of partial derivatives, which leads to a very convenient way to derive a Newton scheme for shape optimization. In particular, this approach also allows an interpretation of the Sobolev gradient descent methods as an approximation of the shape Hessian when the perimeter is used as a regularizer. The methodology quite naturally leads to a variational formulation to compute the 2nd order shape updates and opens the door to novel numerical schemes, not only in classical aerodynamic shape optimization, but also to construct particularly fast solvers for free surface flows and microfluids.

Stephan Schmidt
Humboldt-Universität zu Berlin
s.schmidt@hu-berlin.de

MS233

On a Linear View on Shape Hessians

Shapes do not define a linear space. Therefore, the classical shape Hessian is not symmetric and does not admit a Taylor series, causing significant computational complications. This talk discusses the linear structure of deformations as a representation of shapes. This transforms shape optimization to a variant of optimal control. The numerical opportunities of this point of view are highlighted in particular in the case of large shape deformations, thus leading to fast algorithms of shape Newton type. In this way, so far often heuristic omissions of non-symmetric terms in the

shape Hessian obtain a theoretical foundation. Numerical results demonstrate the efficacy of this approach.

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

Stephan Schmidt
Humboldt-Universität zu Berlin
s.schmidt@hu-berlin.de

MS234

Efficient Spectral Method with Filters and Reconstruction

Spectral methods utilize global basis functions, and thus achieve exponential convergence, when the solution is smooth. For nonsmooth problems involving discontinuities or singularities, the accuracy degenerates. With filters and/or post-processing methods, the accuracy can be improved. In this talk, we will discuss the efficient spectral methods using filters and reconstruction techniques.

Zheng Chen
University of Massachusetts, Dartmouth
zchen2@umassd.edu

MS234

Three Applications of High-Order Polynomial Kernels That Approximate the Dirac Delta with Vanishing Moments

Necessary properties of a high-order, polynomial approximation of the Dirac delta function with compact support of the form [Jean-Piero et al., SIAM SISC, 2014]

$$\delta_\epsilon^{m,k}(x) = \begin{cases} \epsilon^{-1} P^{m,k}(x/\epsilon), & x \in [-\epsilon, \epsilon], \\ 0, & \text{otherwise.} \end{cases}$$

where $\epsilon > 0$ is the support width, include, (1) that on the compactly supported interval the regularized delta function integrates to unity. Secondly, the kernel $P^{m,k}$ must be designed to vanish the first up to the m^{th} moments and k continuous derivatives at the endpoints of the compact support. To vanish the moments the regularized delta can have negative values. The vanishing moments ensure that the regularized Dirac delta kernel converges to the exact Dirac delta function at $O(\epsilon^{m+1})$ rate. This moment property is necessary for the construction of high-order approximations of singular Dirac delta source terms in spectral approximations of PDEs [Jean-Piero et al., SIAM SISC, 2014]. To preserve high-order spatial accuracy it was further shown that the optimal value for the compact support must be $\epsilon = N_x^{-k/(m+k+2)}$. We will discuss applications of this kernel to solve PDEs with high-order order accuracy including PDEs with singular sources, the regularized solution of PDEs that admit singular solutions, and the approximation of deterministic initial conditions of governing PDEs for probability density functions.

Gustaaf Jacobs
Department of Aerospace Engineering
San Diego State University
gjacobs@mail.sdsu.edu

MS234

Beyond Uniform Meshes: Filtering and Enhance-

ment with LSIAC-MRA

Multiresolution Analysis (MRA) is a useful technique for mesh adaptivity when numerically solving PDEs with discontinuities or other interesting features. It allows for a direct relation of approximations on a coarse mesh with approximations on a fine mesh, enabling adaptive and sparse solvers for PDEs. The critical concept is interpreting a fine mesh approximation as a coarse mesh approximation with additional finer details. Using this hierarchy, it is simple to move an approximation from a fine mesh to a coarse mesh by truncation of detail coefficients. Doing the reverse is more challenging and requires a reconstruction procedure to generate fine mesh data from coarse information. It is for this purpose we consider the Line Smoothness-Increasing Accuracy-Conserving MRA (LSIAC-MRA) procedure for obtaining detail coefficients. This technique utilizes filtering methods together with multi-wavelet MRA principles to enable the transition of data from coarse to fine resolutions, simultaneously decreasing errors in the fine mesh approximation. Initially developed for uniform meshes (Picklo and Ryan, "Enhance Multiresolution Analysis..." 2022), in this talk we will demonstrate the applicability of this filter-based reconstruction technique for error reduction on nonuniform meshes.

Matthew Picklo, Jennifer K. Ryan
Colorado School of Mines
mpicklo@mines.edu, jryan@kth.se

MS234

MSIAC Package: a Filtering Tool for General Applications

The Smoothness-Increasing Accuracy-Conserving filter (SIAC) is a post-processor that extracts accuracy from DG and FEM solutions exploiting the natural superconvergence of these methods. It reduces oscillations in the errors and increases convergence rates, leading in general, to lower errors. These filters have been applied to aeroacoustic problems, visualization of streamlines and isosurfaces, adaptivity and shock regularization. Designing and implementing an effective tool for general applications implies dealing with the underlying mesh and domain boundaries, selecting suitable filter parameters, and addressing computational times. This talk presents a numerical framework and a filtering software package showing applications in DG-FEM data.

Xulia Docampo
Barcelona Supercomputing Center
xulia.docampo@indominus.eu

Jennifer K. Ryan
Colorado School of Mines
jryan@kth.se

MS235

On the Dynamical Low-Rank Numerical Method for Kinetic Equations

Solving inverse problems for kinetic equations requires an efficient solution of the forward problem. Usually, the computational effort of numerically solving kinetic equations is relatively high. We consider the concept of dynamical low-rank approximation that has recently been used for an efficient and accurate numerical solution of kinetic equations. This concept is applied to the thermal radiative transport equation and an energy stable and mass conservative dy-

namical low-rank algorithm is presented. This is joint work with Lukas Einkemmer, Christian Klingenberg and Jonas Kusch.

Lena Baumann
Wuerzburg University
lena.baumann@mathematik.uni-wuerzburg.de

MS235

Asymptotic Preserving Dynamical Low-Rank Algorithms for Kinetic Theory

Solving kinetic equations, owing to the up to 6-dimensional phase space, is extremely expensive from a computational point of view. Recently, using dynamical low-rank approximation has emerged as an attractive complexity reduction technique that, in many cases, succeeds in drastically reducing the time required to run such simulations. As kinetic equations are often posed in a context where either a fluid or diffusive limit is of interest, it is of importance to derive numerical methods that are able to cope with these limits. In this talk we will report on recent advances of such asymptotic preserving dynamical low-rank algorithms both from a practical and theoretical point of view. In particular, we consider a way to treat the compressible fluid limit of the Boltzmann-BGK equation and present a mathematical analysis based on energy estimates for the radiative transfer equation.

Lukas Einkemmer
University of Innsbruck
lukas.einkemmer@gmail.com

MS235

Numerics of Inverse Problems for Kinetic Models and Their Corresponding Macroscopic PDEs

The viewpoint proposed in this mini-symposium is that uncertainty quantification for partial differential equations (PDE) does not always represent the viewpoint of the experimentalists. Instead they want to determine the uncertain coefficient in a PDE by measuring data of the experiment corresponding to the PDE. In other words experimentalists are interested in solving an inverse problem. To this end we shall have lectures on a range of mathematical techniques needed to approach this problem. Our underlying tenet is to approach the inherent illposedness of inverse problems for macroscopic PDEs by considering inverse problems for the corresponding kinetic equations. This is well illustrated by considering models from mathematical biology, namely the motion of cells, as described by the kinetic chemotaxis equations. The corresponding macroscopic Keller-Segel type model will be a diffusion equation. The aim is to study the inverse problems in these two settings. We will hear about numerical approaches that need to be developed in this context.

Christian F. Klingenberg
Wuerzburg University, Dept. of Mathematics
Germany
klingen@mathematik.uni-wuerzburg.de

MS236

Multivariate Interpolation in Non-Tensorial Nodes May Lift the Curse of Dimensionality for Trefethen Functions

We extend 1D Newton and Lagrange interpolation to arbitrary dimensions while maintaining their numerical stabil-

ity and computational efficiency. Our generalisation relies on a proper choice of non-tensorial unisolvent interpolation nodes, whose number scales sub-exponentially with the space dimension. While we have proven the resulting interpolation scheme to approximate all continuous Sobolev functions, in addition, empirical demonstrations show that we reach the optimal exponential approximation rate for the Runge function. This suggests our interpolation method to be optimal for a class of analytic functions we term Trefethen functions. Combining sub-exponential node counts with exponential approximation rates, we may lift the curse of dimensionality for interpolation problems involving Trefethen functions. We have condensed our findings into implementations that are included in the open source Python package `minterpy`. We further sketch how `minterpy` might impact computational challenges, such as: solving (inverse) PDE problems by variational spectral methods, delivering novel solutions for computational tasks in numerical differential geometry, and realising efficient and stable regression for strongly-varying multivariate functions given on equispaced or arbitrary, scattered data.

Michael Hecht

Center for Advanced Systems Understanding (CASUS)
m.hecht@hzdr.de

Janina Schreiber

Center for Advanced System Understanding
j.schreiber@hzdr.de

Juan E. Suarez

Helmholtz-Zentrum Dresden Rossendorf
juesuarezca@gmail.com

Gentian Zavalani

Center for Advanced Systems Understanding (CASUS)
Technische Universität Dresden, D-01062 Dresden,
Germany
g.zavalani@hzdr.de

Sachin K. Thekke Veettil

TUD, MPI-CBG, CSBD
sthekke@mpi-cbg.de

MS236

Local-to-Global Support Vector Machines (LGSVMs)

Support Vector Machines (SVMs) are one of the most popular and considered classification methods. In this talk we present an efficient scheme, namely Local-to-Global Support Vector Machines (LGSVMs), which significantly reduces the complexity cost of SVMs when a large number of instances is involved. This method is a global scheme constructed by gluing together local SVM contributions via compactly supported weights, and its background lies in the framework of approximation theory and of local kernel-based models, such as the Partition of Unity (PU) scheme. The effectiveness and the efficiency of the proposed approach are proved by a theoretical analysis of its complexity, as well as by extensive numerical experiments carried out by considering benchmark datasets.

Francesco Marchetti

University of Padova (Italy)
francesco.marchetti@math.unipd.it

Emma Perracchione

Polytechnic University of Torino

emma.perracchione@polito.it

MS236

Physics-Informed Neural Networks (PINNs) and Kernel Methods

Physics-informed neural networks (PINNs) are neural networks trained using physical laws as soft constraints; these physical laws typically take the form of residuals from partial differential equations (PDEs). PINNs can be trained to directly solve PDEs or to infer PDE coefficients from solution data, and therefore can serve as meshless PDE solvers or tools for model discovery. However, they are typically challenging to train as PDE-based constraints result in complicated loss landscapes. We present new training techniques for PINNs that leverage high-order numerical differentiation to significantly speed up training without any loss in accuracy. In order to retain the meshless nature of PINNs, we use kernel-based finite difference formulas, i.e., radial basis function finite differences (RBF-FD). Given time, we will also present results on novel PINN architectures that leverage insights from the kernel literature to further speed up training.

Varun Shankar

School of Computing
University of Utah
shankar@cs.utah.edu

MS236

On the Approximability and Curse of Dimensionality of Certain Classes of High-Dimensional Functions

In this talk, I will discuss the approximability of high-dimensional functions that appear, for example, in the context of many body expansions and hdmr (high dimensional model representation). Such functions, though high dimensional, can be represented as finite sums of lower dimensional functions, i.e. by so-called additive models. I will derive sampling inequalities for such functions, give explicit advice on the location of good sampling points and show that such functions do not suffer from the curse of dimensionality.

Holger Wendland

University of Bayreuth
Germany
holger.wendland@uni-bayreuth.de

MS236

Approximation of Multivariate Functions on Sparse Grids by Kernel-Based Quasi-Interpolation

In this presentation, we introduce a new class of quasi-interpolation schemes for the approximation of multivariate functions on sparse grids. Each scheme in this class is based on shifts of kernels constructed from one-dimensional radial basis functions such as multiquadrics. The kernels are modified near the boundaries to prevent deterioration of the fidelity of the approximation. We implement our scheme using the standard single-level method as well as the multilevel technique designed to improve rates of approximation. The advantages of the proposed quasi-interpolation schemes are twofold. First, our sparse approximation attains almost the same level convergence order as the optimal approximation on the full grid related to the Strang-Fix condition, reducing the amount of data

required significantly compared to full grid methods. Second, the single-level approximation performs nearly as well as the multilevel approximation, with much less computation time. We provide a rigorous proof for the approximation orders of our quasi-interpolations. In particular, compared to another quasi-interpolation scheme in the literature based on the Gaussian kernel using the multilevel technique, we show that our methods provide significantly better approximation rates. Finally, some numerical results are presented to demonstrate the performance of the proposed schemes.

Jungho Yoon

Dept. Mathematics, Ewha W. Univ.,
Seoul, S.Korea
yoon@ewha.ac.kr

MS237

Multi-Fidelity Regression Using Artificial Neural Networks: Efficient Approximation of Parameter-Dependent Output Quantities

When evaluating quantities of interest that depend on the solutions to differential equations, we inevitably face the trade-off between accuracy and efficiency. Especially for parametrized, time-dependent problems in engineering computations, it is often the case that acceptable computational budgets limit the availability of high-fidelity, accurate simulation data. Multi-fidelity surrogate modeling has emerged as an effective strategy to overcome this difficulty. Its key idea is to leverage many low-fidelity simulation data, less accurate but much faster to compute, to improve the approximations with limited high-fidelity data. In this work, we show how to use artificial neural networks (e.g., feedforward neural networks) to tackle multi-fidelity regression problems. A first case deals with a parameter-dependent, stationary Helmholtz problem, where finite element and reduced basis methods are adopted as the high- and low-fidelity models, respectively. We then introduce a novel data-driven framework of multi-fidelity surrogate modeling for parametrized, time-dependent problems using long short-term memory (LSTM) networks, to enhance output predictions both for unseen parameter values and forward in time simultaneously. Numerical performances are assessed on a variety of engineering problems with high- and low-fidelity data generated through fine versus coarse meshes, small versus large time steps, or finite element full-order versus deep learning reduced-order models.

Paolo Conti

Politecnico di Milano, Italy
paolo.conti@polimi.it

Mengwu Guo

University of Twente
m.guo@utwente.nl

Andrea Manzoni

Politecnico di Milano
andrea1.manzoni@polimi.it

Jan S. Hesthaven

EPFL
jan.hesthaven@epfl.ch

MS237

Reduced Order Data Assimilation for Damage Pa-

rameter Estimation in Composite Structures

In order to identify damages in fiber metal laminates, we utilize techniques of Bayesian analysis and data assimilation for (damage) parameter estimation. The underlying physical model describes effects of wave propagation in fiber metal laminates and its interaction with damages, where the damage is characterized through parameters. As further source of information, we utilize (synthetic) sensor data measuring the out-of-plane displacement. Since many forward simulations are required for Bayesian analysis in order to estimate the unknown damage parameters, we utilize reduced order models in order to speed up the computational times. Numerical experiments illustrate the concepts.

Nanda Kishore Bellam Muralidhar

Technische Universität Braunschweig
Institut für Analysis und Algebra, Technische Universität Br
n.bellam-muralidhar@tu-braunschweig.de

Carmen Graessle

Technische Universität Braunschweig
Braunschweig, Germany
c.graessle@tu-braunschweig.de

Natalie Rauter, Rolf Lammering, Andrey Mikhaylenko

Institute of Mechanics
Helmut-Schmidt-University
natalie.rauter@hsu-hh.de, rolf.lammering@hsu-hh.de,
mikhayla@hsu-hh.de

Dirk Lorenz

TU Braunschweig
d.lorenz@tu-braunschweig.de

MS237

Benchmarking Machine Learning Methods for Reduced-Order Modeling on Large-Scale Datasets

There are many recent advances in using machine learning for reduced-order modeling. However, these methods are often demonstrated on dynamical systems that are small enough that the research can be done on a single processor. In this talk, we will explore how well some of these methods scale to complex systems requiring many GPUs. Challenges include having limited but high-dimensional training data and writing software that efficiently communicates between simulations and machine learning code.

Bethany Lusch

Argonne National Laboratory
blusch@anl.gov

MS237

Data-Driven Optimal and Predictive Control with Guarantees

We present novel approximation error bounds for Koopman-based surrogate modeling of control systems by means of extended dynamic mode decomposition. To this end, the approximation error is split into a deterministic projection error, depending on richness of the chosen dictionary, and a probabilistic estimation error depending on the number of samples. We apply our results in an optimal control context, providing bounds on the prediction of cost and state constraints along optimal solutions in order to derive surrogate-based predictive control schemes with

stability guarantees.

Feliks Nüske
Max Planck Institute for Dynamics of
Complex Technical Systems, Magdeburg, Germany
nueske@mpi-magdeburg.mpg.de

Sebastian Peitz
Paderborn University
sebastian.peitz@upb.de

Friedrich Philipp
Technische Universität Ilmenau
friedrich.philipp@tu-ilmenau.de

Manuel Schaller
Technische Universität Ilmenau, Germany
manuel.schaller@tu-ilmenau.de

Karl Worthmann
University of Bayreuth, Germany
karl.worthmann@tu-ilmenau.de

MS238

Quantum Circuits for Hamiltonian Simulation with Quantum Signal Processing

Hamiltonian simulation via quantum signal processing is known to be optimal. In recent years, abstract oracle circuit structures for these types of algorithms have been proposed. However, the best practice for constructing these oracles when using a universal gate set is relatively unclear. In this talk we propose compilation strategies for qubitization which minimise the number of one and two qubit gates. Where we make use of Quantinuum's pytket and show how we can build up layers of abstraction using circuit boxes to construct the high level circuit oracles. Typically these circuit primitives require many ancillary qubits and hence we use a tensor network quantum simulator to simulate these circuits. Where we show some examples of Hamiltonian simulation on the Hubbard model treated as a molecular wire.

Nathan Fitzpatrick
Quantinuum
nfitzpatrickchem@gmail.com

MS239

Parameter Estimation with Convolutional Neural Networks Applied to An ODE and a Maximum Likelihood Problem

Many inverse problems include poorly behaved objective functions or computationally infeasible likelihood functions such as for multidimensional extremes, thus making traditional approaches for computational recovery of model parameters intractable. We propose a deep learning framework to estimate parameters of two models respectively governed by ordinary differential equations and statistical models for multidimensional extremes. In both cases, we use data from model simulations as input to train deep neural networks and learn the output ODE/statistical parameters. Our neural network-based method provides a competitive alternative to existing methods such as pairwise likelihood for multidimensional extremes, as demonstrated by considerable accuracy and computational time improvements.

Julie Bessac

Argonne National Laboratory
jbessac@anl.gov

MS239

Riemannian Multifidelity Covariance Estimation from Statistically Coupled Observations

We present a framework for multifidelity covariance estimation from statistically coupled pairs of high- and low-fidelity random variables. We define a multifidelity covariance estimator as the solution to a regression problem on tangent spaces to product manifolds of symmetric positive definite (SPD) matrices. Given a set of high- and low-fidelity sample covariance matrices, viewed as a sample of a product-manifold-valued random variable, we estimate the underlying true covariance matrices by minimizing a notion of squared Mahalanobis distance between the data and a model for its variation about its mean. This estimator can be employed using any Riemannian geometry for the SPD manifold and reduces to control variates in the chosen geometry when all low-fidelity optimization variables are fixed. We particularly focus on the affine-invariant SPD manifold geometry, under which multifidelity covariance regression estimates are guaranteeably positive definite and the Mahalanobis distance has desirable properties. We demonstrate under three geometries that our estimator can provide significant reductions in MSE relative to single-fidelity covariance estimators. This framework is relevant to the data assimilation setting, which often requires repeated evaluation of expensive high-fidelity forecast models, and particularly applicable to the ensemble Kalman filter, in which covariances must be computed from forecast model evaluations in order to perform the analysis step.

Aimee Maurais, Youssef M. Marzouk
Massachusetts Institute of Technology
maurais@mit.edu, ymarz@mit.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

Terrence Alsup
New York University
Courant Institute of Mathematical Sciences
alsup@cims.nyu.edu

MS239

A Density Estimating Filter Based on Marginal Coupling

In many DA problems, the number of observations is much lower than the number of states. Assuming that the observation is just a subset of states, certain filtering methods accurately sample the marginal posterior of the observed dimensions, but lack a good mechanism to sample from the unobserved dimensions. In this work, we investigate using copulas to sample the unobserved dimension to obtain samples consistent with the joint Bayesian posterior.

Amit Subrahmanya
Virginia Tech
amitns@vt.edu

Julie Bessac
Argonne National Laboratory
jbessac@anl.gov

Andrey A. Popov
University of Texas at Austin
apopov@vt.edu

Adrian Sandu
Virginia Polytechnic Institute and State University
Computational Science Laboratory
sandu@cs.vt.edu

MS239

Unbiased Filtering and Smoothing for Very High-Dimensional Geophysical Systems: Stochastic Particle Flows

Particle Flow Filters and Smoothers allow for sampling posterior probability density functions (pdf) in very high-dimensional spaces. They are based on iterative minimization of the KL-divergence (or other distance measures) between the pdf represented by the particles and the posterior pdf. The methodology can be seen as an ensemble of 3Dvars for a filter, and an ensemble of 4Dvars for a smoother, in which the particles interact during the minimization. Unfortunately, deterministic particle flows are biased for finite ensemble sizes. Recently it has been realized that this bias disappears for any ensemble size when a stochastic version of the particle flow is used. The idea is to consider the flow as a Langevin dynamics over the configuration space spanned by the particles. This leads to an extra stochastic term in the minimization equations, which turns out to be relatively simple to implement. A further problem is that particle flow filters need the gradient of the log of the prior pdf, which is only known by its samples. Hence calculating its log gradient is problematic. We developed a practical solution to this problem of estimating the gradient of the log of the prior in high-dimensional sequential applications, using kernel methods and localization procedures. In this talk I will briefly explain the theory and present results of particle flow filters and smoothers for high-dimensional geophysical systems, comparing stochastic and deterministic versions.

Peter Jan van Leeuwen
Colorado State University and University of Reading
peter.vanleeuwen@colostate.edu

MS240

Fast Methods in Volumetric Potential Theory for PDEs and Applications

This talk will provide an overview of past and ongoing work in fast, high-order accurate methods for solving volumetric problems in conjunction with integral equation techniques. We will also outline a few of our own contributions in the area of high-order accurate kernel-independent methods for evaluation of the Newtonian potential based on numerical quadrature. The methods are based on corrections to (FMM-accelerated) quadratures for smooth functions, with rapidly-generated pre-computed corrections that can be efficiently applied using orthogonal polynomials, achieving > 99% of fast algorithm performance. We will also demonstrate adaptive solvers using quadtree and box-code based fast algorithms. Some of this work is joint with Shravan Veerapaneni (University of Michigan) and Hai Zhu (Flatiron Institute).

Thomas G. Anderson
University of Michigan, Ann Arbor

tganders@umich.edu

MS240

Computation of Negative Fractional Sobolev Norms on Arbitrary Domains: a Volume Potential Approach

The evaluation of weak (negative Sobolev) norms of functions with arbitrary bounded support $D \subset \mathbb{R}^d$ arises e.g. for quantifying mixing unevenness of concentrations advected by fluid flows. In this work, we address the evaluation of the $H^{-r}(D)$ norm of a function $c \in L^2(D)$, with $r \in]1/2, 1]$. To this end, the Riesz representative $u[c] \in H_0^r(D)$ of c is sought, and we then have $\|c\|_{H^{-r}(D)} = (u[c], c)_{L^2(D)}$. Computing $u[c]$ by domain discretization methods (e.g. FEM) for fractional r would entail either the costly accurate evaluation of sufficiently many Dirichlet Laplace eigenfunctions of D or multiple elliptic problem solves. By contrast, the computational treatment proposed here rests on one volume integral operator (VIO) evaluation with density c followed by solving one single-layer integral equation (SLIE) on ∂D (ensuring $\gamma u[c] = 0$). The kernel used in both the VIO and the SLIE depends in closed form on r and the physical coordinates (and has a singularity of order $2r - d$); in fact, the VIO is directly connected with the $H^{-r}(\mathbb{R}^d)$ Fourier-Bessel norm. We prove that the SLIE in weak form is coercive on $H^{1/2-r}(\partial D)$ for any $r \in]1/2, 1]$, making it well-posed. Our procedure thus naturally extends to fractional r the method solving in $H_0^1(D)$ the Poisson problem with given density c . We demonstrate the proposed approach on computational examples with available reference solutions.

Marc Bonnet
POems (CNRS-ENSTA-INRIA), Appl. Math. Dept.
Paris, France
marc.bonnet@ensta.fr

Thomas G. Anderson
University of Michigan, Ann Arbor
tganders@umich.edu

Shravan Veerapaneni
Department of Mathematics
University of Michigan
shravan@umich.edu

MS240

Volume Integral Equations and Single-Trace Formulations for Acoustic Wave Scattering in an Inhomogeneous Medium

We study frequency domain acoustic scattering at a bounded, penetrable, and inhomogeneous obstacle. By defining constant reference coefficients, a representation formula for the pressure field is derived. It contains a volume integral operator, related to the one in the Lippmann-Schwinger equation. Besides, it features integral operators defined on the boundary of the obstacle and closely related to boundary integral equations of single-trace formulation (STF) for transmission problems with piecewise constant coefficients. We show well-posedness of the continuous variational formulation and asymptotic convergence of Galerkin discretizations. Numerical experiments in 2D and 3D validate our expected convergence rates.

Ignacio Labarca
ETH Zurich

Seminar for Applied Mathematics
ignacio.labarca@sam.math.ethz.ch

Ralf Hiptmair
ETH Zurich
ralf.hiptmair@sam.math.ethz.ch

MS240

Density Interpolation Methods for Volume Potentials

We present a simple numerical method for the high-order numerical evaluation of volume integral operators associated with the Poisson and Helmholtz equations in two spatial dimensions. Following the ideas of the density interpolation method for boundary integral operators, the proposed methodology leverages Green's third identity and a local polynomial interpolation of the source function to recast the volume potential as a sum of single- and double-layer potentials and a volume integral with a regularized (bounded or smoother) integrand. The layer potentials can be accurately and efficiently evaluated everywhere by means of existing methods (e.g. the density interpolation method), while the regularized volume integral can be accurately evaluated by applying elementary quadrature rules. Numerical examples based on non-overlapping quadrilateral patch representations of the domain in conjunction with Chebyshev-grid discretizations demonstrate the effectiveness of the proposed approach.

Carlos Perez-Arancibia
University of Twente
Department of Applied Mathematics
c.a.perezarancibia@utwente.edu

Thomas G. Anderson
University of Michigan, Ann Arbor
tganders@umich.edu

Marc Bonnet, Luiz Faria
ENSTA Paris
marc.bonnet@ensta-paris.fr, luiz.maltez-faria@inria.fr

MS240

A Phase Retrieval Approach to Characterize Underwater Transducers Using Time-Harmonic Near-Field Data

Active acoustic devices have been used profusely in monitoring activities in coastal and oceanic marine environments for fishery purposes and the health evaluation of the seabed and its related biological ecosystems. The accurate numerical characterization of those active transducers, which insonifies the underwater fluid media, plays a crucial role in the subsequent numerical simulations of complex marine environments. The transducers can be quickly characterized experimentally using in-house laboratory equipment where their acoustic response can be obtained from time-harmonic near-field phaseless data. Obviously, the extrapolation of those near-field data to predict the time-harmonic far-field generated pressure cannot be computed straightforwardly and requires estimating the directivity pattern associated with the transducer. Different numerical approaches based on an integral representation of the pressure field are analyzed in this work. Since the phaseless experimental near-field data is known in a reduced number of locations, different polynomial discretizations have been utilized to discretize the directivity pat-

tern taking into account low-order spherical harmonics and standard finite element polynomial approximations. Both approaches are compared in realistic scenarios with available closed-form solutions such as omnidirectional sources and end-fire arrays. The impact of using different phase retrieval and quadrature methods is also quantified.

Andrés Prieto
University of A Coruña
Galician Center for Mathematical Research and Technology
andres.prieto@udc.es

Sara Recondo
Marine Instruments
University of A Coruña
sara.recondo@udc.es

MS241

A Stable Mimetic Finite-Different Discretization for Convection-Dominated Diffusion Equations

Convection-diffusion equations arise in a variety of applications such as particle transport, electromagnetics, and magnetohydrodynamics. Simulation of the convection-dominated case, even with high-fidelity techniques, is particularly challenging due to sharp boundary layers and shocks causing jumps and discontinuities in the solution, and numerical issues such as loss of the maximum principle in the discretization. These complications cause instabilities, admitting large oscillations in the numerical solutions when using traditional methods. Drawing connections to the simplex-averaged finite element method (S. Wu and J. Xu, 2020), we develop a mimetic finite-difference (MFD) discretization using exponentially averaged coefficients to guarantee monotonicity of the scheme and stability of the solution as the diffusion coefficient approaches zero. The finite-element framework allows for transparent analysis of the MFD, such as proving well-posedness and deriving error estimates from the finite-element setting. Numerical tests are presented confirming the stability of the method and verifying the error estimates.

Casey Cavanaugh
Louisiana State University
ccavanaugh@lsu.edu

Xiaozhe Hu, James H. Adler
Tufts University
xiaozhe.hu@tufts.edu, james.adler@tufts.edu

Nathaniel Trask, Andy Huang
Sandia National Laboratories
natrask@sandia.gov, ahuang@sandia.gov

MS241

High Order Quasi-Conservative Discontinuous Galerkin Schemes for Hyperbolic Systems in Primitive Variables

In this talk, we present a novel quasi-conservative high order Discontinuous Galerkin (DG) method able to capture contact discontinuities avoiding any spurious numerical artifacts, thanks to the PDE evolution in primitive variables, while at the same time being strongly conservative on shocks, thanks to a conservative a posteriori subcell Finite Volume (FV) limiter. In particular, in this work we consider ADER-DG schemes on Voronoi meshes, and we treat in primitive variables both the predictor step (that

produces a local high order accurate space-time reconstruction) and the corrector step which takes care of the final solution update. This is possible because the algorithm is then supplemented with an a posteriori FV limiter, where the solution is checked for physical admissibility, and on the cells judged to be troubled, it is locally recomputed via a robust FV scheme. At this point, we distinguish, by using a criterion based on the divergence of the velocity field, between troubled zone caused by shock discontinuities and all the rest and only on shocks we now apply a FV acting on the conservative version of the PDE. To prove the capabilities of our approach, after having verified its correctness on classical benchmarks for the single-fluid Euler equations, we show its improved reliability on the multi-fluid Euler system on examples, like the interaction of a shock with a helium bubble, for which we are able to avoid the development of any spurious oscillations.

Elena Gaburro
University of Trento
elena.gaburro@inria.fr

Walter Boscheri
University of Ferrara
Italy
walter.boscheri@unife.it

Simone Chiocchetti
University of Trento
simone.chiocchetti@unitn.it

Mario Ricchiuto
INRIA
mario.ricchiuto@inria.fr

MS241 Robust and Efficient Approximation of the Compressible Navier-Stokes Equations

Structure preserving numerical methods provide theoretical guarantees of reliability for situations where ad-hoc stabilization techniques can fail. In this talk we present fully discrete approximation techniques for the compressible Euler and Navier-Stokes equations that is second-order accurate in time and space and guaranteed to be invariant domain preserving. This means the method maintains important physical invariants and is guaranteed to be stable without the use of ad-hoc tuning parameters. We discuss the underlying algebraic discretization technique based on collocation and convex limiting.

Matthias Maier
Texas A&M University
maier@tamu.edu

MS241 A Stable and Adaptive FE Method for Convection-Dominated Diffusion Problems

We present a stabilized finite element (FE) method for linear convection-diffusion boundary value problems (BVPs). The stabilization is based on the use of specially constructed test spaces and due to a first-order system representation of the underlying BVP, the resulting FE approximations rely on classical trial spaces such as C^0 polynomials and Raviart-Thomas functions. In addition to the FE approximations, the method also allows straightforward computation of a residual-based error representation function. This function leads to a posteriori error esti-

mates which restriction to each element is an error indicator which we employ in adaptive mesh refinements. We present numerical verifications of the method as well as adaptive mesh refinement algorithms based on existing element marking strategies for problems in which convection is strongly dominating.

Eirik Vålseth
Oden Institute, UT Austin
eirik@utexas.edu

Clint Dawson
University of Texas at Austin
clint.dawson@austin.utexas.edu

Albert Romkes
South Dakota School of Mines and Technology
albert.romkes@sdsmt.edu

MS242 A First Approach to the Design of Fully Well-Balanced Lagrange-Projection Schemes for the Shallow Water Model

We will first propose a method to design numerical approximations of the shallow-water equations that are both implicit and exactly well-balanced for water at rest steady states. This is achieved through a two-step process, where the shallow-water system is first solved in Lagrangian coordinates, called the Lagrangian step, and then the results are projected into Eulerian coordinates, called the Projection step. The Lagrangian step can be done implicitly, while the Projection step is always done explicitly. This approach separates the acoustic and the transport phenomena, allowing for larger time steps with restrictions based on slower transport waves instead of acoustic ones. Furthermore, we are currently working on developing implicit fully well-balanced schemes that preserve all the steady states. In this sense, we will also be providing an idea of how to develop explicit fully well-balanced schemes, which will help us accomplish our goal.

Celia Caballero-Cárdenas
University of Málaga
Spain
celiacaba@uma.es

Castro-Díaz Manuel Jesus
University of Malaga
mjcastro@uma.es

Tomás Morales de Luna, María Luz Muñoz-Ruiz
University of Málaga
tmorales@uma.es, mlmunoz@uma.es

MS242 An Asymptotic Preserving Scheme for the Two-Dimensional Shallow Water Equations with Coriolis Forces

We consider the two-dimensional Saint-Venant system of shallow water equations with Coriolis forces. We focus on the case of a low Froude number, in which the system is stiff and conventional explicit numerical methods are extremely inefficient and often impractical. Our goal is to design an asymptotic preserving (AP) scheme, which is uniformly asymptotically consistent and stable for a broad range of (low) Froude numbers. The goal is achieved using the flux

splitting. We split the flux into the stiff and nonstiff parts and then use an implicit-explicit approach: apply an explicit hyperbolic solver to the nonstiff part of the system while treating the stiff part of it implicitly. Moreover, the stiff part of the flux is linear and therefore we reduce the implicit stage of the proposed method to solving a Poisson-type elliptic equation, which is discretized using a standard second-order central difference scheme. We conduct a series of numerical experiments, which demonstrate that the developed AP scheme achieves the theoretical second-order rate of convergence and the time-step stability restriction is independent of the Froude number. This makes the proposed AP scheme an efficient and robust alternative to fully explicit numerical methods.

Alina Chertock

North Carolina State University
Department of Mathematics
chertock@math.ncsu.edu

MS242

Implicit and Semi-Implicit High-Order Well-Balanced Methods for One-Dimensional Systems of Balance Laws

The goal of this work is to design one-dimensional implicit and semi-implicit high-order well-balanced numerical schemes for general systems of balance laws with stiff numerical flux and/or source term. When the relaxation parameter is very small, the relaxation terms become very strong and highly stiff, and numerical schemes may produce spurious results. The well-balanced reconstruction procedure introduced by two of the authors in some previous work is applied, based on the development of well-balanced reconstruction operators. The well-balanced property is preserved when quadrature formulas are used to approximate the averages and integral of the source term in the cells. This technique is combined with a time discretization method for the time evolution of type RK-IMEX or RK-implicit. The methodology will be applied to several systems of balance laws, such as the Burgers equation or the shallow water equations with Manning friction.

Irene Gómez-Bueno

University of Málaga
Spain
igomezbueno@uma.es

Sebastiano Boscarino
University of Catania
boscarino@dmf.unict.it

Castro-Díaz Manuel Jesus
University of Malaga
mjcastro@uma.es

Carlos Parés
University of Málaga. Spain
pares@anamat.cie.uma.es

Giovanni Russo
University of Catania
Department of Mathematics and Computer Science
russo@dmf.unict.it

MS242

Steady States and Well-Balanced Schemes for Shal-

low Water Moment Equations with Topography and Wet-Dry Areas

We consider a new Shallow Water Moment Equations (SWME) model based on the ones described in [Koellermeier, J. and Rominger, M.; Analysis and Numerical Simulation of Hyperbolic Shallow Water Moment Equations; Communications in Computational Physics; 2020] and [Kowalski, J. and Torrilhon, M.; Moment Approximations and Model Cascades for Shallow Flow; Communications in Computational Physics; 2019]. This new system developed in [Koellermeier, J. and Pimentel-García, E.; Steady States and Well-Balanced Schemes for Shallow Water Moment Equations with Topography; Applied Mathematics and Computations; 2022] is a nonconservative hyperbolic system from which we can get some information about its steady state solutions. This is mandatory when using the high-order well-balanced methodology developed in [Castro, M.J., Pars, C.; Well-balanced high-order finite volume methods for systems of balance laws; Journal of Scientific Computing; 2020]. We focus on the first- and second-order schemes and we solve the different difficulties that we find. Some numerical tests are shown in order to verify the well-balanced property of the methods and how they deal when wet-dry fronts appear in the solution.

Ernesto Pimentel-García

University of Málaga
erpigar@uma.es

Julian Koellermeier
University of Groningen
j.koellermeier@rug.nl

MS242

Asymptotic Preserving IMEX Schemes for the Shallow Water Equations on Unstructured Meshes

We propose a novel numerical method for the solution of the shallow water equations in different regimes of the Froude number making use of general polygonal meshes. The fluxes of the governing equations are split such that advection and acoustic-gravity sub-systems are derived, hence separating slow and fast phenomena. This splitting allows the nonlinear convective fluxes to be discretized explicitly in time, while retaining an implicit time marching for the acoustic-gravity terms. Consequently, the novel schemes are particularly well suited in the low Froude limit of the model, since no numerical viscosity is added in the implicit solver. Besides, stability follows from a milder CFL condition which is based only on the advection speed and not on the celerity. High order time accuracy is achieved using the family of semi-implicit IMEX Runge-Kutta schemes, while high order in space is granted relying on two discretizations: (i) a cell-centered finite volume (FV) scheme for the nonlinear convective contribution on the polygonal cells; (ii) a staggered discontinuous Galerkin (DG) scheme for the solution of the linear system associated to the implicit discretization of the pressure sub-system. The novel schemes are proved to be Asymptotic Preserving (AP) and well-balanced by construction. Accuracy and robustness are then validated against a set of benchmark test cases with Froude numbers ranging in the interval $\approx [10^{-6}; 5]$.

Maurizio Tavelli

University of Bozen
maurizio.tavelli@unibz.it

Walter Boscheri

University of Ferrara
Italy
walter.boscheri@unife.it

MS243

Power and Performance Oriented Analysis of Stencils

Finite-difference stencils have been always one of the key numerical schemes to solve PDE such as wave propagation for seismic computing, but also in other scientific domains, such as heat transfer, fluid dynamic, electromagnetic and structural mechanics problems. Common performance analysis focuses on the arithmetic intensity to demonstrate the efficiency of algorithms by using well known analysis such as multi-level rooflines when using caches for both spatial and temporal blocking. It demonstrates that in some cases, finite differences-based algorithms can easily extend from high memory bandwidth demand, up to strong compute demand. As energy-efficiency becomes crucial in HPC in general, we focus on the energy-to-solution extra dimension. We develop ad hoc tools based on collection of CPU specific performance counters, for instance instruction retired, frequencies and power/energy. Gathering this altogether, we extract the pJ/FP_{op} FoM for various algorithms and different optimization levels and technics. As the various algorithms have quite different memory bandwidth or compute demands, we collect the performance and power data at different CPU core and uncore frequency data points. Our goal is to characterize different seismic computing algorithms from both the time-to-solution and the energy-to-solution standpoints and adapt the both code optimization and CPU settings to get the best Perf/Watt.

Sauge Ludovic, Cedric andreolli, Mrabet Amine, Thierry Philippe
Intel
ludovic.sauge@intel.com, cedric.andreolli@intel.com,
amine.mrabet@intel.com, philippe.thierry@intel.com

MS243

Automated Temporal Blocking in the Devito Compiler

The finite-difference (FD) method often solves PDEs, resulting in stencil kernels. Our work is motivated by accelerating stencil kernels through performance optimisations that exploit memory hierarchy. We focus on temporal blocking (aka time-tiling), which has long been proven beneficial for performance. Temporal blocking reduces the required memory bandwidth of stencil computations by re-using data from the cache for multiple time steps. However, applying temporal blocking to practical applications' stencils remains challenging. These computations are more complex as they include sparsely located operators not aligned with the computational FD grid, like sources and receivers that perform scatter and gather operations, boundary and initial conditions, complicated physics, data dependencies and others. Like other stencil optimisations, temporal blocking is error-prone and hard to implement by hand, requiring careful software engineering. We present a fully-automated temporal blocking code generation through Devito, demonstrated in various applications of significant industrial interest. Devito offers a DSL to model simulations from a high-level symbolic abstraction and a compiler framework to automatically generate optimised stencil solvers. Our implementation works in tandem with other domain-specific and compiler opti-

misations. We evaluate the performance of our scheme on modern CPUs against highly-optimised stencil kernels generated by the Devito framework.

George Bisbas
Imperial College London
g.bisbas18@imperial.ac.uk

Fabio Luporini
Devito
fabio@devitocodes.com

Mathias Louboutin
Georgia Institute of Technology
mlouboutin3@gatech.edu

Rhodri Nelson
Imperial College, United Kingdom
rhodri.nelson@imperial.ac.uk

Gerard J. Gorman, Paul Kelly
Imperial College London
g.gorman@imperial.ac.uk, p.kelly@imperial.ac.uk

MS243

Stencil Operator Inference Using Regression Techniques

Due to rapid developments in the field of machine learning and data science, the derivation of mathematical relationships by learning ODEs or PDEs from field data has gained much attraction in recent years. While this directly yields a (more or less) continuous description of the underlying problem, this leaves the application experts still with the burden of discretizing the problem when a numerical simulation is anticipated. I present recent work on the inference of stencil expressions from prescribed field data, i.e. learning directly a discrete relationship from the data. Focussing on the use of regression techniques, I discuss challenges such as the handling of noise in the data, stability of the derived stencil expressions or the treatment of heterogeneous coefficients in space or time. I will base my discussion on various well-known first- and second-order differential problems and their solutions which are used to generate the field data.

Philipp Neumann
Helmut Schmidt University, Germany
philipp.neumann@hsu-hh.de

Yannis Schumann
Helmut-Schmidt-University Hamburg
Chair for High Performance Computing
yannis.schumann@hsu-hh.de

MS243

Leveraging Stencil Computation Performance with Temporal Blocking

Stencil computations constitute the core kernel in many temporally explicit approaches found in structured grid finite-difference, finite-volume, and finite-element discretizations of partial differential equation conservation laws. Using various blocking dimensions, the Spatial Blocking (SB) approach enables data reuse within multiple cache levels. Introduced in GIRIH, the Multi-core Wavefront Diamond blocking (MWD) method further optimizes memory accesses of stencil algorithms by combining

the concepts of diamond tiling and cache-aware wavefront temporal blocking, leading to a significant increase in data reuse and locality. We evaluate the performance of MWD on a variety of recent multi-core architectures.

Long Qu
King Abdullah University of Science and Technology
(KAUST)
long.qu@kaust.edu.sa

Hatem Ltaief
KAUST, Saudia Arabia
hatem.ltaief@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology
(KAUST)
david.keyes@kaust.edu.sa

MS244

Weak form System Identification: Computational Efficiency and Application to an Ideal MHD Fluid

In recent years, data-driven modeling methods have received substantial attention for their ability to learn equations directly from data. These approaches consider a candidate library of terms and use sparse regression to prune the model to find the model which best describes the data. Our Weak form Sparse Identification of Nonlinear Dynamics (WSINDy) method transforms a vast number of computationally challenging system identification problems into a simple weighted least squares solve. These systems can be identified in seconds (on a laptop) and for many equations in the presence of greater than 100% noise levels. In this talk, I will present an overview of WSINDy, highlighting the features that give WSINDy its computational efficiencies. I will also present an illustration of applying WSINDy to several challenging examples, including recovering the equations of an ideal magneto-hydrodynamic fluid in a simple test problem. I will also discuss potential future improvement efforts.

David M. Bortz
University of Colorado
Department of Applied Mathematics
dmbortz@colorado.edu

Dan Messenger
Department of Applied Mathematics
University of Colorado - Boulder
daniel.messenger@colorado.edu

Brian O'Shea
Dept. of Computational Mathematics, Science and Engineering
Michigan State University
oshea@msu.edu

Gina Vasey
Michigan State University
vaseygin@msu.edu

Andrew J. Christlieb
Michigan State University
Dept. of Comp. Math., Sci & Engr.

christli@msu.edu

MS244

Structure Preserving Machine Learning, a Symplectic Nn

In this talk, I will discuss symplectic neural networks and neural asymptotic-preserving surrogates. For the first part I will use field-line Poincare plots as an example and for the second part I will use a recently developed nearly-periodic symplectic neural network as an example. The latter can be applied to guiding center dynamics of plasma.

Joshua Burby
Los Alamos National Laboratory
maruchanil1@gmail.com

Qi Tang
Department of Mathematical Sciences
Rensselaer Polytechnic Institute
tangqimath@gmail.com

MS244

Non-Equilibrium Transport and the Curse of Dimensionality, the Need for Advanced Methods

Even with exascale computing, modeling non-equilibrium transport for plasmas has major challenges. At the heart of these challenge is the curse of dimensionality (CoD). The CoD acknowledges that computational complexity grows exponentially with the number of dimensions, making many problems (specially in non-equilibrium regimes) intractable. In this talk we will outline the challenges and give a brief over view of a new effort, the center for hierarchical and robust modeling of non-equilibrium transport (CHaRMNET). CHaRMNET overview: The curse of dimensionality acknowledges that computational complexity grows exponentially with the number of dimensions, making many problems (specially in non-equilibrium regimes) intractable. The CHaRMNET team is focused on developing a family of hierarchical approaches for problems in plasmas, related to fusion energy, that leverage novel reduced representations. CHaRMNET brings together traditional hierarchical models in plasma science, structure preserving algorithms (including structure preserving ML), ML methods for model discovery, blended algorithms that merge ML with traditional scientific computing, ML low rank approximations, sparse-grid approximations, and new algorithms for non-traditional computing platforms. This multi-fidelity multi-model holistic approach to algorithm development is designed to enable new discoveries and is an essential aspect of CHaRMNET's vision.

Andrew J. Christlieb
Michigan State University
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

MS244

A Local Macroscopic Conservative (LoMaC) Low Rank Tensor Method with the Discontinuous Galerkin Method for the Vlasov Dynamics

In this paper, we propose a novel Local Macroscopic Conservative (LoMaC) low rank tensor method with discontinuous Galerkin (DG) discretization for the physical and phase spaces for simulating the Vlasov-Poisson (VP) system. The LoMaC property refers to the exact local con-

servation of macroscopic mass, momentum and energy at the discrete level. The LoMaC low rank tensor algorithm (recently developed in arXiv:2207.00518) simultaneously evolves the macroscopic conservation laws of mass, momentum and energy using the kinetic flux vector splitting; then the LoMaC property is realized by projecting the low rank kinetic solution onto a subspace that shares the same macroscopic observables. This paper is a generalization of our previous work, but with DG discretization to take advantage of its compactness and flexibility in handling boundary conditions and its superior accuracy in the long term. Extensive numerical results are performed to showcase the efficacy of the method.

Jingmei Qiu
University of Delaware
jingqiu@udel.edu

Wei Gou
Texas Tech University
weimath.guo@ttu.edu

MS245 Massively Parallel Bayesian Optimization

One way to reduce the time of conducting optimization studies is to evaluate designs in parallel rather than just one-at-a-time. For expensive-to-evaluate black-boxes, batch versions of Bayesian optimization have been proposed. They work by building a surrogate model of the black-box that can be used to select the designs to evaluate efficiently via an infill criterion. Still, with higher levels of parallelization becoming available, the strategies that work for a few tens of parallel evaluations become limiting, in particular due to the complexity of selecting more evaluations. It is even more crucial when the black-box is noisy, necessitating more evaluations as well as repeating experiments. Here we propose a scalable strategy that can keep up with massive batching natively, focused on the exploration/exploitation trade-off and a portfolio allocation. We compare the approach with related methods on noisy test functions. The results show similar or better performance than existing methods, while being orders of magnitude faster. Then we illustrate the potential on the optimization of a large-scale epidemiological simulator.

Mickael Binois
Inria
mickael.binois@inria.fr

Nicholson Collier, Jonathan Ozik
Argonne National Laboratory
nccollier@anl.gov, jozik@anl.gov

MS245 Handling Hidden Constraints in Blackbox Optimization

In the context of blackbox optimization, the values of the objective function and of the constraints are given through industrial experiments of computer simulations. Sometimes, the values cannot be returned because the experiment failed or the simulation crashed. This causes to lose an evaluation in the budget for a point that brings very little information. In such case, it is said that the point hits a 'hidden constraint'. Therefore, it becomes helpful to avoid those hidden constraints and use the information given by the fact that a point is hitting or not a hidden constraint. This presentation offers different ways to handle

those hidden constraints using surrogates and subproblem formulations in the algorithm MADS and shows numerical results.

Stéphane Jacquet
Polytechnique Montreal
stephanejacquet2005@hotmail.fr

MS245 Derivative-Free Variance-Reduced Jacobian Sketching

We consider the setting of derivative-free finite-sum minimization. Motivated in particular by problems in nuclear model calibration, we suppose that each summand function is computationally expensive, and that moreover, the summand functions can be evaluated independently. In this setting, we propose and demonstrate a novel method inspired by variance reduction methods in machine learning; such a method permits us to more judiciously select a subset of summand functions to evaluate on each iteration according to a particular probability distribution. We then combine this methodology with sketching methods, enabling judicious sketches of the problem Jacobian. Numerical results demonstrating the efficiency of our Jacobian sketching method will be presented.

Matt Menickelly
Argonne National Laboratory
mmenickelly@anl.gov

Stefan Wild
Argonne National Laboratory
LANS, Mathematics and Computer Science
wild@lbl.gov

MS245 Derivative-Free Spectral Residual Methods for Nonlinear Systems of Equations and Applications to Wheel-Rail Contact Models in Railway Systems

Spectral residual methods are derivative-free and low-cost per iteration procedures for solving systems of nonlinear equations [La Cruz, W., Martinez, J. M., Raydan, M. 2006 Spectral residual method without gradient information for solving large-scale nonlinear systems of equations. *Math. Comput.* 75, 1429-1448]. They are generally coupled with a nonmonotone linesearch strategy and compare well with Newton-based methods for large nonlinear systems and sequences of nonlinear systems. The residual vector is used as the search direction and the steplength is inspired by the Barzilai Borwein method [Barzilai, J., Borwein, J. 1988 Two point step gradient methods. *IMA J. Numer. Anal.* 8, 141-148]. Analogously to spectral gradient methods for minimization, choosing the steplength has a crucial impact on the performance of the procedure. In this talk we address, both theoretically and experimentally, the steplength selection and provide results on a real application such as a rolling contact problem modelling a wheel-rail contact in railway systems.

Enrico Meli
University of Firenze
Dipartimento di Ingegneria Industriale
enrico.meli@unifi.it

Benedetta Morini
Dipartimento di Ingegneria Industriale
Universita' di Firenze

benedetta.morini@unifi.it

Margherita Porcelli
Dipartimento di Matematica
Universita' di Bologna
margherita.porcelli@unibo.it

Cristina Sgattoni
Institute for Applied Mathematics Mauro Picone (IAC)
CNR
cristina.sgattoni@fi.iac.cnr.it

MS245

Stochastic Blackbox Optimization Methods in the Presence of Dynamical Constraints

In simulation-based optimization, one is often faced with constraints that stem from physical processes, and are expressed under the form of differential equations. Recent interest in machine learning architectures based on differential equations has generated renewed interest for this class of problems. Indeed, it gave rise to several complex optimization formulations where the dynamics play a prominent role and the objective function can be viewed as the result of an expensive procedure, typically not directly available to the optimizer. In this talk, we investigate constrained optimization problems where the objective function is the result of a blackbox simulation, but the dynamics expressed in the constraints are available as a white box. We provide an algorithmic framework that is equipped with theoretical guarantees, even when the objective function cannot be accessed directly, and stochastic estimates are available instead. We also illustrate the performance of our algorithm on task involving neural architectures inspired by differential equations.

Clement W. Royer
Université Paris Dauphine-PSL
clement.royer@lamsade.dauphine.fr

MS246

Enabling An Environment for a Sustainable Software Development on HPC Systems

Research software is an increasingly important pillar of the scientific environment. This comes with an increase in importance for a sustainable and efficient development and development workflows and RSE practices to be applied for research software projects. An important cornerstone for the efficient and sustainable development of these are continuous integration and testing or similar techniques. Commonly available solutions rely on shared, public infrastructures and services. While these offer good accessibility and usability for many projects, they are not well suited for all demands of software targeting usage on HPC clusters. Such projects need to be developed, tested and verified with the targeted cluster, architecture and potential usage of accelerators in mind. In addition, infrastructure offering a stable and reliable environment, especially in terms of performance, are an essential component for deriving and comparing benchmark results in a continuous manner. For supporting research software projects targeting HPC clusters, an environment for a sustainable development, which offers a similar usage-experience as the previously mentioned services, needs to be adopted, integrated and offered on HPC clusters.

Rene Caspart
Karlsruhe Institute of Technology (KIT)

rene.caspart@kit.edu

MS246

On the Path Towards Better Sustainability and Productivity for Research Software

This talk will introduce the minisymposium and discuss the ongoing challenges and solutions that will take the research software development and user communities on a path towards improved software sustainability and productivity in scientific computing. We will discuss broad challenges in these areas, both from the technical perspective (how different tools, techniques and better practices help) as well as from the social perspective (how do people/teams/communities that are made of personnel such as research software engineers (RSEs) help). This talk will inspire attendees to reflect on the roadblocks they face in their teams, organizations, and communities; how to realign our focus on meeting these challenges; and how to encourage culture change so as to build better scientific software.

Rinku Gupta
Argonne National Laboratory
rgupta@anl.gov

Simon Hettrick
University of Southampton
s.hettrick@software.ac.uk

Daniel S. Katz
University of Illinois at Urbana-Champaign
d.katz@ieee.org

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich GmbH
r.speck@fz-juelich.de

MS246

Tools to Rapidly Develop Sophisticated HPC Software Libraries

Applications run on supercomputers having heterogeneous nodes require sophisticated HPC software libraries, e.g., Kokkos, VT, and these HPC software libraries provide intelligent capabilities involving, e.g., load balancing, resilience, support for performance profiling and support for debugging. These sophisticated HPC software libraries use sophisticated heuristics or AI to guide its decisions during application execution. For rapid development of such HPC software, tools in particular low-level mechanisms that sit just above compiler and runtime for parallelization - need to be enhanced. In this talk, we first explain novel mechanisms implemented within MPIch and clang/LLVM OpenMP, specifically (1) tunable locality-sensitive loop scheduling for clang/LLVM's OpenMP and (2) the use of asynchronous communication through support for MPI continuations implemented in OpenMPI. We provide theoretical analyses and show experimental results of these features to demonstrate the effectiveness of the features use within HPC software on current and future supercomputing platforms. Finally, we discuss how these features can be used as a tool within HPC software libraries enable rapid development sophisticated strategies within the libraries.

Vivek Kale
Brookhaven National Laboratory

vivek.lkale@gmail.com

MS246**Producing Software for Science with Class**

The Computer Languages and Systems Software (CLaSS) Group at Berkeley Lab researches and develops programming models, languages, libraries, and applications for parallel and quantum computing. The open-source software under development in CLaSS includes the GASNet-EX networking middleware, the UPC++ partitioned global address space (PGAS) template library, the Berkeley Quantum Synthesis Toolkit (BQSKit), and the MetaHipMer metagenome assembler. This talk will start with an overview of CLaSS software and the software sustainability practices commonly employed across the group. The talk will then dive more deeply into the our burgeoning contributions to the ecosystem supporting modern Fortran, including our test development for the LLVM Flang Fortran compiler. This presentation will demonstrate how agile software development techniques are helping to ensure robust front-end support for standard Fortran 2018 parallel programming features. The talk will also present several key insights that inspired our design and development of the CoArray Fortran Framework of Efficient Interfaces to Network Environments (Caffeine) parallel runtime library, emphasizing the design choices that help to ensure sustainability. Lastly, the talk will demonstrate the productivity benefits associated with the first Caffeine application in Motility Analysis of T-Cell Histories in Activation (Matcha).

Damian W. RousonLawrence Berkeley National Laboratory, U.S.
rouson@lbl.gov**MS246****Automatically Exploring GPU Program Design Spaces for Increased Productivity and Sustainability**

Graphics Processing Units (GPUs) have revolutionized the computing landscape in the past decade and are seen as one of enabling factors in recent breakthroughs in Artificial Intelligence. However, it is very difficult to unlock to full computational power of the GPU. This is because there are many degrees of freedom in GPU programming, and typically only a handful of specific combinations of optimizations and parameter choices result in near-optimal performance. To obtain such highly-efficient kernels it is required to search vast and discontinuous program design spaces, which is infeasible for programmers to do by hand. Moreover, this search process would have to be repeated for different hardware and for different input problems, leading to productivity and sustainability issues with GPU applications. This talk gives a brief introduction to Kernel Tuner, a tool that allows programmers to create tunable applications that can be automatically optimized for any combination of hardware and input problems. Using runtime kernel selection and compilation based on auto-tuning results, programmers can create sustainable applications that can achieve near optimal performance on a wide variety of hardware and inputs.

Ben van Werkhoven

Netherlands eScience Center

b.vanwerkhoven@esciencecenter.nl

MS247**Continuum Limit of the Eikonal Equation on Graphs**

In this work, we study a graph approximation of the time-dependent (local) Eikonal equation with Dirichlet-type boundary conditions, where the kernel in the non-local problem is properly scaled. Based on the theory of viscosity solutions, we prove existence and uniqueness of the viscosity solutions of both the local and non-local problems, as well as regularity properties of these solutions in time and space. We then derive error bounds between the solution to the graph problem and its continuum limit both in continuous-time (gradient flow) and Euler time discretization. In particular, we establish that if the kernel scale parameter decreases at an appropriate rate as n grows, then almost surely, the solution of the problem on graphs converges uniformly to the viscosity solution of the local problem as the time step vanishes and the number vertices n grows large.

Jalal FadiliCNRS-ENSICAEN-Université Caen
Jalal.Fadili@ensicaen.fr**MS247****Inverse Reinforcement Learning via Variational System Identification of Fokker-Planck Equation**

Inverse Reinforcement Learning (IRL) is an inverse problem of Reinforcement Learning (RL) that uncovers the reward function of a Markov decision process from observed trajectories of an agent. It is an appealing way to quantify the rationale behind the behavior of humans or any living agents. Like other inverse problems, IRL is ill-posed and becomes especially challenging when the transition model is unknown or inaccessible for sampling because the evolution of states and actions depends both on the transition model and the agent's policy, which is derived from the transition model as well together with the reward function. Current IRL approaches either assume the transition model is accessible or utilize data-driven methods to estimate the transition model beforehand without considering physics. We propose a novel physics-aware IRL algorithm that can simultaneously infer the reward function and transition probability function. We first build a connection between the value function in RL and the potential function in the Fokker-Planck equation and then employ a variational system identification method to infer the potential function. The reward function and transition probability function can then be evaluated instantaneously with the built connection. We demonstrate our new approach using a Gridworld benchmark problem and a biological problem of cancer cells.

Chengyang Huang, Siddhartha SrivastavaUniversity of Michigan
chengyah@umich.edu, sidsriva@umich.eduKrishna Garikipati
Mechanical Engineering
University of Michigan
krishna@umich.edu

Xun Huan

University of Michigan

xhuan@umich.edu

MS247

Monotone Discretizations of Levelset Convex Geometric PDEs

We present a new algorithm that converges to level set convex viscosity solutions of high-dimensional Hamilton-Jacobi equations. The algorithm can be used to solve a wide class of curvature motion PDEs, as well as a recent Hamilton-Jacobi equation for the Tukey depth, which is a statistical depth measure of data points. The algorithm is based on monotone schemes that involve partial derivatives in directions orthogonal to the gradient. We provide the convergence analysis of the algorithm on regular Cartesian grids and unstructured point clouds in any dimensions, and numerical experiments that demonstrate approximated solutions of affine flows in 2D and Tukey depth measure of high dimensional data, such as MNIST and FashionMNIST datasets.

Wonjun Lee

Institute for Mathematics and its Applications
University of Minnesota
lee01273@umn.edu

MS248

Implementation of Coupled Simulations in the ForestClaw Library

We present an implementation of the coupling of two distinct adaptive simulations, run on a separate hierarchy of adaptive Cartesian meshes within the ForestClaw library. Both simulations have their own set of options, and own meshes derived from p4est quadtree/octree meshes, and solver characteristics. Through coupling, however, both can be run in the same executable. This work goes beyond the typical coupling of hyperbolic solvers with elliptic and/or parabolic solvers in that each simulation may involve its own operator split approach and runs on its own mesh. The target codes are solvers MAGICForest (Model for Acoustic Gravity wave Interactions and Coupling, J. Snively Embry-Riddle Aeronautic Univ. (ERAU)) with the ionosphere code TreesGEMINI (Geospace Environment Model of Ion-Neutral Interactions, M. Zettergren, ERAU). MAGICForest currently runs within the ForestClaw environment, and progress is being made to extend ForestClaw with the GEMINI solvers. We will report on the progress towards full coupling of these two solvers, and underlying challenges that must be overcome to allow multiple simulations to run and communicate through a single executable.

Scott Aiton, Donna Calhoun

Boise State University
scottaiton@boisestate.edu,
donna.calhoun@boisestate.edu

donnacal-

MS248

A Numerical Method of Self-Similar Solutions to Ideal Magnetohydrodynamics by Using P4est

We present a numerical method to obtain self-similar solutions to the ideal magnetohydrodynamics (MHD) equations using p4est. In this effort, the initial-value problem (IVP) is converted into a boundary-value problem (BVP) by eliminating time and transforming the system into self-similar coordinates. The self-similar solution to the BVP is then solved using an iterative method embedded in an

adaptive mesh refinement (AMR) framework based on the p4est methodology. P4est is a parallel cell-based AMR mesh generation and management software with a 2:1 balancing on computational domains composed of multiple connected two-dimensional quadtrees or three-dimensional octrees. Existing Riemann solvers (e.g., Roe, HLLD, etc.) can be modified in a relatively straightforward manner and used on the present problem. P4est effectively refines the mesh around discontinuities and improves computational efficiency. Extensive numerical tests illustrate that the obtained self-similar solution to the BVP exhibits sharper discontinuities than the corresponding IVP solution. In problems where vortex sheets occur, the self-similar approach results in a well-posed system, when compared with the IVP which experiences convergence loss with mesh refinement.

Fang Chen

Kaust

fang.chen@kaust.edu.sa

Ravi Samtaney

KAUST

ravi.samtaney@kaust.edu.sa

MS248

Improved Parallelization of the Radiosity Method for Visualization

The radiosity method is an approach to simulate viewpoint-independent global illumination in contrast to viewpoint-dependent methods like ray tracing. Radiosity works by computing the pairwise light transfer between surface patches in a scene, which are determined by enforcing strict local radiation balance. In this talk, we present our approach to revive the radiosity method by using a distributed octree data structure, where the Morton space-filling curve and the octree's inherent recursivity lend themselves to handle spatial data in parallel. This research is an initial step towards parallel adaptive in-situ visualization that utilizes a widely known simulation data structure. The MPI parallelization of the radiosity method is a challenging task since the calculations to detect occluding surfaces are highly non-local, and this non-locality of relevant data requires communication schemes that limit the communication runtime overhead by asynchronous point-to-point MPI messages. Using this approach, we improved the strong scaling of the radiosity method. We present the current state of our research, an adaptive, parallel and natively supported radiosity solver operating on a distributed octree data structure and focus on the parallelization and our scalability results.

Tim Griesbach

Universität Bonn, Germany
Institute for Numerical Simulation
tim.griesbach@uni-bonn.de

Carsten Burstedde

Universität Bonn

burstedde@ins.uni-bonn.de

MS248

The Power of Modular Tree-Based AMR – Resolving Hanging Nodes and Cutting Holes

In this talk I will describe t8code, a new C/C++ library we have developed to manage parallel adaptive meshes (AMR) with various element types. t8code uses a collection (a for-

est) of multiple connected adaptive space-trees in parallel and scales to at least one million MPI ranks and over 1 Trillion mesh elements. t8code extends the well-known and efficient space-filling curve (SFC) technology from quads and hexes to all commonly used element shapes. I will include examples from several earth system modelling simulation codes that use t8code.

Johannes Holke

German Aerospace Center (DLR)
johannes.holke@dlr.de

Johannes Markert
German Aerospace Center
johannes.markert@dlr.de

MS248

Tensor-Product Space-Time Goal-Oriented Error Control and Adaptivity with Partition-of-Unity Dual-Weighted Residuals for Nonstationary Flow Problems

In this presentation, the dual-weighted residual method is applied to a space-time formulation of nonstationary Stokes and Navier-Stokes flow. Tensor-product space-time finite elements are being used to discretize the variational formulation with discontinuous Galerkin finite elements in time and inf-sup stable Taylor-Hood finite element pairs in space. To estimate the error in a quantity of interest and drive adaptive refinement in time and space, we demonstrate how the dual-weighted residual method for incompressible flow can be extended to a partition of unity based error localization. We derive the space-time Newton method for the Navier-Stokes equations and substantiate our methodology on 2D benchmark problems from computational fluid mechanics.

Julian Roth, Jan Philipp Thiele

Leibniz Universität Hannover
roth@ifam.uni-hannover.de, thiele@ifam.uni-hannover.de

Thomas Wick
Leibniz Universität Hannover
Université Paris-Saclay
wick@ifam.uni-hannover.de

Uwe Köcher
Helmut Schmidt University
mail@uwe.koecher.cc

MS249

Linear and Nonlinear Output Representations in Operator Networks

Supervised learning in function spaces is an emerging area of machine learning research with applications to the prediction of complex physical systems such as fluid flows, solid mechanics, and climate modeling. By directly learning maps (operators) between infinite dimensional function spaces, these models are able to learn discretization invariant representations of target functions. A common approach is to represent such target functions as linear combinations of basis elements learned from data. However, there are simple scenarios where, even though the target functions form a low dimensional submanifold, a very large number of basis elements is needed for an accurate linear representation. Here we present NOMAD, a novel operator learning framework with a nonlinear decoder map capable of learning finite dimensional representations of nonlinear

submanifolds in function spaces. We show this method is able to accurately learn low dimensional representations of solution manifolds to partial differential equations while outperforming linear models of larger size. Additionally, we compare to state-of-the-art operator learning methods on a complex fluid dynamics benchmark and achieve competitive performance with a significantly smaller model size and training cost.

Jacob Seidman, George Kissas, George J. Pappas, Paris Perdikaris

University of Pennsylvania
seidj@sas.upenn.edu, gkissas@seas.upenn.edu, pappasg@seas.upenn.edu, pgp@seas.upenn.edu

MS249

Learning Moment-Preserving Filters for Discontinuous Data

Filters are used to enhance accuracy, remove noise, and reduce spurious oscillations in a numerical approximation. This is especially useful when unphysical artifacts such as aliasing error or Gibbs phenomena arise, or when extracting higher accuracy in a numerical solution. As part of this, filters should increase the decay rate of the coefficients away from a discontinuity. Near a shock, filters reduce spurious oscillations. The motivation of this work is to minimize the size of the region with order one numerical error typically found near a discontinuity, which in turn contributes to enhancing the overall accuracy of a numerical approximation. We propose a novel way to learn filters for discontinuous data while preserving moments. The learned filter is a convolutional neural network, which learns from time-evolved, linearly advected data of different initial condition profiles that include C^0 and C^1 discontinuities. We enforce a consistency condition as a hard constraint and maintain polynomial reproduction through a penalty term in the training. In this talk, we will present construction and training aspects of our learned filter along with benchmark tests of the effectiveness of the filter.

Soraya Terrab, Samy Wu Fung, Jennifer K. Ryan
Colorado School of Mines
sterrab@mines.edu, swufung@mines.edu, jryan@kth.se

MS250

Quantum Algorithms for Scientific Computing: An Introduction

In this talk I will provide an overview of the opportunities and challenges in quantum algorithms for scientific computing, starting with a brief introduction to quantum computing. As such no background in quantum computing is assumed and the talk is appropriate for any computational scientist. By the end of the talk you will know the key results and obstacles in the design of both "NISQ" [Quantum 2, 79] and error-corrected quantum algorithms, and have the necessary background for upcoming talks in this minisymposium.

Ryan LaRose

EPFL
ryan.larose@epfl.ch

MS250

A Classical-Classical Type Approximation to Quantum States

Quantum computation and related disciplines are opening

a new area in computer science. Researchers are eager to understand certain behavior and exploit powerful functions to make quantum computation more powerful and applicable. Quantum correlation contributes remarkably to speeding up algorithms by its non-negligible capability in processing massive data. Surprisingly, not only do entanglements have quantum correlations but separable states. However, quantifying the quantum correlation has been proven as an NP-hard problem, which drives us to approximate it alternatively. We consider a more general situation that turns quantifying quantum correlation into a particular case. Through our talk, we will consider the Classical-Classical (CC) formation to approximate quantum states in a bipartite system. The approximation problem is cast into an optimization problem that constraints on Stiefel manifolds. A Riemannian gradient-driven descent flow is established to tackle the problem. We will finally show the objective function will decrease along the flow and successfully recover a synthetic CC state. Moreover, the consistency and convergence of the proposed method will be demonstrated via numerical results.

Bingze Lu, YuChen Shu
National Cheng Kung University
Dept. Mathematics
l18081028@gs.ncku.edu.tw, ycshu@mail.ncku.edu.tw

Matthew M. Lin
Department of Mathematics
National Cheng Kung University
mhlin@mail.ncku.edu.tw

MS250

Near-Term Quantum Algorithms for Optimization

One of the most prominent applications of quantum computers is to solving hard constraint satisfaction and optimization problems. In this talk, I will discuss recent work on the applicability of the near-term quantum algorithm QAOA (the Quantum Approximate Optimization Algorithm) to problems in this domain. First I will discuss theoretical and numerical results on the ability of QAOA to solve the fundamental boolean satisfiability problem, in the form of random k -SAT. In this setting, based on these results, QAOA may be able to outperform leading classical algorithms. Second, I will discuss whether QAOA could be applied to solve hard protein folding problems. In this case, the situation seems more challenging and QAOA may not significantly outperform existing methods. This talk is based on the papers arXiv:2208.06909 (joint work with Sami Boulebnane) and arXiv:2204.01821 (joint work with Sami Boulebnane, Xavier Lucas, Agnes Meyder, and Stanislaw Adaszewski).

Ashley Montanaro
University of Bristol
Phasecraft
ashley.montanaro@bristol.ac.uk

MS250

Constrained Combinatorial Optimization with Multi-Qubit Gates

Quantum algorithms for unconstrained optimization problems, such as the Quantum Approximate Optimization Algorithm (QAOA), have become increasingly popular. Recent work has shown that the QAOA can also be applied to constrained combinatorial optimization problems by incorporating the problem constraints within the design of

the quantum circuit ansatz. This work investigates the resource requirements of QAOA when targeting instances of the Maximum Independent Set problem. We consider three variants of QAOA each of which make different tradeoffs between the amount of quantum and classical resources required. Additionally, we consider the quantum cost of decomposing the QAOA circuits for hardware which may support different qubit technologies and native gate sets.

Teague Tomesh
princeton university
ttomesh@princeton.edu

Nicholas Allen
Princeton University
ndallen@princeton.edu

Zain Saleem
Argonne National Laboratory
zsaleem@anl.gov

MS250

A Theory of Quantum Differential Equation Solvers: Limitations and Fast-Forwarding

We study limitations and fast-forwarding of quantum algorithms for solving linear ordinary differential equation (ODE) systems with the focus on non-quantum dynamics, where the coefficient matrix in the ODE is not anti-Hermitian or the ODE is inhomogeneous. On the one hand, for homogeneous linear ODEs, by proving worst-case lower bounds, we show that generic quantum algorithms suffer from computational overheads due to two types of “non-quantumness”: real part gap and non-normality of the coefficient matrix. We then reach the conclusion that quantum algorithms are most efficient at solving quantum dynamics. To obtain these lower bounds, we propose a general framework for proving lower bounds of quantum algorithms that can be regarded as amplifiers (i.e., can amplify the difference between a pair of quantum states). On the other hand, improved asymptotic complexity scalings and fast-forwarding are possible for special cases of ODEs. We obtain quadratic to exponential improvements in terms of the evolution time and the spectral norm of the coefficient matrix in several cases, including inhomogeneous ODEs with negative definite coefficient matrix, inhomogeneous ODEs with restricted eigenvalues and eigenstates, and spatially discretized inhomogeneous heat equation and advection-diffusion equation. Our fast-forwarding algorithms are conceptually different from existing ones in the sense that neither time discretization nor solving high-dimensional linear systems is required.

Qi Zhao
University of Hong Kong
zhaqiqi@cs.hku.hk

Dong An, Jin-Peng Liu, Daochen Wang
University of Maryland
dongan@umd.edu, jliu1219@terpmail.umd.edu,
wdaochen@gmail.com

MS251

A Numerical Study of a Micro-Macro Model-Reduced 3-Level Parareal Method for Scale-Separated SDEs

Time-parallel methods can reduce the wall clock time required for the accurate numerical solution of a Stochas-

tic Differential Equation (SDE) by parallelizing across the time-dimension. In this talk, we study the convergence behavior of a multiscale, micro-macro version of a 3-level Parareal method. In our method, the fine propagator of the SDE is based on a high-dimensional slow-fast microscopic model; a hierarchy of coarser propagators are model-reduced versions of the latter, that capture the low-dimensional, model-reduced, effective dynamics at slower time scales. We investigate numerically how the model errors of the approximate models influence the convergence of the micro-macro 3-level Parareal algorithm.

Ignace Bossuyt, Giovanni Samaey
 KU Leuven
 ignace.bossuyt1@kuleuven.be,
 giovanni.samaey@kuleuven.be

Stefan Vandewalle
 Department of Computer Science
 KU Leuven - University of Leuven
 stefan.vandewalle@kuleuven.be

MS251

Multigrid-Reduction-in-Time for Large Eddy Simulation

Simulations of turbulent flow present challenges in terms of accuracy and affordability on modern highly-parallel computer architectures. A multigrid-reduction-in-time algorithm is used to provide a framework for separately evolving different scales of turbulence and for parallelizing the temporal domain, thereby increasing the concurrency. It is hypothesized that the space-time locality of the small scales of turbulence can be used to circumvent difficulties in applying temporal multigrid to flows dominated by inertial physics. For algorithms that fall well short of spectral accuracy (fourth-order is used in this work) attention must be paid to the accuracy of features on scales transferred between multigrid levels. Results from applying the approach to an infinite-Reynolds number Taylor-Green flow provide strong evidence that the approach has merit. The multigrid-reduction-in-time framework can be used to parallelize the temporal domain of a high-Reynolds-number turbulent flow and permit independent convergence of different scales.

Xinfeng Gao, Stephen Guzik
 Colorado State University
 xinfeng.gao@colostate.edu, stephen.guzik@colostate.edu

Robert D. Falgout
 Lawrence Livermore National Laboratory
 falgout2@llnl.gov

Jacob B. Schroder
 Department of Mathematics and Statistics
 University of New Mexico
 jbschroder@unm.edu

MS251

Multigrid Reduction-in-Time for Linear Advection Problems

In recent decades, many iterative parallel-in-time methods have been proposed for efficiently solving diffusion-dominated partial differential equations (PDEs). However, most of these approaches are not efficient when applied to hyperbolic PDEs, or advection-dominated PDEs more broadly. In this talk, we consider the specific parallel-in-

time algorithms of multigrid reduction-in-time (MGRIT) and Parareal. We discuss recent advances we have made in solving linear hyperbolic PDEs with these methods. This includes improved theoretical understanding of the difficulties that arise when applying these methods to hyperbolic PDEs, and new and improved coarse-grid operators that use semi-Lagrangian discretization techniques.

Oliver A. Krzysik
 University of Waterloo
 okrzysik@uwaterloo.ca

Hans De Sterck
 University of Waterloo, Canada
 hdsterck@uwaterloo.ca

Robert D. Falgout
 Lawrence Livermore National Laboratory
 falgout2@llnl.gov

Jacob B. Schroder
 Department of Mathematics and Statistics
 University of New Mexico
 jbschroder@unm.edu

MS251

Multigrid Reduction in Time for Chaotic Problems

Multigrid Reduction in Time (MGRIT) extends the well known two-level Parareal algorithm to a full multigrid method, and has demonstrated optimal scaling for parabolic problems. However, chaotic problems have proved more difficult, since chaotic initial value problems are inherently ill-conditioned. MGRIT relies on a hierarchy of successively coarser time-grids to iteratively correct the solution on the finest time-grid, but due to the nature of chaotic systems, subtle inaccuracies on the coarser levels can lead to poor coarse-grid corrections. We propose a modification to nonlinear FAS multigrid, as well as a novel time-coarsening scheme, which together better capture long term behavior on coarse grids and greatly improve convergence of MGRIT for chaotic initial value problems. We present our implementation of the algorithm in the XBraid software package, and provide supporting numerical results for the Lorenz system and demonstrate parallel speedup for the Kuramoto-Sivashinsky equation.

David Vargas
 University of New Mexico
 dvargas2@unm.edu

Robert D. Falgout, Stefanie Guenther
 Lawrence Livermore National Laboratory
 falgout2@llnl.gov, guenther5@llnl.gov

Jacob B. Schroder
 Department of Mathematics and Statistics
 University of New Mexico
 jbschroder@unm.edu

MS252

Efficient Proximal Subproblem Solvers for An Inexact Nonsmooth Trust-Region Method

Minimizing the sum of a smooth and nonsmooth function has been the core structure of many optimization applications in recent years. These include basis pursuit denoise problems in data science, optimal control of partial

differential equations, and, more generally, where inverse problems desire sparsity in the optimized parameters. In previous work, we developed novel trust-region methods that minimize the sum of a smooth, nonconvex and a nonsmooth yet convex function. A question left open by both such work is how to solve the subproblem, which is often a quadratic model of the smooth component added to either the nonsmooth regularizer or a model of the regularizer. This work elaborates on potential trust-region subproblem solvers for various assumptions on the nonsmooth component. We extend our inexactness criteria to the nonsmooth regularizer and the evaluation of proximal operators, which form the core of the trust region routine. For convex nonsmooth regularizers, we describe variations of a spectral projected gradient algorithm, a nonlinear conjugate gradient routine, and a dogleg method. Finally, we illustrate performance on numerical examples in PDE-constrained optimization.

Robert Baraldi
Sandia National Laboratories
Optimization & Uncertainty Quantification
rjbaral@sandia.gov

MS252

An Interior Proximal Gradient Method for Nonconvex Optimization

In this talk we consider composite minimization problems subject to smooth inequality constraints and present an algorithm that combines interior point (IP) and proximal gradient schemes. While traditional IP methods cannot cope with nonsmooth objective functions and proximal algorithms cannot handle complicated constraints, their combined usage is shown to successfully compensate the respective shortcomings. We provide a theoretical characterization of the algorithm and its asymptotic properties, deriving convergence results for fully nonconvex problems, thus bridging the gap with previous works that successfully addressed the convex case. Our interior proximal gradient algorithm benefits from warm starting, generates strictly feasible iterates with decreasing objective value, and returns after finitely many iterations a primal-dual pair approximately satisfying suitable optimality conditions.

Alberto De Marchi
Universität der Bundeswehr München
alberto.demarchi@unibw.de

Andreas Themelis
Kyushu University)
andreas.themelis@ees.kyushu-u.ac.jp

MS252

Proximal Quasi-Newton Methods for Nonsmooth Regularized Bound-Constrained Optimization

We develop an interior-point method for nonsmooth regularized bound-constrained optimization problems of the form $x \geq 0$. Steps are computed by minimizing the sum of models of the smooth f and nonsmooth h terms, with the smooth term containing the barrier information. This subproblem minimization is done via a first-order method, such as proximal gradient descent. Both f and h may be nonconvex. We establish global convergence to a first-order stationary point under the assumptions that f has Lipschitz-continuous gradients and h is proper and lower-semicontinuous. We comment on the complexity of our approach to bring a measure of stationarity below $\epsilon \in (0, 1)$.

Dominique Orban
GERAD and Dept. Mathematics and Industrial Engineering
Ecole Polytechnique de Montreal
dominique.orban@gerad.ca

MS252

Global Convergence of Operator Splitting Methods in Absence of Monotonicity

Splitting methods have seen widespread use in a variety of applications in both optimization and for solving variational inequalities. Despite this popularity, convergence results of these methods have been largely limited to the convex/monotone setting. In this work, we study two such methods in the absence of the monotonicity assumption and without requiring smoothness, namely Douglas-Rachford splitting (DRS) and the Chambolle-Pock (CP) algorithm. To this end, we introduce the concept of semimonotonicity and provide sufficient conditions for global convergence of DRS and CP involving the sum of two semimonotone operators. Most notably, it is shown that DRS and CP converge even when the sum of the involved operators (or of their inverses) is nonmonotone. Our analysis relies on establishing a connection between both methods and the preconditioned proximal point algorithm applied to a corresponding underlying operator satisfying an oblique weak Minty assumption.

Panagiotis Patrinos, Brecht Evens, Pieter Pas, Puya Latafat
KU Leuven
panos.patrinos@esat.kuleuven.be,
brecht.evens@kuleuven.be, pieter.pas@kuleuven.be,
puya.latafat@kuleuven.be

MS253

Frequency Sweep for the Helmholtz Equation: Preconditioning and Efficient Solves

Boundary integral equation discretizations for the scalar Helmholtz equation lead to large dense linear systems. Efficient boundary element method (BEM) solvers, such as the fast multipole method (FMM) and H-matrix based methods, focus on structured low-rank approximations of subblocks in these systems. The ranks of these subblocks increase linearly with the wavenumber. Recently, in [S. Dirckx, D. Huybrechs & K. Meerbergen, "Frequency extraction for BEM-matrices arising from the 3D scalar Helmholtz equation", (to appear in SISC)], we explored a data compact representation of the wavenumber dependence of these BEM matrices, based on a technique called 'frequency extraction'. This allows for the efficient, quadrature-free construction of all relevant BEM-matrices in a pre-selected wavenumber range. Now we explore preconditioning techniques and efficient frequency sweeps. The preconditioning is two-fold. Firstly, we can generate so-called 'operator preconditioners' quadrature-free from our data-compact representation, resulting in mesh-independent condition numbers of the BEM-systems. Secondly, Krylov subspace recycling is used to accelerate convergence of iterative solvers (e.g. GMRES), by collecting strategically chosen deflation subspaces. Our data-compact representation is augmented by these deflation spaces, and can then be used to accelerate solves at ar-

bitrary wavenumbers in the selected range.

Simon Dirckx
KU Leuven
simon.dirckx@kuleuven.be

Daan Huybrechs
Department of Computer Science
KU Leuven
daan.huybrechs@cs.kuleuven.be

Karl Meerbergen
KU Leuven
karl.meerbergen@kuleuven.be

MS253

An OSRC Preconditioner for the EFIE

The Electric Field Integral Equation (EFIE) is a well-established tool to solve scattering problems, but the development of efficient and easy to implement preconditioners for this equation remains an active research area. In recent years, operator preconditioning approaches for the EFIE, where the electric field operator is regularised by multiplication with another convenient operator, have become popular. A particular operator that is likely to be considered as a preconditioner is the exact Magnetic-to-Electric (MtE) operator. However, evaluating this operator is as expensive as solving the original EFIE. In [El Bouajaji, M., Antoine, X., & Geuzaine, C. *Approximate local magnetic-to-electric surface operators for time-harmonic Maxwell's equations*. Journal of Computational Physics], approximate local Magnetic-to-Electric surface operators for the time-harmonic Maxwell equation were proposed that can be efficiently evaluated through the solution of sparse surface problems. In [Fierro-Piccardo, I., & Betcke, T. *An OSRC Preconditioner for the EFIE*] and this presentation we show the preconditioning properties of the approximate MtE operator for the EFIE. We describe the implementation and present a number of numerical comparisons against other preconditioning techniques for the EFIE to demonstrate the effectiveness of this new technique.

Ignacia Fierro Piccardo
University College London
maria.piccardo.19@ucl.ac.uk

MS253

Accurate Numerical Quadrature for 2D Stokes Flow Quantities in a Boundary Integral Equation Framework

Highly accurate quadrature schemes for computing the velocity, pressure, stress, vorticity, pressure gradient, velocity gradient and their averages, of a 2D Stokes flow are presented. There are many benefits of using a boundary integral method such as this one to numerically solve Stokes equations. An advantage is that the dimensionality is reduced by one as the partial differential equation is rewritten into a boundary integral equation that must only be solved over the boundary of the domain in order to determine a so-called layer density. As a post-processing step, different quantities of the flow can be found via integrals of the layer density, which become difficult to accurately compute when evaluated at or close to the boundary of the domain due to the Greens function in the integrand. Periodic problems are also considered, where so-called Ewald decompositions of the aforementioned quantities are derived to allow for

rapid evaluation of the arising sums using FFTs.

David Krantz, Lukas Bystricky
KTH Royal Institute of Technology
davkra@kth.se, lukasby@kth.se

Anna-Karin Tornberg
KTH
akto@kth.se

MS253

Computing Singular and Near-Singular Integrals over Curved Boundary Elements

In this talk, we present algorithms for computing singular and near-singular integrals arising when solving the 3D Helmholtz equation with curved boundary elements. These are based on the computation of the preimage of the singularity in the reference element's space using Newton's method, singularity subtraction, the continuation approach, and transplanted Gauss quadrature. We demonstrate the accuracy of our method for quadratic basis functions and quadratic triangles with several numerical experiments, including the scattering by two half-spheres.

Hadrien Montanelli
Institut Polytechnique de Paris
hadrien.montanelli@gmail.com

MS253

A High-Order Close Evaluation Scheme of Stokes Layer Potentials in 3D

The close evaluation problem arises when using boundary integral methods to simulate dense suspensions of particles in Stokes flow. In this talk, we present a high-order evaluation scheme of Stokes layer potentials to handle arbitrary geometries in three dimensions, which is crucial to model the hydrodynamic interactions and detailed flow structure. The scheme relies on expressing Stokes layer potentials in terms of Laplace potentials and a recently introduced differential geometry approach to convert nearly singular surface integral into several line integrals. We then demonstrate the application of this method through several numerical examples.

Hai Zhu
Flatiron Institute
hzhu@flatironinstitute.org

Shidong Jiang
Flatiron Institute, Simons Foundation
Center for Computational Mathematics
sjiang@flatironinstitute.org

Manas N. Rachh
Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

Shravan Veerapaneni
Department of Mathematics
University of Michigan
shravan@umich.edu

MS254

Continuous Testing and Benchmarking in a HPC

Sparse Linear Algebra Software

With a strong reliance on research software projects in both industry and for scientific simulations, research software sustainability is increasingly becoming a major point of contention. A necessary but nonsufficient aspect of software sustainability is Continuous Integration and Benchmarking (CI/CB/Cx). In addition, a mathematical and HPC software's testing strategies can be complex due to the different hardware behavior, and the need to ensure numerical accuracy. In this talk, we will showcase the Cx practices of the Ginkgo sparse linear algebra framework, which was designed with unit testing, good software design techniques as well as Cx practices as centerpieces. Because of HPC software constraints, access to HPC clusters for both CI and CB is essential to ensure correctness. We will showcase our testing and design strategies as well as our configurable solution for running Cx on HPC clusters.

Terry Cojean
Karlsruhe Institute of Technology
terry.cojean@kit.edu

Hartwig Anzt
University of Tennessee, U.S.
hanzt@icl.utk.edu

MS254

Ensuring Reliability in Next-Generation Linear Algebra Libraries

The SLATE project is a GPU-accelerated, distributed memory, dense linear algebra, comparable in scope to ScaLAPACK. To ensure its reliability over a diverse set of platforms, including CUDA, ROCm, and oneMKL, we built an extensive test suite with continuous testing to execute a standard set of acceptance tests before merging in code. Tests and experiments with SLATE require exploring values for a wide variety of parameters – matrix dimensions, block sizes, transposition operations, storage format, CPU or GPU computation, MPI grid dimensions, algorithm choice, test matrices, etc. For this purpose, we developed the TestSweeper library to perform parameter sweeps, which provides significant flexibility and ease of use compared to LAPACK's robust test suite. For instance, TestSweeper parameters are specified on the command line, as individual values or ranges. New parameters are trivial to add with several lines of code. Output is in tabular form to be both human readable or parsed by scripts, such as by Python's pandas for analysis and plotting. SLATE includes an extensive test matrix generation code, with a flexible string description of the matrix types. We also touch on testing SLATE, LAPACK, and BLAS for consistent propagation of Inf and NaN exceptional values, which the existing reference BLAS and many vendor BLAS do not handle consistently.

Mark Gates
Computer Science Department
University of Tennessee
mgates3@icl.utk.edu

Asim YarKhan
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
yarkhan@eecs.utk.edu

Dalal Sukkari
KAUST

sukkari@icl.utk.edu

Kadir Akbudak
King Abdullah University of Science and Technology
kadir.akbudak@kaust.edu.sa

Sebastien Cayrols
University of Tennessee, Knoxville
sebastien.cayrols@icl.utk.edu

Daniel R. Bielich
University of Colorado at Denver
drbieli@icl.utk.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS254

Testing and Debugging for Numerical Exceptions in GPU Scientific Applications

Testing scientific applications for numerical correctness and reproducibility in GPU systems is a challenging task. Detecting and mitigating floating-point exceptions, such as division by zero or not a number (NaN) operations, is crucial in any testing campaign since such operations can lead to unreliable numerical answers or reproducibility issues. Commonly used GPUs, such as NVIDIA GPUs, provide no support to detect floating-point exceptions, which makes it tricky for developers to mitigate such errors. In this talk, we present techniques and tools to mitigate different classes of numerical exceptions and anomalies, with a particular focus on GPUs. We present FPChecker and BinFPE, two tools based on LLVM instrumentation and binary instrumentation to detect the location of code affected by floating-point exceptions in the CPU and GPU; we also present Xscope, a Bayesian optimization approach to find inputs that trigger such exceptions in the GPU.

Ignacio Laguna
Lawrence Livermore National Laboratory
ilaguna@llnl.gov

MS254

Constrained-Based Testing for Floating-Point Code: Challenges and Opportunities

Computational science code makes intense use of floating-point numbers. Unfortunately, powerful constraint-based analysis techniques rarely support floating point, or they scale poorly to this domain. In this talk, I will discuss our experience developing techniques that can reason about floating-point constraints. These include methods that make use of the theory of floating point offered by some SMT solvers, solving floating-point constraints via fuzzing, and approximating floating point via fixed point. Based on joint work with Daniel Liew, Daniel Schemmel, Alastair Donaldson, Rafael Zhl, Klaus Wehrle, Thom Hughes, Martin Nowack, J. Ryan Stinnet.

Cristian Cadar, Daniel Schemmel
Imperial College, London

c.cadar@imperial.ac.uk, d.schemmel@imperial.ac.uk

MS255

Deploying a Tensor Network Simulator on Modern HPC Architectures

We present QXTools, a framework for simulating quantum circuits using tensor network methods. QXTools is written in Julia and is designed to run on large distributed compute clusters and to support GPU accelerators. We discuss the deployment of QXTools on different HPC systems, available on the LRZ Bavarian Energy, Architecture, and Software Testbed (BEAST) and including ARM ThunderX2 (THX2), Fujitsu A64FX, AMD Rome and Intel Xeon Skylake. The application can run within containers on such systems, without compromising security, stability, and performance, thus ensuring a simple workflow for HPC users.

Xiaolong Deng
Leibniz Rechenzentrum
LRZ
xiaolong.deng@lrz.de

Luigi Iapichino
Leibniz Supercomputing Centre (LRZ), Garching
luigi.iapichino@lrz.de

MS255

Tensor Networks From Quantum Mechanics to Artificial Intelligence

Tensor Networks (TNs) are a numerical tool initially designed to simulate quantum many-body systems on a classical computer. Since their first success story over three decades ago, TNs have been applied to investigate various quantum systems, such as Lattice Gauge Theories modelling the fundamental particles of our universe or condensed matter systems underlying the upcoming technology of Quantum Computers. However, in recent years it turned out that TNs can be seen as a more general tool for representing information with applications in domains, such as applied mathematics, Informatics or medical physics. In fact, we at *Tensor Solutions* are a start-up developing AI-Technology based on this quantum-inspired method. In this talk, I will present the concept of Tensor Networks and their journey from addressing problems in Quantum physics up to solving Machine Learning problems in Artificial Intelligence.

Timo Felser
Tensor Solutions, Germany
timo.felser@tensor-solutions.com

Marco Trenti
Tensor AI Solutions GmbH
marco.trenti@tensor-solutions.com

Simone Montangero
Università di Padova
Dipartimento di Fisica e Astronomia G.Galilei
simone.montangero@unipd.it

MS255

Simulation of Quantum Circuits with Tensor Networks

In this talk, I will introduce two tensor-network methods

targeting the sampling problem of the Sycamore circuits. The first method computes amplitudes and probabilities for a large number of correlated bitstrings. The obtained results verify the Porter-Thomas distribution of the large and deep quantum circuits of Google and can be used for spoofing the linear cross entropy benchmark of quantum supremacy using a post-sampling approach. The second method can be used to generate one million uncorrelated bitstrings which are sampled from a final distribution of the Sycamore circuit with 53 qubits and 20 cycles. With the approximate state having fidelity of roughly 0.0037. The whole computation has cost about 15 hours on a computational cluster with 512 GPUs. If our algorithm could be implemented with high efficiency on a modern super-computer with ExaFLOPS performance, we estimate that ideally, the simulation would cost a few dozens of seconds, which is faster than Google's quantum hardware.

Pan Zhang
Chinese Academy of Sciences
panzhang@itp.ac.cn

MS256

A Discontinuous Galerkin-Finite-Difference Hybrid Method for Computational Astrophysics

Conservative finite difference and finite volume methods have proven extremely robust and reliable for magnetohydrodynamics simulations of binary neutron star mergers, core collapse supernova, and accretion onto a black hole. However, finite difference methods are generally less accurate and efficient than spectral methods when the solution is smooth, e.g. away from the stellar surfaces and shocks. The attractiveness of spectral methods has been demonstrated by thousands of long and highly accurate binary black hole simulations. Discontinuous Galerkin methods seek to provide the accuracy of spectral methods while also robustly capturing shocks in hydrodynamics simulations. I will give an overview of a discontinuous Galerkin-finite-difference hybrid method that inherits the best properties of both spectral and finite difference methods. I will show longterm simulations of magnetized and rotating neutron stars, as well as test problems to demonstrate the reliability and accuracy of the hybrid method.

Nils Deppe, Yoonsoo Kim
California Institute of Technology
ndeppe@caltech.edu, ykim7@caltech.edu

MS256

A Two-Moment Neutrino Transport Method for Nonrelativistic Fluid Velocities

When modeling neutrino transport through a moving fluid, the use of momentum space coordinates associated with a frame of reference comoving with the fluid greatly simplifies the inclusion of neutrino-matter interactions. However, this choice of momentum space coordinates complicates the discretization of the phase-space advection operator due to the appearance of velocity-dependent terms. In this context, we consider a multidimensional, spectral two-moment model — closely related to that promoted by Lowrie, Mihalas, and Morel (2001; JQSRT, 69, 291-304) — which includes special relativistic effects to order v/c , where v is the fluid velocity and c is the speed of light. This model maintains wave speeds bounded by c and is consistent with conservation of energy and momentum in the laboratory frame. Here we present a numerical method for evolving this model. The method uses the discontin-

uous Galerkin method for phase-space discretization and implicit-explicit time-stepping. The method is carefully constructed to maintain physical radiation moments during evolution. We also consider convergence of the implicit neutrino-matter coupling solver and the simultaneous conservation of particle number and energy. Numerical results that demonstrate the properties of the proposed method will be presented.

Eirik Endeve, Paul Laiu, Austin Harris
Oak Ridge National Laboratory
endevee@ornl.gov, laiump@ornl.gov, harrisja@ornl.gov

Zachary Elledge
University of Tennessee Knoxville
zelledge@vols.utk.edu

Anthony Mezzacappa
University of Tennessee
mezz@utk.edu

MS256

Neutrino Radiation Hydrodynamics with Thornado and Weaklib

Neutrino-matter coupling via weak interactions is one of the most important physical mechanism in the evolution of core-collapse supernovae (CCSN). In this talk, we will describe recent changes to the toolkit for high-order neutrino-radiation hydrodynamics (thornado), a code to simulate neutrino radiation hydrodynamics in a spectral two-moment model using the discontinuous Galerkin method, as well as WeakLib, a library that provides the weak interaction opacity tables used in thornado. We will describe the update of WeakLib from the canonical neutrino opacity set of Bruenn (1985) to now include nucleon-nucleon Bremsstrahlung and electron capture on heavy nuclei using a table of electron capture rates and the corresponding spectra of emitted electron neutrinos, as well as the necessary changes to thornado to use these new opacities. We will present results from detailed comparisons of CCSN simulations using the canonical Bruenn 1985 opacity set and the updated opacity tables in WeakLib.

Vassilios Mewes
Oak Ridge National Laboratory
mewesv@ornl.gov

Ran Chu
Department of Physics and Astronomy,
University of Tennessee Knoxville
rchu@vols.utk.edu

Eirik Endeve, Austin Harris, William Raphael Hix, Paul Laiu, Eric Lentz
Oak Ridge National Laboratory
endevee@ornl.gov, harrisja@ornl.gov, raph@ornl.gov, laiump@ornl.gov, lentzej@ornl.gov

Bronson Messer
Oak Ridge National Laboratory
& University of Tennessee
bronson@ornl.gov

Anthony Mezzacappa
University of Tennessee

mezz@utk.edu

MS256

Multi-Dimensional MHD Core Collapse from the Onset of the Explosion to Nucleosynthesis

Core-collapse supernova form a diverse class of explosions produced at the end of the lives of massive stars. Studying these events is heavily dependent on supercomputing due to the complexity of the turbulent flows and the neutrino transport in dense matter. We rely on a high resolution code in a hybrid parallelisation model (OpenMP/MPI) to solve the coupled equations of magnetohydrodynamics and two-moment neutrino transport. Our three-dimensional models including all relevant physics and covering long simulation times demonstrate how the subset of progenitors with strong magnetic fields and high rotational energies can be responsible for producing particularly high explosion energies and for playing an important role in the chemical enrichment of early galaxies with heavy elements in ways that are not accessible to ordinary, primarily neutrino-driven explosions.

Martin Obergaulinger, Miguel Angel Aloy
University of Valencia
martin.obergaulinger@uv.es, miguel.a.aloy@uv.es

MS257

Studies of Transition to Turbulence in Oscillatory Cerebrospinal Fluid Flow

Presence of turbulence like flow fluctuations in physiologic flows has been known since a long time. In this work we study oscillatory flow in stenosed geometries to quantify and characterize the onset of turbulence in zero mean oscillatory flows. Simulations on a stenosed pipe in axisymmetric and eccentric configurations were conducted using the LBM solver Musubi with various Reynolds numbers and pulsation frequencies. Stenosis with area reduction of 75%, 60%, 50% and 25% were studied in both axisymmetric and eccentric configurations. Meshes of up to 2.8 billion cells were created and simulations were conducted on 300000 CPU cores of the SuperMUC-NG petascale system in Munich, GERMANY. Main findings include: 1. The flow transitions only in higher degrees of stenosis namely 50%, 60% and 75%, where Re 1800 is the approximate threshold for transition. 2. The flow reversal stabilizes the flow field. 3. A higher pulsation frequency leads to earlier breakdown of flow a phenomenon that is seen mostly for lower stenoses degrees. 4. The eccentricity of the stenosis is one of the major factors for flow transition. The results advocate that transition to turbulence is a possibility in physiologic flows, and the geometry of the conduit is the most prominent factor that results in the onset of turbulence.

Kartik Jain
University of Siegen, Germany
k.jain@utwente.nl

MS257

Refactoring Legacy Fortran Applications to Leverage Modern Heterogeneous Architectures in Extreme-Scale CFD

Recent trends and advancement in including more diverse and heterogeneous hardware in High-Performance Computing is challenging software developers in their pursuit

for good performance and numerical stability. The well-known maxim "software outlives hardware" may no longer necessarily hold true, and developers are today forced to re-factor their codebases to leverage these powerful new heterogeneous systems. In this talk, we present Neko - a portable framework for high-order spectral element flow simulations. Unlike prior works, Neko adopts a modern object-oriented Fortran 2008 approach, allowing multi-tier abstractions of the solver stack and facilitating various hardware backends ranging from general-purpose processors, accelerators down to exotic vector processors and Field-Programmable Gate Arrays (FPGAs) via Neko's device abstraction layer. Focusing on Neko's performance, scalability, and accuracy, we present performance measurements on a wide range of accelerated computing platforms, including the EuroHPC pre-exascale system LUMI, for performing large-scale direct numerical simulation (DNS) of turbulent fluid flow.

Niclas Jansson, Martin Karp
KTH Royal Institute of Technology
njansson@kth.se, makarp@kth.se

Stefano Markidis
KTH, Sweden
markidis@kth.se

Philipp Schlatter
KTH Royal Institute of Technology
pschlatt@mech.kth.se

MS257

Numerical Precision in High-Fidelity Computational Fluid Dynamics

The impact of numerical precision, both with regard to accuracy and potential performance improvements, has garnered increased interest in several different domains of computational science. However, it is an open question what numerical precision is required for a specific simulation in computational fluid dynamics (CFD). While we have conventionally relied on double-precision arithmetics, several factors affect what numerical precision is required to obtain a resolved flow field. In this presentation, we talk about our work to assess how the numerical precision requirements depend on the fluid scales and simulation size, both in theory and practice. We investigate how numerical precision affects a series of different flow cases and try to isolate the impact of numerical precision and connect it to the dynamics and properties of the considered flow cases. To do this we perform experiments with different discretization schemes and perform a series of computational experiments. We focus on direct numerical simulations and shed some light on what machine epsilon is required in the best case to perform a resolved CFD simulation. Finally, we address whether even higher numerical precision will be required in order to perform direct numerical simulations on exascale and post-exascale parallel systems.

Martin Karp, Niclas Jansson, Philipp Schlatter
KTH Royal Institute of Technology
makarp@kth.se, njansson@kth.se, pschlatt@mech.kth.se

Stefano Markidis
KTH, Sweden

markidis@kth.se

MS257

Walberla: a Multi-Physics Open-Source Software Framework for Scalable and Efficient CFD

waLberla is a modern open-source massively parallel multiphysics simulation framework with a focus on CFD applications. It uses the lattice Boltzmann method (LBM), which is an alternative to classical Navier-Stokes solvers for computational fluid dynamics simulations. waLberla scales on some of the top clusters in the world due to carefully designed distributed data structures. As a showcase, we use particle-resolved sediment transport simulations using LBM and the discrete element method (DEM). In this scenario, the LBM simulation dominates the computational cost. In our implementation, the LBM simulation runs on the GPU and the DEM simulation on the CPU. We analyze the performance on a heterogeneous compute cluster.

Samuel Kemmler
FAU Erlangen
samuel.kemmler@fau.de

MS257

Improving Parallel Scalability of High-Order Preconditioners on GPUs

Matrix-free methods for the solution of linear systems arising from high-order finite element discretizations of partial differential equations, such as the Poisson equation, require robust preconditioners. Two classes of preconditioners prove most effective: multigrid and low-order finite element methods. For multigrid preconditioners built on geometric p -multigrid on GPU architectures, the communication dominated coarse grid solve hinders the parallel scalability of the preconditioner. We propose a strategy to mitigate the increase coarse grid solve cost at scale by treating the coarse grid solve as an additive, rather than multiplicative, correction. This allows overlapping the coarse grid solve, which utilizes the CPU, and the remainder of the multigrid cycle, which utilizes the GPU. A hybrid p -multigrid and low-order finite element preconditioner is also proposed. We demonstrate the effectiveness of these two novel approaches on a variety of problems arising from the spectral element discretization of the Navier-Stokes equations, spanning up to $P = 6144$ NVIDIA V100 GPUs on Summit.

Malachi Phillips
University of Illinois, Urbana-Champaign
malachi2@illinois.edu

Stefan Kerkemeier
Argonne National Laboratory
stgeke@gmail.com

Fischer Paul
University of Illinois, Urbana-Champaign
fischerp@illinois.edu

MS258

Modeling of Chemical Reactions in Rarefied Gas Flows by the Fokker-Planck Method

Flows encountered in space applications, like atmospheric reentry or thruster plume expansion and interactions, can

be described by the well-known Boltzmann equation. A common approach to numerically solve it is the Direct Simulation Monte-Carlo (DSMC). The method is very efficient for high Knudsen numbers but becomes increasingly computationally intensive when approaching the continuum limit. In this regime the Boltzmann equation can be numerically solved using the recently proposed kinetic Fokker-Planck method which, like DSMC, relies on simulated particles to transport mass, momentum and energy through the flow domain, but does not resolve individual collisions. While this approach has been extended to mixtures and diatomic gases, the modeling of chemical reactions remains a research topic. Many problems arise from the fact that properties of the reactants need to be considered to evaluate a chemical reaction. This contradicts the Fokker-Planck approach that achieves its linear scaling in terms of particle numbers by forgoing the creation of interaction pairs. In this paper we propose a novel approach to model chemical reactions in Fokker-Planck which aims to conserve mass, energy and momentum while retaining the performance advantages in areas of high density.

Leo Basov

German Aerospace Center
leo.basov@dlr.de

MS258

A Low-Variance Particle Collision Scheme for Variable Weight DSMC Simulations

Numerical simulation of rarefied plasmas often makes use of Particle-in-Cell-Direct Simulation Monte Carlo (PIC-DSMC) codes [Birdsall, C. K., Particle-in-cell charged-particle simulations, plus Monte Carlo collisions with neutral atoms, PIC-MCC. *Trans. Plasma Sci.*, 19(2), 1991], which simulate the motion of a large number of computational particles, each of which represents a large number of actual molecules, atoms, or electrons. Particle-particle and particle-surface collisions in such codes are performed in a stochastic manner: a single process is simulated for each collision, with the process being randomly chosen based on the probability of it taking place. In the present work a new collision and boundary condition simulation scheme, dubbed "event splitting" [Oblapenko, G. et al, Hedging direct simulation Monte Carlo bets via event splitting, *J. Comput. Phys.*, 466, 2022], is investigated. It simulates multiple events simultaneously via splitting of the particles involved in the collision and has been shown to significantly reduce the level of stochastic noise in unsteady breakdown simulations. The scheme is applied to 1-D modelling of discharges in rarefied argon and helium plasmas, and compared to standard PIC-DSMC collision routines in terms of computational efficiency and numerical noise. Georgii Oblapenko acknowledges the funding provided by the Alexander von Humboldt Foundation for his stay as a guest researcher at the German Aerospace Agency (DLR).

Georgii Oblapenko

German Aerospace Center
georgii.oblapenko@dlr.de

MS258

Explicit Modal Discontinuous Galerkin Method for the Vlasov-Poisson System

In plasma physics, the Vlasov-Poisson system characterizes the behavior of a collisionless plasma subject to electrostatic effects in terms of the corresponding distribu-

tion function in phase space. Due to high dimensionality (3D space + 3D velocity + 1D time) and their intrinsic physical properties (conservation and positivity), the construction of an efficient numerical method represents a challenge and requires a careful balance between accuracy and computational complexity. In this research, the development of a numerical method for solving the multi-dimensional Vlasov-Poisson system is demonstrated. A modal discontinuous Galerkin (DG) scheme based on rectangular/hexahedral meshes is employed for the numerical solution of the multi-dimensional Vlasov-Poisson system. For the spatial discretization, the polynomial solutions are represented by scaled Legendre basis functions, and the numerical flux based on an upwind scheme is used for the solution stability. All the integrals appearing in the DG formulation are approximated with the Gauss-Legendre quadrature rule. A third-order explicit SSP-RK-based temporal discretization scheme is used for the resulting semi-discrete ordinary differential equation. We study the convergence analysis of the proposed scheme. The performance of the proposed scheme is assessed by solving some benchmarks in the simulations of plasma physics, including two-stream instability, weak and strong Landau damping, etc.

Satyvir Singh, Manuel Torrilhon

RWTH Aachen

singh@acom.rwth-aachen.de, mt@acom.rwth-aachen.de

MS258

The Numerical Flow Iteration for the Vlasov-Poisson Equation

The Vlasov-Poisson equation is a high-dimensional partial differential equation used to model the behaviour of plasmas in the collision-less limit. The six-dimensional phase-space and the turbulent motion of plasmas together with formation of strong filamentation in the solution lead to significant complications in the implementation of solvers. While purely Lagrangian schemes are in theory capable to capture the dynamics correctly, the inherent numerical noise makes it hard to run long-time accurate simulations and requires large numbers of particles as well as remeshing. Eulerian and Semi-Lagrangian schemes produce more accurate results, however, generation and management of meshes are prohibitively expensive for simulations in more than four dimensions. In addition, when constructing high-order schemes one often loses conservation properties of the exact solution. We present a novel approach, the numerical flow iteration (NuFI), which iteratively reconstructs the phase flow and then using the method of characteristics can evaluate the distribution function to arbitrary precision, i.e., up to error in time and Poisson discretisation. This approach naturally preserves non-negativity, does not struggle with over-shoots and conserves all L^p norms exactly. From the implementation perspective, a high flop-byte ratio as well as a structure which allows for easy parallelization of this method, lead to efficient portability to modern clusters.

Rostislav-Paul Wilhelm

RWTH Aachen

wilhelm@acom.rwth-aachen.de

Matthias Kirchhart

MathCCES

RWTH Aachen University

kirchhart@acom.rwth-aachen.de

MS258

A Quantum-Inspired Method for Solving the Vlasov Equation

Recently, it has been proposed that matrix product states (MPS), a quantum-inspired but classical algorithm, can be used to approximately solve partial differential equations with exponential speed-up, provided that the PDE is compressible or low-rank so that it can be efficiently represented in the MPS framework. Here, we investigate the utility of MPS methods for solving the Vlasov equation, a 6-dimensional nonlinear set of PDEs that provide an ab-initio description of collisionless plasmas. We report results of select test cases and demonstrate that we are able to capture important features of the dynamics even with significant amounts of compression.

Erika YeMassachusetts Institute of Technology
erikaye@mit.edu

MS259

A Solution Technique for Darcy Flow in Fractured Porous Media that Ensures Local Mass Conservation

Constructing fast solution schemes often involves deciding which errors are acceptable and which approximations can be made for the sake of computational efficiency. Herein, we consider a mixed formulation of Darcy flow and take the perspective that the physical law of mass conservation is significantly more important than the constitutive relationship, i.e. Darcy's law. Within this point of view, we propose a three-step solution technique that guarantees local mass conservation. In the first step, an initial flux field is obtained by using a locally conservative method such as the TPFA Finite Volume Method. Although this scheme is computationally efficient, it lacks consistency and therefore requires a suitable correction. Since this correction is divergence-free, the Helmholtz decomposition ensures that it is given by the curl of a potential field. The second step therefore employs an $H(\text{curl})$ -conforming discretization to compute the correction potential and update the flux field. The pressure field is computed in the final step by using the same TPFA system from the first step. The procedure guarantees local mass conservation regardless of the quality of the computed correction. Thus, we relax this computation using tools from reduced order modeling. We introduce a reduced basis method that is capable of rapidly producing a potential field for given permeability fields. By applying the curl to this field, we ensure that the correction is divergence-free and mass conservation is not impacted. Finally, we extend the method to solving Darcy flow in fractured porous media. We rewrite the equations in terms of mixed-dimensional differential operators and identify the problem as a mixed-dimensional Darcy flow system. In turn, the proposed three-step solution procedure directly applies using the mixed-dimensional curl to ensure local mass conservation.

Wietse M. Boon, Alessio Fumagalli, Alessio Fumagalli
Politecnico di Milano

wietsemarijn.boon@polimi.it, alessio.fumagalli@polimi.it,

alessio.fumagalli@polimi.it

MS259

Multilevel Methods for Nearly-Singular Problems in Mixed Dimensions

We consider parameter-robust algebraic multigrid methods for a class of operators that appear in preconditioning of monolithic multiphysics systems on domains of mixed dimensions. The choice of the coarse spaces and smoothers is modified to take into account the near-kernel of the system operator imposed by the coupling terms. The theoretical results show convergence of the two-level method independently of the discretization and model parameters. We also demonstrate the efficacy and robustness of the method in practice on numerical examples of mixed-dimensional problems arising in brain biomechanics.

Ana BudisaSimula Research Laboratory
ana@simula.no

Xiaozhe Hu

Tufts University
xiaozhe.hu@tufts.edu

Miroslav Kuchta

Simula Research Laboratory
Oslo, Norway
miroslav@simula.no

Kent-Andre Mardal

Department of Mathematics, University of Oslo
Simula Research Laboratory
kent-and@simula.no

Ludmil Zikatanov

Pennsylvania State University
ludmil@psu.edu

MS259

Unfitted Space-Time Finite Element Methods for Coupled Surface-Bulk Problems

Coupled surface-bulk problems represent an important range of application for numerical simulation techniques in computational sciences, as they model a variety of processes in fluid dynamics and biology. In this talk, we focus in particular on a model of soluble surfactants involving convection and diffusion. To solve this model problem, we consider unfitted space-time finite element simulations, which come with the benefit of a flexible handling of complex geometries. The overall coupled problem can be decomposed into several components, which can be studied individually first and finally taken together, such as the mere surface and the mere bulk problem. In relation to the important reference work "A cut finite element method for coupled bulk-surface problems on time-dependent domains", P. Hansbo, M.G. Larson, S. Zahedi (CMAME 307, 96-116) we present some variant methods for several of these steps. Specifically, we investigate a stabilisation variant which we call direct Ghost penalty stabilisation. Moreover, we study the applicability of the higher order geometry description developed for the merely spatial case in "High order unfitted finite element methods on level set domains using isoparametric mappings", C. Lehrenfeld (CMAME

300, 716-733).

Fabian Heimann, Christoph Lehrenfeld
University of Göttingen
f.heimann@math.uni-goettingen.de, lehrenfeld@math.uni-goettingen.de

MS259

Derivation of the Effective Boundary Condition on a Porous Boundary

We derive the new effective boundary condition for the fluid flow in domain with porous boundary. Starting from the Newtonian fluid flow through a domain with an array of small holes on the boundary, using the homogenization and the boundary layers, we find an effective law in the form of generalized Darcy law. If the pores geometry is isotropic, then the condition splits in Beavers-Joseph type condition for the tangential flow and the standard Darcy condition for the normal flow.

Eduard Marusic-Paloka
Department of Mathematics,
University of Zagreb, Croatia
emarusic@math.hr

MS259

Boundary Integral Formulation of the Microscopic Bidomain Model of Cardiac Electrophysiology

At microscopic scale, the cardiac tissue is composed by tightly coupled myocytes, embedded into an extracellular matrix and surrounded by other cells. The microscopic bidomain (μ -bidomain) model is a system of partial differential equations describing the evolution of the electric potential within each cell and in the extracellular space. The model accounts for the cellular membrane dynamics and intercellular coupling. From a numerical perspective, the μ -bidomain model poses a number of difficulties: the cellular subdomains are separated by a sharp interface, that is the cellular membrane; the temporal dynamics is confined on such interfaces; and the membrane dynamics is typically multiscale, nonlinear and very stiff in time. In this work, we first recast the problem to a boundary integral formulation and confine the equations on the cellular membrane, on which the temporal dynamics take place. Then, we present some numerical experiments to compare the boundary integral formulation to a more standard finite element approach. We also discuss the feasibility of the approach towards the application to large scale models.

Giacomo Rosilho De Souza
Universita' della Svizzera italiana
giacomo.rosilhodesouza@usi.ch

Simone Pezzuto
Università della Svizzera italiana
simone.pezzuto@usi.ch

Rolf Krause
Università della Svizzera italiana, Switzerland
rolf.krause@usi.ch

MS260

Prediction of Severe Thunderstorm Events as a Warning Machine with Deep Learning and Radar

Data

One of the most interesting problems in weather forecasting is the prediction of extreme rainfall events such as severe thunderstorms possibly leading to flash floods. The problem of nowcasting extreme weather events can be addressed by applying either numerical methods for the solution of dynamic model equations or data-driven artificial intelligence techniques. In this talk, we focus on this second approach. In detail, we present a novel method based on ensemble deep learning that exploits a deep neural network that takes as input time series of multichannel radar reflectivity images and provides as output the probability that an extreme event occurs, where the extreme event is defined on the basis of a certain level of precipitation and lightning density. One of the main novelty is the use of value-weighted skill scores in the ensemble procedure and for assessing the forecasting performance: such value-weighted skill scores allow ranking prediction errors on the basis of their distribution along time, preferring to show up a warning well in advance of the actual occurrence of an event rather than not to show it at all. The results of this study is a data-driven warning system for supporting the decision-making in the case of extreme rainfall events tailored for the Liguria region in Italy.

Sabrina Guastavino
Dipartimento di Matematica
Università di Genova
guastavino@dima.unige.it

MS260

An Observing System Simulation Experiment (OSSE) Using the Particle Flow Filter (PFF) in a High-Dimensional Atmospheric Model in the Data Assimilation Research Testbed (DART)

Representing the true state of the atmosphere is a very difficult task. In the modern weather forecasting, satellite based observations are assimilated into the numerical model in order to obtain a more realistic representation of the atmosphere in the model. This process is called data assimilation (DA). DA plays an important role in constraining the model heavily with observations. However, due to the high dimension of the atmospheric model, many DA methods are developed based on either Gaussian or linear assumptions, which are not applicable for some types of observations. Recently, a more general DA method, the Particle Flow Filter (PFF), is proposed to deal with the nonlinear observations and has been shown to work effectively in high dimensional models. With some modifications of the algorithm, the PFF can be adapted to run efficiently in the Data Assimilation Research Testbed (DART) framework and is now able to run in parallel. We will demonstrate an observing system simulation experiment (OSSE) using an atmospheric model which has nonlinear interactions between dynamics and thermodynamics. Specifically, the performance of the DA using the PFF will be compared with using the Ensemble Adjustment Kalman Filter (EAKF), which is a widely used DA scheme based on linear and Gaussian assumptions. Discussions will be focused on how the PFF and the EAKF deal with the nonlinear observations differently, and how these differences can affect the model behavior.

Chih-Chi Hu, Peter Jan van Leeuwen
Colorado State University
chihchi.hu@colostate.edu, peter.vanleeuwen@colostate.edu

Jeffrey Anderson

National Center for Atmospheric Research
Institute for Math Applied to Geophysics
jla@ucar.edu

MS260

Data-Driven Fitting for First Hard X-Ray Imaging Results by Solar Orbiter Stix

The Spectrometer/Telescope for Imaging X-rays (STIX) is one of 6 remote sensing instruments on-board Solar Orbiter. It provides hard X-ray imaging spectroscopy of solar flares by sampling the Fourier transform of the incoming photon flux. With the first real observations, called visibilities, we analyze two imaging methods: the first one is an enhanced interpolation/extrapolation algorithm based on data-driven fitting, while the second one is a parametric imaging approach.

Emma Perracchione

Polytechnic University of Torino
emma.perracchione@polito.it

MS260

Machine Learning Methods For Numerical Homogenization

The numerical solution to partial differential equations with multiscale data poses significant computational difficulties. Choosing a mesh scale that is too fine can lead to slow convergence or an intractable problem size. Alternatively, too coarse a mesh scale and features of the PDE will go unresolved, yielding an incorrect solution. Methods of numerical homogenization address these by creating an effective model defined at the macro scale by propagating relevant information from the fine space to the coarse space. Unfortunately, state-of-the-art homogenization approaches are sensitive to individual PDE configurations. We propose machine learning methods to construct effective models that are more robust and efficient. Our approach investigates linear and nonlinear models with L_1 sparsity enforcing penalties on heterogeneous PDEs without relying on separable scales or periodicity.

Alexander Stone

University at Buffalo
awstone@buffalo.edu

Matthew G. Knepley

University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

MS261

A Parameter Uniform Domain Decomposition Method for Solving 4th Order Singularly Perturbed Ordinary Differential Equation

In this work, a fourth order singularly perturbed ordinary differential equation is considered where the highest order derivative component is multiplied by a small positive parameter. In order to solve the problem numerically, the differential equation is transformed into a coupled system of singularly perturbed differential equations associated with appropriate boundary conditions. The original domain is divided into three subdomains that overlap: two layer subdomains and one regular subdomain. On each subdomain we consider backward Euler scheme on a uniform mesh in time and the standard central difference scheme on a uni-

form mesh in space. Further, we demonstrate that the proposed scheme is almost second order uniformly convergent. The numerical results also provide to assure the efficiency of the proposed numerical method.

Aakansha Aakansha

Indian Institute of Technology (BHU), Varanasi
aakansha.rs.mat18@itbhu.ac.in

Sunil Kumar

Indian Institute of Technology (BHU)
Varanasi, India
skumar.mat@iitbhu.ac.in

MS261

A New Framework for Easy Derivation of Stiffness-Resilient Exponential Methods

In this talk, we introduce a new theoretical framework for deriving exponential methods. This framework addresses the complexity of deriving exponential schemes and enforces advantageous properties. We show that the dominant error terms of stiff problems are canceled for the methods derived with this framework. This has the effect of increasing the accuracy of the methods and makes them more resilient to the stiffness of the problem. A wide range of exponential methods can be easily derived using this framework. We present the derivation of high order exponential Runge-Kutta, exponential multi-step, and a new class of exponential multi-values methods.

Valentin Dallerit

University of California Merced
vdallerit@ucmerced.edu

MS261

Efficient General Linear Methods for Stiff ODEs

In this work we show how to use the theoretical framework of General Linear Methods (GLMs) to analyze and generalize some classes of existing methods, in order to derive efficient numerical schemes for stiff ordinary differential equations. In particular we focus on multistep methods such as *Backward Differentiation Formulae* (BDF) and *Modified Extended BDF* [J.R. Cash, *The integration of stiff initial value problems in ODEs using modified extended backward differentiation formulae*, Comput. Math. Appl. 9 (5), 645–657 (1983)], to show how to derive the class of *Generalized Linear Multistep Methods* (GLMMs) [G. Izzo, Z. Jackiewicz, *Generalized linear multistep methods for ordinary differential equations*, Appl. Numer. Math., 114, 165–178 (2017)]. Within this latter class of multi-step multi-stage methods, we present new implicit schemes with order up to $p = 9$ and high stage order q (i.e. $q = p - 1$ or $q = p$). This last feature turns out to be crucial to avoid the order reduction phenomenon. Our analysis indicates that the proposed methods have better accuracy and stability properties than the original ones, and the reported numerical experiments confirm that they are competitive and can have better performance on mildly stiff and stiff problems.

Giuseppe Izzo

The University of Naples Federico II
Department of Mathematics and Applications
giuseppe.izzo@unina.it

Zdzislaw Jackiewicz

Arizona State University

jackiewicz@asu.edu

MS261**Weak Stage Order Theory and Order Barriers for Runge-Kutta Methods**

Runge-Kutta (RK) methods may exhibit order reduction when applied to stiff problems. For linear problems with time-independent operators, order reduction can be avoided if the method satisfies certain weak stage order (WSO) conditions. WSO is less restrictive than traditional stage order conditions and compatible with a DIRK structure. This talk will present the first general order barrier bounds relating the WSO of a scheme to its classical order and number of stages. These bounds characterize the fundamental accuracy limit of RK methods applied to stiff problems. New necessary conditions are also established on how one needs to split the spectrum of the Butcher matrix A to devise schemes with WSO which we use to construct new families of high WSO schemes. The key mathematical ideas are to recast WSO into a pair of orthogonal invariant subspaces and perform calculations modulo minimal polynomials. We also provide new formulas for the RK stability function.

David ShirokoffNew Jersey Institute of Technology
david.g.shirokoff@njit.edu

Abhijit Biswas

King Abdullah University of Science and Technology
abhijit.biswas@kaust.edu.sa

David I. Ketcheson

CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Benjamin Seibold

Temple University
seibold@temple.edu**MS261****Construction of Implicit-Exponential Partitioned Integrators**

Over the past several decades a number of new numerical tools were introduced to solve large scale stiff systems of differential equations. A traditional approach to solving such problems is implicit integration. More recently exponential integration ideas produced several classes of methods that can yield computational savings compared to other state-of-the-art techniques. In this talk we discuss how various approaches such as implicit, exponential and partitioned integration can be combined or used separately to design efficient schemes for stiff problems. We will particularly present new implicit-exponential (IMEXP) partitioned methods and demonstrate their advantages using a suite of test problems. These IMEXP schemes are built for problems where forcing term can be additively partitioned into two nonlinear components with stiffness carried by both, but only one possessing an efficient preconditioner. We introduce a new ansatz for constructing this type of methods and outline what advantages this new form can offer for stiff systems of this type.

Mayya TokmanUniversity of California, Merced
School of Natural Sciences

mtokman@ucmerced.edu

MS262**Multigrid Methods for Structured Matrices at ExaScale**

Not only with the advent of exascale systems the specific requirements of these machines have become apparent. Developers of simulation codes and numerical methods are confronted not just with a tremendous amount of parallelism but also with a high degree of heterogeneity, hierarchies in networking and memory and challenging programming models. To face these problems we focus on algebraic multigrid methods for structured matrices. In these we use the level hierarchy naturally present in the system matrices while using advanced results from numerical linear algebra to design optimal components for these matrices. The presence of structure not only allows us to obtain the hierarchy straightforwardly, it also eases the efficient implementation on modern computer architectures, including accelerators like GPUs. Nevertheless, even with efficient implementations, the methods' performance is still memory bound. Certain approaches like higher order discretization, but also techniques from linear algebra like block smoothers or polynomial smoothers, can increase the arithmetic intensity and thus reduce the effect of this memory boundness. In the talk an overview over the approach will be given, some recent findings from numerical linear algebra will be presented and the efficient implementation as well as methods to overcome the implications of modern computer architectures will be discussed.

Matthias BoltenUniversity of Wuppertal
bolten@math.uni-wuppertal.de**MS262****Optimizing Hypr Communication with Node Aware Parallelism**

Hypr is an industry-leading, high-performance solver for sparse system of equations. To perform communication during sparse matrix-vector multiplication and sparse matrix-matrix multiplication, Hypr uses MPI in a basic way using only non-blocking sends and receives. Neighborhood collectives are an alternative communication interface in MPI that provide more context than a single message enabling the use of this context to optimize communication within MPI. This work explores optimizing point-to-point communication in Hypr through the use of neighborhood collectives and investigates the effects of locality-aware optimizations.

Gerald CollomUniversity of New Mexico
Albuquerque NM, USA
geraldc@unm.edu**MS262****Chronos: a Gpu-Accelerated General Purpose AMG Solver for Industrial Applications**

Algebraic Multigrid (AMG) is a very popular iterative method used in several applications. This wide diffusion is due to its effectiveness in solving linear systems arising from PDE discretizations. The key feature of AMG is its optimality, i.e., the ability to guarantee a convergence rate independent of the mesh size for different problems. This is

obtained through a good interplay between the smoother and the interpolation. Unfortunately, for difficult problems, such as those arising in industrial applications, standard AMG techniques are not effective and more elaborate strategies and algorithms need to be developed. The implementation of novel AMG methods became even more difficult with the advent of many-core hardware such as GPU accelerators. While AMG application simply consists in matrix by vector products that have been already successfully ported on GPU by a number of authors, the set-up stage needs to be completely re-designed. In this work, we present the AMG preconditioner included in Chronos, tailored for industrial applications from both solid mechanics and CFD and entirely running on GPU. Thanks to some numerical experiments we will show its performance and scalability on real-world problems.

Carlo Janna
M3E S.r.l.
c.janna@m3eweb.it

MS262

Multigrid Method for the Indefinite Helmholtz Equation

Multigrid methods are optimal for many types of elliptic problems. Convergence depends on a complementary principle : few iterations of a smoother capture high frequency information, while the geometrically smooth information, the near-kernel space, is transferred to a direct method on the coarsest level. Restriction and interpolation operators generate a collection of coarse spaces, and transfer information between them. The usual relaxation methods are efficient for capturing oscillatory information, while interpolation rules for the smooth part are generally not a concern. The Helmholtz equation has two major difficulties. Some eigenvalues are negative, requiring an adapted smoothing method and because the near-kernel space is oscillatory, the geometrical smoothness assumption can not be applied to build efficient interpolation rules. While Krylov iterations or polynomial smoothers based on the normal equations can handle indefiniteness, the choice of a sparse interpolator is still an open question. We describe a method converging in a constant number of iterations independent of the problem size. We use an Ideal framework for SPD matrices to correct initial least square minimization interpolator built from accurate candidate vectors approximating the near-kernel space. Benchmarks are presented for a practical Helmholtz equation, but also for hypothetical oscillatory problems where matrices have an almost equal proportion of negative and positive eigenvalues.

Clément Richefort
Lawrence Livermore National Laboratory
French Atomic Energy Commission (CEA), INRIA
richefort.clement@protonmail.com

MS262

Schwarz Smoothers for Pressure Projection in low-Mach Navier-Stokes Combustion Models

PeleLM is an adaptive mesh low Mach Navier-Stokes combustion code developed under DOE's Exascale Computing Program. A key feature of the model is that the fluid density varies considerably across the computational domain. Extremely ill-conditioned problems arise for incompressible and reacting flows in the low Mach flow regime, particularly for cut-cell meshes in complex geometries, where non-covered cells that are cut by the domain boundar have

arbitrarily small volumes and areas. The standard Jacobi and Gauss-Seidel AMG smoothers are less effective in these cases at reducing the error at each level of the AMG V-cycle and may result in very large iteration counts for the GMRES+AMG solver. Prenter (2020) improved convergence rates for cut-cells and conjugate-gradient solvers with AMG preconditioners by employing Schwarz smoothers. We combine ILU smoothers using iterative triangular solves with restricted additive Schwarz (RAS) adapted to hypre for a new iterated Gauss-Seidel formulation of GMRES with AMG preconditioner. The iteration counts tend to remain constant and these smoothers reduce run times on many-core GPU's in the strong-scaling limit.

Arielle K. Carr
Lehigh University, U.S.
arg318@lehigh.edu

Amik St Cyr
Shell
amik.st-cyr@shell.com

Stephen Thomas
Advanced Micro Devices (AMD)
stephethomas@gmail.com

MS263

Mixed Precision Randomized Preconditioners

Support for floating point arithmetic in multiple precisions is becoming increasingly common in emerging architectures. Mixed precision capabilities are already included in a quarter of the machines on the TOP500 list and are expected to be a crucial hardware feature in coming exascale machines. In this talk, we discuss recent work on exploiting mixed precision in randomized algorithms for low-rank approximation, and in particular, the use of such approximations within preconditioners for Krylov subspace methods.

Erin Claire C. Carson, Ieva Dauzickaite
Charles University MFF
carson@karlin.mff.cuni.cz, dauzikaite@karlin.mff.cuni.cz

MS263

Chopblas: Simulating Mixed-Precision and Stochastically Rounded Linear Algebra

In this talk, we present the ChopBLAS MATLAB library for simulating basic linear algebra operations using mixed precision and non-standard floating-point formats such as BFloat16, stochastic rounding, and custom formats. Existing simulation frameworks use MATLAB classes and operator overloading to implement custom precision data types, which makes implementing mixed-precision computations difficult and leads to decreased computational performance due to the overhead of MATLAB classes. Instead, ChopBLAS uses operation-level rounding, where instead of a custom data type, we store all data as double precision floating-point values, use double precision floating-point arithmetic, and then round the data to the desired precision using either the chop or cpfloat rounding function. By operating on double precision values we are able to exploit the built-in vectorization capabilities of MATLAB operations, leading to a speed-up of nearly 90x and 1050x when simulating stochastic rounding using chop and cpfloat, respectively, compared to a reference BLAS implementation

in MATLAB with operation level rounding.

Ian McInerney
The University of Manchester
i.mcinerney17@imperial.ac.uk

Nicholas J. Higham
School of Mathematics
The University of Manchester
nick.higham@manchester.ac.uk

MS263

Monotonicity of Multi-Term Floating-Point Adders

We demonstrate that standard techniques for performing multi-term addition with four or more arguments, without normalization of intermediate quantities, can be non-monotonic. Since summation is part of dot product and matrix multiplication operations, routines that have increasingly started appearing in the hardware of supercomputers, exploring monotonicity or nonmonotonicity is of wide interest. We provide details on what features multi-term addition algorithms need in order to avoid non-monotonic behaviour. We also look at algorithms that utilize multi-term summation and vector-vector, matrix-vector, or matrix-matrix operations that can have non-monotonic behaviour. We suggest that nonmonotonicity of multi-term summation, in some available hardware devices from large-scale chip companies, is a feature that may have appeared unintentionally as a consequence of design choices, and neither the fact of appearance in these devices nor that nonmonotonicity can happen in multi-term floating-point adders in general has been analysed in literature before, to the best of our knowledge.

Mantas Mikaitis
The University of Manchester
m.mikaitis@leeds.ac.uk

MS263

Solving Total Least Squares Problems Using Mixed Precision

With the recent emergence of mixed precision hardware, there has been a renewed interest in its use for solving numerical linear algebra problems fast and accurately. The solution of total least squares problems, i.e., solving $\min_{E,f} \|[E, f]\|_F$ subject to $(A + E)x = b + f$, arises in numerous application areas. This requires finding the smallest singular value and corresponding right singular vector of $[A, b]$, which is challenging when A is large and sparse. An efficient algorithm for this case due to Björck et al. is based on Rayleigh quotient iteration coupled with conjugate gradient preconditioned via Cholesky factors. In this talk, we present a mixed precision variant of this algorithm and several numerical experiments.

Eda Oktay
Charles University
oktay@karlin.mff.cuni.cz

Erin Claire C. Carson
Charles University MFF
carson@karlin.mff.cuni.cz

MS263

A New Mixed-Precision Benchmark for HP Com-

puters

We present a new benchmark for high-performance (HP) computers. Similar to HPCG, the new benchmark is designed to rank computers based on how fast they can solve a sparse linear system of equations. The main novelty of the new benchmark is that it is based on GMRES (combined with Geometric Multi-Grid preconditioner with Gauss Seidel smoother) and provides the flexibility to utilize lower precision arithmetic. We present our initial design of the new benchmark, its reference implementation, and its performance on current top-ranked architectures. We also discuss challenges of designing such a benchmark, along with our preliminary numerical results of using 16-bit numerical values (half and bfloat precisions) for solving a sparse linear system of equations.

Ichi Yamazaki
Sandia National Laboratory
iyamaza@sandia.gov

Jennifer A. Loe
Sandia National Laboratories
jloe@sandia.gov

Piotr Luszczek
University of Tennessee
luszczek@icl.utk.edu

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory,
USA
dongarra@icl.utk.edu

MS264

A Partitioned Method for Reduced Order Model-Finite Element Model (ROM-FEM) and ROM-ROM Couplings with Separate Reduced Bases for Interior and Interface Variables

Partitioned methods allow for the solution of multiphysics or multiscale problems in a way that considers the specific mathematical or physical properties of the sub-models and allows for the use of solution tools designed for their specific solution. Introduction of a projection-based reduced order model (ROM) for one or more of the sub-models can significantly decrease the on-line computational cost of the system. In this talk, a heterogeneous sub-model case will be considered, having two non-overlapping sub-domains. We present a partitioned method that provides for the coupling of a ROM on one subdomain with either a ROM or a full-order model on the other. At the crux of this method is a dual Schur complement system which implicitly expresses a Lagrange multiplier (LM), representing the interfacial flux, in terms of the state variables. In order to ensure the non-singularity of this system, we develop the ROM by creating separate reduced bases for the interior and interfacial degrees of freedom as well as providing a reduction for the LM space. The solution of the Schur complement system and the application of an explicit time-stepping scheme allow for the two subdomain equations to be decoupled and independently solved at each time step. We show numerical results demonstrating the capability of

the method to represent each type of coupling.

Amy De Castro
SIAM
agdecas@sandia.gov

Paul Kuberry, Irina K. Tezaur
Sandia National Laboratories
pakuber@sandia.gov, ikalash@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

MS264

Some Examples of Linear Or Nonlinear, Intrusive Or Non-Intrusive Reduced Models Based on Convex Displacement Interpolation

The presentation focuses on examples of parametric problems governed by partial differential equations in which the linear representation of the reduced space fails. We introduce a nonlinear approximation technique called convex displacement interpolation based on a solution mapping as a function of the parameters, via optimal transportation. We discuss the advantages and disadvantages in this framework of linear or nonlinear, intrusive or non-intrusive model reduction approaches.

Angelo Iollo
Institut de Mathématiques de Bordeaux
angelo.iollo@math.u-bordeaux1.fr

Simona Cucchiara
Optimad Engineering
simona.cucchiara@optimad.it

Tommaso Taddei
INRIA Bordeaux South-West
tommaso.taddei@inria.fr

Haysam Telib
Optimad
haysam.telib@optimad.it

MS264

Real-Time Simulation of Parameter-Dependent Fluid Flows Through Deep Learning-Based Reduced Order Models

Simulating fluid flows in different virtual scenarios is of key importance towards the design of predictive digital twins in computational fluid dynamics. Compared to high-fidelity, full-order models, reduced order models (ROMs) relying, e.g., on proper orthogonal decomposition (POD) provide reliable approximations to parameter-dependent fluid dynamics problems in rapid times. However, they might require expensive hyper-reduction strategies for handling nonlinear terms, and enriched reduced spaces if a mixed velocity-pressure formulation is considered, possibly hampering the evaluation of reliable solutions in real-time. Dealing with fluid-structure interactions entails even greater difficulties. We propose a new class of deep learning (DL)-based ROMs that overcome all these limitations by learning, in a nonintrusive way, both the nonlinear trial manifold and the reduced dynamics. To do so, they rely on deep neural networks, such as convolutional autoencoders

and deep feedforward neural networks, after performing a former dimensionality reduction through POD that enhances training times substantially. Moreover, including long short-term memory (LSTM) cells enables long-term prediction of complex systems evolution, with respect to the training window, for unseen input parameter values. Numerical results show the capability of POD-DL-ROMs to provide accurate results, in real-time, in multiple contexts, ranging from haemodynamics to fluid-structure interaction problems.

Andrea Manzoni, Stefania Fresca
Politecnico di Milano
andrea.l.manzoni@polimi.it, stefania.fresca@polimi.it

Federico Fatone
Politecnico di Milano, Milan, Italy
federico.fatone@polimi.it

MS264

Reduced Order Modeling for Real-Time Shape Optimization of the Total Cavopulmonary Connection in Fontan Surgical Planning

The Fontan operation is a surgery for children born with ventricular defects, where the inferior vena cava (IVC) is routed to the pulmonary arteries (PA) to form the total cavopulmonary connection (TCPC). This represents a shape optimization problem, where differences in the anatomy result in unique hemodynamics. However, the costs of simulating such cases via computational fluid dynamics (CFD) becomes extremely expensive, especially for testing various TCPC geometries in a trial-and-error fashion. Here, we introduce a reduced order modeling (ROM) framework for real-time shape optimization of the TCPC. In the offline phase of the ROM, the IVC offset with respect to the PAs was parameterized resulting in 100 transient CFD simulations being performed, from which a snapshot library was created. The library was recycled in the online stage for a new set of parameters using the Proper Orthogonal Decomposition approach. A shape optimization problem was implemented using the framework, where the optimization criteria was based on the hepatic flow distribution (HFD). ROM solutions were in strong agreement with full order model (FOM) solutions with regards to pressure, velocity, and flow dynamics. The FOM required an average of 153.5 seconds while the ROM reduced this to 27.5 seconds, an 82.1% decrease in computation time. Using this framework, the optimal TCPC geometry for a given HFD value was rapidly calculated, requiring less than 15 minutes for finding the optimal solution.

Imran Shah
Emory University
ishah3@emory.edu

Huang Chen, Lakshmi Dasi
Wallace H. Coulter Department of Biomedical Engineering, Geo
huang.chen@gatech.edu, lakshmi.dasi@gatech.edu

Ajit P. Yoganathan
Department of Biomedical Engineering
Georgia Institute of Technology
ajit.yoganathan@bme.gatech.edu

Zhenglun Wei
Department of Biomedical Engineering, University of Massachusetts

zhenglunalan_wei@uml.edu

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

Omer San
Mechanical and Aerospace Engineering
Oklahoma State University
osan@okstate.edu

Alessandro Veneziani
MathCS, Emory University, Atlanta, GA
ale@mathcs.emory.edu

MS265

An Approximate Bayesian Computation Approach to Goal-Oriented Bayesian Experimental Design

Conventional Bayesian experimental design techniques aim to find the optimal design by maximising the expected information gain (EIG) on the model parameters. However, the EIG may not always be a suitable criterion for experimenters as it is often hard to interpret and may be an indirect way of understanding the experiment. This calls for goal oriented experimental designs that directly addresses the goal of the experimenter. An added advantage of the goal oriented formulation lies in eliminating the need to work on high dimensional parameter spaces. In this paper, we consider the EIG as a function of the goal directly and thus aim to choose a design that maximises the reduction of the entropy from the marginalised prior of the goal function to the posterior of the goal function over all possible experimental outcomes. To address the computationally challenging task of efficiently computing the EIG across all design points, we propose an approach using Approximate Bayesian Computation (ABC) that utilises simulations from the model in place of likelihood evaluations. We demonstrate the improved performance of our method over traditional Nested Monte Carlo approaches in several examples.

Atlanta Chakraborty, Xun Huan
University of Michigan
atlantac@umich.edu, xhuan@umich.edu

Thomas A. Catanach
Sandia National Laboratories
tacatan@sandia.gov

MS265

Using Local Lipschitz Stability of Posteriors for Model Error-Aware Experimental Design

We consider experimental design for Bayesian inverse problems, where the parameter-to-observable map or 'forward model' is the composition of a parameter-to-solution map and a solution-to-observable map, which we call the 'solution map' and 'observation map' respectively. Experimental design often involves choosing the observation map to minimise uncertainty in the parameter estimate. However, for some inverse problems, error in the solution map may arise due to inaccurate mathematical models or the use of surrogates. This error may propagate to error in the forward model and thus lead to biased parameter estimates. If the bias in the estimate is large, then minimising uncertainty about the estimate is undesirable, and the need

arises for experimental design that takes into account error in the solution map. We consider some natural approaches for modeling the error in the solution map, and use theoretical results about local Lipschitz stability of the posterior in order to suggest how one may choose the observation map for each approach.

Han Cheng Lie
University of Potsdam
hanlie@uni-potsdam.de

Harshit Bansal
Eindhoven University of Technology
h.bansal@tue.nl

Nada Cvetkovic, Karen Veroy-Grepl
TU/e Eindhoven University of Technology
n.cvetkovic@tue.nl, k.p.veroy@tue.nl

MS265

Gibbs Optimal Design of Experiments

Gibbs (or generalised Bayesian) inference is a generalisation of Bayesian inference made by replacing the log-likelihood in Bayes' theorem by a (negative) loss function. The loss function identifies desirable parameter values for given responses. The advantage of Gibbs inference over traditional Bayesian inference is that it does not require the specification of a probabilistic data-generating process and, therefore, should be less sensitive to this process. This talk proposes Gibbs optimal design of experiments for this inferential framework, extended decision-theoretic Bayesian optimal design. The challenge is that the decision-theoretic approach relies on a probabilistic data-generating process that is notably absent from Gibbs inference. This is circumvented by assuming a designer model: a probabilistic data-generating process which is only used to find a design rather than in the ensuing inference. Because of this, the designer model can encapsulate very general data-generating processes with the aim of introducing robustness into the design procedure. The proposed Gibbs optimal design framework is demonstrated on several illustrative examples.

Antony Overstall
University of Southampton
a.m.overstall@soton.ac.uk

MS265

Disturbance Rejection of Interval Type-2 Fuzzy Stochastic Systems : A Disturbance Observer Based Approach

Disturbance rejection of interval type fuzzy nonlinear stochastic systems with asynchronous premise variables is proposed. The disturbance observer based on disturbance attenuation control structure is constructed to estimate disturbance which is generated by exogeneous system. The construction of a fuzzy feedback controller with asynchronous premise variables ensures that the membership functions of the premise are not necessarily the same as in interval-type fuzzy systems. To guarantee the stochastic stability of the closed loop fuzzy system, a new adequate condition is constructed through the method of linear matrix inequalities by integrating the infinitesimal operator and the selection of appropriate Lyapunov-Krasovskii functional candidate. The efficiency of the provided theory is then demonstrated with a numerical illustration for

Mass Spring Dashpot system.

Aarthi Subramanian, Marshal Anthoni
B
Anna University Regional Campus Coimbatore
aarthi15.au@gmail.com, smarshalanthoni@gmail.com

MS266

Finite Element Methods for the Stokes-Onsager-Stefan-Maxwell Equations of Multicomponent Flow

The Onsager framework for linear irreversible thermodynamics provides a thermodynamically consistent model of mass transport in a phase consisting of multiple species, via the Stefan-Maxwell equations, but a complete description of the overall transport problem necessitates also solving the momentum equations for the flow velocity of the medium. We derive a novel nonlinear variational formulation of this coupling, called the (Navier)Stokes-Onsager-Stefan-Maxwell system, which governs molecular diffusion and convection within a non-ideal, single-phase fluid composed of multiple species, in the regime of low Reynolds number in the steady state. We propose an appropriate Picard linearization posed in a novel Sobolev space relating to the diffusional driving forces, and prove convergence of a structure-preserving finite element discretization. This represents some of the first rigorous numerics for the coupling of multicomponent molecular diffusion with compressible convective flow. The broad applicability of our theory is illustrated with simulations of the centrifugal separation of noble gases and the microfluidic mixing of hydrocarbons.

Francis Aznaran
University of Oxford
aznaran@maths.ox.ac.uk

Alexander Van-Brunt
Bosch USA
alexander.van-brunt@us.bosch.com

MS266

On the Finite Element Approximation of Incompressible Implicitly Constituted Fluids

In this talk I will provide a general overview of the finite element approximation of systems describing incompressible implicitly constituted fluids. Given the highly nonlinear nature of the models, and possible non-uniqueness of solutions, the results available so far focus mostly on proving (weak) convergence of the numerical approximations to minimal regularity solutions. I will go through the main ideas in the convergence proof, highlighting new results obtained in collaboration with Alex Kaltenbach, which focus on nonconforming discretisations in both space and time.

Pablo Alexei Gazca Orozco, Alex Kaltenbach
Albert-Ludwigs-Universität Freiburg
alexei.gazca@mathematik.uni-freiburg.de,
alex.kaltenbach@mathematik.uni-freiburg.de

MS266

A Multilevel Substructuring Method for Fast Vibration Response Calculations of Finite Periodic Structures

Periodic metamaterial structures are attracting increasing interest in the scientific community because of their unique properties that go beyond the ones of natural materials.

Due to the often detailed nature of the unit cells (UCs) in these architected structures, the UC finite element models can become large, which would rapidly render finite structure models comprised of UC assemblies computationally unaffordable. To tackle this problem, conventional substructuring approaches approximate the dynamics of the single UC with the subspace and assembly the reduced UCs to obtain the reduced order model. However, the size of the reduced order model grows linearly with respect to the number of UCs. Hence, a natural enhancement is to add more levels of decomposition as the automated multilevel substructuring (AMLS) method. To accelerate this approach even further, we present a multilevel substructuring method combined with the interface reduction and subspace recycling for the forced vibration response calculations of finite periodic structures. We analyze the scaling property of the method numerically comparing with conventional substructuring approaches. The work has been funded by the project "MOR4MDesign", which is part of the MacroModelMat (M3) research program coordinated by Siemens (Siemens Digital Industries Software, Belgium) and funded by SIM (Strategic Initiative Materials in Flanders) and VLAIO (Flanders Innovation & Entrepreneurship Agency).

Fei Qu
KU Leuven, Department of Mechanical Engineering,
SIM M3 program, Technologiepark 48, 9052 Zwijnaarde,
Belgium
fei.qu@kuleuven.be

Lucas Van Belle
KU Leuven, Department of Mechanical Engineering,
DMMS Core lab, Flanders Make.
lucas.vanbelle@kuleuven.be

Wim Desmet
KU Leuven, Department of Mechanical Engineering
wim.desmet@kuleuven.be

Elke Deckers
KU Leuven, Department of Mechanical Engineering,
DMMS Core lab, Flanders Make.
elke.deckers@kuleuven.be

MS266

Mathematical and Numerical Modeling of Multi-Component Diffusion

An often used model for multi-component diffusion in mixtures are the Stefan-Maxwell equations. These equations can be written as a linear system for the diffusion velocities with a singular coefficient matrix, the so-called friction matrix, and a right hand side vector containing the driving forces for diffusion. To express the diffusion velocities in terms of mass fraction gradients, we first have to regularize the friction matrix, taking into account the mass flux constraint, and subsequently convert mole fractions to mass fractions. The resulting expressions are substituted in the continuity equations, leading to a coupled system of conservation laws, where the diffusion is expressed in terms of a diffusion matrix. To compute numerical solutions of this system, we employ the finite volume method in combination with the system version of the complete flux scheme. The basic idea of this scheme is to compute the numerical flux vector at a cell interface from a local system boundary value problem, thus including the coupling between the constituent equations in the discretization. The scheme has favorable properties, such as conservation of

total mass. We will demonstrate the performance of the scheme for some examples.

Jan Ten Thije Boonkkamp
Department of Mathematics and Computer Science
Eindhoven University of Technology
j.h.m.tenthijeboonkkamp@tue.nl

Jan van Dijk
Department of Applied Physics
Eindhoven University of Technology
j.v.dijk@tue.nl

MS266

Thermodynamically Compatible Non-Isothermal Phenomenological Model for Mullins Effect: Implementation and Simulations of Selected Problems

This talk builds on the talk by V. Prua: "A Thermodynamic Framework for Non-Isothermal Phenomenological Models of Mullins Effect". We derive a thermodynamically compatible model for the Mullins effect including a consistent evolution equation for the temperature. The system of balance equations together with an implicit constitutive relation for the Cauchy stress tensor is formulated in the Eulerian setting. For convenience, the equations are transformed to the Lagrangian setting and used for simulations of several different problems such as the investigation of the interplay between the Mullins effect and the Gough-Joule effect or the study of the Mullins effect for a cylinder subjected to combined extension and torsion. The model is implemented with Firedrake finite element code that provides automatic differentiation and is tightly connected with PETSc solver library, which enables us to use effective iterative solvers. The non-linear problem is solved with the Newton method. Consequent linear sub-problems are solved with GMRES with geometric multigrid used as a preconditioner.

Karel Tuma
Charles University, Prague, Czech Republic
ktuma@karlin.mff.cuni.cz

Vit Prusa
Charles University, Faculty of Mathematics and Physics
Sokolovská 83, 18675 Prague 8, Czech Republic
prusv@karlin.mff.cuni.cz

Pablo Alexei Gazca Orozco
Albert-Ludwigs-Universität Freiburg
alexei.gazca@mathematik.uni-freiburg.de

David Cichra
Charles University, Prague
Czech Republic
cichra.david@gmail.com

MS267

Numerical Simulation of Ice Crystal Trajectories and Fragmentation Dynamics Around the Xrfl Fuselage Nose

Since the early 1920s, icing has been reported to appear on airplane surfaces, thus on probes, or in aircraft engine compressors. It is a significant safety concern for the aeronautic industry since it can lead to dramatic events. HAIC and recently MUSIC-haic models are developed in CEDRE solver, an ONERA's CFD code, to help the aeronautic in-

dustry to assess ice crystal icing on new products and future generations of aircraft engines and meet certification requirements. This study considers the numerical simulation of ice crystal trajectories around an airplane fuselage nose. First, the aerodynamics airflow field is investigated. Then, trajectories and the associated fragmentation dynamics on walls are considered. Numerical simulation with a full deposition model is compared to numerical simulations with two different fragmentation models: HAIC and MUSIC-haic models. The comparison between the fragmentation models aims at assessing the capabilities of the recently proposed ice crystal fragmentation model (MUSIC-haic) with respect to the state-of-the-art model (HAIC). Both models show significant deviation from the full deposition model, therefore stating the relative importance of particle fragmentation on particle concentration along the fuselage nose. Although the two fragmentation models show some agreement, they also exhibit a large discrepancy in the fragmentation process.

Moussa Diop, Tristan Soubrié
ANDHEO
moussa.diop@andheo.fr, tristan.soubrie@andheo.fr

Julien Cliquet
Airbus
julien.cliquet@airbus.com

Jean-Mathieu Senoner, Philippe Villedieu
ONERA
jean-mathieu.senoner@onera.fr,
philippe.villedieu@onera.fr

MS267

Kinetic-Based Two-Phase Flow Model: a Reduced-Order Model of Polydisperse Oscillating Droplets with Geometrical Variables

We consider two-phase compressible flows involving both separate regime where the interface dynamics is "smooth" and disperse regimes i.e. sprays of droplets. We introduced multi-fluid two-scale models whose large-scale is adapted to separate regime while the small-scale is for the disperse regime. The modeling of the latter relies on a kinetic modeling of the small-scale disperse phase. Originally proposed by [Williams, 1958], a kinetic-based model fits well the disperse configuration when the droplets are sufficiently distant from each other. However, the separate regime is not suited for such a description and averaged descriptions of the dynamics have been privileged. [Pope, 1988] proposed a probabilistic framework for separate regime, and it has been successfully extended to disperse flows by [Essadki, 2019] for spherical droplets. These works also underlined the key role of geometrical variables in the construction of a unified model for all regimes. This work proposes a polydisperse flow model of oscillating droplets extending Essadki's work. First, we propose a new set of geometrical variables along with the associated Williams-Boltzmann equation, recovering the work of [ORourke & Amsden, 1987]. Second, the moments compatible with Popes framework are identified. Finally, we propose realizable schemes [Ait Ameur, 2022] for simulations of the spray with both monodisperse and polydisperse closures.

Arthur Loison
école polytechnique

arthur.loison@polytechnique.edu

MS267

Fredholm Theory for the Linearized Boltzmann Operator for Polyatomic Gases

This talk is devoted to the Fredholm property of the linearized Boltzmann operator for a mixture of polyatomic gases with a continuous internal energy variable. The collision operator is constructed according to a Borgnakke-Larsen procedure. The linearized operator around equilibrium is written as a perturbation of a multiplication operator. We prove that the perturbation operator is a Hilbert-Schmidt operator and the multiplication operator is coercive. This yields the Fredholm property of the linearized Boltzmann operator.

Stephane Brull
Université de Bordeaux
stephane.brull@math.u-bordeaux.fr

Marwa Shahine
University of Bordeaux
marwa.shahine@math.u-bordeaux.fr

Philippe Thieullen
Université de Bordeaux
philippe.thieullen@math.u-bordeaux.fr

MS267

An Entropic Asymptotic Preserving Scheme for the M1 Model Based on the Unified Gas Kinetic Scheme

Kinetic equations appears in many field of study (radiative transfer, rarefied-gas dynamics, etc) to model particles transport. The multi-scale nature of those equations leads to numerical challenges in terms of asymptotic behaviour to recover all Knudsen number flows. Considering this challenge, the UGKS (Unified Gaz Kinetic scheme) has been developed by K. Xu for the BGK model. This characteristic based scheme correctly captures both the continuum and rarefied flow regimes, with a time step much larger than the collision time. However, the problems dimension induces a huge computational cost. Under certain assumptions, moments models of those kinetic equations allow to reduce the velocity space dimension by assuming the distribution function shape. Elaborating robust asymptotic preserving schemes for those models is also challenging. Following L. Mieussens work on an adaptation of UGKS to a linear kinetic model of radiative transfer theory, we propose a new numerical scheme based on UGKS for the associated M1 model. We prove that this new scheme preserves the moments achievability and ensures the entropy decreases.

Thomas Vigier
University of Bordeaux
thomas.vigier@u-bordeaux.fr

Luc Mieussens
Institut de Mathematiques de Bordeaux
luc.mieussens@math.u-bordeaux.fr

Julien Mathiaud, Jean-Luc Feugeas
CELIA, University of Bordeaux, CNRS, CEA, UMR 5107
F-33405 Talence, France
julien.mathiaud@u-bordeaux.fr, jean-luc.feugeas@u-

bordeaux.fr

MS267

Symmetric Regularized 13-Moment Equations in the Linear Regime

We develop the symmetric regularized 13-moment equations in the rarefied gas dynamics for a general class of linearized collision models in the linear regime. By reconstructing the basis functions, the distribution function is expanded to a finite number of terms with $O(Kn^3)$ accuracy where Kn is the Knudsen number. The resulting system of moment equations based on such expansion is shown to have a symmetric structure. We also propose the Onsager boundary conditions for the moment equations. The model is then tested for one-dimensional stationary channel problems, and numerical experiments are carried out to show the performance of our model.

Siyao Yang, Zhenning Cai
National University of Singapore
matsiya@nus.edu.sg, matcz@nus.edu.sg

Manuel Torrilhon
Schinkelstrasse 2
RWTH Aachen University
torrilhon@rwth-aachen.de

MS268

Algorithmic Challenges in Topology Optimization of District Heating Networks

District heating networks have the potential to decarbonize heat supply. By designing them for low supply temperatures, the potential of renewables and waste heat can be fully exploited. The high investment costs, however, hamper their wide implementation. Over the last decades, several attempts have been made to directly optimize the network topology for maximal return on investment. Most of these studies successfully apply mixed-integer linear or nonlinear programming methods to decide on the optimal piping connections in the network. However, while the linear methods are limited in applicability, the nonlinear methods suffer from an exponential computational cost scaling with the network size. In this contribution, we present how adjoint-based topology optimization methods can be adopted for the topology optimization of district heating networks. We describe the key steps to adapt density-based topology optimization methods to achieve discrete solutions for the diameters of the heating pipes. Using an augmented Lagrangian method, large amounts of engineering constraints are efficiently accounted for in the design. Moreover, a benchmark with a mixed-integer nonlinear programming method confirms the advantageous scaling of the presented topology optimization method. Thus, we are able to showcase first topology optimization results for large-scale district heating networks.

Maarten Blommaert
KU Leuven
maarten.blommaert@kuleuven.be

Yannick Wack
KU Leuven, Department of MEchanical Engineering,
Celestijnenlaan 300 B-3001 Heverlee (Leuven) Belgium
yannick.wack@kuleuven.be

Robbe Salenbien
Vito

robbe.salenbien@vito.be

Martine Baelmans
 KU Leuven, Department of Mechanical Engineering, 3001
 Leuven
 martine.baelmans@kuleuven.be

MS268

Shape Optimization for Parabolic Problems on Time-Dependent Domains

This talk is concerned with the solution of time-dependent shape optimization problems. Specifically, we consider the heat equation in a domain which might change over time. We compute Hadamards shape gradient in case of both, domain integrals and boundary integrals. As particular examples, we consider the one-phase Stefan problem and the detection of a time-dependent inclusion. We discuss the numerical solution of these problems and present respective results.

Rahel Brügger, Helmut Harbrecht
 University of Basel
 Department of Mathematics and Computer Science
 ra.bruegger@unibas.ch, helmut.harbrecht@unibas.ch

Johannes Tausch
 Southern Methodist University
 Department of Mathematics
 tausch@smu.edu

MS268

A Novel Density Based Approach for Topology Optimization

A new method for performing density based topology optimization is presented. We discuss the approach in the setting of Stokes flow. It is based on classical topology optimization and phase field approaches, and introduces a different way to relax the underlying infinite-dimensional mixed integer problem. We compare the novel approach to state-of-the-art density based and phase field approaches, and approaches based on the topological derivative. Moreover, we give a theoretically founded choice of the relaxed problems. The density is modeled on a space that allows for jumps along hypersurfaces, such as BV or fractional order Sobolev spaces. Degrees of freedoms in modeling the optimization problem are chosen based on theoretical considerations, i.e., by discussing the arising optimization problems concerning existence, differentiability and convergence towards solutions of the unrelaxed problem. We present some numerical results. Moreover, in order to show the potential of the new approach we do a comparison to a classical approach.

Johannes Haubner
 Simula Research Laboratory
 haubnerj@simula.no

Michael Ulbrich
 Technical University of Munich
 Chair of Mathematical Optimization
 mulbrich@ma.tum.de

Franziska Neumann
 Technical University of Munich
 Institute of Automotive Technology

franziska.neumann@tum.de

MS268

Topology Optimization of Superhydrophobic Surfaces to Delay Laminar-Turbulent Transition

The microstructure of SuperHydrophobic Surfaces (SHSs) can trap a layer of gas when submerged in a liquid. This phenomena reduces the skin friction of the overlying fluid and has been shown to delay laminar-turbulent transition. Research into the macroscopic layout of SHSs as a means of passive flow control, however, remains relatively under investigated. Here, topology optimization is used to design an inhomogeneous SHS that can further delay spatially developing modal transition in channel flows; a highly non-linear and non-convex optimization problem. The influence of the SHS is modelled using a spatially varying slip length on the walls, forming a 2D optimization domain. The spectral element solver Nek5000 is used to first establish a steady baseflow in a wall-bounded channel, and then conduct unsteady direct numerical simulations of the transition to turbulence. The optimization procedure uses the adjoint-variable method to compute sensitivities and uses a checkpointing strategy to overcome the data storage requirements of the unsteady adjoint simulation. This methodology has been successfully applied to produce a surfaces that moves the transition location 29% further downstream when compared to a homogeneous counterpart by inhibiting the growth of secondary instability modes.

Harrison Nobis
 KTH Royal Institute of Technology
 Stockholm
 nobis@kth.se

Philipp Schlatter
 KTH Royal Institute of Technology
 pschlatt@mech.kth.se

Dan Henningson
 Department of Mechanics,
 Royal Institute of Technology, Stockholm
 henning@mech.kth.se

Martin Berggren
 Department of Computing Science, Umea University
 martin.berggren@cs.umu.se

Eddie Wadbro
 Department of Computing Science
 Umeå University
 eddiew@cs.umu.se

MS269

Relating Abgrall's Correction Schemes to Filters and Artificial Dissipation

As a means of improving both solution accuracy and robustness, methods that are capable of preserving auxiliary properties of the governing equations have been in strong development. In the case of conservation laws, these schemes often employ non-standard flux formulations which can entail intrusive coding modifications. Alternatively, Abgrall introduces a procedure for satisfying additional conservation relations such as entropy [R. Abgrall, *J. Comput. Phys.*, 2018]; the approach utilizes a telescoping correction term that is agnostic to the baseline method. Abgrall's correction term is composed of first and second

moment statistics (i.e., the mean and variance) of the relevant transformation variables (e.g., entropy variables). A generalization of the methodology is thus made possible by relating averaging to filtering. Therefore, the current work re-interprets Abgrall's correction procedure as specialized filter-based artificial dissipation methods [A. K. Edoh, *J. Comput. Phys.*, 2018]. We consider finite difference discretizations within the SBP-SAT framework and study the impacts of filter choice (e.g., filter order, stencil localization, wavenumber attenuation) on the overall performance of the approach. This work has been approved for public release (AFRL Public Affairs Clearance Number AFRL-2022-4517

Ayaboe K. Edoh
ERC, Inc. (AFRL)
ayaboe.edoh@jacobs.com

MS269

High Order Strong Stability Preserving Multi-Derivative Implicit and IMEX Runge–Kutta Methods with Asymptotic Preserving Properties

In this work we present a class of high order unconditionally strong stability preserving (SSP) implicit multi-derivative Runge-Kutta schemes, and SSP implicit-explicit (IMEX) multi-derivative Runge-Kutta schemes where the time-step restriction is independent of the stiff term. The unconditional SSP property for a method of order $p \geq 2$ is unique among SSP methods, and depends on a backward-in-time assumption on the derivative of the operator. We show that this backward derivative condition is satisfied in many relevant cases where SSP IMEX schemes are desired. We devise unconditionally SSP implicit Runge-Kutta schemes of order up to $p=4$, and IMEX Runge-Kutta schemes of order up to $p=3$. For the multi-derivative IMEX schemes, we also derive and present the order conditions, which have not appeared previously. The unconditional SSP condition ensures that these methods are positivity preserving, and we present sufficient conditions under which such methods are also asymptotic preserving when applied to a range of problems, including a hyperbolic relaxation system, the Broadwell model, and the Bhatnagar-Gross-Krook (BGK) kinetic equation. We present numerical results to support the theoretical results, on a variety of problems.

Zachary J. Grant
Oak Ridge National Laboratory
Grantza1@msu.edu

MS269

Scalable, High Order Remapping Methods for Multimodel Problems

Coupled multimodel simulations involve high degree of computationally complex workflows, and achieving consistently accurate solutions is strongly dependent on the choice of spatiotemporal numerical algorithms used to resolve the interacting scales in physical models. Rigorous spatial coupling between components in such systems involves field transformations, and communication of data across multiresolution grids, while preserving key attributes of interest such as global integrals and local features, which is usually referred to as the process of "remapping". Such remap procedures are critical in ensuring stability and accuracy of scientific codes simulating multiphysics problems that typically occur in many different scientific domains such as Fluid-Structure interaction, nuclear reactor analysis, magneto-hydrodynamics modeling,

and Climate and Weather simulations to name a few. Improper treatment of the coupled solution terms in multiphysics simulations can lead to large inaccuracies that can propagate and destroy the overall accuracy and stability of the numerical computation. Since solution transfer between different discretizations of the domain plays a critical role here, we introduce an infrastructure based on the MOAB unstructured mesh library to handle field projections to recover high-order accuracy in a scalable way, while preserving key remapping metrics of interest for various applications.

Vijay S. Mahadevan, Iulian Grindeanu, Robert Jacob
Argonne National Laboratory
mahadevan@anl.gov, iulian@anl.gov, jacob@anl.gov

MS269

Entropy-Stable Positivity Preserving Schemes for Multiphase Flows

The multiphase description of flow phenomena is an attractive option for modelling complex flows as all fluids in the domain can be modelled using a single set of governing equations as opposed to separate sets of equations for each phase. The current work is aimed at studying the problem of an acoustically generated bubble collapse using high accuracy simulations. Typically, high-order methods have stability issues, especially when considering complex physics. However, in the last decade, there has been considerable progress made in developing provably nonlinearly stable schemes for conservative problems by combining summation-by-parts (SBP) operators and symmetric two-point flux functions. The goal of this work is to develop provably-stable high-order schemes for multiphase flows in non conservative form. To achieve this goal, the framework of Renac is leveraged to develop an entropy-stable SBP scheme using a non-symmetric two-point flux function. However, quantities such as density and void fractions of each phase require positivity to be preserved. To this end, the work of Upperman and Yamaleev is utilised to determine the appropriate time step to take to preserve positivity of these quantities. In this presentation, the details of the resulting schemes and numerical tests on smooth flows are presented.

Ben Simpson
University of Waterloo
bj2simps@uwaterloo.ca

David C. Del Rey Fernandez
University of Waterloo
Department of Applied Mathematics
dcdelrey@gmail.com

Sivabal Sivaloganathan
University of Waterloo
Applied Mathematics
ssivaloganathan@uwaterloo.ca

MS270

Large Scale Simulations of Diffusio-Phoretic Suspensions

In the last decades catalytic micron-sized particles have been developed to mimic natural micro-swimmers and for engineering applications like drug delivery or water purification. These suspensions of artificial phoretic particles can form coherent structures due to dual hydro-chemical interactions. The numerical modelling of such systems re-

mains particularly challenging given the coupled nature of the chemical and the hydrodynamical problems and the rapid increase in the required computational resources for large ensembles of particles. In this work, we propose a new framework based on the core ideas of the Force Coupling Method, able to handle both the chemical reaction-diffusion and Stokes flow problems for a large number of particles (Diffusio-phoretic Force Coupling Method). We implement this method in an efficient parallel solver, in order to model the collective motion of diffusio-phoretic suspensions made of Janus particles. We will show how the macroscopic behaviours emergent in these suspensions strongly depend on the chemical surface properties of the individual particles.

Blaise Delmotte

LadHyx, CNRS/Ecole Polytechnique
blaise.delmotte@ladhyx.polytechnique.fr

Francisco Rojas, Sebastien Michelin
LadHyX, Ecole Polytechnique
rojas@ladhyx.polytechnique.fr,
bastien.michelin@ladhyx.polytechnique.fr

MS270

High-Order Boundary Integral Equations on Implicitly Defined Surfaces

We present ongoing work on accurately solving boundary integral equations on implicitly defined surfaces in \mathbb{R}^d . The method relies on combining a dimension-independent technique for generating a high-order surface quadrature on level-set surfaces, with the general-purpose density interpolation method for handling the singular and nearly-singular integrals ubiquitous in boundary integral formulations. The proposed methodology, based on a Nystrom discretization scheme, bypasses the need for generating a body conforming mesh for the implicit surface, allowing in principle for an efficient coupling between a robust dynamic level-set representation of deforming surfaces, and boundary integral equation solvers. Particular attention will be paid to the computation of singular integrals when only a surface quadrature is available (i.e. in the absence of an actual mesh). We believe such techniques could prove useful in applications involving microscopic flows governed by the Stokes equations; in particular, the simulation of micro-swimmers and droplet microfluidics.

Luiz Faria

ENSTA Paris
luiz.maltez-faria@inria.fr

Aline Lefebvre-Lepot
CMAP, CNRS/Ecole Polytechnique
aline.lefebvre@polytechnique.edu

Dongchen He
Ecole Polytechnique and ENSTA
dongchen.he@polytechnique.edu

MS270

Hydrodynamic Interactions in Dense Active Suspensions

We discuss a methodology to model dense active suspensions and present one implementation in the fast Stokesian dynamics framework. The algorithm accounts for full hydrodynamic interactions (both long-range and lubrication) and has an $O(N)$ scaling, where N is the number of par-

ticles. We apply this method to study the dynamics and rheology of a suspension of squirmers under cyclic shear. The relationship between the suspension viscosity and particle diffusivities, as well as the effect of hydrodynamic interactions, will be discussed.

Zhouyang Ge

University of British Columbia
Dept. Mech. Eng.
zhoge@mech.kth.se

Gwynn Elfring

University of British Columbia
gelfring@mech.ubc.ca

MS270

Topological Data Analysis of Spatial Point Patterns

Effectively characterizing spatial point patterns in physical, chemical, and biological systems is often challenging due to the great diversity in possible arrangements. The pair correlation function was thus introduced almost a century ago to succinctly capture the distribution of interparticle distances. Despite its great success in providing deep insight into countless physical systems and mechanisms, this method can be highly sensitive to structurally insignificant differences among patterns and also insensitive to structurally significant ones. We show how topological data analysis provides a natural and robust discrete version of this method that is both simple to compute and easy to analyze, providing insight into ordered and disordered systems alike.

Emanuel A. Lazar

University of Pennsylvania
mlazar@math.biu.ac.il

Vasco Worlitzer

Technische Universität Berlin
Hardenbergstraße 36, 10623 Berlin, Germany
vasco.m.worlitzer@campus.tu-berlin.de

Gil Ariel

Bar-Ilan University
arielg@math.biu.ac.il

MS270

Numerical Simulations of Suspensions: Lubrication Correction Including Fluid Correction

We address the problem of numerical simulation of suspensions of rigid particles in a Stokes flow. We focus on the inclusion of the singular short range interaction effects (lubrication effects) in the simulations when the particles come close one to another. The problem is solved without introducing new hypothesis nor model. The key idea is to decompose the velocity and pressure flows in a sum of a singular and a regular part. The singular part is computed using an explicit asymptotic expansion of the solution when the distance goes to zero. This expansion can be computed for any locally convex (that is the particles have to be convex close to the contact point) and regular shape of particles. We prove that the remaining part is regular in the sense that it is bounded independently of the distance. As a consequence, only a small number of degrees of freedom are necessary to compute numerically this remaining regular part of the solution and obtain accurate results. The method is tested in dimension 2 for clusters of two or three aligned particles with general rigid velocities.

We show that, as expected, the convergence is independent of the distance.

Aline Lefebvre-Lepot
CMAP, CNRS/Ecole Polytechnique
aline.lefebvre@polytechnique.edu

Flore Nabet
CMAP, Ecole Polytechnique, France
flore.nabet@polytechnique.edu

MS271

A Discrete Framework for Hypocoercivity Methods in the Context of Kinetic Equations

We propose a numerical method for a Vlasov-Fokker-Planck model with external applied force field and prove quantitative results ensuring that it is Asymptotic-Preserving for both the macroscopic regime and the long time behavior simultaneously. We illustrate these results with various numerical experiments in which we observe, among others, transition phase between macroscopic and long time behavior as well as other finer phenomena, such as oscillations, which were not predicted by our theoretical analysis.

Alain Blaustein
Université de Toulouse
Institut de Mathématiques de Toulouse
alain.blaustein@math.univ-toulouse.fr

MS271

Multi-Scale Inverse Problems for a Kinetic Chemotaxis Model

The directed motion of particles such as bacteria is governed by their turning rate. We study the inverse problem of determining this turning rate in a mesoscopic run-and-tumble model from macroscopic, directionally averaged data. The lack of directional information in the measurements poses problems in the reconstruction. We consider theoretical reconstructability of this rate in the kinetic and investigate the macroscopic limit behaviour for the inverse problem. Furthermore, we present first numerical results. This is joint work with Christian Klingenberg (Wrzburg, Germany), Qin Li (Madison, Wisc., USA) and Min Tang (Shanghai, China).

Kathrin Hellmuth
Wuerzburg University, Germany
Kathrin.hellmuth@mathematik.uni-wuerzburg.de

MS271

Recent Trends on Nonlinear Filtering for Inverse Problems

In many applications is required to determine a parameters set which is suitable for competing models. To this end, we are interested in ensemble methods to solve multi-objective optimization problems. Here, the ensemble Kalman Filter method is applied and adapted in order to solve coupled inverse nonlinear problems using a weighted function approach. An analysis of the mean field limit of the ensemble method yields an explicit update formula for the weights. Numerical examples show the improved performance of the proposed method.

Michael Herty

RWTH Aachen University
herty@igpm.rwth-aachen.de

MS271

A Hybrid Learning-Inversion Reconstruction Scheme for Inverse Transport Problems

We propose a computational framework to couple machine learning methods with classical model-based computational inversion techniques for nonlinear inverse problems to the radiative transport equation. In a nutshell, for a given configuration of the inverse problem, we train an approximate inverse operator based on the subspace factorization idea and then utilize the approximate inverse to improve model-based inversion. We provide some computational evidence to demonstrate the effectiveness of the method together with a little theoretical understanding of it.

Kui Ren
Columbia University
kr2002@columbia.edu

MS272

On Advantages of Strong-Form Local Meshless Methods for Phase-Field Modelling of Dendritic Solidification

The strong-form local meshless methods are applied to solve phase-field models for dendritic solidification. Dendrites, tree-like branched structures, are the most common microstructure in the solidification processing of metals. The microstructure evolution directly controls the properties of solidified material. Therefore, understanding dendritic solidification plays a crucial role in the quality assurance of the solidified product. The phase-field method represents the preferred technique for predicting dendritic microstructure under various casting conditions. The strong-form meshless methods turn out to be very suitable for solving phase-field models for dendritic growth. The solution of these models is prone to mesh-induced anisotropy, which can be in the meshless methods straightforwardly circumvented by using scattered node distributions. In the present contribution, the radial basis function generated finite difference method and the diffuse approximate method are tested on regular and scattered node distributions. The emphasis is given to the analysis of the growth in arbitrary preferential growth directions. The influence of node distribution type on the accuracy is thoroughly analysed for both methods. Additionally, a space-time adaptive algorithm developed to accelerate meshless calculations is presented. The algorithm dynamically ensures the highest density of computational nodes and the finest time-stepping at the solid-liquid interface.

Tadej Dobravec
Laboratory for Fluid Dynamics and Thermodynamics
University of Ljubljana
tadej.dobravec@fs.uni-lj.si

Boštjan Mavric
Department of Information Technology
Uppsala University
bostjan.mavric@it.uu.se

Boidar Šarler
University of Ljubljana, Faculty of Mechanical Engineering

bozidar.sarler@fs.uni-lj.si

MS272

A New Perspective on the Stability of RBF Methods for Conservation Laws

Radial basis function methods are powerful tools in numerical analysis and have demonstrated good properties in many different simulations. However, only a few stability results are known for RBF methods for hyperbolic conservation laws. In particular, when boundary conditions are included, stability issues frequently occur. We address how provable stability for RBF methods can be obtained. In particular, we discuss a new perspective on the stability of RBF methods using recent developments on summation-by-parts operators, often used in the finite difference and the finite element community.

Jan Glaubitz

Department of Mathematics, Dartmouth College
Jan.Glaubitz@Dartmouth.edu

MS272

Adapting Meshfree Galerkin Schemes for Representing Highly Anisotropic Fields

Methods for representing highly-anisotropic fields are presented, building on the element free Galerkin (EFG) and the flux-coordinate independent (FCI) methods. A mapping function is used to provide information about the local direction of the anisotropy, with one of the global coordinates chosen to parameterize the parallel position along the mapping in a one-to-one manner. Nodes are arranged on successive planes of constant parallel coordinate with fine spacing on each plane necessary to represent the small-scale variations perpendicular to the mapping direction, but with large spacings then possible between these planes because of the small variation along the mapping. This greatly reduces the number of degrees of freedom required to represent highly-anisotropic fields in this space and the associated computational cost of simulations involving such fields. The use of standard unstructured finite element meshes on each plane is explored, as well as fully-meshfree arrangements. No mesh connectivity is defined between planes in either case, and field-aligned basis functions are constructed using the mapping function to extend the bases into the full domain. Integration of the bases is addressed using variationally consistent integration (VCI), and the scheme (as well as other EFG schemes) is shown to be locally conservative under certain conditions. Robust convergence of several test problems is demonstrated.

Samuel A. Maloney, Ben McMillan

University of Warwick
sam.maloney@warwick.ac.uk,
b.f.mcmillan@warwick.ac.uk

MS272

Variably Scaled Kernels and Their Applications

This talk aims to provide an overview of variably scaled kernels (VSKs) and to highlight their applications. We mainly focus on RBF scattered data interpolation. VSKs work quite well in cases spoiled by a considerable instability of the standard method. They can significantly improve the recovery quality by preserving shape properties and particular features of the function underlying the given data. We also show how a proper choice of the scaling func-

tion can be an effective tool to approximate non-smooth functions.

Milvia F. Rossini

DIPARTIMENTO DI MATEMATICA E Applicazioni
University of Milano Bicocca, Italy
milvia.rossini@unimib.it

MS272

A Meshfree Geometric Multilevel Method for Systems Arising from Elliptic Equations on Point Cloud Surfaces

We present a new meshfree geometric multilevel (MGM) method for solving linear systems that arise from discretizing elliptic PDEs on surfaces represented by point clouds. These surface PDEs arise in many areas of science and engineering, including models of atmospheric flows, chemical signaling on cell membranes, morphogenesis, and textures for computer graphics. MGM uses a Poisson disk sampling-type technique for coarsening the point clouds and new meshfree restriction/interpolation operators based on polyharmonic splines for transferring information between the coarsened point clouds. These components are then combined with standard smoothing and operator coarsening methods in a V-cycle iteration. MGM is applicable to discretizations of elliptic PDEs based on various localized meshfree methods, including RBF finite differences (RBF-FD) and generalized finite differences (GFD). We test MGM both as a standalone solver and preconditioner for Krylov subspace methods on several test problems using RBF-FD and GFD, and numerically analyze convergence rates, efficiency, and scaling with increasing point cloud sizes. We also perform a side-by-side comparison to algebraic multigrid (AMG) methods for solving the same systems. Finally, we further demonstrate the effectiveness of MGM by applying it to three challenging applications on complicated surfaces: pattern formation, surface harmonics, and geodesic distance.

Grady B. Wright

Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

MS273

Approximation of Response Surfaces from Partial Information

We discuss how to use dimension reduction methods to organize unstructured ensembles of partial observations of a systems response surface. We assume to have no knowledge about where on the surface the observations are obtained. We do, however, record measurements not only at the point of observation, but also in its immediate vicinity. We demonstrate how such partial and disorganized observation ensembles can be integrated into coherent response surfaces whose dimension and parametrization can be systematically recovered in a data-driven fashion. The approach can be justified through the Whitney and Takens embedding theorems. We demonstrate our approach by organizing unstructured observations of a cusp bifurcation surface for Hydrogen combustion in a Continuous Stirred Tank Reactor. Finally, we discuss more generally which adjustments to common measurement procedures are necessary to obtain useful reconstructions.

Felix Dietrich

Technical University of Munich

felix.dietrich@tum.de

MS273

Generalized Quadratic Embeddings for Nonlinear Dynamical Systems

Dynamical modeling of a process is essential to study its dynamical behavior and perform engineering studies such as control and optimization. With the ease of accessibility of data, learning models directly from the data have recently drawn much attention. Constructing simple and compact models describing complex nonlinear dynamics is also desirable for efficient engineering studies on modest computer hardware. To achieve our goal, we focus on the lifting principles—that is, sufficiently smooth nonlinear systems can be rewritten as quadratic models in an appropriate coordinate system. Therefore, we focus on identifying suitable coordinate systems such that a quadratic model can describe the dynamics in the obtained coordinate system. To determine such a coordinate system, we leverage the powerful expressive capabilities of deep learning, particularly autoencoders. Moreover, in several physical systems where energy preservation is preserved, we focus on identifying the coordinate systems that can also preserve energy. We illustrate the methodologies to learn the desired coordinate systems for nonlinear dynamical models by illustrative examples.

Pawan Goyal
Max Planck Institute for
Dynamics of Complex Technical Systems Magdeburg
goyalp@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical
Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS273

Solving the Electronic Schrödinger Equation with Deep Learning

The electronic Schrödinger equation is a second-order partial differential equation for the wave function in $3N$ dimensions that underlies all quantum chemistry. The search space of its solutions grows exponentially with N , limiting exact methods to only the smallest molecules. Approximate quantum chemistry methods can be considered as reduced-order models that scale only polynomially with N and seek to strike a balance between accuracy and computational cost. Paulinet is one such model based on a physics-inspired deep-learning ansatz for the wave function that is optimized in a data-driven fashion within the framework of variational quantum Monte Carlo. It integrates three ingredients: exact physics as constraints, approximate physics as baselines, and neural networks as universal approximators. Paulinet competes in accuracy with state-of-the-art quantum chemistry methods while offering better scaling with system size.

Jan Hermann
Microsoft Research AI4Science
jan.hermann@microsoft.com

MS273

Model Reduction Based on the Koopman Genera-

tor

The determination of low-dimensional collective variables (CVs) is a critical problem in various application areas, including molecular dynamics. However, finding good CVs, for example in the sense of retaining slow dynamical time scales of the full system, is a highly non-trivial problem. In this talk, I will present recent progress on the determination of CVs using the projection formalism for stochastic dynamics developed by [Legoll and Lelivre, Nonlinearity (2010)] and [Zhang et al, Faraday Discuss. (2016)]. This formalism amounts to projecting the Koopman generator of the full system onto the space of functions which only depend on the CV space. The first result I will show is an error estimate comparing the slow time scales of the projected and the full system. We arrive at a Galerkin-type estimate, bounding the error in terms of the projection error for dominant eigenfunctions [Nske et al, Entropy (2021)]. Then, I will move on to show how a data-driven approximation of the projected generator can be obtained by a technique called generator extended dynamic mode decomposition (gEDMD) [Klus, Nske, et al, Physica D (2020)]. This method can also be formulated if a tensor product basis is used as Galerkin subspace [Lcke and Nske, J. Nonlinear Sci. (2022)], enabling the use of powerful approximation spaces. Finally, I will discuss bounds for the estimation error of the gEDMD method in terms of the amount of simulation data [Nske et al, arXiv:2108.07102 (2021)].

Feliks Nüske
Max Planck Institute for Dynamics of
Complex Technical Systems, Magdeburg, Germany
nueske@mpi-magdeburg.mpg.de

MS273

Data-Driven Prediction of Parametric Flows

Model reduction techniques enable to transform large-scale dynamical systems into systems of much smaller dimension. In general, such reduced systems are achieved by projecting the high fidelity equations onto a reduced subspace obtained by the Proper Orthogonal Decomposition. With varying parameters, such projected models may experience instabilities even if the original system is asymptotically stable. Moreover, the data may originate from experience making hard the application of Galerkin projection. To overcome these issues, we propose a data-driven approach to construct parametric reduced order models (DD/PROM). Regardless of the origin of the data, the DD/PROM is achieved by learning the model through residual assimilation. Then, interpolation on the tangent space of the Grassmann manifold is used to handle parameter variations. We show on the example of the flow past a cylinder, that the PROM/DD succeeds to reproduce the dynamics with a good accuracy, for a wide range of Reynolds number values.

Mourad Oulghelou
Arts et Métiers ParisTech
mourad.oulghelou@ensam.eu

MS274

Matrix Product State Simulation of the Burgers'-Fisher's Equation

The quantum-inspired, matrix product state (MPS) methodology is invoked to simulate the Burgers'-Fisher-KPP one-dimensional advection-diffusion-reaction equa-

tion. The MPS, developed for approximation of quantum many-body systems, exploits structures to capture the flow physics in a low-rank form (i.e., in a massively reduced state space). Here, the governing system is recast in the context of MPS, and is temporally simulated for various degrees of truncation. The generated MPS-reduced order solution is appraised against the results obtained via direct numerical simulation of the original equations. The computational complexity of MPS is assessed to determine its potential for simulating complex multi-scale turbulent reacting flows.

Peyman Givi
University of Pittsburgh, USA
Givi@engr.pitt.edu

Robert Pinkston
University of Pittsburgh
rop56@pitt.edu

Nikita Gourianov
University of Oxford
nikita.gourianov@physics.ox.a

Dieter Jaksch
University of Hamburg
dieter.jaksch@uni-hamburg.de

MS274

Exact and Efficient Lanczos Method on a Quantum Computer

We present an algorithm that uses block encoding on a quantum computer to exactly construct a Krylov space, which can be used as the basis for the Lanczos method to estimate extremal eigenvalues of Hamiltonians. While the classical Lanczos method has exponential cost in the system size to represent the Krylov states for quantum systems, our efficient quantum algorithm achieves this in polynomial time and memory. The construction presented is exact in the sense that the resulting Krylov space is identical to that of the Lanczos method, so the only approximation with respect to the Lanczos method is due to finite sample noise. This is possible because, unlike previous quantum versions of the Lanczos method, our algorithm does not require simulating real or imaginary time evolution. We provide an explicit error bound for the resulting ground state energy estimate in the presence of noise. For this method to be successful, the only requirement on the input problem is that the overlap of the initial state with the true ground state must be $\Omega(1/\text{poly}(n))$ for n qubits.

William Kirby
Tufts University
IBM Quantum
wmkirby1@gmail.com

MS274

Variational Quantum Algorithms for Solving the Diffusion and Burgers' Equation

We present variational quantum algorithms to solve general (linear and nonlinear) partial differential equations (PDEs) based on the FeynmanKitaev Hamiltonian formalism. The method was originally developed for the variational time evolution of quantum systems by Barison et al. [Phys. Rev. Research 4, 043161 (2022)]. It makes use of a clock quantum register to get the full time evolution with a single optimization routine. We generalize the approach to solve

classical PDEs and benchmark its performance on applications including diffusion and the creation of shockwaves in Burgers equation. Our analysis is carried out using a statevector simulator as well as the IBM Q System One quantum computer in Ehningen in Germany. Additionally we corroborate the validity of the variational framework for solving nonlinear PDEs on spacetime grids of various sizes via imaginary time evolution.

Albert J. Pool
German Aerospace Center
albert.pool@dlr.de

Alejandro D. Somoza
German Aerospace Center (DLR)
Germany
alejandro.somoza@dlr.de

Michael Lubasch
Quantinuum
United Kingdom
michael.lubasch@quantinuum.com

Birger Horstmann
German Aerospace Center (DLR)
Germany
birger.horstmann@dlr.de

MS274

Qclab++: Simulating Quantum Circuits on GPUs

Quantum circuit simulation is a powerful approach for rapid prototyping and development of quantum algorithms in the current era of noisy intermediate-scale quantum (NISQ) computers. QCLAB++ is an object-oriented, fully templated C++ package for creating and simulating quantum circuits developed at Lawrence Berkeley National Laboratory. QCLAB++ leverages the state vector approach for quantum simulation where the amplitudes of the quantum wave function are tracked and updated as if they evolve on a noiseless quantum computer. In this talk, we present the latest developments of porting the QCLAB++ code to GPUs and benchmark it on NERSC's Perlmutter system.

Roel Van Beeumen, Daan Camps
Lawrence Berkeley National Laboratory
rvanbeeumen@lbl.gov, dcamps@lbl.gov

MS275

Integrating Pool and Mathematics in Educational Setting is Viable - the Pilot Phase of the SmartPool in Bulgaria

The paper presents the SmartPool Erasmus+ project activities carried out in Bulgaria under the coordination of Georgi Bozhkov - chairman of the Bulgarian National Billiards Federation, and two mathematics education researchers from the Bulgarian Academy of Sciences. The main activities under consideration include: (i) enriching the school mathematics curriculum via tasks related to billiards and adapting the suggested SmartPool learning resources in harmony with the Bulgarian educational system; (ii) developing dynamic geometry models for demonstrations and explorations of mathematical ideas in a billiard context, and (iii) discussing the adapted learning resources with the teachers, selected to carry out the pilot SmartPool sessions with groups of students from the town of Petrich

(Bulgaria), together with leaders and coaches from the local BO&BO Billiard Club.

Toni Chehlarova
IMI-BAS
toni.chehlarova@gmail.com

Evgenia Y. Sendova
Institute of Mathematics and Informatics
Bulgarian Academy of Science
jenny.sendova@gmail.com

Georgi Bozhkov, Radoslav Milkov
BO&BO Billiard Club, Bulgaria
bozhkov84@yahoo.com, quemaster90@gmail.com

MS275

From Doing Math Research at School Age to Doing Math in Your Professional Life

The paper presents the authors long-term experience as a tutor/mentor in two research summer schools - the High School Student Institute of Mathematics and Informatics (HSSI) in Bulgaria [O. Mushkarov, N. Dimitrova, E. Sendova, Math Research at School Age. In: Meeting in Mathematics (Eds. Georgiev et al.), Sofia, 2008, 8193] and the Research Science Institute (RSI), organized jointly by CEE (<https://www.cee.org/>) and MIT (<https://www.mit.edu/>) in the USA [E. Sendova, You do you understand; you explore you invent: the fourth level of the inquiry-based learning. In: Constructionism and Creativity (Eds. G. Futschek, C. Kynigos) Proceedings of the 3d International Constructionism Conference, August 1923, Vienna, Austria, 2014, 103112]. The similarities and the differences are discussed with a focus on the importance of doing research in mathematics at school age and presenting it in different settings for the students to choose mathematics as the main component in their professional life. Case studies of several alumni from both program who have chosen math as their profession will be presented and discussed.

Evgenia Y. Sendova
Institute of Mathematics and Informatics
Bulgarian Academy of Science
jenny.sendova@gmail.com

MS275

The Math Joy in the Smartpool Curriculum

In pursuing a synergy between mathematics and pool billiard a Smartpool curriculum was solicited at the start of 2019 based on the authors experience as a math teacher and former pool player. The presentation deals with the mathematical foundation of the Smartpool curriculum and the philosophy that such a synergy will manifest math joy. An insight in the psychology of stimulating positive engagement in mathematics through playful active explorations will be demonstrated - from investigating where to theoretically hit a ball, through the use of geometry in measuring angles, areas, distances, to the applicability of this knowledge to help practically execute shots in an efficient manner on the table. The psychological benefits of implementing this program in an educational system will be discussed with an emphasis on its effects for reducing math anxiety.

Erik van Haren
Mathplay.eu

mathplay.eu@gmail.com

MS276

Sparse Spectral Methods for Power Law Kernels

We review recently derived sparse spectral methods for integral equations involving power law kernels, including the case of multivariate disks/balls. The key is new recurrence relationships of power law kernels applied to weighted orthogonal polynomials. Applications include equilibrium measures associated with attractive-repulsive dynamics.

Sheehan Olver
Imperial College London
s.olver@imperial.ac.uk

Timon Gutleb, José A. Carrillo
University of Oxford

timon.gutleb@maths.ox.ac.uk, carrillo@maths.ox.ac.uk

MS276

High-Order Solutions of Volumetric Transmission Scattering Problems in Linear Elasticity

We present a numerical scheme based on integral equations for the solution of two dimensional volumetric transmission scattering problems in linear elasticity. Our method relies on the Helmholtz decomposition of elastic fields and on recently introduced polynomial density interpolation methods (DIM) for the high-order discretization of volumetric Helmholtz potentials. The volumetric DIMs, in turn, rely on representations of the scatterer as a union of quadrilateral subdomains and on subdomain/local polynomial solutions of the Helmholtz operator that act as regularizers for the singular volumetric potentials. We present a variety of numerical results that showcase the high-order convergence of our volumetric integral solvers.

Catalin Turc
Department of Mathematical Sciences
New Jersey Institute of Technology
catalin.c.turc@njit.edu

MS277

HpVersion Discontinuous Galerkin Methods on Essentially Arbitrarily-Shaped Elements

We extend the applicability of the popular discontinuous Galerkin method discretizing advection-diffusion-reaction problems to meshes comprising extremely general, essentially arbitrarily-shaped element shapes. In particular, our analysis allows for curved element shapes without the use of non-linear elemental maps. The feasibility of the method relies on the definition of a suitable choice of discontinuity penalization, which turns out to be explicitly dependent on the particular element shape but essentially independent of small shape variations. This is achieved upon proving extensions of classical trace and Markov-type inverse estimates to arbitrary element shapes. A further new $H^1 - L_2$ -type inverse estimate on essentially arbitrary element shapes enables the proof of inf-sup stability of the method in a streamline-diffusion-like norm. These inverse estimates may be of independent interest. A priori error bounds for the resulting method are given under very mild structural assumptions restricting the magnitude of the local curvature of element boundaries. Numerical experiments are also presented, indicating the practicality of the proposed

approach.

Andrea Cangiani
SISSA
andrea.cangiani@sissa.it

Zhaonan Dong
Inria Paris
zhaonan.dong@inria.fr

Emmanuel Georgoulis
University of Leicester
National Technical University of Athens & IACM-FORTH
eg64@leicester.ac.uk

MS277

Space-Time Embedded/hybridized DG for the Navier-Stokes Problem

Partial differential equations posed on moving domains arise in many applications such as air turbine modeling, flow past airplane wings, etc. The time-dependent nature of the flow domain poses an additional challenge when devising numerical methods for the discretization of such problems. One alternative when dealing with time-dependent domains is to pose the problem on a space-time domain and apply, for example, a finite element method in both space and time. These space-time methods can easily handle the time-dependent nature of the domain. In this talk, we present a space-time hybridizable discontinuous Galerkin method for the discretization of the incompressible Navier-Stokes equations on moving domains. This discretization is pointwise mass conserving and pressure robust, even on time-dependent domains. Moreover, high order can be achieved both in space and time. Numerical experiments will demonstrate the capabilities of the method.

Tamas Horvath
Oakland University
Rochester, MI
thorvath@oakland.edu

MS277

A Discontinuous Galerkin Level Set Reinitialization with Subcell Stabilization on Tetrahedral Meshes

In this study, we present a 3D level set reinitialization based on a high-order, local discontinuous Galerkin method on unstructured tetrahedral meshes. The flow of time Eikonal equation is discretized to construct an approximate signed distance function instead of the standard hyperbolic level set reinitialization. Using the Eikonal equation removes the regularization parameter in the standard approach which allows more predictable behavior and faster convergence speeds around the interface. A finite volume-based subcell stabilization is used to improve the nonlinear stability of the approach. A set of numerical experiments with smooth and non-smooth interfaces show that the method is stable for highly distorted interfaces and experimentally achieves design order accuracy.

Ali Karakus
Middle East Technical University

akarakus@metu.edu.tr

MS277

Coupled 1D/2D Discontinuous Galerkin Methods for Modeling Overland and Open-Channel Flows

In this talk, I will present on the continued development and application of a multidimensional (1D/2D), unstructured-mesh model for simulating coupled overland/open-channel flows. While both overland and open-channel flows can be modeled using the full shallow water equations, a simpler approximation to the shallow water equations — namely, the kinematic wave approximation — has been shown to accurately represent overland and open-channel flows under conditions often encountered in natural watersheds. Our modeling approach solves the kinematic wave equations resulting from this approximation making use of discontinuous Galerkin (DG) finite element spatial discretizations of variable polynomial degree p , paired with explicit RungeKutta time steppers, and is supported by advancements made to an automatic mesh generation tool, ADMESH+ , that is used to construct constrained triangulations for channel routing. The developed modeling framework is applied to a number of test cases that demonstrate the accuracy and robustness of our approach and highlight the potential benefits of using p (polynomial) refinement for hydrological simulations.

Ethan Kubatko
Department of Civil, Environmental and Geodetic Engineering
The Ohio State University
kubatko.3@osu.edu

MS277

A Sequential Discontinuous Galerkin Method for Two-Phase Flow in Deformable Porous Media

We formulate a numerical method for solving the two-phase flow poroelasticity equations in three dimensions. The scheme employs the interior penalty discontinuous Galerkin method and a sequential time-stepping method. The numerical results show the accuracy and robustness of the proposed method.

Boqian Shen
KAUST
boqian.shen@kaust.edu.sa

MS278

Code Generation for the Simulation of Ocean Flows on Block-Structured Grids

Stencil computations and their application in various scientific domains have been an active field of research for the past decades. We present our work on utilizing stencil technologies in complex application scenarios transcending simple benchmark problems. One such domain is the simulation of ocean flows which requires support for real-world geometries, advanced numerical components, as well as domain-specific boundary conditions and source terms. Based on an early code version, we first transition from uniform to block-structured grids. They are set up automatically by coarsely meshing a given geometry before applying multiple refinement steps and nodal position optimization. Next, we extend for higher-order quadrature-free discontinuous Galerkin (DG) discretizations. The resulting compute

kernels are more complex than traditional stencil codes but share beneficial characteristics such as vectorizability, consecutive memory accesses, and suitability for GPUs. To tackle performance portability and the high number of possible implementation variants, we employ code generation techniques. The ExaStencils framework serves as an excellent base as it is capable of automatically optimizing stencil-like codes, as well as of parallelizing them with MPI, OpenMP and CUDA. Its Python-based front end GHOD-DESS adds the necessary support for DG discretizations, and further increases performance through domain-specific optimizations on the algorithmic level.

Sebastian Kuckuk

Friedrich-Alexander-Universität Erlangen-Nürnberg
Erlangen National High Performance Computing Center
sebastian.kuckuk@fau.de

MS278

Stencil Computations in Petroleum Industry

To Be Defined..

Ali Momin

Aramco
ali.almomin@aramco.com

MS278

Accelerated Stencil Computations on NVIDIA H100

GPUs are very popular to compute 3D finite difference stencils. Compared to naive implementations, optimized Stencil kernels can provide a significant speedup. Paulius Micikevicius described an optimized implementation on early CUDA GPUs using a mix of register queues and shared memory. Although the method described still apply today, we can apply a few modifications to keep these stencil codes running as fast as possible on the latest architectures. Today's GPUs are significantly more powerful than the NVIDIA Tesla S1070 used in the original paper. The number of compute cores has grown significantly, Tensor Cores are available to accelerate GEMMs, the memory bandwidth numbers have improved by an order of magnitude with the adoption of HBM, the L2 caches are now significantly larger, and newer instructions have been introduced to help move data in and out of the symmetrical multiprocessor (SM) units more efficiently. The memory bandwidth has also grown significantly, and is often the main bottleneck for these problems. Finite difference stencil codes are typically bandwidth-bound, and the goal for stencil codes should be to move the minimum amount of data, at speeds as close as possible to the theoretical peak bandwidth. In this talk we suggest a new method of Stencil Computation with some experimentation on the latest NVIDIA GPU, H100.

Issam Said, Guillaume Thomas-Collignon

Nvidia
isaid@nvidia.com, guillaumet@nvidia.com

MS278

Difference Potentials Method Based on Local Basis Functions for Elliptic Partial Differential Equations in Arbitrary Geometry

We develop efficient and high-order accurate finite difference methods for elliptic partial differential equations in complex geometry in the Difference Potentials framework.

The main novelty in the developed methods is the use of local basis functions defined at near-boundary grid points. The use of local basis functions allow unified numerical treatment of (i) explicitly and implicitly defined geometry; (ii) geometry of more complicated shapes, such as those with corners, multi-connected domain, etc; and (iii) different types of boundary conditions. This geometrically flexible approach with local basis functions is a complement and improvement to the traditional DPM with global basis function approaches, especially in the case where a large number of global basis functions are needed to resolve the boundary data, or where the optimal global basis functions are difficult to obtain. Another feature of this work is that fast Poisson solvers based on FFT can be employed for standard centered finite difference stencils regardless of the designed order of accuracy (2nd, 4th, 6th, etc.). The techniques developed in this paper can also be applied to the numerical solutions of time-dependent partial differential equations, when implicit time stepping is employed.

Qing Xia

KTH Royal Institute of Technology
qingx@kth.se

MS278

Revisiting Stencil Optimizations as the Architecture Gap Between GPUs and CPUs Shrinks

GPUs being traditionally throughput-optimized while CPUs being latency-optimized has been reflected in how stencils are optimized for each architecture. In particular, temporal blocking has been more geared towards CPUs due to the latency penalty of deep temporal blocking on GPUs. On the other hand embarrassingly parallel partial accumulation optimizations have been more suitable for GPUs. The latency of all operations across the spectrum on GPUs drop, while HPC-grade CPUs (ex: A64FX) become equipped with high bandwidth memory and wide vectors. As a results, in some aspects, the gap between GPUs and CPUs is shrinking. In this talk we first describe how GPUs have evolved in the last decade in a way that can make them more accommodating of stencil optimizations that favor low-latency architectures. We also describe the reverse effect: how some HPC-grade CPUs are more accommodating of stencil optimization that favor high-throughput architectures.

Mohamed Wahib

AIST-TokyoTech Open Innovation Lab, Japan
mohamed.attia@riken.jp

Lingqi Zhang

TokyoTech
cslingqizhang@gmail.com

Satoshi Matsuoka

Riken R-CCS and
Tokyo Institute of Technology
matsu@is.titech.ac.jp

MS279

An Implicit, Conservative, Asymptotic-Preserving Electrostatic Particle-in-Cell Algorithm for Arbitrarily Magnetized Plasmas

We discuss a new full-orbit electrostatic particle-in-cell algorithm able to use large timesteps compared to particle gyro-period in arbitrary magnetic fields [Chen et al. JCP,

submitted (2022)]. The algorithm extends earlier electrostatic fully implicit PIC implementations [Chen et al., JCP 230 (2011)] with a new asymptotic-preserving (AP) particle-push scheme [Ricketson et al., JCP 418 (2020)] (recently endowed with a very fast Picard nonlinear solver [Koshkarov et al., JCP 459 (2022)]) that allows timesteps much larger than particle gyroperiods. The AP integrator preserves all the averaged particle drifts in the large-timestep limit, while recovering resolved particle orbits with small timesteps. The scheme allows for a seamless, efficient treatment of particles in coexisting magnetized and unmagnetized regions, conserves energy and charge exactly, and does not spoil implicit solver performance. We demonstrate by numerical experiment with several strongly magnetized problems (diocotron instability, modified two-stream instability, and drift-wave instability) that two orders of magnitude wall-clock-time speedups are possible vs. the standard fully implicit electrostatic PIC algorithm without sacrificing solution quality and while preserving strict charge and energy conservation.

Luis Chacon, Guangye Chen
Los Alamos National Laboratory
chacon@lanl.gov, gchen@lanl.gov

Oleksandr Koshkarov
LANL
koshkarov@lanl.gov

Lee Ricketson
Lawrence Livermore National Laboratory
ricketson1@llnl.gov

MS279

Mixed Model/Precision Methods

Additive Runge-Kutta methods designed for preserving highly accurate solutions in mixed-precision and mixed-model computation were proposed and analyzed in a recent paper by Z. Grant. These specially designed methods use reduced precision for the implicit computations and full precision for the explicit computations. We present a stability analysis and explore convergence, accuracy, runtime, and energy consumption of these methods on test problems. We show that these MP-ARK methods efficiently produce accurate solutions with significant reductions in runtime (and by extension energy consumption).

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

MS279

Approximate Entropy-Based Moment Closures

Despite their elegant mathematical structure, entropy-based moment closures face severe implementation challenges that have limited their wide-spread use. In this talk, I will present two approximations that attempt to address some of these challenges. The first approximation relies on a regularization of the optimization problem that defines the original entropy-based closure. The main advantage of the regularization is that moment vectors need not take on traditional realizable values. However, the resulting equations still retain many important structural features, such as hyperbolicity and an entropy dissipation law. These results reveal the moment entropy as a key tool in constructing approximate closures and motivate a second approx-

imation of the entropy-based closure that is constructed via a convex fit of the moment entropy. We illustrate this approach with numerical examples from a simple linear transport equation.

Cory Hauck
Oak Ridge National Laboratory and University of Tennessee
hauckc@ornl.gov

Graham Alldredge
Berlin
gwak@posteo.net

Martin Frank
Karlsruhe Institute of Technology
martin.frank@kit.edu

Paul Laiu
Oak Ridge National Laboratory
laiump@ornl.gov

William Porteous
University of Texas Austin
afpwilliam@gmail.com

Steffen Schotthöfer
Karlsruhe Institut of Technology
steffen.schotthoef@kit.edu

MS279

Parallel in Time Methods for Plasmas

Kinetic plasma simulations must reckon with many widely varying time-scales and are routinely stability-limited, requiring very many timesteps to resolve the interesting physics. This makes them an ideal candidate for parallel-in-time (PIT) strategies. Unfortunately, due to their extreme cost, these systems are simulated at the minimum threshold of spatial and temporal resolution, thus rendering the usual PIT strategy of spatiotemporal coarsening infeasible. Here we present a parallel-in-time strategy that employs a coarsening of models, rather than meshes, to accelerate time to solution for high-dimensional kinetic plasma models by leveraging their much cheaper, lower dimensional fluid approximations. These fluid models are closed via kinetic data computed in parallel, and are thus applicable far from the fluid limit. We present initial results on the two-stream instability.

Paul Tranquilli
Lawrence Livermore National Lab
tranquilli1@llnl.gov

Lee Ricketson, Terry Haut
Lawrence Livermore National Laboratory
ricketson1@llnl.gov, haut3@llnl.gov

Jeffrey A. Hittinger
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
hittinger1@llnl.gov

MS280

FleCSI: A Programming System for the Future

FleCSI is a task-based programming system designed to address the challenges presented by modern heterogeneous

system architectures. FleCSI supports a rich feature set, including distributed and shared-memory data parallelism; task parallelism; portability across processor architectures and deep memory hierarchies; customizable data structure and topology interfaces; and powerful, high-level mapping capabilities for advanced algorithm development. Given this set of features, applications built on FleCSI are well-prepared to run on the diverse set of supercomputing systems that are being deployed around the world. This presentation will review some of FleCSI's capabilities and how they enable application developers to take advantage of modern computing systems.

Ben Bergen
Los Alamos National Laboratory
bergen@lanl.gov

MS280

Scaling Starpu's Task-Based Sequential Task Flow (STF) Inside Nodes and Between Nodes

To be able to get the best performance out of nowadays' complex architectures comprising GPUs, NICs, disks, etc., the HPC field is leaning more and more to task-based programming instead of e.g. thread-based programming approaches. This approach indeed allows to delegate most of the optimization work to a runtime system: computation scheduling, transferring data concurrently with task execution, data prefetching from disks to GPUs, network communications, etc. The StarPU runtime system has been proposing such a task-based programming paradigm, and promotes the use of the now-quite-common Sequential Task Flow (STF) programming style in which the main code of the application is a loop nest of task submissions which actually looks like a sequential source code. StarPU can then optimize the execution of the resulting task graph inside machines comprising CPUs, GPUs, FPGAs, disks... The StarPU-MPI layer extends this to clusters by automating inter-node communications, allowing to keep the same application source code for both the debuggable sequential version and the scalable distributed heterogeneous version. The obtained performance outperforms the ScaLAPACK implementation, and is on par with the PARSEC performance. Using the example of the SUMMA algorithms variants, we show that we can have both a very simple, easy-to-maintain application code, and flexibility to adapt the algorithm to scale with different applicative situations.

Samuel Thibault
LaBRI
samuel.thibault@u-bordeaux.fr

MS281

A Stochastic Derivative-Free Trust-Region Method in Random Subspaces

This work proposes a framework for large-scale stochastic derivative-free optimization (DFO) by introducing STARS, a trust-region method which achieves scalability using random models in random low-dimensional affine subspaces. STARS significantly reduces per-iteration costs in terms of function evaluations, thus yielding strong performance on large-scale stochastic DFO problems. The user-determined dimension of these subspaces can be chosen via so-called Johnson-Lindenstrauss transforms, and independently of the dimension of the problem. For convergence purposes, both a particular quality of the subspace and the accuracies of random function estimates and models are required to hold with sufficiently high, but fixed, probabilities. Con-

vergence and expected complexity results of STARS are obtained using martingale theory.

Kwassi Joseph Dzahini
Argonne National Laboratory
kdzahini@anl.gov

Stefan Wild
Argonne National Laboratory
LANS, Mathematics and Computer Science
wild@lbl.gov

MS281

Solar: A Solar Thermal Power Plant Simulator for Blackbox Optimization Benchmarking

This work presents SOLAR, a collection of optimization problems provided as a benchmarking tool for blackbox optimization solvers. Each problem optimizes the design of a concentrated solar power plant defined as a blackbox numerical model. The type of variables, dimensionality, and number and type of constraints are different across problems. Optimization may be single or biobjective. The solar plant model considers several subsystems: a heliostats field, a central cavity receiver, a molten salt thermal energy storage, a steam generator and an idealized power block. Benchmark optimization results are provided using the NOMAD software package.

Sébastien Le Digabel
Polytechnique Montréal
Sebastien.Le.Digabel@gerad.ca

MS281

COBYQA A Derivative-Free Trust-Region SQP Method for Nonlinearly Constrained Optimization

This talk introduces COBYQA, a derivative-free trust-region SQP method for general nonlinear optimization problems. The method builds the trust-region quadratic models using Powell's derivative-free symmetric Broyden update. An important feature of COBYQA is that it always respects bound constraints if any. COBYQA is competitive with NEWUOA and BOBYQA while being able to handle more general problems. On linearly constrained problems, COBYQA outperforms LINCOA if the problems also contain bound constraints that cannot be violated. Most importantly, COBYQA outperforms COBYLA on all types of problems, no matter whether bound constraints (if any) can be violated or not. COBYQA is implemented in Python and publicly available at <https://www.cobyqa.com/>. This is joint work with Zaikun Zhang.

Tom M. Ragonneau
The Hong Kong Polytechnic University
tom.ragonneau@polyu.edu.hk

Zaikun Zhang
Hong Kong Polytechnic University
zaikun.zhang@polyu.edu.hk

MS281

Learning Hidden Constraints Using a Stepwise Uncertainty Reduction Strategy with Gaussian Process Classifiers

Design optimization of engineering systems generally in-

volves the use of complex numerical models that take as input design variables and environmental variables. The environmental conditions are generally simulated to assess the reliability of the proposed designs. Hence the simulations can become computationally very expensive. Moreover, some input conditions can lead to simulation failures or instabilities, due, for instance, to convergence issues of the numerical scheme of complex partial derivative equations. Most of the time, the set of inputs corresponding to failures is not known a priori and corresponds to a hidden constraint, also called crash constraint in this special case. Since the observation of a simulation failure might be as costly as a feasible simulation, we seek to learn the feasible set of inputs and thus target areas without simulation failure during the optimization process. Therefore, we propose a Gaussian Process Classifiers (GPC) active learning method to learn the feasible domain. The proposed methodology is an adaptation of Stepwise Uncertainty Reduction strategies usually used with Gaussian Process Regression, in the classification setting with GPC models. The performance of this strategy on toy problems and on a wind turbine design application will be presented.

Morgane Menz
IFP Energies nouvelles
morgane.menz@ifpen.fr

Miguel Munoz Zuniga
IFP Energies Nouvelles
miguel.munoz-zuniga@ifpen.fr

Delphine Sinoquet
IFP Energies nouvelles
delphine.sinoquet@ifpen.fr

MS282

Building on Communities to Further Software Sustainability

Software sustainability is increasingly recognised as a must in scientific research. Well-established tools for source code management, for software package (and dependency) management and/or application-level virtualisation (containerisation) are often instrumental in the development of software. No matter how good these tools alone can hardly ensure long-term sustainability. Teams and communities are equally or even more important than any technical solution chosen. Building community practices to ensure the long-term sustainability of software is rarely given priority in the context of scientific projects which heavily rely on them nonetheless. Co-developing best practices with communities proves to lead to a wider/better adoption, more efficient work/developments, higher quality documentation and training material/tutorials, a much flatter learning curve which makes it easier onboarding newcomers, and all that translates into a better take up by researchers. In this talk, the Nordic Infrastructure Community for Earth System Modelling Tools will explain how they have placed people at the centre of the software development cycle to ensure its long-term sustainability. Successful examples will be presented as well as failures, to highlight what went well or wrong, respectively, and why. It will be discussed in particular the importance to go beyond its own team and community, and to think about a scientific project in a wider context e.g. Open Science.

Anne Fouilloux
Simula Research Laboratory

annef@simula.no

MS282

How RSE Community Identity Leads to Improved Research Software Practices

Research Software Engineers (RSEs) have developed a strong sense of community identity. This identity includes a culture and set of practices that help define the community and its membership. Over time, some members of the community changes practices (to new best practices), and the community identity encourages these practices to be taken up by others in the community fairly quickly. People outside the community see these changes, and then also take up these practices. This talk will discuss this process, including examples such as practices around documentation and testing, and explain how these lead to more sustainable research software.

Daniel S. Katz
University of Illinois at Urbana-Champaign
d.katz@ieee.org

MS282

Challenges of Integrating Workflow Software with Multifactor Authentication: How can RSEs Work with System Administrators and Other Stakeholders to Enable Unattended Remote Access?

Workflow tools (e.g. AiiDA) are becoming more and more popular with researchers working on large-scale HPC systems due to their ability to simplify the management of more complex research requirements and improve the productivity of researchers. However, the increased focus of security on academic HPC systems is leading to multifactor authentication (MFA) processes, usually employing a time-based one-time password (TOPT), becoming commonplace. This causes a large problem for both researchers who wish to use workflow tools which need unattended access to HPC facilities, and for HPC service providers who want to enable research but improve system security. In this presentation I will briefly set the context for this issue before presenting some approaches that have been taken to enable such functionality on HPC services with MFA. I will conclude with showing why solving the problem sustainably requires both workflow software and HPC services to be designed with input from all parties; and why RSEs are the key personnel to enabling long-term, sustainable solutions to enable complex research workflows in a secure manner and improve researcher and HPC service productivity.

Andrew Turner
EPCC, UK
a.turner@epcc.ed.ac.uk

MS282

Are You An RSE?

Did you start off as a researcher who spent time developing software to progress your research and then realised you enjoyed this work more than research? Or maybe you started off from a more conventional software-development background and are drawn to research by the challenge of using software to further research? A growing number of people in academia combine expertise in programming with an intricate understanding of research. Although this combination of skills is extremely valuable, these people lack a

formal place in the academic system. Without a job title, it is difficult for people to rally around a cause and there is no easy way to recognise their contribution, to reward them, or to represent their views. So in 2012, the term Research Software Engineer was created in 2012 and the community is now working to raise awareness of the role, bring RSEs together and advocate for more appropriate career recognition and promotion. Regardless of your formal job title, the attraction of being a Research Software Engineer is similar to that of being a researcher: the role attracts people who want to know how the world works. The Research Software Engineer works with researchers to gain an understanding of the problems they face, and then develops, maintains and extends software to provide the answers. This talk will cover the history of the RSE movement, what it has achieved, what's coming up in 2023 and how you can get involved.

Claire Wyatt
Jülich Supercomputing Centre
clairelouise_wyatt@yahoo.co.uk

MS283

Coupled Parameter and Data Dimension Reduction for Bayesian Inference

In this talk, we introduce a new method to jointly reduce the dimension of the parameter and the data space of high-dimensional Bayesian inverse problems. Recent dimension reduction methods commonly focus on the parameter space, yet many geophysical applications equally necessitate dimension reduction in the data space, which is done in practice through an independent upstream procedure. A recent work [1] considers dimension reduction for both data and parameters albeit treating the two spaces separately. However, choosing a low-dimensional informed parameter subspace influences which data subspace is informative and vice versa. We thus propose a coupled dimension reduction method based on the gradient of the generalised forward operator that maps between the parameter and data spaces. Our method computes in an offline phase a projector for a ridge approximation of the likelihood function to accelerate online multi-query posterior evaluations. We also show how our method can aid experimental design by localising effective sensor placements. Numerical experiments on a 2D shallow water model illustrate the benefits of our proposed method.

Qiao Chen, Olivier Zahm
Univ. Grenoble Alpes, Inria
qiao.chen@inria.fr, olivier.zahm@inria.fr

Elise Arnaud
Grenobles Alpes university
INRIA
elise.arnaud@univ-grenoble-alpes.fr

Ricardo Baptista
MIT
rsb@caltech.edu

MS283

Grassmann Stein Variational Gradient Descent

Stein variational gradient descent (SVGD) is a deterministic particle inference algorithm that provides an efficient alternative to Markov chain Monte Carlo. However, SVGD has been found to suffer from variance underestimation when the dimensionality of the target distribution is high.

Recent developments have advocated projecting both the score function and the data onto real lines to sidestep this issue, although this can severely overestimate the model uncertainty. In this talk, we propose Grassmann Stein variational gradient descent (GSVGD) as an alternative approach, which permits projections onto arbitrary dimensional subspaces and which promotes efficient state-space exploration in high-dimensional problems that have an intrinsic low-dimensional structure.

Andrew Duncan, Xing Liu, Harrison Zhu
Imperial College London
a.duncan@imperial.ac.uk, xing.liu16@imperial.ac.uk,
hbz15@ic.ac.uk

Jean-François Ton
University of Oxford
jean-francois.ton@stats.ox.ac.uk

George Wynne
Imperial College London
g.wynne18@ic.ac.uk

MS283

Influence of Monte-Carlo Sampling on the Convergence Rates of Greedy Algorithms for Reduced-Basis Methods

In this talk will be presented recent results about the mathematical study of a variance reduction technique for the computation of parameter-dependent expectations using a reduced basis paradigm. We study the effect of Monte-Carlo sampling on the theoretical properties of greedy algorithms. In particular, using concentration inequalities for the empirical measure in Wasserstein distance, we provide sufficient conditions on the number of samples used for the computation of empirical variances at each iteration of the greedy procedure to guarantee that the resulting method algorithm is a weak greedy algorithm with high probability. These theoretical results are not fully practical and we therefore propose a heuristic procedure to choose the number of Monte-Carlo samples at each iteration, inspired from this theoretical study, which provides satisfactory results on several numerical test cases.

Virginie Ehrlacher
CERMICS - ENPC, Paris
virginie.ehrlacher@enpc.fr

Mohamed-Raed Blel
CERMICS - ENPC, Paris
mohamed-raed.blel@enpc.fr

Tony Lelievre
Ecole des Ponts ParisTech
lelievre@cermics.enpc.fr

MS284

Exploiting Hierarchical Low-Rank Structure in Ice Sheet Inverse Problems

In this talk, we discuss the problem of estimating solutions of large-scale inverse problems governed by partial differential equation (PDE) based models and the associated uncertainties. The ice flow is modeled as a three-dimensional, viscous, incompressible fluid via the nonlinear (first-order) Stokes equations. Solving inverse problems with expensive forward models and high-dimensional parameters is com-

putationally challenging. Here, we exploit the local sensitivity of model predictions to parameters which suggests that the Hessian has off-diagonal low-rank structure, and so therefore build hierarchical matrix Hessian approximations. We demonstrate that for problems with sufficiently-informative data, hierarchical matrices provide for a more computationally efficient approximation format than the related global low-rank format. This motivates the use of hierarchical matrices for continental-scale problems. Presented computational results are generated using PyAlbany, a Python interface to the C++ multiphysics library Albany. Time permitting, we outline a framework that utilizes Hessian impulse responses of the Hessian and a physically motivated impulse response interpolation scheme, to further reduce the computational cost of generating hierarchical matrix Hessian approximations. Said framework is employed to generate inexact Newton-CG preconditioners, which ultimately reduce the number of PDE solves needed to estimate the inverse problem solution.

Tucker Hartland

University of California, Merced
hartland1@llnl.gov

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Kim Liegeois
Sandia National Laboratories
knliege@sandia.gov

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

Nicholas Alger
The University of Texas at Austin
Oden Institute
nalger225@gmail.com

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

MS284

Computational Efficient Estimation of the Extreme Event Probability of the Mass Loss of Antarctic Ice Sheets

The mass loss of the Antarctic ice sheets is one of the major causes of the sea level rise. Future projections of sea level rise have a large variability due to uncertainties in the climate models and socioeconomic factors, and it is important to study low-probability but high-risk scenarios. Standard Monte Carlo methods may require an infeasible large number of expensive forward simulations to obtain useful estimate of extreme events probability. In this presentation, we explore forward uncertainty quantification methods that rely on optimization strategies to estimate these probabilities. The proposed UQ methods build on ideas from PDE-constrained optimization and the probability estimation of tail probabilities. We illustrate that this approach is computationally more efficient than us-

ing standard Monte Carlo methods with larger number of expensive forward simulations when estimating the probability of extreme mass loss of the Antarctic ice sheets.

Kim Liegeois

Sandia National Laboratories
knliege@sandia.gov

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS284

Designing Iterative Ice Flow Solvers for GPU Supercomputers

The current climate change triggers a critical acceleration of ice loss at global scale, may it happen in Antarctica and Greenland or in mountain regions such as the Alps. This situation ultimately poses a threat for coastal regions affected by sea-level rise or in valley-glaciers subject to surges and collapses of ice masses. Numerical modelling permits to understand and predict the evolution of these complex natural systems resolving the underlying physical processes in three dimensions. Moreover, high spatial and temporal resolution is crucial to capture rapid changes in the system leading to the formation of, e.g., ice streams or collapse features. We here present innovative software tools which provide a way forward in ice dynamics and computational Earth sciences by exploiting massively parallel supercomputing on graphics processing units (GPUs) and scalable iterative solvers. We use Julia language because it features high-level capabilities without compromising high performance and portability amongst multiple backends (e.g., multi-core CPUs, and NVIDIA and AMD GPUs). We will compare different Julia multi-GPU full-Stokes solvers for ice flow with thermo-mechanical coupling and discuss the performance of various discretisation and boundary conditions approaches as well as compare the efficiency of the solvers.

Ludovic Räss, Ivan Utkin

Lab of Hydraulics, Hydrology and Glaciology (VAW)
ETH Zurich
luraess@ethz.ch, iutkin@ethz.ch

MS284

Performance Portable Extensions to Ice-Sheet Modeling in Mali

The road to exascale computing has introduced a need to redesign our scientific software for new computing architectures that require a fundamentally different programming model. In an effort to avoid architecture specific programming and maintain productivity across platforms, the ice-sheet modeling code known as MALI uses high level abstractions to integrate Trilinos libraries and the Kokkos programming model for performance portable code across a variety of different architectures. By utilizing packages within Trilinos and the Kokkos programming model, a significant amount of progress has been made towards developing a framework that supports high resolution simulations of ice sheets on pre-exascale systems and beyond. This presentation focuses on the key performance devel-

opments and future performance goals of developing and maintaining a modern, robust and scalable land ice solver.

Jerry Watkins, Max Carlson, Irina K. Tezaur
Sandia National Laboratories
jwatkin@sandia.gov, maxcarl@sandia.gov,
ikalash@sandia.gov

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

MS285

Modeling Stochasticity In Coupled Dynamical Systems

From space weather, to climate, considerable efforts are being made to improve the accuracy of the physics models. These models have complex interactions between physical phenomena evolving at different scales. Reliability of such multi-physics models remains a grand challenge. When physical phenomena interact, the predictions from such models get polluted and the inherent uncertainties are compounded, thus reducing their reliability. To tackle this problem, we look at a novel approach for uncertainty quantification (UQ) in coupled systems by learning the mathematical operators, called the Perron-Frobenius and Koopman operators, that describe the propagation of uncertainties in such models. While scale bridging of multi-scale physics has received considerable attention, UQ of multi-physics remain a challenge. Existing methods for UQ don't quantify the stochastic nature of interaction in multi-physics models. The key scientific contributions of our work is the development of a novel UQ method that 1) quantifies the stochastic nature of coupling, 2) is non-intrusive and, 3) provides both forward and backward UQ in coupled systems. We achieve this by leveraging the Perron-Frobenius and Koopman Operators that act on functions of the model state and describe the evolution of probability densities and observables forward and backward in time respectively.

Jehanzeb H. Chaudhary
Department of Mathematics
University of New Mexico
jehanzebh@hotmail.com

Nishant Panda
Los Alamos National Lab
nishpan@lanl.gov

MS285

Tractable Learning in Dynamical Systems Excited by Colored Non-White Inputs

Learning the structure of a network from time-series data is of significant interest in many disciplines such as power grids, biology, finance. We consider a networked linear dynamical system with p agents/nodes. We study the problem of learning the underlying graph of interactions/dependencies from observations of the nodal trajectories over a time-interval T . We present a regularized non-casual consistent estimator for this problem and analyze its sample complexity over two regimes: (a) where the interval T consists of n iid observation windows of length T/n (restart and record), and (b) where T is one continuous observation window (consecutive). Using the theory of M-estimators, we show that the estimator recovers the

underlying interactions, in either regime, in a time-interval that is logarithmic in the system size p . To the best of our knowledge, this is the first work to analyze the sample complexity of learning linear dynamical systems driven by unobserved not-white wide-sense stationary (WSS) inputs. Extensions to the setting with cyclo-stationary inputs are provided.

Deepjyoti Deka
Theory Division and the Center for Nonlinear Studies
Los Alamos National Laboratory
deepjyoti@lanl.gov

MS285

Data Assimilation for Strongly Coupled Systems

We investigate the effects of the strength of coupling across scales on the effectiveness of data assimilation in the two-scale Lorenz 96 system. Only the large scale, slow variables are assimilated and a large ensemble is considered to minimize sampling errors. We find that, given similar observational constraints, the effectiveness of data assimilation progressively deteriorates as the coupling strength increases and then fails more abruptly. We explore the dynamical changes that underlie the progressive deterioration and ultimate failure of data assimilation.

Humberto C. Godinez
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
hgodinez@lanl.gov

MS285

Scalable Inference of Non-Newtonian Rheology Parameters in Earth's Mantle on HPC Platforms

Estimating parameters in mantle flow and plate tectonics from surface observations results in an optimization problem governed by a highly nonlinear, heterogeneous, and incompressible Stokes equation. Solving this governing equation is a major challenge by itself (Our computational methods for this forward problem include adaptive mesh refinement and inexact Newton-Krylov with BFBT and multigrid preconditioning for saddle point linear systems.) Estimation of rheological parameters, for which we start with a Bayesian formulation, adds substantially to the solver challenges. To efficiently estimate parameters in the constitutive relation, we use adjoint Stokes equations and a Newton method for the resulting PDE-constrained optimization problem. The Hessian during the Newton steps is approximated with BFGS, while a Gauss-Newton Hessian at the numerical minimum is used for quantifying uncertainties in the estimated parameters. We demonstrate successful inference on large cross-sectional models and explore how the inference can be made computationally feasible at a global scale encompassing the entire Earth.

Johann Rudi
Department of Mathematics
Virginia Tech
jrudi@vt.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Jiashun Hu
Southern University of Science and Technology

hujs@sustech.edu.cn

Michael Gurnis
Caltech
gurnis@gps.caltech.edu

MS286

Neural Implicit Flow: a Mesh-Agnostic Dimensionality Reduction Paradigm of Spatio-Temporal Data

Fluid dynamics exhibits complex, multi-scale spatial structure, chaotic dynamics in time, and bifurcation in the relevant parameters. Among these challenges, spatial complexity is the major barrier for modeling and control of fluid dynamics, which motivates the need of dimensionality reduction. Existing paradigms, such as proper orthogonal decomposition or convolutional autoencoders, both struggle to accurately and efficiently represent flow structures for problems requiring variable geometry, non-uniform grid resolution (e.g., wall-bounded flows, flow phenomenon induced by small geometry features), adaptive mesh refinement, or parameter-dependent meshes. To resolve these difficulties, we propose a general framework called Neural Implicit Flow (NIF) that enables a compact and flexible dimension reduction of large-scale, parametric, spatio-temporal data into mesh-agnostic fixed-length representations. We apply our mesh-agnostic approach to several fluid flows, including flow past a cylinder, sea surface temperature data, 3D homogeneous isotropic turbulence, and a transonic 3D flow over ONERA M6 wing to show the effectiveness of parametric surrogate modeling, efficient differentiable query in space, learning non-linear manifolds, and the interpretable low-rank decomposition of fluid flow data.

Shaowu Pan, Steve Brunton
University of Washington
pans2@rpi.edu, sbrunton@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS286

Neural Galerkin Schemes for Evolution Equations

Neural Galerkin schemes build on Dirac-Frenkel variational methods to propagate forward in time nonlinear, low-dimensional parametrizations for efficiently predicting nonlinear latent dynamics of transport-dominated problems. The nonlinear parameters are learned by integrating systems of differential equations given by the physics of the problem. This is different to collocation methods that globally fit parameters by minimizing the residual based on samples from the space and time domain. Numerical results demonstrate that the proposed approach requires few degrees of freedom to accurately describe and predict transport-dominated dynamics. Furthermore, the approach is compatible with adaptive sampling schemes to efficiently approximate high-dimensional problems with spatially local features.

Yuxiao Wen, Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
yw3210@nyu.edu, pehersto@cims.nyu.edu

Eric Vanden Eijnden

Courant Institute
New York University
eve2@cims.nyu.edu

Joan Bruna
New York University
Courant Institute
bruna@cims.nyu.edu

MS286

Dynamical Weights: Learning Smooth Latent Trajectories for Time-Dependent Problems

We present a data-driven, space-time continuous framework to learn surrogate models for complex physical systems described by advection-dominated partial differential equations. Those systems have slow-decaying Kolmogorov n -width that hinders standard methods, including reduced order modeling, from producing high-fidelity simulations at low cost. In this work, we construct hypernetwork-based latent dynamical models directly on the parameter space of a compact representation network. We leverage the expressive power of the network and a specially designed consistency-inducing regularization to obtain latent trajectories that are both low-dimensional and smooth. These properties render our surrogate models highly efficient at inference time. We show the efficacy of our framework by learning models that generate accurate multi-step roll-out predictions at much faster inference speed compared to competitors, for several challenging examples.

Zhong Yi Wan
MIT
wanzy@google.com

MS287

Parallel-in-Time Solution of PDE-Constrained Optimization Problems with Hyperbolic PDEs

Solving optimization problems with hyperbolic partial differential equations as constraints can be extremely challenging. Furthermore, the necessity of solving both forward and adjoint problem at each step of the optimization algorithm leads to high computational costs. Parallel-in-time algorithms like Parareal can be used to speed up computations. Beside the simple approach of applying Parareal to the forward as well as the adjoint problem, we will also explore approaches for a closer integration of Parareal into the optimization algorithm. In this talk, we will investigate the performance of a combination of the gradient descent method with Parareal and discuss the influence of artificial and physical diffusion on our method.

Judith Angel
TU Hamburg
judith.angel@tuhh.de

Daniel Ruprecht
Institute of Mathematics
Hamburg University of Technology
ruprecht@tuhh.de

Sebastian Götschel
Hamburg University of Technology

sebastian.goetschel@tuhh.de

MS287

Diagonalisation-Based ParaOpt Preconditioners

The recently proposed ParaOpt method [1] constitutes a first step towards the parallel-in-time (PinT) solution of optimality systems arising from non-linear optimal control problems. Solving Newton-like linear systems in the coarse-grid correction (CGC) step inhibits good parallel scaling, however: the non-preconditioned iterative solver requires more iterations when increasing time-parallelism. Inspired by developments for the Parareal PinT algorithm for initial-value problems, we propose using diagonalisation-based preconditioners to speed up CGC and improve scaling. These preconditioners build upon those used in the ParaDiag class of algorithms [2], extended to support both tracking and terminal-cost objectives and having undergone improvements and convergence analyses. We formulate preconditioned ParaOpt for linear problems and demonstrate good weak scaling in numerical tests, supported theoretically by linking to our novel ParaDiag results. In addition, we lay out a clear pathway towards non-linear preconditioned ParaOpt, based on non-linear extensions of ParaDiag. [1]: M. J. Gander, F. Kwok, and J. Salomon. PARAOPT: A Parareal algorithm for optimality systems. SIAM Journal on Scientific Computing, 42(5):A2773A2802, Jan 2020. [2]: M. J. Gander, J. Liu, S.-L. Wu, X. Yue, and T. Zhou. ParaDiag: parallel-in-time algorithms based on the diagonalization technique. arXiv:2005.09158 [cs, math], Apr 2021. arXiv: 2005.09158.

Arne Bouillon, Ignace Bossuyt, Giovanni Samaey, Karl Meerbergen
 KU Leuven
 arne.bouillon@student.kuleuven.be,
 ignace.bossuyt1@kuleuven.be,
 giovanni.samaey@kuleuven.be,
 karl.meerbergen@kuleuven.be

ig-
 gio-

MS287

A Single-Level Parallel-in-Time Method for Hyperbolic Problems

A single-level Parallel-in-Time method for hyperbolic problems In this talk, we propose an iterative method for solving composite collocation problems parallel-in-time, across multiple steps, by firstly forming the all-at-once system. These problems are typically the basis for the multi-level PFASST algorithm. In contrast to PFASSTs multi-level approach, we use single-level preconditioned Richardson iterations with an α -circulant preconditioner which can be diagonalized using fast Fourier transformations (FFT) in time, making this method a class of ParaDiag algorithms. Using FFT in time on the preconditioner, the communication becomes very efficient and there is no need to worry about coarsening strategies, which may be a source of error when solving hyperbolic problems. As a consequence, we obtain time-interval decoupled problems in each iteration which we propose to solve directly in parallel, across the collocation nodes, using a diagonalization of the slightly modified quadrature matrix. This brings us to a doubly parallel-in-time method that solves linear problems very efficiently, from where an extension to nonlinear problems is made. We present theoretical and practical convergence and performance results as well as the strengths and weaknesses of this method.

Gayatri Caklovic

Forschungszentrum Juelich
 gayatri.caklovic@kit.edu

MS287

TriMGRIT - An Extension of Multigrid Reduction in Time for Constrained Optimization

Time-dependent simulations traditionally involve sequential time stepping. In the linear setting, this is equivalent to a forward solve of a block lower-bidiagonal space-time system. Multigrid reduction in time (MGRIT) is an approximate block cyclic reduction method for solving these systems in parallel, applicable also to nonlinear problems. One important goal of MGRIT is non-intrusiveness, where ideally users only wrap their existing code for taking a time step. TriMGRIT extends the MGRIT idea to solve time-dependent constrained optimization and optimal control problems, where the systems have block tridiagonal structure and are coupled both forward and backward in time. In this talk, we will describe the TriMGRIT algorithm and present parallel results and experiences for several model problems.

Robert D. Falgout
 Lawrence Livermore National Laboratory
 falgout2@llnl.gov

MS288

Neural Networks for Fast Simulation of Multiphysics Interactions of Electromagnetic Fields

This paper presents neural network architectures that can be efficiently trained to replace one or more solvers in multiphysics simulations involving electromagnetic fields. We focus on electromagnetic/thermal simulations, where Maxwells equations are jointly solved with the heat equation in lossy media. In particular, neural networks are used to replace the heat equation solver and directly produce the transient and steady state temperature distribution within the medium, using the material properties and the electric field components, computed by a Finite-Difference Time-Domain (FDTD) simulation. Hence, the multiphysics simulation is greatly simplified, as it relies on a single-physics (electromagnetic) solver, combined with a neural network. The training of the network is based on an unsupervised strategy, where the heat equation is directly implemented as a loss function of the network. This approach is compared to the conventional supervised strategy of using the mean square error (MSE) between ground truth data (produced by a finite difference solver for the heat equation) and the network prediction for the temperature. We demonstrate the accuracy and computational efficiency of this approach for general multiphysics problems and its applications to the modeling of microwave hyperthermia scenarios in biological tissues.

Shutong Qi
 Department of Electrical and Computer Engineering
 University of Toronto
 st.qi@mail.utoronto.ca

Costas Sarris
 Electrical and Computer Engineering, University of Toronto
 costas.sarris@utoronto.ca

MS288

Kernel-Based Regressions for Vector-Valued Prob-

lems in Microwave Applications

This talk presents a modeling framework for vector-valued functions based on the multi-output kernel Ridge regression (KRR). The proposed formulation of the KRR relies on a generalized definition of the reproducing kernel Hilbert space (RKHS) to the case of vector-valued functions, thus bridging the gap between multi-output Artificial Neural Network (NN) structures and standard scalar kernel-based approaches. The above multi-output formulation of the KRR can be suitably adopted to train a vector-output surrogate model able to predict the frequency responses of parametric microwave structures affected by several parameters with large variability. Its performance, in terms of parametric and stochastic analysis, is compared with the ones provided by two state-of-the-art techniques, such as the combination of the principal components analysis (PCA) and the least-squares support vector machine (LS-SVM) regression and a multi-output feed-forward NN structure.

Nastaran Soleimani
Dept. Electronics and Telecommunications
Politecnico di Torino, Torino, Italy
nastaran.soleimani@polito.it

Riccardo Trincherò
Dept. Electronics and Telecommunications
Politecnico di Torino
riccardo.trincherò@polito.it

Flavio Canavero
Dept. Electronics and Telecommunications
Politecnico di Torino, Torino, Italy
flavio.canavero@polito.it

MS288

AI-Driven Design and Optimization of Future Electronics

Artificial intelligence (AI) technologies have incentivized paradigm shift for a wide range of domain science areas. Specifically, in the field of electronics devices and systems, by embedding AI technologies, we can significantly reduce the time and cost to the design and optimization process. In this mini-symposium, we will introduce several state-of-the-art works on facilitating AI in wave and material modeling. The presentations will vary across physics-informed neural networks (PINN), kernel-based regressions, quantum machine learning, machine learning and variable-fidelity approaches. We will also address the revolution in microelectronics and high-frequency device design with AI acceleration.

Zhi Jackie Yao
Lawrence Berkeley National Laboratory
jackie_zhiyao@lbl.gov

MS289

SpecSolve: Fast Spectral Methods for Spectral Measures

Spectral measures arise in numerous applications such as quantum mechanics, signal processing, resonances, and fluid analysis. Unlike its matrix counterpart, the spectral measure of a self-adjoint operator may have an absolutely continuous component and an associated density function, e.g., in applications posed on unbounded domains. However, there has been no method able to compute spectral

measures of general infinite-dimensional self-adjoint operators, with previous efforts focused on classes of operators that carry a lot of structure. The `SpecSolve` package [Colbrook, Horning, Townsend, *SIAM Rev.*, 63(3) (2021), pp. 489524] computes spectral measures by combining state-of-the-art adaptive spectral methods with an efficient resolvent-based strategy. By dealing with operators directly, as opposed to previous "discretise-then-solve" techniques, we construct methods with arbitrarily large rates of convergence that recover the spectral measures of general self-adjoint ODEs, PDEs, integral operators and discrete operators. These algorithms are embarrassingly parallelisable. Numerical examples will be given, demonstrating efficiency, and tackling difficult problems taken from mathematics and other fields such as chemistry and physics. A highlight is the computation of hundreds of eigenvalues of Dirac operators to machine precision and without spectral pollution.

Matthew Colbrook
University of Cambridge
mjc249@cam.ac.uk

Andrew Horning, Alex J. Townsend
Cornell University
horninga@mit.edu, townsend@cornell.edu

MS289

Sparse Spectral Methods for Equilibrium Measures

Equilibrium measure problems appear in the continuous limit of particle swarm problems in which particle behaviour is modeled by attractive-repulsive forces. We introduce a banded spectral method using weighted Jacobi polynomials which allows the efficient computation of equilibrium measures for power law kernels. The derivation of the method involves new Jacobi polynomial recurrence relationships for Riemann-Liouville integrals, Riesz potentials and the fractional Laplacian. We present numerical experiments which agree with special case analytic results as well as independent particle swarm simulations in one and higher dimensions.

Timon Gutleb
University of Oxford
timon.gutleb@maths.ox.ac.uk

MS289

Highly Accurate Multi-Domain Multivariate Spectral Collocation Method for (2+1) Dimensional Partial Differential Equations

The novelty of this work rests upon the use of domain partitioning technique in time variable when discretizing the domain of solution in spectral collocation algorithm. The single domain multivariate spectral collocation methods have been proven to be effective in solving time-dependent PDEs defined over small time domains. However, there is a significant loss of accuracy as time domain proliferates and also when the number of grid nodes approaches a definite particular number. Therefore, the establishment of the new innovative multi-domain multivariate spectral quasilinearisation method (MDMV-SQLM) is described for the purpose of solving (2+1) dimensional nonlinear PDEs defined on large time intervals. The main output of this study is confirmation that minimizing the size of time computational domain at each subinterval assures sufficiently accurate results that are attained using minimal number of nodal points and computational time. To highlight the effi-

ciency and accuracy of the MDMV-SQLM, error estimates, condition numbers and CPU time are presented for nonlinear Burgers PDEs. The adoption of the domain decomposition technique is efficacious in suppressing the numerical challenges linked to large matrices and ill-conditioned nature of the resulting coefficient matrix. Also, the obtained results confirm that the numerical scheme is computationally cheap, fast and yield extremely accurate and stable results using fewer number of grid points for large time domains.

Musawenkhosi P. Mkhathshwa
University of South Africa
patsonmkhathshwa@gmail.com

MS289

Sparse Spectral Methods for Fractional PDEs

Fractional partial differential equations (FPDEs) model nonlocal processes such as wave absorption in the brain, long-range geophysical effects, and Levy flights. Many methods require a truncation of the domain to discretize. Aside from the expense, the nonlocal nature of FPDEs and the slow decay of the solutions often cause numerical artefacts to appear even with large truncations. Moreover, the solutions may exhibit singularities that must be accounted for to build methods that converge quickly. We discuss a class of FPDEs that include sqrt-Laplacian, derivative, Hilbert transform and identity operators. By constructing a sum space consisting of weighted Chebyshev polynomials of the second kind and their Hilbert transform counterparts, we precisely identify the action of the fractional operator on our approximation space. This allows us to derive sparse relationships between the sum space that we expand our solution in, and the right-hand side which is expanded in the so-called dual sum space. Affine transformations of the sum space can be combined to together in the approximation where, remarkably, the action of the fractional operator decouples. Hence, we solve over each interval separately and in parallel. We apply our spectral method to problems with smooth data, where we observe spectral convergence, as well as a fractional wave propagation problem.

Ioannis P. Papadopoulos, Sheehan Olver
Imperial College London
ioannis.papadopoulos13@imperial.ac.uk,
s.olver@imperial.ac.uk

MS290

Parallel QR Factorization of Block Low-Rank Matrices

Block Low-Rank (BLR) matrices are a compact representation that can reduce the storage and computational costs of factorizing dense matrices arising from a wide range of scientific applications. We present two new algorithms to perform QR factorization of BLR-matrices: one that performs blocked Householder QR, and another that is based on the fine-grained tiled Householder QR. We show how the blocked algorithm exploits BLR structure to achieve arithmetic complexity of $\mathcal{O}(mn)$, while the tiled algorithm exhibits $\mathcal{O}(mn^{1.5})$ complexity. We also present the parallelization of these algorithms and show that the tiled algorithm allows for highly-parallel task-based execution thanks to its finer task granularity. We compare our algorithms with an existing BLR-QR algorithm based on Modified Gram-Schmidt (MGS) iteration and a state-of-the-art vendor-optimized dense Householder QR in Intel MKL. For

a matrix of size $131k \times 65k$, our methods are more than an order of magnitude faster than the dense QR in MKL. Our methods are also robust to ill-conditioning compared to the existing MGS-based method that suffers from numerical instability. On a CPU with 64 cores, our parallel tiled Householder and blocked Householder algorithms show a speedup of 50 and 37 times, respectively.

Muhammad Ridwan Apriansyah
Tokyo Institute of Technology
Tokyo, Japan
ridwan@rio.gsic.titech.ac.jp

Rio Yokota
Tokyo Institute of Technology
rioyokota@gasic.titech.ac.jp

MS290

Combining Binary Compression with Low-Rank Arithmetic

Recently, a hybrid data compression scheme based on low-rank compression and optimized binary representation was used to compress simulation data. There, the low-rank approximation was further optimized by using binary compression, permitting optimized storage while maintaining error boundaries. We apply such a scheme in the context of low-rank arithmetic and investigate the benefits and costs associated with this, while different methods of representing the binary data are evaluated.

Ronald Kriemann
Max Planck Institute for Mathematics i.t.S.
rok@mis.mpg.de

MS290

A Fast Solver for Linear Systems with Tensor Product Structure via Low-Rank Updates

In this talk we consider the numerical solution of linear systems with a tensor product structure of the form

$$A_1 \otimes I \otimes \cdots \otimes I + \cdots + I \otimes \cdots \otimes I \otimes A_d)x = b$$

where the matrices $A_i \in \mathbb{C}^{n \times n}$ are positive definite and belong to the class of hierarchically semiseparable matrices (HSS). Such linear systems arise in the discretization of Laplace-type differential equations on tensorized grids and in the dimensionality reduction of linear dynamical control systems. We propose and analyze a divide-and-conquer scheme based on the technology of low-rank updates that attains the computational cost $\mathcal{O}(n^d \log(n))$ that is quasi-optimal in the case of unstructured right-hand-sides.

Stefano Massei
Technische Universiteit Eindhoven
stefano.massei@unipi.it

MS291

Deep Importance Sampling Using Tensor Approximations

We propose a deep importance sampling method that is suitable in particular for estimating rare event probabilities in high-dimensional problems. We approximate the optimal importance distribution in a general importance sampling problem as the pushforward of a reference distribution under a composition of order-preserving transformations, in which each transformation is formed by a

squared tensor-train decomposition of a ratio of unnormalized bridging density functions, such as tempered or smoothed versions of the target density. The use of composition of maps moving along a sequence of bridging densities alleviates the difficulty of directly approximating a concentrated target density function. To compute expectations over unnormalized probability distributions, we design a ratio estimator that estimates the normalizing constant using a separate importance distribution, again constructed via a composition of transformations in tensor-train format. This offers better theoretical variance reduction compared with self-normalized importance sampling, and thus opens the door to efficient computation of rare event probabilities in Bayesian inference problems. Numerical experiments on problems constrained by differential equations show little to no increase in the computational complexity with the event probability going to zero.

Sergey Dolgov
University of Bath
Department of Mathematical Sciences
s.dolgov@bath.ac.uk

Tiangang Cui
Monash University
tiangang.cui@monash.edu

Robert Scheichl
Universitaet Heidelberg
r.scheichl@uni-heidelberg.de

MS291

Bayesian Multifidelity Inverse Analysis for Expensive, Non-Differentiable, Physics-Based Solvers in High Dimensions

Computationally expensive, large-scale, numerical models with high-dimensional (stochastic) model input pose some of the biggest challenges for (Bayesian) inverse analysis. The solution process is further impeded when model derivatives are inaccessible, as is often the case in legacy codes. We propose Bayesian multifidelity inverse analysis (BMFIA), which overcomes the difficulties mentioned above by employing computationally inexpensive, lower-fidelity models and constructing a multifidelity likelihood function. The approach builds upon previous developments of the authors in the field of uncertainty quantification. The multifidelity likelihood is learned robustly, and potentially adaptively, from a small number of high-fidelity simulations. It incorporates the epistemic uncertainty that arises from the small amount of high-fidelity data and can also serve as an objective function for deterministic optimization schemes. We enrich the multifidelity dependence with low-dimensional, informative input features, which increases the formulations accuracy. The inference process can be performed using state-of-the-art sampling-based or variational methods for Bayesian inverse analysis and requires only evaluations of the lower-fidelity model(s), which, when appropriately selected, can also furnish derivatives. We demonstrate our approach on large-scale biomechanical and coupled multi-physics problems and compare it with state-of-the-art single- and multifidelity methods.

Jonas Nitzler, Wolfgang Wall
Technical University of Munich
Institute for Computational Mechanics (LNM)
jonas.nitzler@tum.de, wolfgang.a.wall@tum.de

Phaedon-Stelios Koutsourelakis
Technical University of Munich
p.s.koutsourelakis@tum.de

MS291

Concentration Analysis of Multivariate Elliptic Diffusions

We prove concentration inequalities and associated PAC bounds for continuous- and discrete-time additive functionals for possibly unbounded functions of multivariate, nonreversible diffusion processes. Our analysis relies on an approach via the Poisson equation allowing us to consider a very broad class of subexponentially ergodic processes. These results add to existing concentration inequalities for additive functionals of diffusion processes which have so far been only available for either bounded functions or for unbounded functions of processes from a significantly smaller class. We demonstrate the usefulness of the results by applying them in the context of high-dimensional drift estimation and Langevin MCMC for moderately heavy-tailed target densities. This is joint work with Cathrine Aeckerle-Willems and Claudia Strauch.

Lukas Trottner
Aarhus University
trottner@math.au.dk

MS291

Transport-Based Sampling Using Polynomial Density Surrogates

Generating samples from arbitrary probability distributions is an integral task in various areas of modern applied mathematics such as parameter inference and uncertainty quantification. In this talk, we describe a sampling algorithm based on the Knothe-Rosenblatt transport, that can be used to approximately sample from target distributions on the d -dimensional unit cube $[0, 1]^d$ under mild assumptions on the target density. The method is based on the use of polynomial density surrogates, which allows to explicitly and analytically construct the corresponding transport. We discuss efficient implementation schemes and derive error convergence rates for target densities belonging to different smoothness classes.

Josephine Westermann
Universität Heidelberg
josephine.westermann@uni-heidelberg.de

Jakob Zech
Heidelberg University
jakob.zech@uni-heidelberg.de

MS291

Breaking Out of Local Minima: Enhancing the Basin Hopping Algorithm with the Skipping Mechanism

The Basin Hopping algorithm is a widely used stochastic method for solving complex global optimization problems, but it can be hindered by getting trapped in local minima. In this talk, I will present a new variant of this algorithm, named Basin Hopping with Skipping (BH-S), in which the classical perturbation step is replaced with a so-called skipping mechanism from rare-event sampling. This new perturbation strategy makes it easier to escape local minima and enhances non-local exploration, leading to a more effi-

cient and effective optimization process. In this talk, I will provide an in-depth overview of the skipping mechanism, describe its integration into the basin-hopping algorithm, and present some empirical results on benchmark global optimization problems.

Alessandro Zocca

Vrije Universiteit Amsterdam
a.zocca@vu.nl

John M. Moriarty, Maldron Goodridge
Queen Mary University of London
j.moriarty@qmul.ac.uk, m.p.goodridge@qmul.ac.uk

Jure Vogrinc
University of Warwick
jure.vogrinc@warwick.ac.uk

MS292

Multi-Information Source Bayesian Optimization of Culture Media for Cellular Agriculture

Culture media used in industrial bioprocessing and the emerging field of cellular agriculture is difficult to optimize due to the lack of rigorous mathematical models of cell growth and culture conditions, as well as the complexity of the design space. Rapid growth assays are inaccurate yet convenient, while robust measures of cell number can be time-consuming to the point of limiting experimentation. In this talk, we describe the use of multi-information source Bayesian optimization to optimize a cell culture media with 14 components. Following this approach, we were able to design media with 181% more cells than a common commercial variant with a similar economic cost while doing so in 38% fewer experiments than an efficient design-of-experiments method. Furthermore, the optimal medium generalized well to long-term growth up to four passages of C2C12 cells, indicating the multi-information source assay improved measurement robustness relative to rapid growth assays alone.

Zachary Cosenza
Department of Chemical Engineering
University of California Davis
zacosenza@ucdavis.edu

Raul Astudillo
Cornell University
rastudil@caltech.edu

Peter I. Frazier
Dept. of Operations Research and Information
Engineering
Cornell University
pf98@cornell.edu

Keith Baar
Department of Physiology and Membrane Biology
University of California Davis
kbaar@ucdavis.edu

David Block
Department of Chemical Engineering
University of California Davis
deblock@ucdavis.edu

MS292

Overview of Multi-Objective Bayesian Optimiza-

tion for Detector Design at the Electron Ion Collider Detector(s): ePIC and Beyond

The Electron Ion Collider (EIC), the future ultimate machine to study the strong force, is a large-scale experiment with an integrated detector that covers the central, far-forward, and far-backward regions. EIC is utilizing adaptive experimentation strategies during the design and R&D phases, in order to deal with a complex design problem characterized by compute intensive simulations, multiple sub-detectors — each characterized by multiple design parameters — and multiple design criteria or objectives. In this context, we explored state of the art solutions to design the experiment in a more efficient way. This talk provides an overview of the recent progress made during the EIC detector proposal; it will also cover how this work could further progress in the near future.

Cristiano Fanelli

College of William and Mary
cfanelli@wm.edu

MS292

Bayesian Optimization for Peptide Design: Controlling Ice, Binding Metals, and Making Proteins Glow

Peptides are a powerful class of molecules that can serve as drugs, materials and sensors in biological applications. They are challenging to design, in part because experimental evaluation of a peptide's quality in particular application requires time-consuming and expensive chemical synthesis and characterization, and because a peptide's quality is hard to predict. We discuss the application of Bayesian optimization to this challenge, focusing on three testbed problems: peptides that prevent ice recrystallization, with applications to preserving organs for transplantation; and two ways of using peptides as sensors, one focused on binding metals, and the other on supporting enzymes that attach fluorescent chemical subgroups. All involve real experiments and face common challenges: incorporating domain expertise from chemists, making accurate predictions sequences with appropriately quantified uncertainty, optimizing acquisition functions over discrete sets of molecules, and appropriately leveraging time-consuming computational chemistry simulations.

Peter I. Frazier

Dept. of Operations Research and Information
Engineering
Cornell University
pf98@cornell.edu

MS292

Polynomial-Model-Based Optimization for Black-box Problems

We deliver a novel algorithm, termed Polynomial-Model-Based Optimization (PMBO), addressing a class of black-box optimization tasks, for which sampling of the underlying unknown system is extremely costly. These scenarios occur for instance for complex interaction-simulations or hyper-parameter search for machine learning applications. While classic Bayesian Optimization is iteratively updating the surrogate model of the black-box function according to the acquisition function Expected Improvement, PMBO delivers a black-box model by fitting a polynomial surrogate to the objective function that enables fast optimization of the surrogate. Though due to the No Free Lunch

Theorem there is no universal best optimizer, by benchmarking PMBO against other state-of-the-art algorithms for a set of artificial, analytical functions and a biological Tumor Response Model, its performance results to be superior up to dimension 6. Apart from its good performance, the novelty of the PMBO relies on providing a search space parametrization on the fly. This additional information enables post-processing analysis which cannot be achieved by prior established methods.

Janina Schreiber

Center for Advanced System Understanding
j.schreiber@hzdr.de

Uwe Hernandez Acosta
CASUS
u.hernandez@hzdr.de

Damar Wicaksono, Michael Hecht
CASUS - Center for Advanced Systems Understanding,
Helmholtz
d.wicaksono@hzdr.de, m,hecht@hzdr.de

MS293

An Algorithm to Mesh 2D Regions in Segmented Images

Certain applications in medicine, material science, engineering require creating meshes of the structures defined by their segmented regions in images. These meshes can then be used to compute high fidelity simulations of associated physics with the finite element method. In this work, we develop a practical algorithm to create high quality triangulations of regions in segmented images. For this, we identify the boundaries between labeled pixels of the image, smooth and coarsen the boundaries, and use them to guide the triangulation. We demonstrate the effectiveness of the algorithm with several real images from material science.

Gunay Dogan
National Institute of Standards and Technology
gunay.dogan@nist.gov

MS293

Conforming Tetrahedral Meshing of Non-Convex Domains with Sharp Features and Narrow Regions Based on the VoroCrust Algorithm

VoroCrust is the first automated algorithm for conforming Voronoi meshing. It was presented in SIGGRAPH 2020. We present a recent variation, VoroCrust-T, to generate conforming sliver-free tetrahedral meshes with guaranteed quality. Our novel approach utilizes the foundational sphere packing methods in VoroCrust which simultaneously generates a quality triangular mesh of the surface of the input Piecewise Linear Complex (PLC) model and decomposes the enclosed volume by weighted Delaunay simplices with good angles conforming to the surface mesh. VoroCrust has an embedded sizing function that capture the curvature of the model, and robustly represents sharp features and narrow regions that might be associated with the input model. Our new variation VoroCrust-T, similar to VoroCrust, can also handle non-manifold and non-watertight inputs. We illustrate the robustness and output quality of VoroCrust through a collection of models of varying complexity. In this talk we also present the recently released VoroCrust software and discuss its performance in

practice.

Mohamed S. Ebeida
Sandia National Laboratories
msebeid@sandia.gov

MS293

Geometry and Topology Analysis of Characterization and Discovery of Porous Material Structures

Nanoporous materials such have been of growing importance as materials for adsorption-based applications such as CO₂ capture, hydrogen and methane storage, as well as for many others applications, for example, as catalysts. A vast majority of nanoporous materials are based on modular chemistry, which enables construction of chemical structures using various building blocks. As the result, a huge number of material structures can be obtained. Novel tools are required to sieve through these possible structures to identify ones with the desired properties for a given application. Our group has been developing such tools by exploiting computational geometry and topology approaches such as Voronoi tessellation and persistence homology. We have gained the ability to detect and characterize porosity features across various length scales important for adsorption processes, and we successfully used these features to build machine learning structure-property models. This presentation will summarize these developments and hint towards new applications in the analysis of porous material structures.

Maciej Haranczyk
Lawrence Berkeley National Laboratory
maciej.haranczyk@imdea.org

MS293

An Extension to VORO++ for Multithreaded Computation of Voronoi Cells

VORO++ is a software library written in C++ for computing the Voronoi tessellation, a technique in computational geometry that is widely used for analyzing systems of particles. VORO++ was released in 2009 and is based on computing the Voronoi cell for each particle individually. Here, we take advantage of modern computer hardware, and extend the original serial version to allow for multithreaded computation of Voronoi cells via the OpenMP application programming interface. We test the performance of the code, and demonstrate that we can achieve parallel efficiencies greater than 95% in many cases. The multithreaded extension follows standard OpenMP programming paradigms, allowing it to be incorporated into other programs. We provide an example of this using the Voro-Top software library, performing a multithreaded Voronoi cell topology analysis of up to 102.4 million particles.

Jiayin Lu
Harvard University
jiayin.lu@g.harvard.edu

Emanuel A. Lazar
University of Pennsylvania
mlazar@math.biu.ac.il

Chris H. Rycroft
University of WisconsinMadison

rycroft@wisc.edu

MS293

Mapping the Structural Landscapes of Ordered and Disordered Particle Systems Through Voronoi Topology

Group theory provides a clear method of classifying crystalline structures according to their symmetry groups, though doesn't provide direct insight into relationships between different systems. Likewise, although there is great understanding of many different kinds of disordered particle systems, and although many tools provide some insight into similarities and differences between them, no unified setting exists for describing relationships between different kinds of disorder. We show that Voronoi topology can be used to map out the structural landscape of both ordered and disordered systems by considering the distributions of Voronoi cell topologies that can arise in the different systems. This approach provides a new way to think about structural relationships between particle systems, ordered and disordered. A new multithreaded version of VoroTop is used to analyze a large database of materials.

Deborah Schwarcz

Bar-Ilan University (Israel)
deborah.schwarcz@biu.ac.il

Emanuel A. Lazar
University of Pennsylvania
mlazar@math.biu.ac.il

MS294

A New Finite Element Method for Radiation Transport on Spherical Geodesic Grids

We present a new numerical method to solve the Boltzmann equation based on a finite element approach in angle, that combines the advantages of the discrete ordinate method (S_N) and the filtered spherical harmonics method (FP_N) while avoiding their disadvantages. The scheme uses a finite element basis to represent functions of angular variables on a geodesic grid along with a simple limiting strategy to ensure non-negativity of solutions. We then compare this new method with S_N and FP_N schemes using four test problems and find it to perform well when one of the other methods fail.

Maitraya K. Bhattacharyya, David Radice
Pennsylvania State University
mbb6217@psu.edu, dur566@psu.edu

MS294

GR-Athena++: Puncture Evolutions on Vertex-Centered Oct-Tree AMR

'GR-Athena++' is a general-relativistic, high-order, vertex-centered solver that extends the oct-tree, adaptive mesh refinement capabilities of the astrophysical (radiation) magnetohydrodynamics code 'Athena++'. Dynamical spacetimes are simulated using Z4c coupled to the moving puncture gauge with robust and accurate binary black hole (BBH) mergers demonstrated. 'GR-Athena++' leverages the task-based parallelism paradigm of 'Athena++' to achieve excellent scalability. Strong scaling efficiencies above 95% for up to 1.21e4 CPUs and excellent weak scaling up to 1e5 CPUs are measured for production BBH runs. 'GR-Athena++' thus allows for the robust

simulation of compact binary coalescences and offers a viable path towards numerical relativity at exascale. In this talk a general overview of features outlined above will be provided together with our recent development efforts.

Boris Daszuta

Friedrich-Schiller-Universität Jena
boris.daszuta@uni-jena.de

MS294

Handing-Off the Outcome of Binary Neutron Star Mergers for Accurate and Long-Term Post-Merger Simulations

We perform binary neutron star (BNS) merger simulations in full dynamical general relativity on a Cartesian grid with adaptive-mesh refinement. After the remnant black hole (BH) has become nearly stationary, we find that the evolution of the surrounding accretion disk over long timescales ($\sim 1s$) is suboptimal, as Cartesian coordinates over-resolve the angular coordinates at large distances, and the accreting plasma flows obliquely across coordinate lines dissipating angular momentum artificially from the disk. To address this, we present the HandOff, a set of computational tools that enables the transfer of general relativistic magnetohydrodynamic (GRMHD) and spacetime data to a GRMHD code that specializes in modeling BH accretion disks in static spacetimes over long timescales making use of general coordinate systems with spherical topology. We demonstrate that the HandOff allows for a smooth and reliable transition of GRMHD fields and spacetime data, enabling us to efficiently and reliably evolve BNS dynamics well beyond merger.

Federico G. Lopez Armengol

Rochester Institute of Technology
fedelopezar@gmail.com

Zachariah B. Etienne
West Virginia University
zachetie@gmail.com

Scott Noble
ASA Goddard Space Flight Center
scott.c.noble@nasa.gov

Bernard Kelly
NASA Goddard Space Flight Center
bernard.j.kelly@nasa.gov

Leonardo Werneck
University of Idaho
leonardo@uidaho.edu

Brendan Drachler
Rochester Institute of Technology
bcd3735@rit.edu

Manuela Campanelli
Rochester Institute of Technology
Center for Computational Relativity and Gravitation
manuela@astro.rit.edu

Federico Cipolletta
Rochester Institute of Technology
fxcsm@rit.edu

Yosef Zlochower

Rochester Institute of Technology
Center for Computational Relativity and Gravitation
yrzsma@rit.edu

Ariadna Murguia-Berthier
University of California
armurgui@ucsc.edu

Lorenzo Enoggi
Rochester Institute of Technology
le8016@rit.edu

MS294

Modeling Magnetic Fields, Jets, and Turbulence in the Multimessenger Era

Magnetic fields, turbulence, and jet-driven outflows play a critical role in core-collapse supernovae and compact-object mergers. These transients belong to the most luminous and energetic events observed in the universe and are key targets for time-domain astronomy surveys. I will discuss the unique challenges in both input physics and computational modelling for these systems and highlight recent breakthroughs in full 3D simulations. I will pay particular attention to challenges in performing these simulations on modern high-performance computational systems and how to extract as much scientific insight as possible from the produced datasets.

Philipp Moesta, Pablo Bosch
University of Amsterdam
p.moesta@uva.nl, p.boschomez@uva.nl

Sanjana Curtis
University of Chicago
sanjanacurtis@uchicago.edu

Swapnil Shankar, Sebastiaan de Haas, Uddipta Bhardwaj
University of Amsterdam
sshankar@uva.nl, sebastiaan.dehaas@student.uva.nl,
u.bhardwaj@uva.nl

MS295

Models and Numerical Simulations of Collective Cellular Motion

The cells of incipient cancer tumors need oxygen to grow in tissues. They issue growth factors that are felt in nearby blood vessels and stimulate the formation of new blood vessels (angiogenesis) that carry the needed oxygen and nutrients to the tumor. Epithelial cells in move collectively to close a wound or collectively push an epithelial monolayer of a different tissue. Aspects of these multiscale processes in biomedicine can be studied by mathematical and computational models. We discuss angiogenesis at supracellular scales by stochastic differential equations and birth-death processes, derive continuum descriptions of the density of moving blood vessels and analyse a soliton-like attractor. Numerical simulations of a cellular Potts model including Notch signaling between cells describe the growth of blood vessels at cellular and subcellular levels. Collective cellular motion of epithelial cells can be described using simpler active vertex models, in which cells are polygons forming a Voronoi tessellation of the tissue. The dual Delaunay triangulation comprises cells centers that evolve according to dynamics including collective inertia and active forces. Numerical simulations qualitatively agree with experiments on wound healing assays, formation of fingers, and antago-

nistic migration assays of a precancerous cellular aggregate invading the space of a healthy tissue and displacing its cells.

Luis Bonilla, Manuel Carretero
Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience
&IndMath
bonilla@ing.uc3m.es, manuel.carretero@uc3m.es

Filippo Terragni
Universidad Carlos III de Madrid, Spain
fterragn@ing.uc3m.es

Ana Carpio
Universidad Complutense de Madrid
Departamento de Matematica Aplicada
carpio@mat.ucm.es

Carolina Trenado
Universidad Carlos III de Madrid
ctrenado@ing.uc3m.es

Rocio Vega
Universidad Carlos III de Madrid
Millan Institute of Fluid Dynamics, Nanoscience
&IndMath
rvega@ing.uc3m.es

MS295

Optimization in Radiotherapy

In this talk I will provide an overview of models and algorithms in the field of radiotherapy from optimization and feasibility aspects. The results are focused on convex and non-convex settings. The algorithms to be discussed are projection methods and the superiorization methodology.

Aviv Gibali
Braude College
P.O. Box 78, Karmiel 2161002, Israel
avivg@braude.ac.il

MS295

Viral Transmission: Determining the Infection Risk and Comparing Mitigations

SARS-CoV-2, the virus responsible for the COVID-19 pandemic, has claimed more than six million lives by now. In collaboration with many other scientists, we have worked intensively with policy-makers to fight the pandemic. I will present an overview of our pandemic-related research: a) A new modelling framework (Lau et al, Proc. R. Soc. A, 2022) for the airborne transmission of viruses that quickly determines the infection risk in an indoor space, as ventilation changes. The model agrees well with real superspreading events during the COVID-19 pandemic and could be used in future epidemics. b) An agent-based modelling framework (Moore et al, COVID, 2021) that allows quick decision-making and comparison of mitigations in educational settings. We find that, although frequent testing can be used to reduce the number of infected students, investment in masks and ventilation is more effective. c) Modelling air purifiers. Air purifiers can be used to improve ventilation and reduce infection risk from SARS-CoV-2 and other pathogens in indoor spaces. It is, however, frequently not clear where to place them and how strong they should be. Hence, deploying them to schools, hospitals and other

spaces poses interesting challenges for policy-makers and managers of spaces. We provide solutions to some of these challenges. In collaboration with: R. Gonzalez-Faria, T.C. Dale. Funding: Welsh Government, Royal Society, Cardiff University, University of Oxford.

Katerina Kaouri

School of Mathematics
Cardiff University, UK
kaourik@cardiff.ac.uk

Ian M. Griffiths
Mathematical Institute
University of Oxford
ian.griffiths@maths.ox.ac.uk

Zechariah Lau, Alexander Pretty
Cardiff University
zechlau@gmail.com, prettya@cardiff.ac.uk

Attila Kovacs
Mathematical Institute
University of Oxford
kovacs@maths.ox.ac.uk

Aaron English
University of Manchester
aaron.english@manchester.ac.uk

Thomas E. Woolley, Joshua W. Moore
Cardiff University
woolleyt1@cardiff.ac.uk, moorej16@cardiff.ac.uk

MS295

Modelling the Removal of Environmental Contaminants by Adsorption

Tackling environmental challenges is this generations defining task (EC Green Deal 2020). The goal of holding global warming to well below 2C (IPCC 2018; Paris Agreement 2015), can now only be achieved through the extraction of atmospheric greenhouse gases in tandem with emission reductions and a range of other measures. Similarly the goal of a toxic free environment requires the removal of a multitude of existing contaminants. The focus of this talk will be the development and solution of mathematical models for the removal of contaminants by column sorption. Column sorption involves passing a contaminated fluid through a porous media contained in a column. The contaminants are adsorbed by the porous media. This is a common process, capable of removing a multitude of contaminants. However, industrial use is restricted due to a large increase in running costs. Further, laboratory experiments on small scale equipment, a necessary precursor to building the full size version, do not scale up in the predicted manner. The mathematical description of column sorption involves coupling advection-diffusion equations to sink (removal) equations, all applied over an evolving domain. In the talk we will describe our modelling, analytical and numerical solution techniques as well as experiments currently being employed to help improve and optimise this process.

Tim G. Myers

Centre de Recerca Matemàtica
Spain
tmyers@crm.cat

Francesc Font

Centre de Recerca Matemàtica
font@crm.cat

MS296

Variational Multiscale Moment Methods for the Boltzmann Equation

Integrating the Boltzmann equation with the collision invariants leads to conservation equations in mass, momentum and energy. These equations require closure relations for to balance the number of unknowns with the number of equations. Such closure may be attained by the Chapman-Enskog method which assumes that the distribution takes the form of a perturbation series in powers of the Knudsen number scaled by a Maxwellian. By using a scaled Boltzmann equation to solve for successive orders of the perturbation series, we obtain familiar fluid dynamic equations: Euler, Navier-Stokes-Fourier, Burnett, and so on. Whilst considering the linearized Boltzmann equation, we propose a different ansatz. Instead of a perturbation series, we use the variational multiscale (VMS) method to split our unknown into a coarse scale term from the space of collision invariants and a fine scale term from the orthogonal complement to the collision invariants. We show that the Chapman-Enskog method can be viewed as particular way of approximating the fine scale term under the VMS formulation. By approximating the fine scale term differently, we are able to derive an alternative to the Burnett equations (i.e. a correction that is a step above Navier-Stokes-Fourier) for which we can show entropy stability. To show some benchmark results, these alternative equations necessitate extra boundary conditions which we motivate by considering the space-velocity weak form of the linearized Boltzmann.

Michael Abdelmalik

Eindhoven University of Technology
m.abdel.malik@tue.nl

Frimpong Baidoo
University of Texas at Austin
fabaidoo@utexas.edu

Irene M. Gamba
Department of Mathematics and ICES
University of Texas
gamba@math.utexas.edu

Luis Caffarelli
University of Texas at Austin
Department of Mathematics
caffarel@math.utexas.edu

Thomas Hughes
UT Austin
hughes@oden.utexas.edu

MS296

Generalised Moment Approximation for the Boltzmann Equation

Non-equilibrium gas flow process occurs naturally in several real world applications, e.g MEM devices, vacuum pumps. Developing efficient yet accurate mathematical models to describe these non-equilibrium flow problems is of interest to researchers work in the field of rarified gas dynamics. In kinetic gas theory, moment method is a technique to reduce complexity of an underlying kinetic equa-

tion by suitable approximation. In this talk, we discuss derivation of general moment systems from the Boltzmann equation for mono-atomic gases and present a generalized numerical framework based on finite-element-method to solve well-defined arbitrary order of moment systems. We highlight the benefits of our proposed framework by discussing the mathematical properties such as symmetry, generality, linearity, and more. This talk also covers several example problems involving non-equilibrium scenarios.

Edilbert Christhraj
RWTH Aachen University
Applied and Computational Mathematics
Christhraj@acom.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

MS296

Goal-Oriented Adaptive MLMC for Elliptic PDEs with Lognormal Random Inputs

We propose our Adaptive Multilevel Monte Carlo (AMLMC) [Beck, Joakim, et al. "Goal-oriented adaptive finite element multilevel Monte Carlo with convergence rates." Computer Methods in Applied Mechanics and Engineering (2022)] method to solve an elliptic partial differential equation (PDE) with lognormal random input data, where the PDE model is subject to geometry-induced singularity. The previous work [Moon, K-S., et al. "Convergence rates for an adaptive dual weighted residual finite element algorithm." BIT Numerical Mathematics 46.2 (2006)] developed convergence rates for a goal-oriented adaptive algorithm based on isoparametric d-linear quadrilateral finite element approximations and the dual weighted residual error representation in the deterministic setting. This algorithm refines the mesh based on the error contribution to the QoI. This work aims to combine MLMC and the adaptive finite element solver. Contrary to the standard Multilevel Monte Carlo methods, where samples are computed using a discretization-based numerical method, whose resolution is linked to the level, our AMLMC algorithm uses a sequence of tolerances as the levels. Specifically, for a given realization of the input coefficient and a given accuracy level, the AMLMC constructs its approximate sample as the ones using the first mesh in the sequence of deterministic, non-uniform meshes generated by the above-mentioned adaptive algorithm that satisfies the sample-dependent bias constraint.

Joakim Beck
King Abdullah University of Science and Technology (KAUST)
joakim.beck@kaust.edu.sa

Yang Liu
King Abdullah University of Science and Technology
yang.liu.3@kaust.edu.sa

Erik von Schwerin
King Abdullah University of Science and Technology
Department of Applied Mathematics and Computational Science
erik.vonschwerin@kaust.edu.sa

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology

raul.tempone@kaust.edu.sa

MS296

Continuum Approaches for Strongly Non-Equilibrium Flows

The kinetic theory finds its application in an extremely wide range of problems, ranging from environmental problems to the planetary entry. It is used to describe non-equilibrium flows for completely different conditions. However, there is still no single universal, accurate and computationally inexpensive approach suitable for various degrees of gas rarefaction, characterized by the Knudsen number, Kn. When solving problems, it may be necessary to simulate at relatively low values of Kn corresponding to the slip flow regime, where the often-used DSMC approach becomes computationally too expensive. In such conditions, continuum approaches become relevant. The talk is devoted to actively developing continuum approaches using generalized Chapman-Enskog method. One-temperature, multi-temperature and state-to-state descriptions will be considered. The methods used to speed up calculations will be highlighted, including the possibility of combining different approaches. Particular attention will be paid to validation, which is necessary before using detailed and resource-intensive approaches.

Maksim Melnik
St.Petersburg University
melnik.mxm@gmail.com

MS296

Efficient Solution of the Linear Boltzmann Equation for Radiation Therapy Using the Dynamical Low-Rank Approximation

Computational methods in radiation therapy require a discretization of the six-dimensional phase space, composed of energy, spatial position and direction of flight. This prohibits fine numerical discretizations which are essential for the construction of accurate treatment plans. We use a dynamical low-rank approximation (DLRA) of the particle density to evolve the solution of the continuous slowing down approximation to the Boltzmann transport equation on a low-rank manifold in time. Here the energy is treated as a pseudo-time and a rank adaptive integrator is chosen to dynamically adapt the rank in energy. The transport equation is split into collided and uncollided particles through a collision source method. Uncollided particles are described analytically, guaranteeing low computational costs, whereas collided particles are represented by a low-rank solution. The resulting method is more efficient concerning both run-time and memory than conventional radiation transport solvers and agrees well with conventional deterministic solvers as well as state-of-the-art Monte Carlo codes. We further extend our model from electron to the more clinically relevant proton transport, which typically requires finer discretizations due to the extremely forward-peaked scattering and sharp peak in the energy deposition curve. We show that DLRA is well-suited for proton dose computations, already at significantly lower rank than required for electrons.

Pia Stammer
Karlsruhe Institute of Technology
pia.stammer@kit.edu

Jonas Kusch
University of Innsbruck

jonas.kusch@uibk.ac.at

MS297

Fast Boundary Element Methods to Simulate Underwater Explosions and Their Interactions with Submarines

Assessing the impact of underwater explosions on submerged structures (submarines) is an important naval engineering problem. An underwater explosion mainly induces two distinct phenomena: a "shock wave" (fast acoustic perturbation) followed by an oscillating bubble of gas (slow incompressible flow). Our goal is to create an efficient numerical method that accounts for the effects of both phenomena on submerged structures. Due to the unbounded nature of the ocean and the complex mechanical behavior of the submarine we want to take into account, it is natural to consider a Boundary Element Method/Finite Element Method (BEM/FEM) coupling. I will present how we can take advantage of fast BEMs in the frequency domain to model this time-domain problem, all the necessary improvements to be able to consider realistic configurations and the consequences on the convergence of the FEM/BEM coupling.

Stephanie Chaillat

Laboratoire POEMS (CNRS-ENSTA-INRIA)
stephanie.chaillat@ensta-paris.fr

Marc Bonnet

POems (CNRS-ENSTA-INRIA), Appl. Math. Dept.
Paris, France
marc.bonnet@ensta.fr

Alice Nassor

POEMS
alice.nassor@ensta.fr

MS297

Fast Boundary Integral Methods for Confined Suspension Flows

Particulate flows are ubiquitous in the study of self-assembly of biological structures and the design of soft materials. In many examples across application fields, it has become apparent that resolving confinement, short term forces (e.g. contact) and long-term forces (e.g. hydrodynamics) is crucial to capture key features of these phenomena. This talk will focus on the current state-of-the-art on the BIE-based methods for fluid suspensions simulation, as well as ongoing challenges in this area. I will discuss some of our recent work applying this boundary integral framework to simulate Janus particle systems, that is, of particles whose surfaces exhibit two distinct physical properties.

Eduardo Corona

University of Colorado at Boulder
eduardo.corona@colorado.edu

MS297

Scattering from 2d Open Waveguides

We present a new method to solve the scattering problem defined by two planar, rectangular, semi-infinite open wave guides that meet along a common perpendicular line. We use the method of fundamental solutions to reduce the scattering problem to a system of Fredholm integral equa-

tions along a line. This method lends itself to numerical implementation.

Charles Epstein

Department of Mathematics, University of Pennsylvania
cepstein@flatironinstitute.org

Manas N. Rachh

Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS297

Scattering Matrix Computations for Integrated Photonics

We present an accurate and efficient numerical method for large-scale simulations of integrated photonic devices. The governing equations for light propagating through these devices are the Helmholtz equation in two dimensions, and Maxwells equations in three dimensions. The problem is numerically challenging to simulate owing to the size of the object as measured in wavelengths of the propagating modes – the devices are typically of millimeter scale, while the wavelength of the propagating modes is on the micrometer scale. This difficulty results in a large number of degrees of freedom required to discretize the problem which in turn leads to prohibitively large simulation times for each iteration of the design loop. We present a divide-and-conquer approach to enable computer-aided design of photonic devices which relies on the devices consisting of simple function modules with straight waveguides albeit with different numbers of input and output channels. Since each component only supports a small number of propagating modes, we construct scattering matrices which completely characterize the input/output behavior of these components using state-of-art boundary integral equation methods coupled with high-order quadrature rules, and the recursively compressed inverse preconditioning method.

Kosuke Sugita

New Jersey Institute of Technology
ks823@njit.edu

Shidong Jiang

Flatiron Institute, Simons Foundation
Center for Computational Mathematics
sjiang@flatironinstitute.org

Manas N. Rachh

Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS298

Heterogeneity-Induced Mixing and Dissolution Hotspots in Coastal Aquifers Facilitate Karst Propagation

The freshwater-seawater mixing zone is characterized by enhanced chemical activity. However, the the impact of spatial heterogeneity on the dynamics of mixing and calcite dissolution is not well understood. These processes play a key role for the understanding of the development of karst formations. We analyze the impact of different heterogeneity structures and strengths on the mixing and dissolution rates across the saltwater freshwater mixing zone. We observe that the initial heterogeneity structure significantly impacts observed dissolution and mixing patterns,

which sheds some new light on karst propagation in coastal aquifers.

Marco Dentz, Maria Pool, Kevin de Vriendt
IDAEA, Spanish National Research Council (CSIC)
marco.dentz@csic.es, mpoolr@gmail.com, kdvg-geo@idaea.csic.es

MS298

Morphological Attractors in Natural Convective Dissolution

Ever-changing geological features on this planet never fail to capture our imagination and inspire new scientific advances. Among them, the formation of stone forests is one striking geomorphology caused by dissolution and fluid-structure interactions in nature. Recent experiments demonstrate how a soluble body placed in a fluid spontaneously forms a dissolution pinnacle—a slender, upward pointing shape that resembles naturally occurring karst pinnacles found in stone forests. This unique shape results from the interplay between interface motion and the natural convective flows driven by the descent of relatively heavy solute. In this talk, we will discuss a class of exact solutions that act as attractors for the shape dynamics in two and three dimensions. Intriguingly, the solutions exhibit large but finite tip curvature without any regularization, and they agree remarkably well with experimental measurements. The relationship between the dimensions of the initial shape and the final state of dissolution may offer a principle for estimating the age and environmental conditions of geological structures.

Jinzi Mac Huang
University of California, San Diego
machuang@ucsd.edu

MS298

How Fluid-Mechanical Erosion Creates Anisotropic Porous Media

When a porous medium erodes, microscopic changes of the grain morphology give rise to larger-scale features such as channelization. Using a boundary integral formulation, we characterize these changes by simulating erosion of porous media. A Cauchy-integral formulation and associated quadrature formulas enable us to resolve dense configurations of nearly contacting bodies. We observe that substantial anisotropy develops over the course of erosion; that is, the configurations that result from erosion permit flow in the longitudinal direction more easily than in the transverse direction by up to a factor of six. These results suggest that the erosion of solid material from groundwater flows may contribute to previously observed anisotropy of natural porous media.

Bryan D. Quai
Scientific Computing
Florida State University
bquai@fsu.edu

Nick Moore
Colgate University
nickmoore83@gmail.com

Jake Cherry
Florida State University

jrc16g@my.fsu.edu

MS298

Data-Driven Analysis and Modeling of Erosion Networks

Channel formation and branching is widely seen in physical systems where movement of fluid through a porous structure causes the erosion of the medium. We introduce a simple model to capture this feedback mechanism in a multiphase model of flow through a frangible porous medium with dynamic permeability. We explore the model through simulations and examine how the channel morphology changes with different parameters. In addition, given a channel network, we develop a data-driven approach to predict the flow characteristics that led to the network formation.

Chris H. Rycroft
University of WisconsinMadison
rycroft@wisc.edu

MS298

Dissolution in Potential Flow: Maps and Shapes

If we put a dissolving object in a flow, its shape will continuously change. Tracking of the evolving shape requires the solution of coupled flow and transport equation, in an evolving geometry around the shrinking object. If the system is quasi-two-dimensional and the flow-potential, one can use conformal mapping techniques, such as Polubarinova-Galin equation, to track the evolving shapes. Two problems of this kind will be discussed: a disk immersed in a flow and dissolved from the sides and a dissolution finger, propagating into the porous rock. The analytical calculations and numerical simulation results will be compared with the results of microfluidic experiments, with a gypsum block inserted in between two polycarbonate plates, dissolved by inflowing water. Finally, the implications for natural shapes will be briefly discussed.

Piotr Szymczak
Institute of Theoretical Physics
University of Warsaw
piotrek@fuw.edu.pl

MS299

An Acoustic/Transport Splitting Method for the Isentropic Baer-Nunziato Two-Phase Flow Model

We are interested in the computation of compressible two-phase flows with the isentropic Baer-Nunziato two-phase flow model [Baer and Nunziato. A two-phase mixture theory for the deflagration-to-detonation transition (DDT) in reactive granular materials.(1986)] by using a finite volume methods that decouples acoustic and transport phenomena. This approach is well-suited to many industrial applications like the simulation of flows in nuclear reactors when the material velocities are low compared to the sound velocities. For the Euler system, the classic Lagrange-Projection method allows to perform such decoupling by solving the flow equations in Lagrangian coordinates. Unfortunately, the Baer-Nunziato model involves three different material waves that do not allow to exhibit a simple expression of the system using Lagrangian coordinates. Using similar ideas as in [Chalons, Kokh and Spillane. Large time-step numerical scheme for the seven-equation model of compressible two-phase flows.(2011)], we consider an operator splitting strategy that allows to separate an acoustic sub-

system and a transport subsystem for the Baer-Nunziato model. This operator splitting allows to design a numerical scheme with a time-implicit treatment of the (fast) acoustic waves, in order to get rid of a too restrictive CFL condition, and an explicit treatment of the (slow) material waves in order to preserve accuracy, for the isentropic Baer-Nunziato model.

Katia Ait-Ameur

Ecole Polytechnique
katia.ait-ameur@polytechnique.edu

Samuel Kokh

CEA - Maison de la Simulation
samuel.kokh@cea.fr

Marc Massot

CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau Cedex, France
marc.massot@polytechnique.edu

Marica Pelanti

IMSIA, ENSTA, Institut Polytechnique de Paris
marica.pelanti@ensta-paris.fr

Teddy Pichard

École polytechnique
teddy.pichard@polytechnique.edu

MS299

Implementation and Stability Analysis of Self-Adjusting Implicit Multirate Methods

A novel strategy for the implementation of self adjusting multirate approaches is proposed, based on the use of optimal time step estimation also for the locally refined time step. The proposed strategy is applied to efficient DIRK methods and the resulting combination is tested on a number of relevant test cases. Furthermore, a general approach for the stability analysis of multirate methods is presented, which allows to study stability on physically relevant linear model problems. Both results generalize previous work by the author and have been obtained in joint work with F. Casella (Politecnico di Milano), B. Bachmann (Bielefeld University of Applied Sciences), M. Gomez and S. Fernandez (University of Sevilla).

Luca Bonaventura

Politecnico di Milano, Italy
luca.bonaventura@polimi.it

MS299

A Scalable Semi-Implicit Solver for the Barotropic System in a Mode-Split Ocean Model

Most modern ocean models have adopted the barotropic/baroclinic mode-splitting technique for more computationally efficient time integration. Since the two-dimensional barotropic system has a gravity-induced wave speed which is around two orders of magnitude faster than the three-dimensional advection-dominated baroclinic mode, numerical stability and model performance are strongly dependent on numerical methods solving the barotropic system. A scalable semi-implicit barotropic mode solver for the Model for Prediction Across Scale-Ocean (MPAS-O) has been implemented as a competitor of the default explicit-subcycling solver.

The barotropic system is discretized in time with the Crank-Nicolson method. As a Krylov subspace solver, we use the single synchronization biconjugate gradient stabilized solver that has one global synchronization point per solver iteration. To accelerate the convergence of this solver, the restricted additive Schwarz (RAS) preconditioner is introduced. Before the construction of the RAS preconditioning matrix, the global coefficient matrix is reordered to make a near-block diagonal form to obtain optimal performance. Several numerical experiments demonstrate that the semi-implicit barotropic mode solver has faster computational performance up to two times compared to the default solver while maintaining the numerical accuracy. Moreover, GPU acceleration of this semi-implicit solver is underway. We will briefly introduce this work at the minisymposium.

Hyun Kang, Katherine J. Evans

Oak Ridge National Laboratory
kangh@ornl.gov, evanskj@ornl.gov

Philip Jones

Los Alamos National Laboratory
pwjones@lanl.gov

Mark R. Petersen

Los Alamos National Laboratory
Computational and Computer Science Division
mpetersen@lanl.gov

Andrew Salinger

CSRI
Sandia National Labs
agsalin@sandia.gov

Raymond Tuminaro

Sandia National Laboratories
rstumin@sandia.gov

MS299

Operator Splitting in Physics-Dynamics Coupling of the Met Office LFRic Weather and Climate Model

The UK Met Office provides weather forecasts and climate predictions at a range of timescales, for different domains of interest and at varying levels of complexity. Global deterministic model forecasts are currently carried out 4 times a day with ~ 10 km horizontal resolution, while more detailed short-to-medium-range regional forecasts over the UK are run more frequently using a 1.5km resolution. Seasonal and climate configurations use lower resolutions, but with additional complexity in the physics packages and coupling to ocean models. The challenges presented by future generations of computer architecture require a redesign of our software and algorithms, so the Met Office is developing LFRic - its next generation atmosphere model. To span the scales for weather and climate efficiently and accurately, a mixed-finite element non-hydrostatic dynamical core is used alongside finite volume transport and finite difference sub-grid physical parametrizations (e.g., solar radiation, clouds) all combined through a semi-implicit time-step. Achieving greater levels of parallelism and more efficiently targeting spatial resolution will be a key feature to obtain efficient time-to-solution with the greatest level of accuracy. We will discuss two aspects: accessing greater task parallelism through operator splitting in the semi-implicit timestep and refining/coarsening spatial resolution between different components (e.g., dynamical core,

sub-grid physics) of the system.

Ben Shipway
Met Office

ben.shipway@metoffice.gov.uk

MS299

Leveraging a Semi-Discrete Error Analysis Framework to Improve Coupling Between Atmosphere Processes in E3SM

The Energy Exascale Earth System Model (E3SM), like other global climate models, seeks to solve a truly multiphysics problem that includes processes such as fluid flow, heat exchange, chemical reactions, and phase changes. Those processes are partitioned into individual models that, along with the associated numerical discretization, are each developed more-or-less independently of the others, which has resulted in process coupling that remains essentially a sequential splitting approach. Thus, this work leverages a semi-discrete error analysis framework to investigate the potential both of alternative coupling approaches in the literature and of new coupling approaches developed specifically to the software architecture of the E3SM atmosphere model (EAM). The analysis framework isolates the coupling error from that of the temporal and spatial discretization, avoiding the need to rederive error terms for each combination of coupling approach and temporal and/or spatial discretization present in EAM. By leveraging such analysis, an alternative coupling approach for aerosol processes has been developed that substantially improves dust lifetime simulation, which is crucial to reliable climate prediction. Both the improvement in aerosol process coupling and the progress on improving coupling between other atmosphere processes will be presented. Prepared by LLNL under Contract DE-AC52-07NA27344. LLNL-ABS-840173.

Christopher J. Vogl

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

Hui Wan

Pacific Northwest National Laboratory
Richland, WA 99354
hui.wan@pnnl.gov

Carol S. Woodward, Steven Roberts

Lawrence Livermore National Laboratory
woodward6@llnl.gov, roberts115@llnl.gov

Christopher J. Vogl

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

MS300

High-Order Diagonally Implicit Runge-Kutta Schemes for Stiff and Oscillatory Problems

Fully implicit Runge-Kutta (RK) schemes possess superior stability properties over explicit RK schemes. Thus, they are employed to accurately integrate stiff or oscillatory systems of ordinary differential equations or semidiscretizations of partial differential equations. However, the computational cost of implementing implicit RK methods often limits their applicability, especially when combined with large-scale and high-resolution numerical simulations. As

an efficient alternative, diagonally implicit Runge-Kutta (DIRK) schemes are sought. We present an optimization approach for designing DIRK schemes of order six and higher that are A-stable and stiffly accurate. The accuracy and robustness of these schemes are shown via extensive numerical tests on diverse problems. This optimization approach can be adopted for designing numerical integrators with other desired properties.

Yousef Alamri

King Abdullah University of Science and Technology
yousef.alamri@kaust.edu.sa

MS300

Avoiding Order Reduction with Explicit Runge-Kutta Exponential Methods in Nonlinear Initial Boundary Value Problems

A technique will be shown to recover the classical order of the method when explicit exponential Runge-Kutta methods integrate reaction-diffusion problems. In the literature, methods of high enough stiff order for problems with vanishing boundary conditions have been constructed, but that implies restricting the coefficients and thus, quite probably, the computational cost. In contrast, the technique being suggested is cheaper because it just needs, for any method, to add some terms with information on the boundaries. Moreover, time-dependent boundary conditions will be directly tackled.

Begoña Cano

Universidad de Valladolid
bcano@uva.es

MS300

Emerging Time Integration Methods for Atmospheric Modeling

Exponential time integrators are a very accurate approach for stiff systems of differential equations. In recent years, many advances have been made to improve their computational efficiency. In numerical weather prediction models, it is common to use implicit or semi-implicit integrators with large time steps to reduce the cost of the solution. Exponential methods have the advantage of allowing time steps as large as the implicit schemes, while correctly simulating all relevant wave dispersions. In this presentation, we investigate the potential of different exponential schemes in the context of atmospheric modelling using idealised test cases of the shallow water and Euler equations.

Stéphane Gaudreault

Environment Canada
stephane.gaudreault@ec.gc.ca

MS300

Applications of Exponential Integrators to Chemical Combustion Simulations

Exponential integrators have been shown as an effective approach to solve stiff systems of differential equations. These methods possess good stability properties all while offering computational savings over implicit methods for some problems where efficient preconditioners are not available. We will discuss the application of variable time stepping exponential methods of EPIRK-type to chemical combustion problems. These variable time-stepping schemes leverage the KIOPS algorithm to quickly and accurately approximate the exponential-like φ -functions which constitute

the main computational expense of exponential integrators. Combustion problems are notorious for stiffness that arises from both the chemical kinetics portion of the models as well as transport mechanisms driving the process. We demonstrate that even straight-forward, i.e. out-of-the-box, application of exponential methods to these systems yields performance comparable to the implicit methods of CVODE (LLNL) and TrBDF2 (SNL). We carry out a number of experiments using GRI, NHeptane, NButane, NDecane, and Hydrogen combustion mechanisms to compare performance of exponential and implicit methods over the course of ignition at a single point. Our numerical studies include both overall comparison of the performance of the integrators as well as detailed analysis of underlying computations with Krylov-based techniques.

Jared Stewart

University of California Merced
jstewart23@ucmerced.edu

MS301

Building a Knowledge Graph for CSE: From Prototype to Community Project

Mathematical computing knowledge is produced at an immense, and seemingly ever increasing, speed. Very little of it is organised in meaningful ways, making its discovery, insight and discussion harder every year. Following new developments in a given field is time consuming even for experts. Entering a new specialisation is daunting for students. We will show how building a knowledge graph for scientific computing can address these issues. We have created an ontology that semantically links mathematical problems with algorithms, publications and implementations. This ontology encodes possible relationships between the entities in the graph and these connections can be explored using our web-based query frontend. This query frontend enables non-experts to quickly gain an overview of available methods and software for numerical problems in their scientific work. The frontend makes variations of essentially the same algorithms easily discoverable. It allows tracking new publications or software implementations connected to a certain mathematical problem. We will be inviting feedback for our plans to grow this knowledge graph into a community-driven platform with an open, freely accessible API. Our efforts are part of the scientific computing task area in the Mathematical Research Data Initiative (MaRDI), a consortium in the German National Research Data Infrastructure (NFDI).

René Fritze, Frank Wübbeling, Hendrik Kleikamp, Mario Ohlberger, Stephan Rave
University of Münster
rene.fritze@wwu.de, frank.wuebbeling@uni-muenster.de,
hendrik.kleikamp@wwu.de, mario.ohlberger@wwu.de,
stephan.rave@wwu.de

MS301

MaRDI - The Mathematical Research Data Initiative within the German National Research Data Infrastructure (NFDI)

Mathematical research data is vast, complex and multifaceted. It emerges within mathematics but also in other scientific disciplines such as physics, chemistry, or engineering. The growing amount of research data challenges an old requirement in science: its reproducibility and the reusability of results. In an attempt to answer this challenge at current level, the FAIR principles have been formulated.

Yet, despite the existence of special solutions a comprehensive infrastructure for research data that supports the research process and implements the FAIR principles in science or in mathematics is missing. In 2020 Germany started the National Research Data Infrastructure (NFDI) which will be created along the needs of all scientific disciplines and aims at FAIR use of research data. Within the NFDI the Mathematical Research Data Initiative (MaRDI) is the consortium for developing a research data infrastructure for mathematics. Starting with the areas of computer algebra, scientific computing, statistics and machine learning, MaRDI will develop standards for confirmable workflows and certifiable mathematical results and provide new services that assist the research cycle up to peer-review in the publication process. Standardized formats, data interoperability and application programming interfaces need to be established to ensure the ease of use of data across broad disciplines. In this talk, we present concepts and ideas behind MaRDI.

Thomas Koprucki, Karsten Tabelow

Weierstrass Institute for Applied Analysis and Stochastics
koprucki@wias-berlin.de, tabelow@wias-berlin.de

MS301

From Dark Data to Fair Data in Computational Science, Engineering and Mathematics

Dark data are data that are poorly managed and they are becoming a major problem for science and industry. They are diametrically opposed to FAIR data because their epistemic status is unclear, and they are neither findable, accessible, interoperable, nor reusable. As such, research data may be uncurated, unavailable, unannotated, biased, or incomplete. The talk will explore representative examples of dark data in the fields of scientific computing and in statistics. Dark data as a negative category serve well as an analytical tool to work out how the data management processes can be improved and aligned with the FAIR principles with respect to reproducible research. This can be enabled by documenting epistemic metadata, i.e., information about the knowledge status of data. Epistemic metadata include, besides documenting pure process metadata, also the reflection and integration of scientific underpinnings of the research. This in turn means the evaluation of disciplinary scientific practices and conventions in detail. It is here that NFDI consortia, such as the Mathematical Research Data Initiative (MaRDI), can provide dedicated support by encouraging academic communities to engage in community-driven metadata standardization. In case of mathematics this means setting up semantic technology (metadata, ontologies, knowledge graphs) for problems, models, algorithms and software. The effort will be highlighted in the talk.

Björn Schembera

University of Stuttgart
bjoern.schembera@mathematik.uni-stuttgart.de

Martin Thomas Horsch
Norwegian University of Life Science
martin.thomas.horsch@nmbu.no

Dorothea Iglezakis
University Library
University of Stuttgart
dorothea.iglezakis@ub.uni-stuttgart.de

Cedric Riethmüller, Dominik Goeddeke

Institute of Applied Analysis and Numerical Simulation
University of Stuttgart
cedric.riethmueller@mathematik.uni-stuttgart.de,
dominik.goeddeke@mathematik.uni-stuttgart.de

MS301

Design of Cse Workflows Using Fair Principles: A Framework for Abstracting Meta-Data into Multi-Layered Descriptions

Numerical algorithms and computational tools are essential for managing and analyzing complex data processing tasks. With the availability of meta-data and parameter-driven simulations, the demand for automated workflows to reproduce computational experiments across platforms has grown. In general, a computational workflow is defined as a step-by-step description for accomplishing a scientific objective, expressed in terms of tasks and their data dependencies. Characterized through their input-output relation, workflows are designed such that the associated meta-data can be used interchangeably and redundantly. As a part of the MaRDI consortium [www.mardi4nfdi.com] on research data management in numerical mathematics, we develop a novel computational framework that focuses on automation of abstracting meta-data embedded in an ontology of mathematical objects while negating the underlying execution and environment dependencies into multi-layered descriptions. Herein, we consider various computational experiments as case studies, and systematically incorporate them into our workflow tool and data provenance framework. Furthermore, we show how best to apply the FAIR principles to computational workflows, such that abstracted components are Findable, Accessible, Interoperable, and Reusable in nature. Going forward, we also plan to provide the developed FAIR CSE workflow in electronic lab notebooks (ELNs) with user-friendly guidelines.

Pavan Veluvali

MPI for Dynamics of Complex Technical Systems
veluvali@mpi-magdeburg.mpg.de

Jan Heiland, Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
heiland@mpi-magdeburg.mpg.de, benner@mpi-magdeburg.mpg.de

MS302

Pde Analysis of Real-World Problem - Groundwater Modelling

A new academic-business model involving a strategic collaborative effort between a mathematical science-based research centre and a business-entrepreneurial entity is suggested. This model is framed within a Quadruple Helix (QH) structure, with four main distinct clusters (government, industry, academia and community user). This academic-business model is then implemented via a proposed mathematical/industrial mathematics modelling activity, which is essentially rested on the knowledge transfer mechanism and entrepreneurial process. This underlying activity apparently would fulfil our sequence of steps in the modelling of a real-world problem within QH framework. This programme is then applied to a real-world industrial problem. This problem would link up to the QH clusters and is described by a presumed workable industrial problem, incorporating PDE analysis and Green function method on Surface Water (SW), Riverbank Filtration (RBF) and Groundwater (GW) modelling. The

main goal is to generate an integrated water management scheme which would ensure an optimal conjunctive use of SW, RBF and GW sources of treated clean water, thus preventing unscheduled disruption of clean water supply to Water Authority account holders. The final outcome is to achieve and sustain a reasonably practicable and optimal conjunctive/integrated water management system for a community of Water Authority users.

Zainal Aziz

UTM Centre for Industrial and Applied Mathematics
Universiti Teknologi Malaysia
zainalaz@utm.my

MS302

Generation of Aesthetic Curves and Surfaces by Integrable Geometry

We first consider the log-aesthetic curves (LAC), which is a family of planar curves developed and used in the area of industrial design as shape elements with built-in aesthetic nature. The LAC has been originally proposed by extracting the common properties from thousands of plane curves that car designers regard as aesthetic, which includes various spirals as special cases, such as the circle involute, the clothoid, the Nielsen's spiral, and the logarithmic spiral. We present a new mathematical framework of the LAC on the theory of integrable systems and the similarity geometry, a Klein geometry associated with the similarity transformation group. Using this framework, we show that the similarity curvature is governed by the stationary Burgers equation and that the LAC is characterized by the variational principle. Namely, the LAC can be regarded as the similarity geometry analogue of the celebrated Eulers elasticae in the Euclidean geometry. Based on this result, we present some generalizations of LAC: (i) integrable discretization, (ii) space curves, and (iii) surfaces. We expect that those curves and surfaces may be useful for generating aesthetic shapes.

Kenji Kajiwara

Institute of Mathematics for Industry
Kyushu University, Japan
kaji@imi.kyushu-u.ac.jp

MS302

Effective Modelling for PCBA Warpage Simulation

A printed circuit board (PCB) is a multi-layered composite which includes many layers of copper circuits and dielectric materials. It is used as a backbone to carry and connect various electronic components, i.e. PCBA, to achieve certain functions. However, because the components are decreasing in size, the reliability of PCBA have become a critical issue. One of the serious problems is the warpage, which is induced by the thermal mismatch between unevenly-distributed circuits, dielectric layers and assembled components during the manufacturing process which experiences a temperature variation. This may cause defects in the assemblies. It is helpful to simulate in advance to evaluate. However, since the circuit on a PCB is very tiny compared with the size of the PCB or the components, it is too difficult to build a finite element model to include all the details of circuits and components. Hence, a feasible way is needed. Here, a new effective modelling approach which adopts equivalent material properties is proposed for PCBA warpage simulation. In this approach, the multi-layered PCB with connected components (PCBA) are meshed and connected properly, and the model is not

too large while the circuits effect is still considered. With the proposed approaches, this method can be useful and efficient to simulate the PCBA manufacturing process to evaluate the warpage of the PCBA. Even those cases with irregular-shaped PCBs and shielding frames can be also treated efficiently.

Ming-Hsiao Lee, Jiunn-Horng Lee, Jen-Gaw Lee,
Chih-Min Yao
National Center for High-performance Computing
9303103@narlabs.org.tw, jhlee@nchc.narl.org.tw,
jglee@nchc.narl.org.tw, 9403116@narlabs.org.tw

MS302

A Predictive System for Assessment the Right and Wrong Parties in a Car Accident from Traces of the Damage

The problem of jamming in Bangkok is a top issue that still needs to be addressed. There are several factors that can cause disruption. And road accidents are one of the major causes. This research was requested by the Thai General Insurance Association for this problem. This research creates a system for assessing the damage effects of a car as if it were an expert in the investigation. The model predicts the severity and damaged parts of the vehicle, reducing the amount of time and procedure for determining damage in the event of an accident by surveyors. The technology was used to analyze vehicle damage images together with the principles of machine learning in analysis. This is coupled with the use of correlation principles and text analysis to determine the causality of the chronology to be used for analysis based on the information recorded from both parties. Coupled with the system's ability to identify the location and environment of the accident scene, it is used through a central system that stores images and text. The diagnosis uses principles consistent with the database to store patterns of accidents in various ways in accordance with the Disaster Survey Manual under the Land Traffic Act for monitoring the damage. The system can initially determine whether a party is right or wrong, and it can also reduce the process of assessing damages for claims, and reporting an accident for convenience and speed, as well as reducing obstructing traffic on the road.

Busayamas Pimpunchat
Department of Mathematics
King Mongkuts Institute of Technology Ladkrabang
busayamas.pi@kmitl.ac.th

MS302

Heat Kernels of Quantum Interaction Models for Quantum Computers

We discuss the derivation of analytical formulas for the heat kernel, propagator and partition functions of the quantum Rabi model and asymmetric quantum Rabi model, which are the most fundamental models of quantum optics. Although these were introduced as theoretical models, it is known that they have predicted the experimental measurements of the realization of quantum entangled states using superconducting artificial atoms. We also refer to explicit propagators (corresponding to path integrals) and quantum computation.

Masato Wakayama
NTT Institute for Fundamental Mathematics

masato.wakayama.hp@hco.ntt.co.jp

MS303

Mixed Precision in Pivoting Avoiding QR

Rank-deficiency plagues least-squares problems. While the QR factorization with Column Pivoting (QRCP) offers a robust solution, it suffers from scaling issues due to its communication-bound nature. The Pivoting Avoiding QR factorization (PAQR) is an alternative that benefits from a scaling similar to that of the original QR factorization but at the expense of fewer theoretical guarantees compared to QRCP. We show how the application in different precisions aids in retaining the speed of a traditional QR factorization, while still selecting an appropriate column order. Within PAQR, there exists a tolerance parameter (used to detect any possible deficiency) that influences the algorithm's behavior, this parameter can be selected so as to fit the accuracy of FP32 or FP16 arithmetic. In this talk we investigate the use of lower-precisions and mixed precisions using PAQR in the context of deficient least-squares problems.

Daniel R. Bielich
University of Colorado at Denver
drbieli@icl.utk.edu

MS303

An Attempt of Exploiting Low Precision Computing in the GMRES(m) Method

Recently, the use of low precision computing such as FP16 has attracted much attention. We have studied approaches of accelerating the computations in numerical linear algebra, e.g., solution process of sparse and large liner systems. In this presentation, we show our attempt of introducing low precision computing into the GMRES(m) method, which is one of well-known Krylov subspace methods for non-symmetric problems. First, we present numerical results of a mixed precision variant of the GMRES(m) method using FP64 and FP32. We discuss the behavior of the mixed precision GMRES(m) method and its potential to accelerate the solution process. Second, we consider introducing lower precision computing than FP32 into the GMRES(m) method. We implement different variants of the mixed precision GMRES(m) method and examine their behavior. Through the numerical experiments, we investigate the possibility of aggressively using the lower precision computing in the GMRES(m) method and discuss issues for further performance improvement.

Takeshi Fukaya, Zhao Yingqi, Takeshi Iwashita
Hokkaido University
fukaya@iic.hokudai.ac.jp, yqzhao2016@hotmail.com,
iwashita@iic.hokudai.ac.jp

MS303

Adaptive Precision Sparse Iterative Solvers

We introduce a mixed precision algorithm for computing sparse matrix-vector products and use it to accelerate the solution of sparse linear systems by iterative methods. Our approach is based on the idea of adapting the precision of each matrix element to their magnitude: we split the elements into buckets and use progressively lower precisions for the buckets of progressively smaller elements. We carry out a rounding error analysis of this algorithm that provides us with an explicit rule to decide which element goes

into which bucket and allows us to rigorously control the accuracy of the algorithm. We implement the algorithm on a multicore computer and obtain significant speedups (up to a factor $7\times$) with respect to uniform precision algorithms, without loss of accuracy, on a range of sparse matrices from real-life applications. We showcase the effectiveness of our algorithm by plugging it into a GMRES solver for sparse linear systems and observe that the convergence of the solution is essentially unaffected by the use of adaptive precision.

Roméo Molina

Sorbonne Université, CNRS, LIP6 and Université Paris-Saclay
romeomolina@gmx.fr

Stef Graillat

University Pierre and Marie Curie (UPMC)
LIP6
stef.graillat@lip6.fr

Fabienne Jezequel

LIP6, Sorbonne Université
fabienne.jezequel@lip6.fr

Theo Mary

Sorbonne Université, CNRS, LIP6
theo.mary@lip6.fr

MS303

Mixed Precision Iterative Refinement for Low-Rank Matrix and Tensor Approximations

We present a new mixed precision algorithm to compute low-rank matrix and tensor approximations, a fundamental task in numerous applications in scientific computing and data analysis. Our algorithm is reminiscent of the iterative refinement framework for linear systems : we first compute a low-rank approximation in low precision and then refine its accuracy by iteratively updating it. We carry out an error analysis of our algorithm which proves that we can reach a high accuracy while performing most of the operations in low precision. We measure the cost of the algorithm depending on the numerical rank of the input as well as the speed ratio between low and high precision arithmetic. We identify two situations where our method has a strong potential : when the hardware provides fast low precision matrix-multiply-accumulate units, and when the numerical rank of the input is small at low accuracy levels. We assess the potential of our algorithm for computing various low-rank matrix and tensor decompositions such as SVD, QR, Tucker, hierarchical Tucker, and tensor-train.

Matthieu Robeyns, Marc Baboulin, Oguz Kaya

Université Paris-Saclay
matthieu.robeyns@universite-paris-saclay.fr,
marc.baboulin@lri.fr, oguz.kaya@universite-paris-saclay.fr

Theo Mary

Sorbonne Université, CNRS, LIP6
theo.mary@lip6.fr

MS303

Mixed Precision Algebraic Multigrid on GPUs

We present the first GPU-native platform-portable algebraic multigrid (AMG) implementation that allows the

user to use different precision formats for the distinct multigrid levels. The AMG we present uses an aggregation size 2 parallel graph match as the AMG coarsening strategy. The implementation provides a high level of flexibility in terms of configuring the bottom-level solver and the precision format for the distinct levels. We present convergence and performance results on the GPUs from AMD, Intel, and NVIDIA, and compare against corresponding functionality available in other libraries.

Yuhsiang M. Tsai

Karlsruhe Institute of Technology
yu-hsiang.tsai@kit.edu

Natalie N. Beams

University of Tennessee
nbeams@icl.utk.edu

Hartwig Anzt

University of Tennessee, U.S.
hantz@icl.utk.edu

MS304

Model Reduction Using Moment Models

The high dimensionality of many mathematical models in science and engineering leads to prohibitively large computational cost for even an approximate numerical solution. However, an accurate solution of the full model is often not even necessary, as only a small set of variables suffices to characterize the main behavior of the solution. This poses the question of model reduction: How can we efficiently reduce the complexity of the model and arrive at a reduced model, that is both sufficiently accurate and computationally feasible? In this talk, we first briefly discuss different model reduction techniques for kinetic equations. We then focus on moment models as one way to reduce the full model to a set of analytical, lower dimensional equations. The benefits of moment models are the mathematically sound derivation, their hierarchical structure, and the possibility to assess analytical properties of the model from the resulting equations. We consider, among others, examples from rarefied gases and free-surface flows and illustrate those by numerical simulations.

Julian Koellermeier

University of Groningen
j.koellermeier@rug.nl

MS304

Hybrid Low-Rank Discretization of the Vlasov Equation

The Vlasov-Maxwell system describes the evolution of the phase-space distribution of a plasma in its self-consistent fields. This kinetic description of a plasma is computationally very challenging due to the dimensionality of the phase space. Low-rank solutions to the Vlasov-Maxwell equations have the potential to compress the distribution function which enables a fast and accurate solution. However, most simulations so far have been for relatively simple test cases which exhibit a strong low-rank structure. Parallelizing low-rank solvers on a large scale is hard which hampers the applicability in state-of-the-art physics simulations. In joint work with Allmann-Rahn and Grauer, we have recently shown that more complex test cases can be tackled with a compression that is only confined to the velocity space. Another way of exploiting low-rank solutions in large-scale simulations is to use a hybrid simulation that

combines a low-rank solution with a particle-in-cell simulation. In this talk, we consider a particle-in-cell simulation with a low-rank solution as a dynamic control variate. For benchmark test cases, it is shown that the low-rank control variate—in contrast to a time-independent control variate—continues to improve on the accuracy of the particle solution when nonlinear effects become dominant.

Katharina Kormann
Ruhr University Bochum
k.kormann@rub.de

MS304

A Two-phases Hybrid Method for Modelling Neutral Particles in the Plasma Edge of a Fusion Device

In nuclear fusion reactors, neutral particles play an important role in shielding the reactor walls and in particular the divertor targets from the hot plasma in which the fusion reactions take place. The neutral particle behaviour is governed by a kinetic equation that is difficult to solve due to the high dimensionality of the phase space and the high collisionality between the neutral particles and the plasma. High dimensionality is typically resolved by resorting to Monte Carlo methods, but in high collisional regimes a pure Monte Carlo solver becomes intractable. To reduce the computational cost, recent research focuses on hybrid fluid/particle methods that exploit the high-collisionality of the system. In this talk, we elaborate on a new hybrid fluid/particle method that splits the particle distribution in two phases: a fluid phase and a particle phase. The particle phase is solved using a Monte Carlo scheme, the fluid phase follows a reduced fluid model that is derived from the kinetic equation in the high collisional limit. The two phases co-exist in the whole domain, but the amount of mass in each phase is determined by the local collisionality. This way, the amount of particles treated with Monte Carlo is reduced in high collisional regimes, where the cost of Monte Carlo is high, but the reduced fluid model is accurate. In low collisional regimes where the reduced fluid model is not accurate, but Monte Carlo is cheap, most of the mass is in the particle phase.

Vince Maes, Wouter Dekeyser
KU Leuven
vince.maes@kuleuven.be, wouter.dekeyser@kuleuven.be

Julian Koellermeier
University of Groningen
j.koellermeier@rug.nl

Martine Baelmans, Giovanni Samaey
KU Leuven
tine.baelmans@kuleuven.be,
vanni.samaey@kuleuven.be

MS304

Regularized, Structure Preserving, Neural Network Based Entropy Closures for the Boltzmann Moment System

The main challenge of large scale numerical simulation of radiation transport is the high memory and computation time requirements of discretization methods for the Boltzmann equation. In this work, we derive an approximation to the entropy closure method to accurately compute the solution of the multi-dimensional Boltzmann moment system with low memory footprint and competitive computational time. We extend structure preserving, neural net-

work based approximations of the closure of the moment system to the context of regularized entropy closures. The main idea is to interpret structure preserving neural network approximations of the regularized entropy closure as an two stage approximation to the original entropy closure. We conduct numerical analysis of this approximation and investigate optimal parameter choices. Our numerical experiments demonstrate that the method has a much lower memory footprint than traditional methods with competitive computation times and simulation accuracy.

Steffen Schotthöfer
Karlsruhe Institut of Technology
steffen.schotthoefer@kit.edu

MS305

Closely Interacting Rigid Particles in Stokes Flow with Contact Handling and Local Corrections

The multiblob method is an inexpensive and approximate technique for simulations of particles in Stokes flow. We improve on its accuracy for rigid axisymmetric particles in 3D free space by reducing two different error contributions: self-interaction and pair-interaction errors, dominant in the far-field and near-field respectively. The later reduction is done by local pair-corrections as inspired by Stokesian dynamics. As we study the particles dynamically, more challenges are introduced – despite the more accurate description – something that holds true for any choice of Stokes solver. Rigid particles in a Stokesian fluid can physically not overlap, as a thin layer of fluid always separates a particle pair, exerting increasingly strong repulsive forces on the bodies for decreasing separations. Numerically, resolving these lubrication forces comes at an intractably large cost and particle collisions and overlaps do occur. In this talk, non-overlap constraints, in terms of the Euclidean distance between boundary points on the particles, are represented via a barrier energy. We solve for the minimum magnitudes of repelling contact forces between any particle pair in contact to correct for overlaps, by enforcing the barrier energy to be zero at the next time step. The method is tested on suspensions of spheres, rods and boomerang shaped particles where collision free configurations can be obtained at all instances in time.

Anna Broms
KTH Royal Institute of Technology
annabrom@kth.se

Anna-Karin Tornberg
KTH
akto@kth.se

MS305

Dynamics of Torque-Dipolar Micro-Swimmers

We present a model and simulations for micro-swimmers that take into account the counter-rotation of the body and flagella, as seen in motile bacteria or spermatozoa. The disturbance fluid flow of one such swimmer now contains a torque-dipole singularity in addition to the well-known force-dipolar singularity. The linear analysis of the coarse-grained model shows an instability for a range of parameters, which we summarize in a phase diagram. Lastly, we show large-scale simulations of torque-dipolar micro-swimmers and illustrate the collective behavior in the regions of parameter space indicated by the stability analysis.

Enkeleida Lushi

New Jersey Institute of Technology
lushi@njit.edu

MS305

Sedimentation of Flexible Fibers in Structured Environments

The motion of flexible fibers often happens in complex media that are structured by obstacles. Examples range from the transport of biofilm streamers through porous media to the design of sorting devices for DNA molecules. For large number of such problems, the dynamics of the fibers result from the complex interplay between internal elastic stresses, contact forces and hydrodynamic interactions with the walls and obstacles. By means of numerical simulations, experiments and analytical predictions, we investigate the dynamics of flexible fibers settling in a viscous fluid embedded with obstacles of arbitrary shapes. We show that, for a single fiber, the fiber's deformation, orientation, velocity and trajectory are strongly altered by its properties and the presence of the obstacle. We introduce a numerical model to capture the aforementioned changes, both qualitatively and quantitatively. We identify and characterize two types of events : trapping and gliding, the former involves nontrivial trapping conformations on sharp obstacles, and the latter, a lateral dispersion. We show how these effects can be leveraged to propose a new strategy to sort fibers based on their size and/or elasticity.

Ursy Makanga

LadHyx, CNRS/Ecole Polytechnique
ursy.makanga@polytechnique.edu

Mohammadreza Sepahi, Camille Duprat
LadHyX, CNRS/Ecole Polytechnique
m.sepahi@aut.ac.ir,
camille.duprat@ladhyx.polytechnique.fr

Blaise Delmotte

LadHyx, CNRS/Ecole Polytechnique
blaise.delmotte@ladhyx.polytechnique.fr

MS305

An Efficient Coupled Lattice-Boltzmann and Discrete Element Method for Simulation of Particulate Flows

Many interesting particulate flow problems can only be studied using efficient numerical methods. We present a method based on a lattice-Boltzmann fluid coupled to point particles that interact with each other via the Discrete Element Method. Two-way coupling between the solid and fluid phases is used. The effect of (locally) large solid volume fractions on the fluid is accounted for by using a lattice-Boltzmann formulation based on the generalized Navier-Stokes equations. We discuss how intricacies of the coupling scheme must be considered to avoid unphysical behavior and instabilities which were reported in previous studies. Through this analysis, we identify under what conditions the model can be expected to yield accurate results. The method is validated by comparison to experiments and established numerical studies. Our method is built upon the LBM solver Musubi and scales well on the Dutch super-computer Snellius. It is especially suitable for simulating flows containing many small particles in complex geome-

tries. To demonstrate this, we apply our method to study the distribution of particles injected from a catheter into an arterial geometry. Various catheter positions are evaluated, showing how the model could be used to optimize treatment in targeted drug delivery applications. In future work we will apply our model to realistic anatomic geometries to study the injection of radioactive particles in radioembolization, a treatment for liver cancer.

Tristan G. Vlogman

University of Twente
t.g.vlogman@utwente.nl

Kartik Jain

University of Siegen, Germany
k.jain@utwente.nl

MS306

Trainability of Quantum Control: Barren Plateaus and Overparameterization

The fate of emerging quantum technologies crucially depends on our ability to devise efficient ways to manipulate, control and scale quantum systems. In this talk, we leverage recent results in the field of variational quantum computing to analyze two key phenomena that critically determine the trainability of quantum optimal control systems. First, we connect the presence or absence of *barren plateaus* (exponentially concentrated gradients in the control landscape) with the dimension of the *dynamical Lie algebra* (DLA) obtained from the generators of the control system. Our results suggest that systems with exponentially large DLAs will exhibit barren plateaus, and thus suffer from severe trainability issues. Next, we analyze the *overparametrization* phenomenon, a complexity phase-transition in the control landscape. Underparametrized systems do not have enough parameters to explore all potential directions in state space, translating into *false* local minima that dominate the landscape. Conversely, when the number of parameters is beyond some critical number, these spurious local minima start disappearing and the performance of the optimizer is greatly improved. Remarkably, we show this transition point is connected with the dimension of the DLA, suggesting that exponential-DLA systems (for example *controllable* ones) will never access the overparametrized regime, suffering from very unfavorable landscapes.

Martin Larocca, Marco Cerezo, Patrick Coles

Los Alamos National Laboratory
larocca@lanl.gov, cerezo@lanl.gov, pcoles@lanl.gov

Nathan Ju

University of Illinois at Urbana-Champaign
34ndju@gmail.com

Kunal Sharma

Louisiana State University
kunals2693@gmail.com

Gopikrishnan Muraleedharan, Piotr Czarnik

Los Alamos National Laboratory
gopu90@unm.edu, pczarnik@lanl.gov

Diego Garcia-Martin

Barcelona Supercomputing Center

diego.garcia@bsc.es

MS306

Realizing Complex Quantum Gates with Minimal Duration Using Intermediate Targets

We consider the optimal control problem of determining electromagnetic pulses for implementing unitary gates in a closed quantum system. Standard gradient-based optimization techniques work well for realizing unitary gates in smaller quantum systems, but can encounter convergence problems for larger agglomerated gates that involves more than two coupled qubits. In this talk we present a continuation technique that parameterizes intermediate gates to gradually morph from trivial free evolution to a complex unitary gate. By successively increasing the complexity of the target gate, the initial gate duration can be small and gradually increased during the continuation, resulting in an accurate estimate of the minimum gate duration. The proposed technique can be used for finding optimal control pulses for driving larger quantum systems, realizing Hamiltonian simulations, and constructing a family of smoothly varying control pulses for realizing parameterized gates.

Anders Petersson

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
petersson1@llnl.gov

Stefanie Guenther

Lawrence Livermore National Laboratory
guenther5@llnl.gov

MS306

Higher Radix Quantum Computation Through Optimal Control Under Experimental Constraints

Current quantum hardware consists of tens or up to few hundreds of qubits, depending on the underlying architecture. Maximizing the utility of these machines involves not only refining the individual qubits but also improving their control. High-fidelity gates enabled by advanced control techniques applied to an increasing number of qubits will be necessary to demonstrate practical applications. I will present some recent work on maximizing quantum resources by breaking the two-state abstraction and using qutrits and ququarts for computation instead. Here optimal control methods allow for faster direct state manipulation, enabling well-performing higher-radix computation architectures. Based on collaborations with experimental groups, I will further address implications of experimental settings on control methods, discussing constraints as well as opportunities for optimal control approaches.

Maximilian Seifert, Fred Chong

University of Chicago
lmseifert@uchicago.edu, chong@cs.uchicago.edu

MS306

Quantum Collocation Algorithms and Iterative Learning Control

Optimization of control pulses for quantum computing hardware has received significant attention in recent years. However, current methods used in this field are primarily based on shooting methods, in which the dynamics of the system driven by the pulse are simulated forward in time, gradients with respect to pulse parameters are computed,

a gradient-based update is applied to the pulse, and the process is repeated until convergence. However, shooting methods suffer from many shortcomings, including poor numerical conditioning and difficulty dealing with many types of constraints. In contrast, we propose the use of collocation methods for these quantum optimization problems, in which both state and control variables are both simultaneously optimized subject to discretized dynamics constraints. Collocation methods offer the ability to handle general nonlinear constraints on states and controls, the ability to leverage dynamically infeasible warm starting for faster convergence, and dramatically improved numerical properties over shooting methods. We demonstrate the capabilities of our quantum collocation method, PICO, on several examples, including a minimum-time gate transition for a multi-level transmon system.

Aaron Trowbridge, Zachary Manchester

Carnegie Mellon University
atrowbri@andrew.cmu.edu, zmanches@andrew.cmu.edu

MS307

Physics-Dynamics-Chemistry Coupling with Components of Different Resolutions in LFRic

Can we save computational resources in NWP and climate models without degrading the quality of the solution, by running different parts of the model at different resolutions? In traditional models, the physics parametrizations, dynamical core and any chemistry components use the same grid. In this talk, I will present an approach to run these components at different resolutions to one another, within the Met Offices new LFRic model, designed for the next generation of supercomputers. The talk will give an overview of a geometric framework for mapping fields between meshes, which preserves certain desirable mathematical properties as the fields are mapped. Only changes of resolution in the horizontal grid are considered, and the meshes are limited to cases in which the cells of the finer mesh are exactly nested within the cells of the coarser mesh. Key to the design of LFRic is the PSyclone code generation system. This should make LFRic portable between different supercomputer architectures, and to allow it to be easily adapted to different hardware.

Thomas M. Bendall

UK MetOffice
thomas.bendall@metoffice.gov.uk

Alex Brown

UK Met Office
alex.brown@metoffice.gov.uk

MS307

GPU Code Generation for Finite Elements with Firedrake

GPUs and accelerators are now common place on large HPC machines and to harness all of the compute power available on such systems these devices can no longer be ignored. However, programming for the GPU remains a difficult task, and optimising for a specific vendor/model is harder still. Using Firedrake we can remove much of this difficulty by using automatic code generation and targeting the GPU directly, which allows users to utilise accelerators for their PDE simulations. For most projects a "rewrite for GPUs" or a "GPU version" can be created, this is not the case for Firedrake. We create a pathway for automatically generating GPU code for any PDE a user can input,

as well as maintaining all existing CPU functionality. By not completely rewriting the codebase we demonstrate the power of the code generation approach. We present the early progress made in offloading Firedrakes functionality on to accelerators. The main focus is on adding GPU assembly for one-forms to enable matrix free operator evaluation and enabling GPU based solvers. We will demonstrate the functionality within Firedrake to optimise for GPUs specifically and how to minimise host to device transfers. These changes utilise the GPU support within loo.py and PETSc to produce efficient and scalable simulations. Some early results are included that show some of the preliminary functionality for GPU platforms.

Jack Betteridge
Department of Mathematics
Imperial College London
j.betteridge@imperial.ac.uk

MS307

The Devito DSL and Compiler Framework: From Symbolic PDEs to HPC Code

Partial differential equations (PDEs) are used in various scientific domains, such as seismic and medical imaging. In several contexts, they are solved numerically by the finite difference (FD) method. When modelling real-world phenomena, these methods are complex to derive, implement and computationally expensive. We present Devito, an open-source framework written in Python that addresses these issues. Devito consists of a domain-specific language to express PDEs symbolically via FD and a compiler capable of lowering symbolic expressions into highly efficient kernels for a variety of architectures. Thanks to its multi-layer architecture, Devito promotes interdisciplinary research and development. Time-consuming, error-prone and handwritten code practices are mitigated: (i) problem modelling, (ii) implementation of stencils and their adjoints, boundary conditions, sparse operations, and other PDE operators, (iii) performance optimisations (e.g., tiling and FLOPs reduction) and (iv) support for vectorisation, shared- and distributed-memory parallelism, GPUs. Devito has been used successfully in academia and industry for several years, especially in the context of wave modelling in anisotropic media for seismic inversion, delivering competitive performance while drastically reducing complexity.

Fabio Luporini
Devito
fabio@devitocodes.com

Mathias Louboutin
Georgia Institute of Technology
mlouboutin3@gatech.edu

Rhodri Nelson
Imperial College, United Kingdom
rhodri.nelson@imperial.ac.uk

George Bisbas, Edward Caunt, Paul Kelly, Gerard J. Gorman
Imperial College London
g.bisbas18@imperial.ac.uk, edward.caunt15@imperial.ac.uk, p.kelly@imperial.ac.uk,

g.gorman@imperial.ac.uk

MS307

Automated Adjoint in a Full-Waveform Inversion Solver Based on Firedrake

In Full-Waveform Inversion (FWI), the adjoint wave equation is essential in evaluating the misfit function gradient with respect to the model parameters without dealing directly with the Frchet derivatives. The adjoint system can be found by hand, either in a continuous or discrete space, using an augmented function, whose constraint is the discrete/continuous forward equation. On considering that in FWI, the forward problem can change, e. g., by modifying the absorption boundary conditions or the misfit functional or the type of wave equation. Thus, the adjoint derivation and implementation by hands need also to be modified. However, new derivations and implementations gained by hand can be avoided if the adjoint differentiation equations are given by an automatic differentiation (AD) algorithm. In Firedrake, the pyadjoint library is a tool that provides the automated adjoint by the employment of AD techniques. Pyadjoint aims to determine the adjoint problem for the case where the forward solver coding is in the Python interface, written in a weak form in finite elements. In the current work, the AD technique (given by the pyadjoint) is employed in an FWI solver based on Firedrake. The goal is to present the numerical accuracy and also computational performance on different realistic geophysical examples, for acoustic and elastic wave equations in a two-dimensional and three-dimensional spatial domain.

Daiiane I. Dolci
Universisty of São Paulo
dolci@usp.br

Alexandre Olender, Thiago Santos, Ernani Volpe, Bruno Carmo
University of São Paulo
olender@usp.br, santos.td@gmail.com, ernvolpe@usp.br, bruno.carmo@usp.br

MS308

Multiplying 2x2 Sub-Blocks Using 4 Multiplications

Fast recursive matrix multiplication algorithms switch to the cubic time classical algorithm on small sub-blocks as the classical algorithm requires fewer operations on small blocks. We obtain a new algorithm that can outperform the classical one, even on small blocks, by trading multiplications with additions. To this end, we introduce commutative algorithms that generalize Winograds folding technique and combine it with fast matrix multiplication algorithms and the fast base change technique. Thus, when a single scalar multiplication requires ρ times more clock cycles than an addition our technique reduces the computation cost of multiplying the small sub-blocks by a factor of $\frac{\rho+3}{2(\rho+1)}$ compared to using the classical algorithm. Our technique also reduces the energy cost of the algorithm, as the ρ values for energy costs are typically even larger than the values for arithmetic costs. Specifically, we obtain an algorithm for multiplying 2x2 blocks using only four multiplications. This algorithm seemingly contradicts the lower bound of Winograd (1971) on multiplying 2x2 matrices. However, we obtain this algorithm by bypassing the implicit assumptions of the lower bound. We provide a new lower bound matching our algorithm for 2x2 block mul-

tiplication, thus showing our technique is optimal. Joint work with Oded Schwartz

Yoav Gross

The Hebrew University
yoav.gross@mail.huji.ac.il

Oded Schwartz
Hebrew University of Jerusalem
odedsc@cs.huji.ac.il

MS308

Design and Evaluation of Batched Sparse Direct Solver for Multi-GPUs

We will present our recent work on a batched sparse direct solver on multi-GPU machines. Our initial implementation assumes that all the matrices in a batch have the same size and sparsity pattern. This allows us to perform sparsity-preserving ordering and symbolic factorization only once. For the numerical phases we designed two parallel schemes: (1) coarse-grained approach by which multiple MPI ranks are assigned one linear system and multiple linear systems are mapped to one GPU; (2) fine-grained approach by which one MPI rank is assigned multiple linear systems, and the batched operations occur in the internal sparse LU and sparse triangular solve. We will show how each approach can be used in different usage scenarios, and the possibility of hybridizing the two approaches.

Sherry Li, Yang Liu
Lawrence Berkeley National Laboratory
xsli@lbl.gov, liuyangzhuan@lbl.gov

Wajih Boukaram
Lawrence Berkeley National Lab
wajih.boukaram@lbl.gov

Tianyi Shi
Cornell University
tianyishi@lbl.gov

MS308

Batched Sparse Linear Algebra Interfaces, Solvers, Libraries, and Preconditioners

Batched sparse linear algebra solvers form the new frontier for algorithmic development and performance engineering. Many applications require simultaneous solutions of small linear systems of equations that are structurally sparse. To move towards high hardware utilization, it is important to provide these applications with appropriate interfaces to efficient batched sparse solvers running on modern hardware accelerators. We present interface designs in use by HPC software libraries supporting batched sparse linear algebra and the development of sparse batched kernel codes for solvers and preconditioners. We also address the potential interoperability opportunities to keep the software portable between the major hardware accelerators. The presented interface specifications includes batched band, sparse iterative, and sparse direct solvers.

Piotr Luszczyk
University of Tennessee
luszczyk@icl.utk.edu

Ahmad Abdelfattah
University of Tennessee, Knoxville
ahmad@icl.utk.edu

Stanimire Tomov
University of Tennessee
tomov@icl.utk.edu

MS308

Batched Solvers and Preconditioners in Ginkgo.

In this talk, we will elaborate on the batched functionality in Ginkgo, a high performance numerical linear algebra library. Many applications which perform grid-wise computations, such as combustion and fusion plasma simulations require solution of independent linear systems at each grid point. This level of embarrassing parallelism is very well suited to the GPU. Traditionally, for these small problems, batched direct methods have been the method of choice. In this talk we will showcase the advantages of using an iterative solver in a batched fashion, and elaborate on the implementation and the challenges. We will show that we can achieve significant speedups compared to the dense and direct methods within our applications. Finally, we will show some techniques for enhancing and accelerating these batched iterative solvers with batched preconditioners and present the results for general matrices as well as matrices from specific applications.

Hartwig Antzt
UTK
hanzt@icl.utk.edu

Pratik Nayak
Karlsruhe Institute of Technology
pratik.nayak@kit.edu

Aditya Kashi
PhD student
Department of Mechanical Engineering, McGill University
aditya.kashi@mail.mcgill.ca

Isha Aggarwal
Karlsruhe Institute of Technology
isha.aggarwal2@kit.edu

MS308

Batched Linear Solvers in Kokkos Kernels

Applications such as multilevel finite element methods, modeling collision in plasma and reentry simulations are reformulating algorithms to match the hierarchical parallel manycore CPUs/GPUs. This reformulation results in a need for linear solvers at different levels of hierarchical parallel algorithms. Several of these linear systems are small or medium size, independent, linear systems that can be solved at the same time. Depending on the problems, some of these linear systems are sparse and others dense. The community has been focused on solving dense linear systems as a batch, or solving multiple small dense linear systems. However, the work has focused primarily on one level of hierarchical parallelism - using an entire accelerator to solve the linear systems. We have been focused on solving batched dense/sparse linear systems, both at the device level and at the team level so applications can use these options effectively. We will present recent results on CPUs/GPUs using our hierarchical parallel batched linear solver implementations.

Siva Rajamanickam, Kim Liegeois, Luc Berger-Vergiat
Sandia National Laboratories
srajama@sandia.gov, knliege@sandia.gov,

lberge@sandia.gov

MS309

Approaches to Moving Boundary Problems in Systems with Multiple Fluid and Solid Interactions.

One of the challenges in modelling thin film and multi-phase flow is that the domain over which the PDE are solved changes dynamically over time. Two approaches to explicitly track the motion of the domain's boundaries are presented, with different motivating applications and requirements. The first uses a reduced-order model to derive interface evolution equations from which the fluid flow is be subsequently reconstructed. This method will be applied to a three-layer fluid flow with the aim to study long-time interactions between free surfaces coupled by inertialess fluids. The equations are solved via a pseudo-spectral method implemented in a concise Matlab library. This approach allows us to study larger domains over an extended duration when compared to direct numerical simulation of the Navier-Stokes equations, and we present nonlinear phenomena exhibited at these later times. The second approach uses finite-elements to solve the Navier-Stokes equations over a mesh which is allowed to deform elastically in response to changes in the free-surface such that the equations there are also satisfied. The higher-order model allows for more localised boundary conditions and we will focus on the point where the interface meets a solid wall. This approach is used to simulate a fluid being withdrawn from a cylinder, with gravity acting against the motion, where the transition from a steadily rising meniscus to more complex dynamics is of interests in industrial applications.

John P. Alexander

University of Manchester
johnpaul.alexander@manchester.ac.uk

Demetrios Papageorgiou
Department of Mathematics
Imperial College London, UK
d.papageorgiou@imperial.ac.uk

Andrew Hazel
The University of Manchester
andrew.hazel@manchester.ac.uk

MS309

Modeling and Computing Thin Film Flows: New Frontiers and Challenges

While modeling and computation of thin fluid films, both deposited on substrates and free standing, have been an active field of research for a while, during recent years significant new developments have taken place in both computational and modeling aspects of the research on thin films. From the computational side, these include establishment of whole computational libraries as well as of GPU-based simulations. These new computational developments have led to more accurate computational results in large scale simulations that could not have been imagined few years ago. From modeling side there has been significant progress on inclusion of various effects including considering non-Newtonian fluid behavior, phase change, flows with particles, bio-focused flows, and many others. In this talk, we will introduce and review some of the novel methods used for simulating thin films, and then follow up with providing few examples illustrating the need for further research and

new developments in the field.

Lou Kondic

Department of Mathematical Sciences, NJIT
University Heights, Newark, NJ 07102
kondic@njit.edu

Linda Cummings
New Jersey Institute of Technology
linda.cummings@njit.edu

MS310

Differentially Private Algorithm for Constrained Federated Learning

Differentially private federated learning (FL) is a privacy-preserving machine learning (ML) that enables learning from the distributed data owned by multiple agents without the need for transferring the data into a central server while guaranteeing differential privacy (DP) against inference attacks. In practice, the agents often need to impose constraints on the model parameters of the training models for explaining decisions suggested by ML models, promoting fair decisions, and observing some physical laws. For such constrained FL models, we develop an inexact alternating direction method of multipliers algorithm with the objective perturbation method that injects the noise into the loss function for DP. In particular, the perturbation method enables the resulting parameters to satisfy the constraints and the outperformance as compared with the state-of-the-art algorithms based on the output perturbation. We further enhance the training algorithm by reducing the communication costs of multiple local update technique. We show the privacy and convergence analyses of the proposed algorithm for convex constrained FL. Using datasets for image classification, we demonstrate that our algorithm significantly reduces the testing error compared with the existing DP algorithms while achieving the same level of data privacy.

Minseok Ryu, Kibaek Kim
Argonne National Laboratory
mryu@anl.gov, kimk@anl.gov

MS311

Multilevel Techniques for the Adaptive Reduction of Scientific Data

While multilevel techniques for the adaptive reduction of data (such as MGARD) have the advantage of being supported by a rich mathematical theory, the practical application of such techniques is limited by the form of the data and, despite being proven to have optimal complexity, they require a large memory overhead. We present some of our recent work in which we seek to develop new, more efficient alternatives to MGARD that retain the advantages of a multilevel approach in terms of the mathematical underpinnings, but which do not suffer from the need to have a large memory overhead. Examples will be presented comparing the new approach with existing techniques. Joint work between Mark Ainsworth and John Ahn (Brown University).

Mark Ainsworth
Division of Applied Mathematics
Brown University

Mark_Ainsworth@brown.edu

MS311

How Much Can We Really Compress Scientific Data?

Lossy compression of scientific data sees a growing interest from different domain sciences. During the past 5-6 years, we have witnessed a significant diversification of use cases ranging from classic visualization to more advanced lossy compression of data in communications or memory. We are also observing a blooming of lossy compression algorithms and, more importantly, continuous progress in compression quality, speed, and ratio. The lossy compression methodology has also progressed drastically with the SDRBench benchmark, the Libpressio unified API, and the Zchecker error analysis tool. However, a fundamental part is missing: the estimation of bounds of lossy compressibility. We have the Shannon entropy to compute the maximum compressibility of a data set from entropy coding. But there is no such result about the lossy compressibility of scientific data. We can formulate the research question as: Can we come up with the equivalent of a roof line of lossy compressibility? Answering this question is fundamental for the research on algorithms for lossy compression of scientific data. This talk will review recent progress in lossy compression for scientific data and in lossy compressibility bound estimation.

Franck Cappello
ANL
cappello@mcs.anl.gov

MS311

Further Investigative Analysis of ZFP Compression

Typically, the error from any lossy compression algorithm is deemed acceptable as the data gathered is noisy, either from simulation error, such as truncation, iteration and round-off error, or observational error, such as finite precision measurements and measurement noise. In many statistical analysis applications, additional assumptions on the error are applied that may not be valid in practice. Some emerging tools provide post-analysis of the compressor for the original dataset; however, this is one instance where assumptions about the compressor may mislead users. The focus of this talk is to provide additional theoretical understanding of the components of the ZFP lossy floating-point compressor, ensuring the assumptions we have empirically studied are indeed valid. Prepared by LLNL under Contract DE-AC52-07NA27344.

Alyson Fox, Peter Lindstrom
Lawrence Livermore National Laboratory
fox33@llnl.gov, lindstrom2@llnl.gov

MS311

High Efficiency Imaging

Compressed sensing is the most sensational topic of scientific research in the past century. The original paper was cited over 30,000 times in only 16 years (2006-2022). It purports to exert data compression at the phase of signal acquisition. However, compressed sensing massively loses signal quality. In theory, it violates fundamental principle of approximation theory, which spawned the enormous success of data compression technologies. We exemplify a new mathematical theory and method for high efficiency

imaging. In compliance with the principle of best n-term approximation and based on discovery of a new logical phenomenon, high efficiency sensing radically rectifies mathematical rationale, while immensely improves technical performance. The core idea is simple yet powerful. High efficiency denotes high quality plus high speed. Demo software and test data are downloadable at our website www.lucidsee.ca.

Xiteng Liu
University of South Carolina
xitengliu@gmail.com

MS311

A General Framework for Progressive Data Compression

The cost of data movement across supercomputer memory hierarchies is becoming an increasingly important bottleneck for scientific simulations and observations in general. Compression techniques that eliminate redundancy and bits of marginal accuracy in floating-point arrays are commonly employed in many scientific applications to deal with this bottleneck. However, the variety of post-processing analyses performed by domain scientists causes regular compression workflows to be conservative and achieve low reduction ratios. In this context, progressive data compression and retrieval arise as a solution for achieving higher overall ratios. Moreover, it allows for the adaptive handling of compressed data according to the needs of a given post-processing task. Currently, a few algorithms natively support progressive data processing. In this work, we present a general framework for adding progressiveness in terms of precision to scientific data compressors. Our framework is based on a multiple-component representation where each new component reduces the error between the original and compressed fields. We have implemented it on top of the ZFP, SZ, MGARD, and SPERR compressors, and, in this presentation, we discuss results obtained with several data fields from the SDRBENCH collection. Numerical results show that our approach is effective in terms of accuracy compared to each of the standalone compressors it builds upon. Prepared by LLNL under Contract DE-AC52-07NA27344.

Victor A. Paludetto Magri, Peter Lindstrom
Lawrence Livermore National Laboratory
paludettomag1@llnl.gov, lindstrom2@llnl.gov

MS312

Robust Optimal Experimental Design for Bayesian Inverse Problems

An optimal design is defined as the one that maximizes a predefined utility function which is formulated in terms of the elements of an inverse problem. An example being optimal sensor placement for parameter identification. This formulation generally overlooks misspecification of the elements of the inverse problem such as the prior or the measurement uncertainties. In this talk, we present efficient recipes for designing optimal experimental design schemes, for Bayesian inverse problems, such that the optimal design is robust with respect to misspecification of elements of the inverse problem.

Ahmed Attia
Argonne National Laboratory
attia@mcs.anl.gov

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

MS312

Experimental Design in the Presence of Model Error

Solving inverse problems for model calibration of parametrized models plays an important role in predictive modelling and relies on the availability of informative data. Optimal experimental design techniques are devised to guide the data acquisition in a way that maximizes the information content in the data about the parameter which is the object of the inversion. In the presence of model error, data additionally must be informative of the model error in order to be able to correct the model and improve models predictive power. In this talk we present different approaches for treating the model error which exist in the literature on Bayesian inference and optimal experimental design. By applying Lipschitz stability results for Bayesian inverse problems we compare the posterior distributions of the presented approaches with the posterior obtained when the model error is ignored. In this way we are able to quantify the effect of including the model error. Furthermore, we are able to characterize the properties of the observation operator that increase the information content in the data about the model error, thereby increasing the models predictive power.

Nada Cvetkovic
TU/e Eindhoven University of Technology
n.cvetkovic@tue.nl

Han Cheng Lie
University of Potsdam
hanlie@uni-potsdam.de

Harshit Bansal
Eindhoven University of Technology
h.bansal@tue.nl

Karen Veroy-Grepl
TU/e Eindhoven University of Technology
k.p.veroy@tue.nl

MS312

Geometric Deep Neural Operators for PDE-Constrained Bayesian Optimal Experimental Design

Bayesian Optimal Experimental Design (BOED) problems governed by PDEs in high or infinite parameter dimensions are often intractable. Efficient evaluation of the parameter-to-observable (PtO) map, which involves solution of the forward model, is key to making BOED problems tractable. Surrogate approximations of PtO maps can greatly accelerate solution of BOED problems, provided an accurate surrogate can be trained with modest numbers of model solves. Unfortunately, constructing such surrogates presents significant challenges when the parameter dimension is high and the forward model is expensive. Deep neural network surrogates of PtO maps, also known

as neural operators, have emerged as leading contenders for overcoming these challenges. Black box application of DNNs for problems with infinite dimensional parameter fields leads to poor results when training data are limited. However, by constructing a network architecture that exploits the geometry of the PtO map—in particular its smoothness, anisotropy, and intrinsic low-dimensionality as revealed through adjoint-PDE-based Gauss-Newton Hessians, one can construct a dimension-independent geometric neural operator with superior generalization properties using only limited training data. We employ this neural operator to make tractable the BOED problem of finding sensor locations that maximize the expected information gain from the data, and demonstrate on an inverse wave scattering problem.

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

Keyi Wu
Meta
keyiwu.w@gmail.com

Thomas O’Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Peng Chen
The University of Texas at Austin
peng@cc.gatech.edu

MS312

Mathematical Model and Characterization of Emotional Disorders Based on Non-Intrusive Monitoring System

According to the WHO mental disorders represent 12.5% of all health problems. Approximately 1 in 4 people will suffer from a mental disorder throughout their lives and between 35% and 50% will receive inadequate treatment, causing up to 800,000 suicides each year. Minimising the time that elapses between the appearance of the first symptoms and the application of the treatment is key to recovery, as it can reduce the severity of the pathology or even slow down its development. In adolescents, especially, if symptoms are not detected and treated in time, it could persist in the adult stage in more severe forms. Each person has different symptomatology and emotional state changes can be suddenly triggered by endogenous and/or exogenous factors, which makes detecting symptoms challenging. One common symptom for many mental illnesses is significant changes in sleep patterns, such as inability to fall asleep, to stay awake, or to get up in the morning. The main objective of this research is to characterize the emotional state of each person in order to predict future changes by analysing their sleeping activity. With that aim, this research presents a mathematical model that allows the characterization of a personalized emotional state. The data used by the model are collected in real time by an alternative and non-intrusive sleep monitoring system that works without the need of supervision or intervention from the monitored person, guaranteeing thus that the data are reliable

Yury A. Jimenez Agudelo
CUNEF UNIVERSIDAD

yury.jimenez@cunef.edu

Maria Espinosa
CUNEF Universidad
msoledad.espinosa@cunef.edu

Victoria López
CUNEF University
Madrid
victoria.lopez@cunef.edu

MS312

Optimal Design of Experiments in Models with Uncertainties

This talk will discuss optimal design of experiments for Bayesian inverse problems governed by uncertain models. When the model uncertainty is irreducible, this amounts to an optimal design problem under uncertainty. When the uncertainty is reducible (i.e., it can be reduced through data), this becomes a probability marginalization problem. I will discuss the differences between these formulation and present numerical examples.

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Karina Koval
Heidelberg University
karina.koval@iwr.uni-heidelberg.de

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Isaac Sunseri
Department of Mathematics
NC State University
ipsunser@ncsu.edu

MS313

Advanced Training Algorithms for Pinn-Reconstruction of Gravity Currents

Inverse problems in geophysical flows are often ill-posed, thus, it is challenging or sometimes even impossible to solve them using traditional methods. Moreover, the generation of simulated data for ill-posed inverse problems can become very costly where simulation needs to be performed multiple times to either discover missing physics in the model or calibrate the free parameters in the model. One possible alternative for solving these problems is through the use of Physics-Informed Neural Networks - PINNs, in which we approximate the problems solution using Neural Networks, while incorporating the known data and physical laws during the training phase Here we show a PINN-based framework for reconstructing gravity currents velocity, concentration, and pressure fields in which the available data can be given in two different ways: a set of scattered measurements for the concentration field, or a set of velocity and concentration measurements that follows the configurations often used in real-life experiments. We discuss

PINNs advanced training algorithms aided by optimal sensor placement and dimensionality reduction.

Rômulo Silva
Federal University of Rio de Janeiro
romulo.silva@nacad.ufrj.br

Alvaro Coutinho
Dept. of Civil Engineering
Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

George E. Karniadakis
Brown University
george.karniadakis@brown.edu

MS313

A New Regularized Reduced Order Model for the Barotropic Vorticity Equation

We develop a ROM for the quasi-geostrophic equations in a computationally efficient finite volume environment. By using the analogy between LES and truncated modal projection, we introduce an eddy viscosity closure approach based on an alpha regularization model to stabilize the resulting surrogate model.

Michele Girfoglio
SISSA, Mathematics Area, mathLab
International School for Advanced Studies
mgirfogl@sissa.it

Annalisa Quaini
Department of Mathematics, University of Houston
quaini@math.uh.edu

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS313

Reduced Order Modeling of the Quasi-Geostrophic Equations

We summarize some recent reduced order model (ROM) developments for the quasi-geostrophic equations (QGE) (also known as the barotropic vorticity equations). The QGE are a simplified model for geophysical flows in which rotation plays a central role, such as wind-driven ocean circulation in mid-latitude ocean basins. Since the QGE represent a practical compromise between efficient numerical simulations of ocean flows and accurate representations of large scale ocean dynamics, these equations have often been used in the testing of new numerical methods for ocean flows. ROMs have also been tested on the QGE for various settings in order to understand their potential in efficient numerical simulations of ocean flows. We survey the ROMs developed for the QGE in order to understand their potential in efficient numerical simulations of more complex ocean flows: We explain how classical numerical methods for the QGE are used to generate the ROM basis functions, we outline the main steps in the construction of projection-based ROMs (with a particular focus on the under-resolved regime, when the closure problem needs to be addressed), we illustrate the ROMs in the numerical simulation of the QGE for various settings, and we present several potential future research avenues in the ROM explo-

ration of the QGE and more complex models of geophysical flows.

Traian Iliescu, Changhong Mou
Department of Mathematics
Virginia Tech
iliescu@vt.edu, cmou3@wisc.edu

Zhu Wang
Department of Mathematics
University of South Carolina
wangzhu@math.sc.edu

David Wells
University of North Carolina at Chapel Hill
drwells@email.unc.edu

Xuping Xie
Courant Institute of Mathematical Sciences,
New York University
xxie@nyu.edu

MS313

Data-Driven Sparse Wasserstein Barycenters for Model Reduction

We present a model order reduction method based on sparse Wasserstein barycenters that can be applied to parametric PDEs posed in the space of measures. Typical examples include conservation laws and gradient flows. Such problems are notoriously difficult to reduce with classical linear approximation methods on Hilbert/Banach spaces due to the transport of shocks and discontinuities. Shifting the point of view to measures spaces allows to better enforce the correct location of shocks and the preservation of discontinuities through the Wasserstein metric.

Olga Mula
Eindhoven University of Technology
o.mula@tue.nl

MS314

Direct Sampling Methods for Recovering Inhomogeneous Inclusions of Different Nature

In this talk, we shall address some effective and robust direct sampling methods for simultaneously recovering inhomogeneous inclusions arising from two different physical coefficients in the typical elliptic PDE model and the electromagnetic system. The essential motivation, detailed algorithm, theoretical justification, and numerical results will be presented. Authors: Yat Tin Chow, University of California, Riverside, yatinc@ucr.edu; Fuqun Han, Chinese University of Hong Kong, fqhan@math.cuhk.edu.hk; Jun Zou, Chinese University of Hong Kong, zou@math.cuhk.edu.hk

Fuqun Han
Department of Mathematics,
The Chinese University of Hong Kong
fqhan@math.cuhk.edu.hk

MS314

Neumann-Cayley Orthogonal Gated Recurrent Units with Some Applications in Bio-molecular

Data

TBA

Edison Mucllari
University of Kentucky
edison.mucllari@uky.edu

MS314

Synchronization of Coupled Oscillators

The emergence of a synchronous behavior of coupled dynamical systems is widely observed in many biological and engineered networks and is essential to understand their mechanism and control them. A plethora of work is done to find conditions that foster synchronization in diffusively coupled systems, in which the coupling vanishes on the synchronization manifold. However, there are few analytical approaches to understanding synchronization behavior in non-diffusively coupled networks. Motivated by neuronal models connected through chemical synapses, we investigate sufficient conditions for non-diffusively coupled oscillators to synchronize globally. Our global stability method follows an analytical contraction-based approach that explicitly relates synchronization to systems' intrinsic dynamics, coupling dynamics, and the underlying network topology. We compare this method to a local or linear stability approach called the master stability function, a numerical method that provides necessary and sufficient conditions for network synchronization. Authors: Fatou Kineh Ndow, University of Iowa, fatou-ndow@uiowa.edu; Zahra Aminzare, University of Iowa, zahra-aminzare@uiowa.edu

Fatou Kineh Ndow
University of Iowa
fatou-ndow@uiowa.edu

MS314

Shape Optimization for Interface Identification in Nonlocal Models

Shape optimization methods have been proven useful for identifying interfaces in models governed by partial differential equations. Here we consider a class of shape optimization problems constrained by nonlocal equations which involve interface-dependent kernels. We derive a novel shape derivative associated to the nonlocal system model and solve the problem by established numerical techniques. Authors: Matthias Schuster, Trier University, schusterm@uni-trier.de; Christian Vollmann, Trier University, vollmann@uni-trier.de; Volker Schulz, Trier University, volker.schulz@uni-trier.de

Matthias Schuster
Trier University
schusterm@uni-trier.de

MS315

From Data to Reduced-Order Models via Generalized Balanced Truncation

This talk discusses a data-driven model reduction approach on the basis of noisy data. Firstly, the concept of data reduction is introduced. In particular, we show that the set of reduced-order models obtained by applying a Petrov-Galerkin projection to all systems explaining the data characterized in a large-dimensional quadratic matrix inequality (QMI) can again be characterized in a lower-dimensional QMI. Next, we develop a data-driven general-

ized balanced truncation method that relies on two steps. First, we provide necessary and sufficient conditions such that systems explaining the data have common generalized Gramians. Second, these common generalized Gramians are used to construct projection matrices that allow to characterize a class of reduced-order models via generalized balanced truncation in terms of a lower-dimensional QMI by applying the data reduction concept. Additionally, we present alternative procedures to compute a priori and a posteriori upper bounds with respect to the true system generating the data. Finally, the proposed techniques are illustrated by means of application to an example of a system of a cart with a double-pendulum.

Azka Muji Burohman
University of Groningen
a.m.burohman@rug.nl

Jacqueline M. Scherpen
Rijksuniversiteit Groningen
j.m.a.scherpen@rug.nl

Bart Besselink
University of Groningen
b.besselink@rug.nl

Kanat Camlibel
University of Groningen
Institute of Mathematics and Computer Science
m.k.camlibel@rug.nl

MS315

Enforcing Uniform Stability and Passivity in Data-Driven Multivariate Model Order Reduction

We present a comprehensive framework for data-driven model order reduction of parameterized Linear and Time-Invariant systems, with explicit certification of model stability and dissipativity. These properties are mandatory when representing passive physical systems, unable to generate energy. Lack of these properties makes the models practically useless in any simulation setting in which they are used as components, as in Computer Aided Design of electrical and electronic systems. The proposed approach builds on

- a particular model structure, which casts the transfer function of the model as a ratio of two rational functions, each with constant predefined “basis” poles and parameter-dependent residues;
- a parameterization of the above residues expressed in a Bernstein polynomial basis, whose partition of unity and positivity are exploited in the derivation of our main result and algorithms;
- availability of sampled responses in frequency and parameter domain, that are fitted by the model through an iterative reweighted linear Least Squares (LS) process;
- the above LS problems is complemented by an explicit, purely algebraic (finite-size) set of constraints that are sufficient for both uniform stability and uniform dissipativity in the parameter space; these constraints provide the key novelty in this contribution.

The proposed approach is illustrated through several electrical CAD application examples.

Stefano Grivet-Talocia, Tommaso Bradde
Politecnico di Torino

stefano.grivet@polito.it, tommaso.bradde@polito.it

MS315

A Realization-Free Balancing-Related Model Reduction Approach from Frequency Response Data

Balancing-related model order reduction methods guarantee stable reduced-order models, and additionally, a priori error bounds have been derived for them. However, the explicit numerical implementation of such methods is often computationally demanding and it is intrusive; this means that it requires explicit access to the matrices scaling the state variables in the original (full-order) dynamical system. For the classical Balanced Truncation (BT) method, and for a few more recent extensions of it, approximate implementations have been derived, which make the application to the large-scale setting more feasible. For time simulation near steady state or for low-frequency applications, the Singular Perturbation Approximation (SPA) method [Fernando/Nicholson, *Internat. J. Control*, 1982] is known to have better approximation properties than BT. In this contribution, we derive an approximate implementation of SPA that alleviates the computational burden for computing the reduced-order models and can be applied also for such cases in which explicit access to the original dynamical system is not granted. Only transfer function evaluations are required in our approach. Specifically, we adapt and extend the purely data-driven (approximate) implementation of BT in [Gosea/Gugercin/Beattie, *SIAM J. Sci. Comput.*, 2022] to SPA, which as such is strongly related to the Loewner framework. The performance of our proposed method is illustrated by several numerical results.

Bjoern Liljegren-Sailer
Trier University
bjoern.sailer@uni-trier.de

Ion Victor Gosea
Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

MS315

Least-Squares Parametric Reduced-Order Modeling

In this talk, we consider reduced-order modeling for (parametric) linear time-invariant systems, based on nonlinear least-squares. We show how it is a special case of a more general \mathcal{L}_2 -optimal parametric reduced-order modeling problem by the choice of a measure space. Based on this, we propose a gradient-based optimization algorithm for finding locally optimal reduced-order models. Then, we discuss its relation to the vector fitting method for least-squares reduced-order modeling of linear time-invariant systems. Furthermore, we present the necessary optimality conditions in the interpolation form, related to interpolatory \mathcal{H}_2 -optimality conditions. Finally, we demonstrate the results on a number of numerical examples.

Petar Mlinaric
Virginia Tech
mlinaric@vt.edu

Serkan Gugercin
Virginia Tech
Department of Mathematics

gugercin@vt.edu

MS316

Bayesian OED for Sensor Placement: Analysis and Optimization of Seismo-Acoustic Monitoring Networks with Bayesian Optimal Experimental Design

Bayesian optimal experimental design (OED) seeks to identify data, sensor configurations, or experiments which can optimally reduce uncertainty. OED formulates the choice of experiment as an optimization problem that maximizes the expected information gain (EIG) about quantities of interest given prior knowledge and models of expected observation data. We use Bayesian OED to find optimal sensor configurations for detecting seismic events as part of a seismic monitoring network. We configure sensor networks by choosing sensor locations, types, and fidelity in order to improve our ability to identify and locate seismic sources. In this work, we develop the framework necessary to use Bayesian OED to optimize a sensor network's ability to locate seismic events from arrival time data of detected seismic phases. As part of this framework we introduce methods for performing importance sampling on difficult prior distributions and methods for placing unusual constraints on our optimization domain to account for real world limits on sensor location. Once we have developed this framework, we can explore many relevant questions to monitoring such as how to trade off sensor fidelity and earth model uncertainty, how choice of prior distribution and domain restrictions affect sensor placement, and how sensor types, number, and locations influence uncertainty.

Thomas A. Catanach
Sandia National Laboratories
tacatan@sandia.gov

Jake Callahan
Sandia National Labs
jpcalla@sandia.gov

Ruben Villarreal
Sandia National Laboratories
rubvill@sandia.gov

Kevin Monogue
Stanford University
kcmonogue@gmail.com

MS316

Robust Expected Information Gain in Optimal Experimental Design

Expected information gain (EIG) is a useful measure of experiment optimality, but it can rarely be computed exactly, and there are drawbacks associated with different approximations. Nested Monte Carlo and related methods are popular approaches to estimating EIG. We note two issues with the use of these estimators. The first is the effect that outliers can have in the estimators when they are undersampled. The second is the sensitivity of experiments to perturbations in the prior: the sensitivity of an experiment's EIG to perturbations in the prior should affect its suitability in a risk averse situation. We propose a quantity called robust expected information gain (REIG) that maximizes the minimum EIG over perturbations in the prior which maps onto a fast post-processing of Monte Carlo estimators, which in some cases can also address the outlier

affect. We show numerical experiments comparing REIG estimators to EIG estimators on simple problems, then on an OED problem modeling experiments with catalysis reactors to determine kinematic reaction coefficients.

Tobin Isaac
Argonne National Lab
tisaac@anl.gov

Jinwoo Go
Georgia Institute of Technology
jgo31@gatech.edu

MS316

Risk-Averse Goal-Oriented Optimal Experimental Design Using Nonlinear Models

Traditionally, nonlinear Bayesian optimal experimental design (OED) uses numerical models to predict when and where to collect data to minimize uncertainty in model parameters. This talk will present novel goal-oriented nonlinear Bayesian OED strategies that minimize uncertainty in predictions. Our risk averse strategies can be used to produce designs that most inform tail statistics in a scalar or vector valued quantity of interest. In the latter situation a specialized case of our general framework provides a nonlinear interpolation between I and G optimality which minimize average and worst case variance respectively. We will demonstrate the utility of our OED strategies for placing sensors that minimize uncertainty in predictions of an advection diffusion model.

John D. Jakeman
Sandia National Labs
jadjakem@sandia.gov

Rebekah White
Sandia National Laboratories
rebwhit@sandia.gov

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS316

Bayesian Optimal Experimental Design in the Presence of Model Uncertainty

We consider optimal experimental design (OED) of nonlinear inverse problems in the Bayesian framework. The problem is further complicated by considering a set up in which the governing equations contain further auxiliary parameters which are also unknown, but are assumed to be of little interest. To account for the additional model uncertainty in the governing equations we employ the so-called Bayesian approximation error (BAE) approach. We show that failure to take into account the additional model uncertainty at the OED stage relates in a sub-optimal design, while failure to take into account at the inference

stage relates in significant over confidence in heavily biased estimates, i.e., an infeasible posterior. In this talk, we consider the optimal design of infinite-dimensional nonlinear Bayesian inverse problems. We derive a marginalised A-optimality criterion and develop an efficient computational approach to solve the associated optimisation problem. We demonstrate the applicability of our approach by considering the estimation of a Robin coefficient for the Poisson problem when the (unknown) internal conductivity is treated as an auxiliary parameter.

Ruanui Nicholson

Department of Engineering Science
University of Auckland
ruanui.nicholson@auckland.ac.nz

Alen Alexanderian

North Carolina State University, U.S.
alexanderian@ncsu.edu

Noemi Petra

University of California, Merced
npetra@ucmerced.edu

MS316

Optimal Design of Validation Experiments Using Sensitivity Indices

The main objective of using physical models is the prediction of quantities of interest (QoIs) under specific regimes called prediction scenarios. These QoIs are usually not observable, as they may refer to future predictions or predictions in conditions that cannot be reproduced in a laboratory. This raises the fundamental issue of validating a model with respect to these QoIs. Several validation processes and metrics have been proposed to assert whether a model is deemed valid. However, they rarely take into account the QoIs, essentially for the reasons mentioned above. One has then to make a leap of faith to use the model, validated w.r.t. given observables, for predicting other QoIs. Our goal is nevertheless to gain confidence in the predictive capabilities of the model by defining validation experiments specifically designed toward this objective. We propose here a methodology to design a validation scenario that relates the QoIs in the prediction scenario with the observables in the validation scenario. It consists in casting the design of the validation scenario as an optimal design problem. We introduce here an optimal design problem whose objective functional minimizes the difference between sensitivity indices (e.g. Sobol indices) of the QoIs in the prediction scenario and those related with the same QoIs in the validation scenario. We will demonstrate the usefulness of our methodology on several examples by showing that it can identify false positives.

Serge Prudhomme

École Polytechnique de Montréal
serge.prudhomme@polymtl.ca

Antonin Paquette-Rufiange

Polytechnique Montréal
antonin.paquette-rufiange@polymtl.ca

Marc Laforest

École Polytechnique de Montréal

marc.laforest@polymtl.ca

MS317

Machine-Learned Stochastic Closures and Regional Superparameterization for Convection

The representation of convection processes and cloud formation in coarse-scale global models of the atmosphere remains a challenge. I will discuss work on two related approaches to tackle this. One is to embed local, convection-resolving high-resolution models in model columns of a global model, with two-way coupling between them. This is the approach of superparameterization, originally proposed by Grabowski and co-workers. We developed a setup to apply superparameterization regionally (in selected model columns) rather than globally, thereby making it computationally feasible to use fully 3-d LES for the local models. In an alternative approach we replaced the local high-res model by a computationally cheap stochastic surrogate model trained on data from the LES or from observations. A discrete Markov chain model was used to this end. We developed a method employing machine-learning assisted resampling to construct a stochastic surrogate, and tested it on the L96 model.

Daan Crommelin

CWI Amsterdam and University of Amsterdam
Daan.Crommelin@cwi.nl

MS317

Mapping Meteorological Conditions to Predator Prey Dynamics

Earth's energy budget is a key quantity for our climate and how it changes. Stratocumulus clouds have a great effect on Earth's energy budget because (i) they can reflect sunlight; and (ii) they cover vast amounts of our planet. It is thus important to understand and model stratocumulus clouds and in this talk, we make connections between two very different models. The first model is a large eddy simulation (LES), which is a cloud resolving 3D atmospheric simulation that is computationally expensive to run. The second model is a scalar delay differential equation (DDE), which is trivial to run and that interprets the interactions of precipitation and cloud as a predator (rain) and prey (cloud). We connect these two models by estimating parameters of the predator prey model from LES that reflect a variety of meteorological conditions. We rely on a feature-based approach to parameter estimation and numerically solve the problem using an affine invariant ensemble sampler. The result of our computations is a map of meteorological conditions to the parameter space of the predator prey model, which is needed if one were to use it to parameterize stratocumulus clouds in Earth system and climate models.

Rebecca Gjini

Scripps Institution of Oceanography
University of California, San Diego
rgjini@ucsd.edu

MS317

Dynamical System Analysis of a Rich Cloud Model and Implications for Cloud Aggregation

Shallow-cumulus evolution and self-organization, or aggregation, into different cloud patterns is a complicated process, with many not well understood elements. To shed more light on the physical processes involved we study the

dynamical behavior of a simple, one-dimensional model of warm cloud evolution, consisting of three prognostic differential equations. Cloud growth is treated as a non-linear combination of vertical motion, cloud droplet concentration, and rain droplet concentration, linked in a system of balance equations. A bifurcation analysis shows that the model has rich mathematical structure, with steady, periodic, and chaotic solutions, and we will discuss the physics and feedbacks of the underlying phases and their transitions, including the route to chaos. If time permits, we will report on a dynamical systems analysis of a two-dimensional field of coupled one-dimensional cloud models.

Michelle Kanipe
Colorado State University
mkkanipe@gmail.com

Peter Jan van Leeuwen
Colorado State University and University of Reading
peter.vanleeuwen@colostate.edu

MS317

Unsupervised Learning of Cloud Microphysical Process Rates

Cloud microphysics schemes employed in atmospheric and climate models have traditionally used moments of the droplet size distribution to parameterize process rates such as collision-coalescence. However, it is not clear that prognostic moments form an optimal basis set for representing microphysical processes across the range of conditions experienced in the atmosphere. In addition, microphysical parameterizations employing prognostic moments require developing closure schemes, which are known to suffer from both parametric and structural uncertainty in their representations of inherently higher dimensional cloud processes. These uncertainties limit model fidelity and lead to forecasting errors. Recently, machine learning methods have been used to develop reduced order models of various physical systems by discovering governing equations directly from high-dimensional observational data. Here we investigate how machine learning methods such as variational auto-encoders can be used to learn optimal predictors for microphysical process rates in an unsupervised manner, by simultaneously learning lower dimensional representations of droplet size distributions and predicting their dynamic evolution. We discuss how these findings could be used to either guide moment-based scheme development, or replace them with alternative formulations.

Kara D. Lamb, Sean P. Santos
Columbia University
kl3231@columbia.edu, sean.santos@pnnl.gov

Hugh Morrison
National Center for Atmospheric Research
morrison@ucar.edu

Marcus Van Lier-Walqui
Center for Climate Systems Research
Columbia University/NASA GISS
mv2525@columbia.edu

MS317

A Novel Stochastic Transport Model for Atmospheric Aerosols

Recently, it has become clear that atmospheric aerosols

have a non-negligible effect on radiative forcing within Earth's climate and the computational models that simulate it [Carslaw, et al., Nature, 2013]. However, aerosol microphysical processes, chemistry, and transport properties are not yet fully understood and are highly parameterized, leading to uncertainty in the results of our models. Thus, we must obtain accurate aerosol models that are also predictive, particularly in a time when aerosol-emitting ships may soon traverse the polar arctic ocean and there is credible talk about climate intervention strategies like stratospheric aerosol injection. This work focuses on the transport properties of atmospheric aerosols, namely their mixing and spreading via turbulent transport in the planetary boundary layer. We propose a novel stochastic model comprised of advection and diffusion parameters, operating on distinct spatial scales, that vary according to their respective machine-learned probability distribution, parameterized to be a function of space, time, and relevant exogenic properties. The model is evaluated using ensembles of particles that advect and diffuse both deterministically and via random walks with a magnitude drawn from the corresponding learned distribution. To verify our model, we employ data sources varying across a wide range of scales from algorithmically-classified satellite imagery of ship-exhaust tracks in the Pacific to a tabletop fog chamber.

Michael J. Schmidt, Pete Bosler, Jake Zenker, Christian Pattyn, Lekha Patel
Sandia National Laboratories
mjschm@sandia.gov, pabosle@sandia.gov,
jpzenke@sandia.gov, capatty@sandia.gov, lpatel@sandia.gov

MS318

Mitgcm-Ad & Sicopolis-Ad V2: Tangent Linear and Adjoint Modeling Frameworks for Oceans and Ice Sheet Modeling Enabled by Automatic Differentiation Tool Tapenade

We present new inverse modeling frameworks for the ocean general circulation model MITgcm and ice sheet model SICOPOLIS that are enabled by source transformation using the open-source Automatic Differentiation (AD) tool Tapenade. Oceans and ice sheets are dynamic entities whose evolution is governed by non-linear Partial Differential Equations (PDEs) that conserve mass, momentum, and energy. Their evolution is driven by their present state and uncertain forcings such as surface temperature, precipitation, basal sliding, geothermal heat flux, etc. These uncertainties propagate to our economically and societally important Quantities of Interest (QoI) such as projections for sea-level rise or the estimated strength of the Atlantic Meridional Overturning Circulation (AMOC), which contributes substantially to the Earth's climate systems. It is thus desirable to evaluate the sensitivities of our QoI to these independent input variables. Furthermore, recently collected data can be used to calibrate model parameters to improve our simulations, which is an exercise in PDE-constrained gradient-based optimization. Posterior uncertainties around these optimal sets of parameters can be quantified elegantly using the Bayesian framework. The gradient for the optimization is evaluated efficiently using the adjoint model generated by the AD tool Tapenade. The frameworks are open source and freely available.

Shreyas S. Gaikwad
The University of Texas at Austin
The University of Texas at Austin
shreyas.gaikwad@utexas.edu

Laurent Hascoet
INRIA, France
laurent.hascoet@inria.fr

Shk Narayanan
Argonne National Laboratory
snarayan@mcs.anl.gov

Liz Curry-Logan
The University of Texas at Austin
liz.curry.logan@gmail.com

Ralf Greve
Hokkaido University, Japan
greve@lowtem.hokudai.ac.jp

Patrick Heimbach
The University of Texas at Austin
heimbach@oden.utexas.edu

MS318

A Framework for Time-Dependent Ice Sheet Uncertainty Quantification, Applied to Three West Antarctic Ice Streams

We seek to characterise the uncertainty in model projections of marine ice sheet loss, which arises from calibration with data. Here, we use a Bayesian approach, to quantify the degree to which observational uncertainty translates to parametric uncertainty (posterior uncertainty of inversions for basal drag and ice stiffness fields) and to uncertainty in projected quantities of interest (QoIs) such as sea level contribution. Our framework implements the Shallow Shelf Approximation (SSA), and a control methods approach to invert for the basal drag and ice stiffness fields. Beginning with a cost function optimization which can allow for either gridded or point-cloud velocities, we generate a low-rank approximation to the posterior covariance of the parameters via the use of the cost function Hessian. In our work, the Hessian is calculated through algorithmic differentiation (AD) using the “complete Hessian rather than the GaussNewton approximation. We then project the covariance on a linearization of the time-dependent ice sheet model to estimate the growth of the QoI uncertainty over time. We apply for the first time this framework to real ice streams in the Amundsen basin, and present several model experiments exploring the impact on uncertainty due to different strengths of priors, sliding laws and velocity inputs. We explore the model capabilities when we applied it to these ice streams and lay out our plans to scale our framework into a larger domain.

Beatriz Recinos, James Maddison, Daniel Goldberg, Joe Todd
University of Edinburgh
beatriz.recinos@ed.ac.uk, j.r.maddison@ed.ac.uk,
dan.goldberg@ed.ac.uk, joe.todd@ed.ac.uk

MS319

Learning Probabilistic Graphical Models of Non-Gaussian Scientific Data

One type of probabilistic graphical model is an undirected graph where the edges capture the conditional dependencies among a collection of random variables. Knowledge of this sparse edge structure is valuable for modeling multivariate distributions and for performing inference and prediction tasks. While several algorithms have been devel-

oped to learn the graph structure for Gaussian data, consistent methods for continuous and non-Gaussian distributions remain limited. In this presentation, we propose a structure learning algorithm based on a non-Gaussian characterization of conditional independence that uses Hessian information of the joint density. For a broad class of distributions, we show that the Hessian is related to conditional mutual information, a classical measure of conditional independence, while being more tractable for computation. To estimate the Hessian, our algorithm estimates the density using a deterministic coupling, known as a transport map, and exploits the sparse structure in the map to reveal the sparsity in the graph. We present results for learning the graph structure of non-Gaussian datasets drawn from chaotic dynamical systems, biology, and multi-disciplinary engineering systems where the graph structure arises from couplings between subdomain models.

Ricardo Baptista
MIT
rsb@caltech.edu

MS319

Deep Learning-Based Reduced Order Models in Fluid Dynamics and Structural Mechanics

The solution of differential problems by means of full order models (FOMs), such as, e.g., the finite element method, entails prohibitive computational costs when it comes to real-time simulations and multi-query routines. The purpose of reduced order modeling is to replace FOMs with reduced order models (ROMs) characterized by much lower complexity but still able to express the physical features of the system under investigation. Conventional ROMs anchored to the assumption of modal linear superimposition, such as proper orthogonal decomposition (POD), may reveal inefficient when dealing with nonlinear time-dependent parametrized PDEs, especially for problems featuring coherent structures propagating over time. To overcome these difficulties, we propose an alternative approach based on deep learning (DL) algorithms used to build an efficient nonlinear surrogate. In the resulting DL-ROM, both the nonlinear trial manifold and the reduced dynamics are learned in a non-intrusive way by relying on DL models trained on a set of FOM snapshots, obtained for different parameter values [Fresca et al, JSC, 2021; Fresca et al, CMAME, 2022]. Accuracy and efficiency of the DL-ROM technique are assessed on parametrized PDE problems in fluid dynamics [Fresca et al, Fluids, 2021] and structural mechanics, by considering, for example, MEMS micromirrors [Fresca et al, JNME, 2022], showing that new queries to the DL-ROM can be computed in real-time.

Stefania Fresca, Andrea Manzoni
Politecnico di Milano
stefania.fresca@polimi.it, andrea1.manzoni@polimi.it

MS319

Learning Physics-Based Reduced-Order Models from Data Using Quadratic Manifolds

Reduced-order models are imperative in making computationally tractable outer-loop applications that require simulating systems for many scenarios with different parameters and under varying inputs. For many physics-based systems linear dimension reduction (which underlies a large class of model reduction techniques) imposes a fundamental limitation to the accuracy that can be achieved using reduced-order models. In this talk I will propose a novel

approach for learning a data-driven quadratic manifold from high-dimensional data, then employing this quadratic manifold to derive efficient physics-based reduced-order models. The key ingredient of the approach is a polynomial mapping between high-dimensional states and a low-dimensional embedding. This mapping consists of two parts: a representation in a linear subspace (computed in this work using the proper orthogonal decomposition) and a quadratic component. Combining the quadratic manifold approximation with the operator inference method for projection-based model reduction leads to a scalable non-intrusive approach for learning reduced-order models of dynamical systems. Applying the new approach to transport-dominated systems of partial differential equations illustrates the gains in efficiency that can be achieved over approximation in a linear subspace. We also discuss several extensions to the proposed framework centered around the use of nonlinear dimension reduction techniques beyond the quadratic case.

Rudy Geelen

Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
rudy.geelen@austin.utexas.edu

Stephen Wright

University of Wisconsin-Madison
Dept. of Computer Sciences
swright@cs.wisc.edu

Karen E. Willcox

UT Austin
kwillcox@oden.utexas.edu

MS320

Globalized Simulation-Driven Miniaturization of High-Frequency Components Using Machine Learning and Variable-Fidelity Computational Models

Miniaturization is an important consideration in the design of contemporary high-frequency components. Small size requirements come in pair with increasing demands on electrical properties, which makes high-frequency design challenging, as reducing physical dimensions is detrimental to performance. Miniaturization-oriented parametric optimization is a heavily constrained task, which is also multimodal due to a parameter redundancy associated with topological modifications introduced in compact structures. Identification of truly minimum-size designs can only be realized through global search. Utilization of nature-inspired algorithms is impractical due to their poor computational efficiency as well as the fact that reliable evaluation of electrical characteristics requires full-wave electromagnetic (EM) simulation. Here, we introduce a machine learning procedure for low-cost global optimization-based size reduction of high-frequency structures. Our methodology involves a parameter space pre-screening, and kriging surrogate models iteratively refined using minimization of the predicted objective function as an infill criterion. Numerical experiments conducted using a number of antenna and microwave structures indicate that the proposed framework is capable of yielding consistent results for multiple algorithm runs, as well as achieving competitive miniaturization rates at low computational cost, as compared to the benchmark.

Slawomir Koziel

School of Science and Engineering

Reykjavik University, Iceland
koziel@ru.is

Anna Pietrenko-Dabrowska

Electronics, Telecommunications and Informatics
Gdansk University of Technology
anna.dabrowska@pg.edu.pl

MS320

AI Empowered Codesign of Beyond Moores Law Microelectronics

The co-design of next-generation hardware is currently a static process that involves human-in-the-loop evolution via repeated experiments, modeling, and design space exploration. Our goal is to develop an automation framework that integrates activities across the co-design process leveraging AI. AI has the potential to accelerate the exploration of strongly correlated parameter spaces such as those involved in co-design across scales and disciplines necessary for revolutionary jumps in computing performance. In this talk, we will outline the vision and then present results in data-driven modeling for device and circuit modeling.

Lavanya Ramakrishnan, Zhi Jackie Yao, Zhe Bai

Lawrence Berkeley National Laboratory
lramakrishnan@lbl.gov, jackie_zhiyao@lbl.gov,
zhebai@lbl.gov

Andrew Tritt

Lawrence Berkeley National Lab
ajtritt@lbl.gov

Ali Zaidi, Hasita Veluri, Dilip Vasudevan, John Shalf

Lawrence Berkeley National Laboratory
azaidi2@lbl.gov, hasitaveluri@lbl.gov, dilipv@lbl.gov,
jshalf@lbl.gov

MS321

CodeRefinery: What We Learned from Teaching Software Engineering Practices to Students and Researchers in Nordics and Beyond

CodeRefinery is a project which supports students and researchers across all academic disciplines by advancing the FAIRness of research software and development practices so that they can collaboratively develop, review, discuss, test, share, and reuse their codes. This initiative teaches the software engineering practices which are needed for modern collaborative computational research but are often missed in a traditional university curriculum. Thus far, over 2000 students and researchers have been trained. The project currently focuses on the Nordic/Baltic countries, but aims to expand beyond this region. CodeRefinery aims to operate as a community project with support from academic organizations. In this contribution we will share our experiences and lessons learned from developing and evolving the lesson curriculum and course format, from organizing over 50 events and from growing a community of instructors and research software engineers with the goal of improving software engineering practices and thus the quality of computational research and engineering. We will summarize which development tools and practices we have found to be essential and offer recommendations on where to start and how to progress.

Radovan Bast

UIT/Code Refinery

radovan.bast@uit.no

MS321

Can Software Metrics Improve Software Quality?

Software quality in computational science and engineering can be interpreted in two ways: does the software implement the algorithm or simulate the phenomenon as expected, and does it perform as expected i.e. does it scale, is it maintainable, is it secure? Software metrics are the degree to which a software system possesses some relevant property. Many software metrics are defined as quantitative measurements based on the analysis of source code. Others, such as those being developed by the CHAOSS initiative, focus on measuring the ability of a project to deliver software that meets various criteria, such as community health or development efficiency. A key question is how useful software metrics are in improving software quality when applied to development in computational science and engineering, which can differ from software engineering in other areas because of evolving or unclear requirements, deployment to large-scale systems and architectures, and focus on performance. I will consider which types of software metric might benefit researcher-developers and research software engineers working in computational science and engineering. Can useful metrics be identified by considering the differences in the way that software is developed these fields?

Neil P. Chue Hong
Software Sustainability Institute
University of Edinburgh
n.chuehong@epcc.ed.ac.uk

MS321

The Pull Request Assistant

A healthy GitHub project changes over time. Developers fix bugs, add new features and rework code to make it more efficient. A well-run project uses a Pull Request (PR) to review changes before they are allowed into the main branch. Our experience working with PR reviewers is that they have a difficult job. Not only do they have to review code changes that are part of a PR, but also review documentation and test-cases related to those changes, which may reside in other files separate from the PR. We are developing a Pull-Request Assistant, called MeerCat, that uses repository mining to make a reviewers job more manageable by showing the PR author (a) potential problems with the PR they have submitted and (b) modifications they can make to address those problems. A PR author is expected to address these problems before the PR moves into formal review stage. In essence, to aid a PR reviewer, MeerCat puts the onus on the PR author to come to review with a well-thought-out set of changes. The types of problems MeerCat currently focuses on are in the areas of documentation and testing: we aid a PR author in aligning both documentation and testing with code changes. In many cases, MeerCat can offer specific directions to bring this alignment about, including bringing missing but related files into the PR and suggesting specific modifications. MeerCat is in beta-testing with several projects at the moment.

Anshu Dubey
Argonne National Laboratory
adubey@anl.gov

Stephen Fickas, Jason Prideaux

University of Oregon
fickas@cs.uoregon.edu, jprideau@cs.uoregon.edu

Pablo Flores
Hop Skip Technologies
jpfloresd.97@gmail.com

Jared O'Neal, Rajeev Jain, Akash V. Dhruv
Argonne National Laboratory
joneal@anl.gov, jain@mcs.anl.gov, adhruv@anl.gov

Klaus Weide
University of Chicago
klaus@flash.uchicago.edu

MS321

Human-Computer Interaction in Open Computational Science

Two important tendencies in scientific research are the increasing importance of computational modelling and the nascent Open Science movement. Together, they force us to revise the way scientists and non-scientists use computational tools and models, as all stakeholders must be able to develop trust in these tools and models at their respective levels of domain knowledge and engagement. Today's software engineering techniques, inherited from an industrial context, focus on small development teams creating computational tools for a large number of users that interact with the tools at a much more abstract level. Software developers thus become gatekeepers of the computational models that are embedded in these tools, which mere users can neither fully understand nor safely modify. Open Science requires reducing this epistemic opacity, allowing all stakeholders to interact with computational models at their level of scientific rather than technical expertise. A strategy that I have been exploring is moving the models out of the tools and into the human-computer interface in the form of specifications. The central user interface of such a future scientific computing environment is an authoring tool for a computational medium that can represent models. Computational tools operate on this medium, reading models and parameters from it and adding back their results. I will conclude my presentation with a demonstration of a prototype implementation.

Konrad Hinsien
CNRS
konrad.hinsien@cnrs.fr

MS322

Towards the Use of Anderson Acceleration in Coupled Transport - Gyrokinetic Turbulence Simulations

Simulations predicting the long-time behavior of magnetically confined fusion plasma requires accurately and efficiently bridging the gap between transport and gyrokinetic turbulence processes evolving on vastly different time scales. The nonlinear transport equation solver, Tango, couples implicit 1D transport equations with a 5D gyrokinetic turbulence model using a relaxed fixed-point iteration method to enable simulations of very long times, in particular, steady state. In this talk we present the application of Anderson acceleration to the iteration scheme in Tango utilizing the KINSOL nonlinear solver package from SUNDIALS. We show that Anderson acceleration offers increased robustness to the choice of relaxation parameter enabling

faster convergence in a nonlinear diffusion test problem. Additionally, we will highlight recent results applying Anderson acceleration to test cases with increased stiffness and noisy transport fluxes.

David J. Gardner, Carol S. Woodward, Lynda LoDestro
Lawrence Livermore National Laboratory
gardner48@llnl.gov, woodward6@llnl.gov, lodestro1@llnl.gov

MS322

Filtering for Anderson Acceleration

Anderson acceleration is a powerful and low-cost method for accelerating the convergence of fixed-point iterations. On any given problem, how successful it is depends on the details of its implementation. In this talk we will review recent results on filtering the columns of the matrix of the least squares problem to control the condition number of the least-squares problem solved at each iteration and improve convergence.

Sara Pollock
University of Florida
s.pollock@ufl.edu

MS323

A Generalized Tau Method for Spectral Boundary Conditions in Multiple Dimensions

It has long been understood how to produce banded spectral discretizations of many PDEs using Petrov-Galerkin methods with orthogonal polynomials. One of the primary remaining challenges to building fully automatic fast spectral solvers is handling arbitrary boundary conditions. For domains with a single bounded dimension (e.g. an annulus or spherical shell), this can be reliably achieved with spectral integration and the classical tau method. For domains with multiple bounded dimensions (e.g. closed squares and cubes), consistency conditions at the corners and shared edges pose a challenge, particularly for mixed boundary conditions. Here we will discuss a generalized tau scheme where analytic perturbations are added to the PDEs and boundary conditions that result in consistent and exactly solvable discrete systems. In particular, we will examine the Poisson equation with generic mixed boundary conditions on the square and cube. This tau-perturbed system can be easily modified to use different test and trial spaces and is entirely isotropic, making it particularly favorable for incorporation into spectral element schemes.

Keaton Burns
Department of Mathematics, MIT
kjburns@mit.edu

Daniel Fortunato
Flatiron Institute
dfortunato@flatironinstitute.org

Geoff Vasil
School of Mathematics and Statistics, Univ. Sydney
gvasil@ed.au.uk

Keith Julien
Applied Mathematics
University of Colorado, Boulder

keith.julien@colorado.edu

MS323

A High-Order Fast Direct Solver for Surface PDEs

We introduce a fast direct solver for variable-coefficient elliptic partial differential equations on surfaces based on the hierarchical Poincaré-Steklov method. The method takes as input a high-order quadrilateral mesh of a surface and discretizes surface differential operators on each element using a high-order spectral collocation scheme. Elemental solution operators and Dirichlet-to-Neumann maps tangent to the surface are precomputed and merged in a pairwise fashion to yield a hierarchy of solution operators that may be applied in $O(N \log N)$ operations for a mesh with N elements. The resulting fast direct solver may be used to accelerate implicit time-stepping schemes, as the precomputed operators can be reused for fast elliptic solves on surfaces. We apply the method to a range of problems on both smooth surfaces and surfaces with sharp corners and edges, including the static Laplace-Beltrami problem, the Hodge decomposition of a tangential vector field, and some time-dependent reaction-diffusion systems.

Daniel Fortunato
Flatiron Institute
dfortunato@flatironinstitute.org

MS323

A Spectral Multistep Method for Time-Dependent PDEs

Krylov subspace spectral (KSS) methods are high-order accurate, explicit time-stepping methods for partial differential equations (PDEs) with stability characteristic of implicit methods. Unlike other time-stepping approaches, KSS methods compute each Fourier coefficient of the solution from an individualized approximation of the solution operator of the PDE. As a result, KSS methods scale effectively to higher spatial resolution. This talk will present explicit and implicit multistep formulations of KSS methods to provide a best-of-both-worlds situation that combines the efficiency and simplicity of multistep methods with the stability and scalability of KSS methods. Convergence analysis will be presented, and the effectiveness of spectral multistep methods will be demonstrated through numerical experiments. It will also be shown that the region of absolute stability exhibits striking behavior that helps explain the effectiveness of these new methods.

James V. Lambers, Bailey Rester
University of Southern Mississippi
School of Mathematics and Natural Sciences
James.Lambers@usm.edu, bailey.rester@usm.edu

MS323

Generalising the Classical Tau Method for Fun and Profit

The classical tau method for differential equations refers to solving an approximate system exactly over polynomials, instead of finding approximate solutions to exact ODEs & PDEs. While very similar, this distinction has practical consequences for the analysis and acquisition of solutions. Classical methods set the residual as the span of the highest numerical basis elements. However, once an appropriate residual space is proscribed, a finite polynomial solution exists independent of any algorithm to obtain it. Cheby-

shev polynomials are excellent numerically, but even lowly monomials (however poorly conditioned in a computer) are fantastic for analysing solutions on paper. This talk will present some data on using a wider range of tau polynomials. We will show the dependence of eigenvalues on the choice of residual space, including the appearance of spurious spectrum and the accuracy of resolved modes. Finally, we will discuss systematic efforts to optimise performance by considering tau-corrections as oscillatory low-rank updates to differential operators.

Geoff Vasil

School of Mathematics and Statistics, Univ. Sydney
gvasil@ed.au.uk

Keaton Burns

Department of Mathematics, MIT
kjburns@mit.edu

Daniel Lecoanet

Engineering Science and Applied Math, Northwestern Univ.
daniel.lecoanet@northwestern.edu

Keith Julien

Applied Mathematics
University of Colorado, Boulder
keith.julien@colorado.edu

MS324

Leveraging Half-Precision in Wireless Communication

Massive Multiple-Input-Multiple-Output is a crucial technology for Next-Generation networks (Next-G). It uses hundreds of antennas at transceivers to exchange data. However, its accurate signal detection relies on solving an NP-hard optimization problem in real-time. In this presentation, we introduce a new GPU-based detection algorithm that demonstrates the positive impact of low-precision arithmetic (FP16, INT16, INT8) and multiple GPUs to achieve next-G latency/scalability/accuracy requirements. Our approach iteratively extends a solution with several symbols representing the best combination out of the aggregated levels. The computation spanning iterations is formulated as a general matrix-matrix multiplication (GEMM) operation to leverage GPU architectures. Preliminary results using A100 GPU show around 2× time complexity improvement by exploiting low-precision GEMM arithmetic over the reference FP32 implementation. Our GEMM-based approach turns out to be oblivious to precision loss by reporting similar accuracy using FP32, FP16, and even INT16/INT8 operations. We demonstrate performance scalability using four A100 GPUs, achieving 2.3× speedup against a single-GPU version.

Adel Dabah

King Abdullah University of Science and Technology (KAUST)
adel.dabah.1@kaust.edu.sa

Hatem Ltaief

KAUST, Saudi Arabia
hatem.ltaief@kaust.edu.sa

Zouheir Rezki

University of California Santa Cruz
zrezki@ucsc.edu

Slim Alouini, David E. Keyes

King Abdullah University of Science and Technology (KAUST)
slim.alouini@kaust.edu.sa, david.keyes@kaust.edu.sa

MS324

Mixed Precision Linear Algebra for High Fidelity Real-Time Wavefront Reconstruction on Giant Optical Telescopes

While the largest ground-based telescopes will soon reach 40m diameter to provide means required to detect faint rocky exoplanets, they must overcome optical distortions induced by atmospheric turbulence. This is done thanks to Adaptive Optics (AO) employing a real-time controller (RTC), operating at high speed (1kHz) to catch up with the rapidly changing optical turbulence, and responsible for reconstructing aberrations measured by wavefront sensors (WFS) to drive deformable mirrors (DM) actuators used to compensate them. In particular Multi-Conjugate AO uses tomographic reconstruction in a linear control scheme: input measurements from several sensors are multiplied by a control matrix to produce an output command vector to drive several DMs. In this context, I will present the various avenues for leveraging approximate computing within the AO RTC. Taking benefit from sensors data coarse grain quantization, mixed precision approximations can be used in the real-time control loop to minimize the servo lag and maximize AO end-to-end performance. Additionally, whilst the tomographic reconstructor must be updated in quasi-real-time, the associated computational burden can be significantly relaxed through mixed precision linear algebra, leveraging intrinsic data sparsity from input matrices, with the goal to either come-up with more complex high-fidelity control schemes or optimize the computer system dimensioning with impact on construction, operation and maintenance costs.

Damien Gratadour

The Australian National University
damien.gratadour@obspm.fr

MS324

Low Rank (Almost) Everywhere

Tile low rank and hierarchical low rank matrices can exploit the data sparsity that is discoverable all across computational science: integral equations, differential equations (Schur complements), spatial statistics (covariances), optimization (Hessians), data compression, RBF-based meshing, non-Fickian diffusion, and in various applications such as seismic redatuming, acoustic scattering, adaptive optics, climate and weather predictions, and more. Exploiting data sparsity can improve performance dramatically by allowing the working set to dwell higher in the memory hierarchy than for dense algorithmic counterparts. We illustrate in large-scale applications and hybridize with similarly motivated mixed precision representations, while featuring recent research with many collaborators, including a multi-institutional 2022 Gordon Bell finalist paper.

David E. Keyes

King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS324

Portable Mixed Precision for the Iterative Solution

of Sparse Linear Systems

The convention in scientific computing is to employ IEEE double-precision (64-bit) arithmetic for all computations involving floating-point data. Nonetheless, appealing benefits from the adoption of mixed precision schemes have been reported for the solution of dense and sparse linear systems, on variety of computer architectures, via iterative refinement. In this talk, we will illustrate the benefits of leveraging mixed precision in terms of execution time and energy efficiency. For this purpose, we will target several case studies arising in the iterative solution of sparse linear systems on GPUs, with codes currently integrated the Ginkgo library (<https://ginkgo-project.github.io>). In some detail, this research effort exploits the fact that, for sparse linear algebra operations, the cost is dominated by the memory accesses while the arithmetic is largely irrelevant. To leverage this property, the Ginkgo solvers store certain parts of the data in reduced precision in memory, but operate in "full" 64-bit precision in order to bound the accumulation of rounding errors. Reduced-precision storage can be then exploited to maintain approximation operators, such as the preconditioner, or in a GMRES-based solver that adjust the precision of the basis to the problem.

José I. Aliaga
Universitat Jaume I
aliaga@icc.uji.es

Hartwig Anzt
University of Tennessee, U.S.
hanzt@icl.utk.edu

Thomas Grützmacher
Karlsruhe Institute of Technology
thomas.gruetzmacher@kit.edu

Enrique S. Quintana-Orti
Universitat Politècnica de València
quintana@disca.upv.es

Andrés Tomás
Universidad Jaume I
tomasan@uji.es

MS325

Global Stochastic Optimization of Stellarator Coils

In the construction of a stellarator, the manufacturing and assembling of the coil system is a dominant cost. These coils need to satisfy strict engineering tolerances, and if those are not met the project could be canceled as in the case of the National Compact Stellarator Experiment (NCSX) project. Therefore, our goal is to find coil configurations that increase construction tolerances without compromising the performance of the magnetic field. In this paper, we develop a gradient-based stochastic optimization model which seeks robust stellarator coil configurations in high dimensions. In particular, we design a two-step method: first, we perform an approximate global search by a sample efficient trust-region Bayesian optimization; second, we refine the minima found in step one with a stochastic local optimizer. To this end, we introduce two stochastic local optimizers: BFGS applied to the Sample Average Approximation and Adam, equipped with a control variate for variance reduction. Numerical experiments performed on a W7-X-like coil configuration demonstrate that our global optimization approach finds a variety of

promising local solutions at less than 0.1% of the cost of previous work, which considered solely local stochastic optimization.

Silke M. Glas
University of Twente
s.m.glas@utwente.nl

Misha Padidar, Ariel Kellison, David Bindel
Cornell University
map454@cornell.edu, ak2485@cornell.edu,
bindel@cornell.edu

MS325

High-Dimensional Multi-Objective Bayesian Optimization for Optical Design

The ability to optimize multiple competing objective functions with high sample efficiency is imperative in many applied problems across science and industry. Multi-objective Bayesian optimization (BO) achieves strong empirical performance on such problems, but even with recent methodological advances, it has been restricted to simple, low-dimensional domains. Most existing BO methods exhibit poor performance on search spaces with more than a few dozen parameters. In this work we propose MORBO, a method for multi-objective Bayesian optimization over high-dimensional search spaces. MORBO performs local Bayesian optimization within multiple trust regions simultaneously, allowing it to explore and identify diverse solutions even when the objective functions are difficult to model globally. We show that MORBO significantly advances the state-of-the-art in sample-efficiency for several high-dimensional synthetic and real-world multi-objective problems, including a vehicle design problem with 222 parameters, demonstrating that MORBO is a practical approach for challenging and important problems that were previously out of reach for BO methods.

David Eriksson
Meta
deriksson@fb.com

MS325

Bayesian Optimization for Hyperparameter and Neural Architecture Search

Hyperparameter and neural architecture search (HNAS) has been arguable one of the biggest success stories of Bayesian optimization. This is manifested in a long list of well-known open-source frameworks and commercial services that provide Bayesian optimization solutions for HNAS problems. However, despite its success, vanilla Bayesian optimization faces several challenges on popular HNAS use-cases: First, Bayesian optimization itself needs to become more robust and automated, to not open another hyperparameter optimization problem on top. Second, with the ever increasing computational demands of machine learning, we have to accelerate the optimization process as much as possible. Third, there is a growing interest to optimize more than just the validation performance, such as inference latency or model size. In this talk, I will present an overview of our recent work that tries to overcome the current short-coming of Bayesian optimization. I argue that these are necessary steps towards fully automated machine learning systems that are able to learn truly end-to-end.

Aaron Klein

Amazon Web Services Berlin
kleiaaro@amazon.com

MS325

Latent Space Bayesian Optimization for Molecular Design

Computational drug discovery can often be cast as a black-box optimization problem, where we aim to design a molecule of interest which maximizes some black-box objective function (i.e. its binding affinity to some target protein). However, optimization over molecules is challenging since the objective function is defined over the discrete and structured space of all possible molecules. Latent space Bayesian optimization (LS-BO) has recently emerged as a promising approach for optimizing over such search spaces. In LS-BO, a deep autoencoder model (DAE) maps molecules into a continuous latent space where familiar Bayesian optimization tools can be more readily applied. While recent work in this area has made rapid progress on the design of better DAEs for this task, relatively little attention has been paid to the optimization component. This is problematic, as the DAEs used typically have high dimensional latent spaces, which significantly degrades the performance of traditional Bayesian optimization algorithms. To address this issue, we develop a novel method to adapt recent work on high-dimensional Bayesian optimization to the LS-BO setting. Our method achieves as much as 20x improvement over state-of-the-art methods across five molecular design tasks.

Natalie Maus

University of Pennsylvania
nmaus@seas.upenn.edu

MS326

A Fast and Arbitrarily High-Order Solver for Highly Oscillatory ODEs

Oscillatory systems are ubiquitous in physics: they arise in celestial and quantum mechanics, electrical circuits, molecular dynamics, and beyond. Yet even in the simplest case, when the frequency of oscillations changes slowly but is large, the vast majority of numerical methods struggle to solve such equations. Methods based on approximating the solution with polynomials are forced to take $O(k)$ timesteps, where k is the characteristic frequency of oscillations. This scaling can generate unacceptable computational costs when the ODE in question needs to be solved billions of times, e.g. as the forward modelling step of Bayesian parameter estimation. In this talk I will introduce an efficient method for solving 2nd order, linear ODEs with highly oscillatory solutions. The solver employs two methods: in regions where the solution varies slowly, it uses a spectral method based on Chebyshev nodes and with an adaptive stepsize, but in the highly oscillatory phase it automatically switches over to an asymptotic method. The asymptotic method constructs a nonoscillatory phase function solution of the Riccati equation associated with the ODE. In the talk I will present how the method fits in the landscape of oscillatory solvers, the theoretical underpinnings of the asymptotic solver, a summary of the switching and stepsize-update algorithms, results from numerical experiments, and a brief error analysis.

Fruzsina Agocs

Center for Computational Mathematics
Flatiron Institute
fagocs@flatironinstitute.org

Alex H. Barnett
Flatiron Institute
Simons Foundation
abarnett@flatironinstitute.org

MS326

Wide-Band Butterfly Networks for Wave-Based Inverse Problems in the Super-Resolution Regime

We propose an end-to-end deep learning framework that comprehensively solves the inverse wave scattering problem across all length scales. Our framework consists of the newly introduced wide-band butterfly network coupled with a simple training procedure which dynamically injects noise during training. While our trained network provides competitive results in classical imaging regimes, most notably it also succeeds in the super-resolution regime where other comparable methods fail. This encompasses both (i) reconstruction of scatterers with sub-wavelength geometric features, and (ii) accurate imaging when two or more scatterers are separated by less than the classical diffraction limit. We demonstrate these properties are retained even in the presence of strong noise and extend to scatterers not previously seen in the training set. In addition, our network is straightforward to train requiring no restarts and has an online runtime that is an order of magnitude faster than optimization-based algorithms. We perform experiments with a variety of wave scattering mediums and we demonstrate that our proposed framework outperforms both classical inversion and competing network architectures that specialize in oscillatory wave scattering data.

Matthew Li

Massachusetts Institute of Technology
mtcli@mit.edu

Leonardo Andrés Zepeda Núñez
University of Wisconsin-Madison
zepedanunez@wisc.edu

Laurent Demanet
Department of Mathematics, MIT
laurent@math.mit.edu

MS326

Butterfly Algorithms: An Overview

The so-called butterfly matrix or butterfly operator is widely known in the signals processing community from its appearance in algorithms such as the Fast Fourier Transform. The basic idea is to split a signal into two pieces, apply separate linear transformations to each piece, and then linearly re-combine the output. The matrix corresponding to the Discrete Fourier Transform can be analytically factored into a product of $\log(N)$ block-diagonal butterfly matrices. Generalizations of such matrix factorizations yield what are now known as butterfly factorizations, and have found applications in PDE, integral equations, signals processing, special function transforms, etc. They often allow for the compression of an $N \times N$ matrix into a product of $\log(N)$ matrices, each with $O(N)$ sparsity. This talk will give a brief overview of such factorizations and applications, modern developments, and open problems as an introduction to the corresponding mini-symposium.

Michael O'Neil

New York University

oneil@cims.nyu.edu

MS326**Efficient Identification of Butterfly Sparse Matrix Factorizations**

Large neural networks perform well in many domains, but they suffer from long training and inference time. Introducing sparsity in these deep models by enforcing most entries in their weight matrices to zero can reduce their complexity. However, learning such sparse networks is difficult. Motivated by recent works showing the expressivity of the butterfly structure for neural network design, a promising approach is to promote sparsity during training by approximating any dense weight matrices with a product of sparse butterfly factors. Our main contribution is to show that any butterfly factorization $\mathbf{Z} = \mathbf{X}^{(1)} \dots \mathbf{X}^{(J)}$ where the factors $\mathbf{X}^{(\ell)}$ of size $N \times N$ follow the butterfly constraint is *essentially unique*, and we provide an efficient hierarchical algorithm that can recover these factors up to unavoidable scaling ambiguities from the product \mathbf{Z} . Our method consists in recursively factorizing the considered matrix into two factors, by relying on a non-trivial application of the singular value decomposition to compute best rank-one approximations of specific submatrices, instead of iterative gradient descent steps. The complexity of our algorithm is only of $O(N^2)$, and it enables fast $O(N \log N)$ matrix-vector multiplication. Hence our work provides a first step toward the construction of an efficient proximal operator associated to the butterfly structure for sparse neural network training.

Léon Zheng

valeo.ai

ENS Lyon

leon.zheng@ens-lyon.fr

Elisa Riccietti

Ecole Normale Supérieure (ENS) Lyon

LIP

elisa.riccietti@ens-lyon.fr

Remi Gribonval

INRIA Rennes - Bretagne Atlantique

remi.gribonval@inria.fr

MS327**Mathematical Challenges in Electric Machine Simulation**

Increasing the efficiency of electric machines got renewed interest, triggered by the fight against climate change. The design of electric machines relies on a combination of behavioural, network and finite-element models of the underlying electromagnetic and thermodynamic partial differential equations, accomplished by nonlinear (and possibly hysteretic) material models, a field-circuit coupling to power-electronics circuits and a mechanical model of the drive train. This talk aims at the construction of mathematical models, including a technical prioritisation of the model features, and a discussion of the contemporary challenges for a numerical solution thereof.

Herbert De Gersem

Institut Theorie Elektromagnetischer Felder, TU

Darmstadt

Graduate School Computational Engineering, TU

Darmstadt

degersem@temf.tu-darmstadt.de

MS327**Energy-Stable Stator-Rotor Coupling and Consistent Torque-Computation by Harmonic Mortaring**

The working principle of electric machines is based on the effective Magneto-mechanic energy conversion, and a key issue in electric machine simulation is the accurate representation of this mechanism after discretization. One particular difficulty in this endeavor is the correct handling of the relative motion between stator and rotor. In this talk, we discuss the efficient coupling of the magnetic field across the air gap by a harmonic mortar method. The use of trigonometric polynomials as Lagrange multipliers enables an efficient and elegant treatment of the relative motion of stator and rotor. Explicit formulas for the torque computation can be derived by energetic considerations, and their evaluation by harmonic mortar finite element and isogeometric analysis discretizations is possible. Numerical tests are presented to illustrate the theoretical results and demonstrate the potential of harmonic mortar methods for the evaluation of cogging torque and torque ripple.

Herbert Egger

Johannes Kepler University Linz

Institute for Computational Mathematics

herbert.egger@jku.at

Melina Merkel, Sebastian Schöps

Computational Electromagnetics Group (CEM), TU

Darmstadt

Centre for Computational Engineering (CCE), TU

Darmstadt

melina.merkel@tu-darmstadt.de,

schoeps@gsc.tu-

darmstadt.de

MS327**Tearing and Interconnecting Approach for $H(\text{curl})$**

In the setting of domain decomposition, Tearing and Interconnecting (TI) methods such as FETI (Finite Element TI) are well-established. They enable the parallel treatment of subdomains to efficiently simulate large and sparse problems. We aim to exploit this property for optimizing 3D models of electrical machines. Consequently, very large systems need to be evaluated multiple times in the optimization process. Hence, an efficient solver is a necessity. As a first step toward reaching this target, combining the discretization scheme Isogeometric Analysis (IGA) with TI methods for $H(\text{curl})$ yields a parallelizable procedure for magnetostatics. This work presents the construction of this solver and explains how tree-cotree gauging is used to handle intricacies arising from the function space $H(\text{curl})$. Additionally, the derived results are verified numerically and the efficiency is examined in praxis.

Mario Mally, Melina Merkel, Sebastian Schöps

Computational Electromagnetics Group (CEM), TU

Darmstadt

Centre for Computational Engineering (CCE), TU

Darmstadt

mario.mally@tu-darmstadt.de,

melina.merkel@tu-darmstadt.de,

schoeps@gsc.tu-darmstadt.de

Olaf Steinbach

TU Graz, Institute of Applied Mathematics

o.steinbach@tugraz.at

MS327

Efficient 3D Finite Element Scheme for Solving Electromagnetic Field Problems Considering Vector Hysteresis

Simulations of nonlinear magnetic field problems often use unique history-independent material laws, such as the commutation- or anhysteretic curve. However, the motivation to optimize, e.g., T-joints in three phase transformers or the field distribution in electric drives, improved and more realistic material models are needed. In this work, an energy based (EB) vector hysteresis model is used as a constitutive law with memory, which shows high potential for physically and thermodynamically consistent modeling of ferromagnetic and ferroelectric materials with high precision and efficiency. The model itself is based on consistent thermodynamic principles and allows an efficient numerical implementation by utilizing the convexity of the Gibbs energy potential, leading to a scalar optimization problem, for which a Newton-Raphson method is used. This material model is then used in a reduced magnetic scalar potential formulation in 2D and 3D for which the nonlinear permeability $\mu(\mathbf{H})$ (as a result of the EB model) can be used, contrary to \mathbf{B} -based formulation, where the reluctivity $\nu(\mathbf{B})$ is needed. This is a major advantage because the dependence on \mathbf{B} would require an inversion of the hysteresis model, which is, to the authors knowledge, acausal. The nonlinear electromagnetic problem is then solved by means of a Quasi-Newton scheme, using Good Broydens method and different line search approaches are compared.

Klaus Roppert

Graz University of Technology
klaus.roppert@tugraz.at

Lukas Domenig, Alexander Sauseng, Andreas Gschwentner, Manfred Kaltenbacher

Institute of Fundamentals and Theory in Electrical Engineering
TU Graz

lukas.domenig@tugraz.at, alexander.sauseng@tugraz.at,
andreas.gschwentner@tugraz.at, manfred.kaltenbacher@tugraz.at

MS328

Mathematics for Potato Cultivation and Breeding

Potato is one of the most versatile food crops with the yearly production around 359 million tons (2020 data). Climate change and ever increasing competition have produced serious challenges for potato breeding, cultivation and processing industries. Of course, this holds for other areas of horticulture and for agriculture, in general, as well. Although agriculture was among the earliest applications of statistics and experimental design, these new challenges often require novel mathematical approaches. For example, large throughput phenotyping is a well-known application for modern image processing methods. In this talk we discuss mathematical and computational problems that arose during our participation in a series of potato-industry-driven projects and share our solutions. We focus on the problems of spatial and temporal effect removal in field trials and on over-parameterized machine learning of predictive models for cultivation and breeding. In both cases one needs to recover a piecewise smooth function of either several or extremely many variables from poorly structured

and sparse data.

Neil Budko

Delft University of Technology
n.v.budko@tudelft.nl

Elisa Atza

TU Delft
e.atza@tudelft.nl

MS328

Two-Scale Topology Optimization for 3D Printing

Subject of my presentation is a novel approach for optimizing both the macroscopic shape and the porous mesoscopic structure of components. The key feature is the introduction of an additional local volume constraint (LVC), which allows to adjust the desired spatial scales. The main novelty is that the radius of the LVC may depend both on space and a local stress measure. This allows for creating optimal topologies with heterogeneous mesostructures enforcing any desired spatial grading and accommodating stress concentrations by stress dependent pore size. I will present some analytical results for the resulting optimal control problem and conclude with numerical simulations showing the versatility of our approach for creating optimal macroscopic designs with tailored mesostructures.

Dietmar Hoemberg

Weierstrass Institute for Applied Analysis and Stochastics
hoemberg@wias-berlin.de

MS328

Large-Scale and Distributed Optimization, Inference, and Learning in the Presence of Heavy-Tail Noise

In both centralized and distributed settings, stochastic gradient descent (SGD) is a frequently used tool to solve stochastic optimization problems. Therein, a noisy version of the cost functions gradient is utilized as a search direction, where the noise originates, e.g., due to data sampling. There have been several recent studies showing that the gradient noise that arises in many scenarios, e.g., when training deep learning models, is heavy-tailed. In distributed environments, an additional source of heavy-tailed noise is the imperfect communication between different computing units, e.g., nodes in a wireless sensor network. In this context, we develop both centralized and distributed SGD-type algorithms that provably work in the presence of heavy-tailed gradient and/or communication noises. The methods are based on introducing a generic nonlinearity in the iterate updates in order to combat the noise. We provide for the methods strong convergence and convergence rate guarantees, including almost sure, mean square rate, and asymptotic normality, in a challenging heavy-tailed setting, when noises are of infinite variance. We present the analysis for two different settings: 1) centralized SGD algorithms that provably work under heavy-tailed gradient noise and generic cost functions; and 2) distributed algorithms for stochastic quadratic costs (corresponding to a distributed estimation scenario) that provably work under heavy-tailed gradient and/or communication noises.

Dusan Jakovetic

University of Novi Sad

dusan.jakovetic@dmi.uns.ac.rs

MS328

Group Equivariant Non-Expansive Operators: a Mathematical Tool for Explainable Artificial Intelligence

Recently the topic of explainability of Artificial Intelligence has been gaining increasing interest from a large part of the scientific community. First, this is due to the desire to investigate the mechanisms of Artificial Intelligence more thoroughly, to better understand its core functioning, so that anomalies, unexpected behaviors and even counterfeits can be prevented. Second, the interest is in part encouraged by the demands of policymakers, who demand techniques that can be understood by humans and relied upon for public decision-making, not only for excellent results but also for inherent transparency. In this context Group Equivariant Non-Expansive Operators (GENEOs) have been developed to establish a mathematical theory of information processing agents. GENEOs are understood as the mathematical formulation of agents who act on the data, and give an opinion. Equivariance is indeed the key property of GENEOs, since it guarantees to filter out the data information that is not relevant for the problem in focus. In this talk we will introduce GENEOs and describe how they have been successfully employed to face some industrial problems, in the field of medicinal chemistry, implants inspection and maintenance and materials science.

Alessandra Micheletti

Department of Mathematics and ADAMSS Centre,
Universita' degli Studi di Milano, Milano, Italy
alessandra.micheletti@unimi.it

Giovanni Bocchi

Università degli Studi di Milano
Department of Environmental Science and Policy
giovanni.bocchi1@unimi.it

Patrizio Frosini

University of Bologna
patrizio.frosini@unibo.it

MS329

Low-Rank Tensor Methods for High-Dimensional Gaussian Processes

We consider the special case of Gaussian process kernel learning where the covariance function is given by a sum of products of RBF kernels. For a given dataset, the parameters of the kernel are learned by minimizing the log marginal likelihood. Computing the log-determinant of the covariance matrix is prohibitive for large datasets or in high dimensions unless one exploits its low-rank tensor structure. We employ a stochastic trace estimation together with a Lanczos algorithm for TT-tensors (Tensor Trains). This allows us to break the curse of dimensionality and to perform a gradient-based optimization even in high dimensions.

Josie König

University of Potsdam
josie.koenig@uni-potsdam.de

Max Pfeffer

TU Chemnitz
max.pfeffer@math.tu-chemnitz.de

Martin Stoll

Chemnitz University of Technology
martin.stoll@mathematik.tu-chemnitz.de

MS329

Preconditioning for Scalable Gaussian Process Hyperparameter Optimization

Gaussian process hyperparameter optimization requires linear solves with, and log-determinants of, large kernel matrices. Iterative numerical techniques are becoming popular to scale to larger datasets, relying on the conjugate gradient method (CG) for the linear solves and stochastic trace estimation for the log-determinant. This work introduces new algorithmic and theoretical insights for preconditioning these computations. While preconditioning is well understood in the context of CG, we demonstrate that it can also accelerate convergence and reduce variance of the estimates for the log-determinant and its derivative. We prove general probabilistic error bounds for the preconditioned computation of the log-determinant, log-marginal likelihood and its derivatives. Additionally, we derive specific rates for a range of kernel-preconditioner combinations, showing that up to exponential convergence can be achieved. Our theoretical results enable provably efficient optimization of kernel hyperparameters, which we validate empirically on large-scale benchmark problems. There our approach accelerates training by up to an order of magnitude.

Jonathan Wenger

University of Tuebingen
jonathan.wenger@uni-tuebingen.de

MS329

Adaptive Factorized Nyström Preconditioner for Gaussian Kernel Matrices

The spectrum of a kernel matrix significantly depends on the parameter values of the kernel function used to define the kernel matrix. This makes it challenging to design a preconditioner for a regularized kernel matrix that is robust across different parameter values. In this talk, we present the Adaptive Factorized Nyström (AFN) preconditioner. The preconditioner is designed for the case where the rank k of the Nyström approximation is large, i.e., for kernel function parameters that lead to kernel matrices with eigenvalues that decay slowly. Other Nyström preconditioners for a regularized kernel matrix use the Sherman–Morrison–Woodbury (SMW) formula and must solve with or compute an eigendecomposition of a k -by- k matrix. In contrast, AFN is a factorized form that avoids the SMW formula. It deliberately chooses a well-conditioned k -by- k matrix to solve with that has a sparse approximate inverse. This makes AFN efficient for large k . AFN also adaptively chooses k to balance accuracy and cost.

Yuanzhe Xi

Department of Mathematics
Emory University
yuanzhe.xi@emory.edu

Edmond Chow

School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

Shifan Zhao

Department of Mathematics

Emory University
shifan.zhao@emory.edu

Tianshi Xu
Department of Computer Science and Engineering
University of Minnesota, Twin Cities
xuxx1180@umn.edu

MS329

Efficient and Scalable Stochastic Variational Gaussian Processes

Gaussian processes (GP) are a popular probabilistic learning framework, especially when inference with uncertainty estimation is necessary. Oftentimes, GPs with derivative information can be useful in many settings where derivative information is available, including numerous Bayesian optimization and regression tasks that arise in the natural sciences. Given the $\mathcal{O}(n^3)$ cost in computation for n training points for GPs, stochastic variational Gaussian processes (SVGP) scale GP inference to large datasets through inducing points and stochastic training. However, the SVGP training process involves hard multimodal optimization, and often suffers from slow and suboptimal convergence when initializing inducing points directly from training data. Moreover, SVGP scales poorly with the number of dimensions when derivative information is incorporated. We split this presentation into two parts to focus on the two issues. In the first part, we discuss an efficient initialization of SVGP to make the training more efficient and in the second part we discuss a fully scalable SVGP method with derivatives, the training cost of which is independent of both the number of training data and the number of input dimensions.

Xinran Zhu, Misha Padidar, Leo Huang
Cornell University
xz584@cornell.edu, map454@cornell.edu,
ah839@cornell.edu

Jacob Gardner
University of Pennsylvania,
acobrg@seas.upenn.edu

David Bindel
Cornell University
bindel@cornell.edu

MS330

On the Robustness of Inverse Scattering for Penetrable, Homogeneous Objects with Complicated Boundary

The acoustic inverse obstacle scattering problem consists of determining the shape of a domain from measurements of the scattered far field due to some set of incident fields (probes). For a penetrable object with known sound speed, this can be accomplished by treating the boundary alone as an unknown curve. Alternatively, one can treat the entire object as unknown and use a more general volumetric representation, without making use of the known sound speed. Both lead to strongly nonlinear and nonconvex optimization problems for which recursive linearization provides a useful framework for numerical analysis. After extending our shape optimization approach developed earlier for impenetrable bodies, we carry out a systematic study of both methods and compare their performance on a variety of examples. Our findings indicate that the volumetric ap-

proach is more robust, even though the number of degrees of freedom is significantly larger. We conclude with a discussion of this phenomenon and potential directions for further research.

Carlos Borges
University of Central Florida
carlos.borges@ucf.edu

Leslie Greengard
Simons Foundation and Courant Institute
New York University
lgreengard@flatironinstitute.org

Manas N. Rachh
Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS330

Fast Algorithms for Certain Simulations in Quantum Optics

In this talk we present a fast and accurate algorithm for the numerical solution of a certain set of nonlocal partial differential equations arising in the study of the process of collective spontaneous emission in a two-level atomic system containing a single photon. The method works by first writing the system as an equivalent set of coupled integro-differential equations for the atomic degrees of freedom. This reformulation gives rise to a history dependence in the system. In order to evaluate the history contribution quickly we employ sum-of-exponentials representations. The performance of this approach is illustrated with several numerical examples.

Jeremy Hoskins
University of Chicago
jeremyhoskins@uchicago.edu

MS330

Fourier Representation of the Diffusion MRI Signal Using Layer Potentials

The diffusion magnetic resonance imaging signal arising from biological tissues can be numerically simulated by solving the Bloch-Torrey partial differential equation. Numerical simulations can facilitate the investigation of the relationship between the diffusion MRI signals and cellular structures. With the rapid advance of available computing power, the diffusion MRI community has begun to employ numerical simulations for model formulation and validation, as well as for imaging sequence optimization. Existing simulation frameworks use the finite difference method, the finite element method, or the Matrix Formalism method to solve the Bloch-Torrey partial differential equation. We propose a new method based on the efficient evaluation of layer potentials. In this paper, the mathematical framework and the numerical implementation of the new method are described. We demonstrate the convergence of our method via numerical experiments and analyze the errors linked to various model and simulation parameters. Since our method provides a Fourier-type representation of the diffusion MRI signal, it can potentially facilitate new physical and biological signal interpretations in the future.

Chengran Fang, Demian Wassermann, Jing-Rebecca Li
INRIA Saclay

chengran.fang@inria.fr, demian.wassermann@inria.fr, jin-grebecca.li@inria.fr

MS330

Deep Learning for the Inverse Obstacle Scattering Problem

The inverse scattering problem has many applications such as sonar and CT. In this short talk, I will introduce our work that combines the traditional methods and deep learning. We use the boundary integral method to generate training data and train a convolutional neural network to give an accurate initialization for the inverse problem. Then, we use a Newton-like traditional inverse algorithm to refine the boundary. Numerical examples show that this hybrid method is able to handle hard problems.

Mo Zhou

Duke University
mo.zhou366@duke.edu

MS332

Multiscale Finite Elements for Transient Atmospheric Flows with Sub-Grid Features

Spatial scale differences are a particular problem in climate and atmospheric modeling, where oftentimes small scale processes interact with larger scales. In numerical representations those processes acting below a resolvable spatial scale are often parameterized by averaging their effect over grid cells. However, such processes may exhibit spatial structures that alter the large scale behavior in a non-uniform manner such that averaging approaches are not suitable. Multiscale Finite Elements (MsFEM) address this shortcoming of traditional parameterization approaches and have been successful in stationary or quasi-stationary settings. We present a semi-Lagrangian reconstruction method to derive a MsFEM method for transient (transport-dominated) flow problems.

Jörn Behrens

Dept. of Mathematics, Universität Hamburg
joern.behrens@uni-hamburg.de

MS332

A Discontinuous-Galerkin-in-Time Framework for Multirate Time Integration of Interface-Coupled Problems

A framework is presented to design multirate time stepping algorithms for two dissipative models with coupling across a physical interface. The coupling takes the form of boundary conditions imposed on the interface relating the solution variables for both models. The multirate aspect arises when numerical time integration is performed with different time step sizes for the components. We describe a unified approach to develop multirate algorithms for these problems. This effort is pursued though the use of discontinuous-Galerkin time stepping methods, acting as a general unified framework, with different time step sizes. The two models are coupled across user-defined intervals of time, called *coupling windows*, using polynomials that are continuous on the window. The coupling method is shown to reproduce the correct interfacial energy dissipation, discrete conservation of fluxes, and asymptotic accuracy. Also, the method could provide a limit for the total communication needed between components, independent of the number of time steps on a coupling window. In prin-

ple, methods of arbitrary order are possible. As a preliminary step, we focus on the presentation and analysis of monolithic methods for advection-diffusion models coupled via generalized Robin-type conditions. The monolithic methods could be computed using a Schur-complement approach. We also discuss some recent progress for partitioned algorithms and different types of coupling conditions.

Jeffrey Connors

University of Connecticut
jeffrey.connors@uconn.edu

MS332

Multirate and Multi-Modeling Frameworks for High-Order Non-Hydrostatic Atmospheric Models

We are exploring process Multirate methods whereby each process in a system of nonlinear partial differential equations (PDEs) uses a time-integrator and time-step commensurate with the wave-speed of that process. We have constructed Multirate methods of any order using extrapolation methods. Along this same idea, we have also developed a multi-modeling framework (MMF, what was previously called super-parameterization) designed to replace the physical parameterizations used in weather/climate models. Our approach is to view the coarse-scale and fine-scale models through the lens of Variational Multi-Scale (VMS) methods in order to give MMF a more rigorous mathematical foundation. Our end goal is to use MMF in order to better resolve the inner core of hurricanes. Our model, NUMA, is a 3D nonhydrostatic atmospheric model that runs on large CPU clusters and on GPUs. We intend to show results for such simulations.

Francis X. Giraldo

Naval Postgraduate School
fxgiraldo@nps.edu

MS332

Tost.II :: A Temporal Operator-Splitting Template Library for deal.II

Temporal OS methods with real coefficients of order greater than two necessarily require backward time steps; however, temporal OS methods of arbitrary order can be implemented by a simple abstract pattern. This presentation describes a simple view of the general OS pattern and presents tost.II, its concrete implementation that is based on data structures in the deal.II finite-element library. The versatility and ease-of-use of general OS methods is demonstrated through several multi-physics examples.

Kevin R. Green

Global Institute for Water Security
University of Saskatchewan
kevin.green@usask.ca

Raymond J. Spiteri

University of Saskatchewan, Canada
Department of Computer Science
spiteri@cs.usask.ca

MS333

Fairification of MSE Simulation Data with Application to Machine Learning

The main targets of the Materials Science & Engineer-

ing (MSE) community are the characterization of materials and the study of the impact of materials processing and manufacturing to design materials with optimized properties. Laboratories develop and use different tools to generate data due to the vast number of different experimental, computational and analytical methods used in the MSE domain. Often, this myriad of tools does not follow standardized schemas and results in heterogeneous data when generated in different laboratories and even within the same laboratory. Heterogeneous data is difficult to exchange, reuse and process, resulting in uncoordinated MSE research. In the Data Analytics in Engineering group, we often store numerical results (e.g., homogenized stresses, stress fields, strain fields, plastic strains, ...) as a function of the loading. Description of the loading, its provenance, and authorship is quintessential to know. This information is important for machine learning tasks, however, metadata is usually missing in the produced datasets. Therefore, we need a common and complete description of the stored experimental results. Ontologies and knowledge graphs are state-of-the-art technology to enable such a unified data view. Thereby, we have designed and developed a knowledge graph; by using standardized World Wide Web Consortium (W3C) technologies and interfaces, the access to MSE experiment data enables compliance with the FAIR principles.

Amir Ladhaar

University of Stuttgart
Institute of Applied Mechanics
amir.laadhar@mib.uni-stuttgart.de

Renan Pereira Alessio, Felix Fritzen
University of Stuttgart
Data Engineering and Analytics
alessio@mib.uni-stuttgart.de,
fritzen@mib.uni-stuttgart.de

Rossella Aversa, Mehmet Soysal
Karlsruhe Institute of Technology
Steinbuch Centre for Computing
rossella.aversa@kit.edu, mehmet.soysal@kit.edu

MS333

Paving the Way for Open and Reproducible Research Within the CRC 1456

We witness an era where unprecedented amounts of data are acquired in experimental research in the natural sciences. While new measurement techniques and instruments keep being devised and improved for inexpensive and efficient data acquisition, the current bottleneck is how to extract meaningful information from the resulting vast amounts of such measurements. The DFG-funded Collaborative Research Center 1456 "Mathematics of Experiments" conducts research to improve mathematical tools for analyzing experimental data. In several subprojects between research groups from the natural sciences and mathematics, experiments and algorithms are developed, improved and analysed. Research activities in this field are very diverse which makes it challenging to define and ensure quality measures for good scientific practice in view of open research and reproducibility. In this talk, we present different actions undertaken as part of central efforts of the CRC's infrastructure project. This includes basic means such as infrastructure and training to ensure reproducibility, but also fostering scientific communication of research results through interactive documentations.

Pedro Costa Klein, Christoph Lehrenfeld, Sebastian

Mohr, Jan-Hendrik Müller, Markus Osterhoff
University of Göttingen
p.klein@math.uni-goettingen.de, lehrenfeld@math.uni-goettingen.de,
sebastian.mohr@stud.uni-goettingen.de,
jan-hendrik.mller@stud.uni-goettingen.de,
mosterh1@gwdg.de

MS333

Offering General Purpose Applications in Scientific Communities

Software interfaces are a central aspect when it comes to the usability of simulation frameworks. The coupling library preCICE is a general purpose software framework, which couples simulation software stemming from different physical domains in order to simulate a global system behavior using the coupled solvers. This talk deals with the inherent challenges arising due to the complexity of coupling different software interfaces employing different concepts as provided by each simulation software. A key aspect to cope with this complexity was the development of software-tailored adapter codes that bridge the gap between preCICE and the coupled application code. Nowadays, the preCICE ecosystem comprises several such adapter codes, corresponding tutorials as well as a broad documentation for each component. The variety of ready-to-use infrastructure in the preCICE ecosystem is one of the main reasons for a steadily growing and vibrant user community.

David Schneider

University of Stuttgart
Institute for Parallel and Distributed Systems
david.schneider@ipvs.uni-stuttgart.de

Benjamin Uekermann
University of Stuttgart
benjamin.uekermann@ipvs.uni-stuttgart.de

MS334

Quadrature-Based Lattice Boltzmann Models for Variable Prandtl Number Flows

The widely-used BGK or single-relaxation-time collision model suffers from the well-known problem that the Prandtl number, representing the ratio of momentum to thermal diffusivity, is fixed at unity. This limitation is overcome in numerous other approaches, of which we focus on the Shakhov and ellipsoidal statistics models. In this contribution, we address the implementation of these models in the spirit of the lattice Boltzmann approach, by discretizing the velocity space based on Gauss quadratures and replacing the equilibrium distribution by a truncated projection with respect to a set of orthogonal polynomials. We discuss an implicit-explicit implementation which allows the ideal hydrodynamics limit to be approached while keeping the time step finite.

Victor E. Ambrus

Department of Physics, West University of Timișoara
victor.ambrus@e-uvt.ro

MS334

The Magic Two-Relaxation-Time Lattice Boltzmann Method As a Macroscopic Finite Difference Scheme with Stress Relaxation

The two-relaxation-time lattice Boltzmann method assigns

different relaxation rates to odd and even moments of the distribution functions. Some "magic" combinations of odd and even relaxation rates offer good qualitative properties, e.g. placing the point of zero tangential velocity exactly half-way between grid points at no-slip boundaries. We present a reformulation using relaxation rates for forwards- and backwards-propagating distributions. The "magic" combination then sets the forwards-propagating distributions to equilibrium. Applying this approach over three time levels eliminates all non-equilibrium contributions from neighboring grid points. We thus derive closed macroscopic evolution equations over three time levels for the fluid density and velocity alone. These discrete equations are naturally interpreted as discretisations of PDEs with second time derivatives. The evolution equation for velocity decomposes into a momentum conservation law, and a separate stress evolution equation that simulates a linear Maxwell fluid with finite stress relaxation. All other kinetic effects decouple. We also derive closed macroscopic boundary conditions from bounce-back boundary conditions. Numerical experiments with flows driven by oscillating walls, and by vortex dipoles colliding with walls, confirm that solutions computed with the underlying lattice Boltzmann scheme and the closed macroscopic finite difference scheme agree to within round-off error in 128-bit arithmetic.

Paul Dellar
Mathematical Institute
University of Oxford
dellar@maths.ox.ac.uk

MS334

The Lattice Boltzmann Deviatoric Stress: Burnett Order Contributions and Their Implication for Boundary Conditions

Most common lattice Boltzmann models are used to compute solutions of the Navier-Stokes equations, for which the deviatoric stress is linear in the velocity gradient and proportional to the Knudsen number. We solve analytically a lattice Boltzmann model for planar channel flow to determine its deviatoric stress tensor. We find that the model contains an additional tangential stress, proportional to the square of the Knudsen number, that matches the stress in the Burnett equations. This additional stress has zero divergence, and hence does not affect the flow directly. However, the neglect of this additional stress by existing boundary conditions for the lattice Boltzmann method triggers spurious oscillations in the computed solutions. We propose new Burnett stress boundary conditions and analyse them theoretically. We show that they completely eliminate the oscillations, and agree precisely with analytical solutions at second order in Knudsen number. We combine these Burnett stress conditions with Navier-Maxwell slip boundary conditions to study rarefied flows in microcavities. We show that our algorithm captures non-equilibrium effects in the pressure that are not captured by the Navier-Stokes equations. The results are in very good agreement with Direct Simulation Monte Carlo (DSMC) solutions of the Boltzmann equation.

Tim Reis
University of Greenwich

t.reis@greenwich.ac.uk

MS335

Machine Learning and Digital Twins

The use of machine learning for digital twins is just beginning. In this talk we will discuss and survey the state of the art and point out the numerous challenges. The presentation will draw from the author's recent book [Asch M., A Toolbox for Digital Twins, SIAM, 2022]. The audience will be encouraged to contribute to the discussion.

Mark A. Asch
LAMFA UMR CNRS 7352
Université de Picardie Jules verne
mark.asch@u-picardie.fr

MS335

Reduced Order Models for Deformable Capsules

In this talk, we present a generic approach of a dynamical data-driven model-order reduction technique for three-dimensional microfluidic fluidstructure interaction problems. A low-order continuous linear differential system is learned from snapshot solutions of a high-fidelity (HF) solver. The reduced-order model uses different ingredients, such as proper orthogonal decomposition, dynamic mode decomposition and Tikhonov-based robust identification techniques. An interpolation method is used to predict the capsule dynamics for any values of the governing non-dimensional parameters that are not in the training database. Then a dynamical system is built from the predicted solution. Numerical evidence shows the ability of the reduced model to predict the time evolution of the capsule deformation from its initial state, whatever the parameter values. Accuracy and stability properties of the resulting low-order dynamical system are analysed numerically. The numerical experiments show very good agreement, measured in terms of modified Hausdorff distance between capsule solutions of the full-order and low-order models, in the case of both confined and unconfined flows. This work is a first milestone to move towards real-time simulation of fluidstructure problems, which can be extended to nonlinear low-order systems to account for strong material and flow nonlinearities. We believe that is a valuable innovative approach for deriving an efficient digital twin in this particular context.

Florian De Vuyst
Université de Technologie de Compiègne
CNRS
florian.de-vuyst@utc.fr

Claire Dupont
Université de Technologie de Compiègne CNRS
claire.dupont@utc.fr

Anne-Virginie Salsac
Université de Technologie de Compiègne, CNRS
anne-virginie.salsac@utc.fr

MS335

ϕ -Fem: Finite Elements on Unfitted Grids to Train Convolutional Neural Networks

ϕ -FEM is a recently proposed finite element method for the efficient numerical solution of partial differential equations posed in domains of complex shapes, using simple struc-

tured meshes, not necessarily fitted to the domain, and achieving the optimal accuracy. The only geometrical input for ϕ -FEM is a level-set function of the domain. In this talk, we shall present a combination of ϕ -FEM with the Convolutional Neural Networks (CNN) allowing for real-time predictions to be computed on any shape, with the ultimate goal of creating digital twins of human organs. We recall that CNN require a particular topological structure, in the form of a regular 2D or 3D grid. In themselves, they are well suited for e.g. processing images. However, in our application domain, where the governing equations should be solved on domains reproducing the geometries of real organs to train the network, the requirement of a structured mesh is challenging for traditional finite element solvers. In this context, ϕ -FEM turns out to be a promising alternative for training a CNN to provide predictions under the varying applied forces and under the varying geometries, conveniently represented by level sets that can be acquired from the medical images.

Stéphane Cotin, Michel Duprez, Sidaty El Hadramy
MIMESIS team, Inria Nancy - Grand Est,
MLMS team, Université de Strasbourg, France
stephane.cotin@inria.fr, michel.duprez@inria.fr,
sidaty.el-hadramy@loria.fr

Vanessa Lleras
IMAG, Univ Montpellier, CNRS, France
vanessa.lleras@umontpellier.fr

Alexei Lozinski
Laboratoire de Mathématiques de Besançon
CNRS et Université de Franche-Comté
alexei.lozinski@univ-fcomte.fr

Vincent Vigon
IRMA, Université de Strasbourg, CNRS, France
vincent.vigon@math.unistra.fr

Killian Vuillemot
MIMESIS team, Inria Nancy - Grand Est,
IMAG, Univ Montpellier, CNRS, France
killian.vuillemot@inria.fr

MS336

Quantifying Uncertainty in Stochastic Programming Parameters for Oil Refinery Planning

Engagement with industry practitioners is one of the important activities or programs in promoting applied and industrial mathematics as they face a lot of uncertainties in carrying out their business. Among other types of uncertainty they faced are knowledge uncertainties, natural uncertainties, and decisions to be made under uncertainties. In an oil refinery, uncertainties in the supply and product demand cover all aspects of uncertainties. This talk will present how stochastic programming parameters be estimated to help oil refinery planners make decisions when the supply and demand in the industry are hit by the rapid fluctuations in oil prices. The fluctuations in oil prices and unstable product demands result in disruption at procurement, production, and inventory stages. The parameters and the types of uncertainties involved in this project were identified through literature and discussion with industry practitioners. Then, an accurate forecast is required to quantify the uncertainties in the oil prices dependent on the behaviour of its volatility. In this talk, the successful implementation of the project will also be

highlighted.

Arifah Bahar
UTM Centre for Industrial and Applied Mathematics
Universiti Teknologi Malaysia
arifah@utm.my

MS336

Mobility and Population Growth Models for Fisheries, with Solutions

Fish population densities vary significantly in space, due to heterogeneous environments and harvesting. Popular spatial models are predominantly reaction-diffusion equations of Fisher-KPP type. However, Brownian motion that leads to a constant diffusion coefficient is a rough first approximation to fish behaviour. For two species of reef snapper, around 5% of the population (rangers) travel much further than the majority (home-stayers). This is not explainable by linear diffusion. Possibilities of improved mobility models include density-dependent mobility, Gaussian mixture models and non-Brownian Levi processes with long-tailed distributions. The resulting densities satisfy non-linear reaction-diffusion equations of fractional or integer order. Exact solutions can be obtained for the interior of a no-take area (NTA) with lethal boundary conditions that represent over-fishing in the exterior. This gives exact expressions for the critical domain size for survival in the NTA.

Philip Broadbridge
Department of Mathematical and Physical Sciences
La Trobe University
broadbridgephil@gmail.com

MS336

Generalization of Reeb Spaces and Application to Data Visualization

In many cases, data sets can be considered to be discrete samples of differentiable maps between manifolds. For a differentiable multivariate function into \mathbf{R}^p with $p \geq 2$, its Reeb space is the space of connected components of its fibers. This is a generalization of the notion of Reeb graphs for univariate functions in the case of $p = 1$. It has been known that Reeb spaces are often very useful for visualizing the given multivariate function. In this talk, we generalize the Reeb space in such a way that it captures more of the topological features of the fibers, not only their connected components. This theoretical part essentially relies on the global singularity theory of differentiable maps between manifolds developed mainly by the author. Such techniques have been used for efficiently visualize large scale data. If time permits, we will also discuss an application to multi-objective optimization problems.

Osamu Saeki
Institute of Mathematics for Industry
Kyushu University
saeki@imi.kyushu-u.ac.jp

MS336

Lattice Basis Reduction and Its Application to Cryptanalysis

The security of modern lattice-based cryptography relies on the hardness of solving lattice problems, such as the shortest vector problem (SVP) and the closest vector problem

(CVP). Lattice basis reduction is a strong tool for solving lattice problems. Given a basis of a lattice, a reduction algorithm is to find a new basis of the same lattice whose basis vectors are relatively short and nearly orthogonal each other. In this talk, I introduce several reduction algorithms, such as the Lenstra-Lenstra-Lovasz (LLL) and the block Korkine-Zolotarev (BKZ) algorithms that are popular in cryptanalysis. I also describe mathematical properties of such reduction algorithms and demonstrate how a reduction algorithm can solve lattice problems.

Masaya Yasuda
Department of Mathematics
Rikkyo University
myasuda@rikkyo.ac.jp

MS337

A Parallel Dynamical Low-Rank Integrator for Radiation Transport

Radiation transport problems are posed in a high-dimensional phase space, limiting the use of finely resolved numerical methods. An emerging tool to efficiently reduce computational costs and memory footprint in such settings is dynamical low-rank approximation (DLRA), which evolves the solution on a low-rank manifold. Since this manifold can exhibit a high curvature, time integrators for DLRA need to be carefully constructed to guarantee a robust error bound while preserving crucial properties of the original problem. This talk discusses a novel rank-adaptive and fully parallel DLRA time integrator. The integrator shares the robust error bound of the projector-splitting and the unconventional integrator while allowing basis functions and coefficients to be updated in parallel. Moreover, it allows for rank adaptivity while not requiring the potentially expensive augmented Galerkin step of the rank-adaptive unconventional integrator. We demonstrate the effectiveness of the proposed integrator for various radiation transport problems including radiation therapy. The presented results have been developed in cooperation with Gianluca Ceruti and Christian Lubich.

Jonas Kusch
University of Innsbruck
jonas.kusch@uibk.ac.at

MS337

A Large-Stepsize Integrator for Charged-Particle Dynamics in a Strong Magnetic Field

Xiao and Qin [Computer Physics Comm., 265:107981, 2021] recently proposed a remarkably simple modification of the Boris algorithm to compute the guiding centre of the highly oscillatory motion of a charged particle with step sizes that are much larger than the period of gyrorotations. They gave strong numerical evidence but no error analysis. In this talk, we provide an analysis of the large-stepsize modified Boris method in a setting that has a strong non-uniform magnetic field and moderately bounded velocities, considered over a fixed finite time interval. The error analysis is based on comparing the modulated Fourier expansions of the exact and numerical solutions, for which the differential equations of the dominant terms are derived explicitly. Numerical experiments illustrate and complement the theoretical results.

Yanyan Shi
University of Tübingen
shi@na.uni-tuebingen.de

Christian Lubich
Mathematisches Institut
Universität Tübingen
lubich@na.uni-tuebingen.de

MS337

Variable Dynamic Mode Decomposition for Model Order Reduction in Multiscale Particle Transport Problems

Variable Dynamic Mode Decomposition (V-DMD), a method for calculating time eigenvalues which does not require a constant time step, is demonstrated as a method for model order reduction in multi-scale problems. V-DMD is shown to be accurate for computing time eigenvalues in systems that were intractable for DMD due to the presence of a large separation of relevant time scales. Time eigenvalues of an infinite medium neutron transport problem with delayed neutrons, and consequently several time scales, are found with V-DMD. Lastly, V-DMD is demonstrated to calculate time eigenvalues with accuracy very similar to the typical DMD approach in systems where that typical approach can be used.

Ethan Smith
Notre Dame University
esmith45@nd.edu

MS338

Quantum Process Engineering using Iterative Learning Control

Quantum optimal control (QOC) promises to implement efficient pulse instructions for repeatable quantum processes such as gates and algorithm subroutines. Typically, QOC objectives are stated in terms of a terminal cost. Alternatively, the QOC objective can include both a running and terminal cost for the purpose of tracking a design trajectory. If existing QOC software is treated as a black box, then a robust, fast, and optimal-bandwidth control pulse can be obtained and used to induce a corresponding best-case design trajectory. This QOC solution can still suffer from model mismatch with the experimental system. In this work, we apply iterative learning control (ILC) to eliminate the tracking error from model mismatch by leveraging a limited number of state tomography experiments along the process trajectory. Many industrial control applications successfully operate by tracking specific trajectories of repetitive tasks across iterations. As the process repeats, any errors along the iteration axis would also repeat; by combining the model and feedback, ILC offers a framework for suppressing these errors. To demonstrate our approach, we apply ILC to realize the necessary pulses to track the design trajectory of representative one and two-qubit gates under quasi-static control errors.

Andy J. Goldschmidt
University of Washington
andygold@uchicago.edu

MS338

Leveraging Quantum Decoherence for Fast Ground State Reset and Pure-State Preparation

A key requirement for building general purpose quantum computers is the ability to perform fast and accurate state preparation, e.g. to initialize quantum algorithms into well defined states, for fault-tolerant quantum memory, or

quantum correction schemes requiring a continuous supply of qubits in a low-entropy state, such as the ground state. Unconditional quantum state preparation requires the system under study to be coupled to a quantum dissipative channel, allowing for entropy exchange between the system of interest and a bath. In this talk, we will present a framework for encouraging these processes through driven controls that are optimally designed using numerical optimal control techniques on classical HPC platforms. We model the underlying quantum dynamics in terms of the density matrix satisfying Lindblad's master equation, driven by control pulses whose frequencies precisely trigger qubit state transitions, and apply gradient-based optimization to design optimal control pulse shapes. Utilizing a specific optimization objective, we show how an ensemble of initial density matrices can be used as a single initial state throughout the optimization process, while achieving the desired target state for any initial qubit state. The scheme enables automated design of optimal state preparation protocols that achieve faster ground state reset than traditional protocols as well as general pure-state preparation targeting decoherence-protected qubits.

Stefanie Guenther

Lawrence Livermore National Laboratory
guenther5@llnl.gov

Anders Petersson

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
petersson1@llnl.gov

MS338

Optimal Control of Open Quantum Systems as a Practical Tool in the Quantum Technologies

Optimal control theory (OCT) is a versatile tool that can be used to identify control strategies in the presence of decoherence, either avoiding or exploiting the environment. In this talk, I will discuss examples for both. First, for qubits encoded in the infinite-dimensional Hilbert space of bosonic modes, I will show how to employ OCT to systematically enhance performance in the implementation of strong and on-demand interactions between the qubits. Second, I will discuss OCT in driven dissipative evolutions. For example, in trapped ion qubits, a protocol derived with OCT allows for the resource-efficient dissipative generation of an entangled state. Finally, I will assess prospects for quantum control in reservoir engineering where the desired dissipation can be realized via quantum non-demolition measurements.

Christiane Koch

FU Berlin
christiane.koch@fu-berlin.de

MS339

Spectrum Aware Coarsening for Graph Partitioning

Graph partitioning is a problem with applications in data mining, parallel computing, and scientific computing. For parallel computations, obtaining a balanced partition is desirable when distributing workloads. Unfortunately well studied methods such as spectral partitioning are too slow to be effective for this step at scale. For this reason, multilevel methods such as METIS coarsen the graph before performing spectral partitioning at the coarsest level and subsequently refining the partition. The quality of the par-

titions yielded from these multilevel methods tend to rely heavily on the quality of the coarsening. This quality is measured by how well the spectral cuts of the coarsened graph represent the true spectral cuts on the finest graph. We investigate the spectral approximation properties of different coarsening methods and implement these methods in parallel on GPU. We show that the spectral approximation properties of different coarsening algorithms vary significantly as do the quality of their output partitions.

Christopher Brissette, George M. Slota

Rensselaer Polytechnic Institute
brissc@rpi.edu, slotag@rpi.edu

MS339

ParMOO: A Python Library for Parallel Multiobjective Simulation Optimization

ParMOO is a Python framework and library of solver components for building and deploying highly customized multiobjective simulation optimization solvers. ParMOO is designed to help engineers, practitioners, and optimization experts exploit available structures in how simulation outputs are used to formulate the objectives for a multiobjective optimization problem. In this talk, we introduce our response surface modeling-based framework for building custom solvers and describe how this enables users to exploit a broad range of structures that commonly arise in multiobjective simulation optimization problems. We then show how ParMOO implements this framework and some additional features that make this a practical and user-friendly solution. ParMOO is pip- and conda-installable (package name: parmoo) and its source code is publicly maintained on GitHub: github.com/parmoo/parmoo

Tyler H. Chang

Argonne National Laboratory
tchang@anl.gov

Stefan Wild

Argonne National Laboratory
LANS, Mathematics and Computer Science
wild@lbl.gov

MS339

Targeted Adaptive Design

Modern advanced manufacturing and advanced materials design often require searches of relatively high-dimensional process control parameter spaces for settings that result in optimal structure, property, and performance parameters. The mapping from the former to the latter must be determined from noisy experiments or from expensive simulations. We abstract this problem to a mathematical framework in which an unknown function from a control space to a design space must be ascertained by means of expensive noisy measurements, which locate optimal control settings generating desired design features within specified tolerances, with quantified uncertainty. We describe targeted adaptive design (TAD), a new algorithm that performs this optimal sampling task. TAD creates a Gaussian process surrogate model of the unknown mapping at each iterative stage, proposing a new batch of control settings to sample experimentally and optimizing the updated log-predictive likelihood of the target design. TAD either stops upon locating a solution with uncertainties that fit inside the tolerance box or uses a measure of expected future information to determine that the search space has been exhausted with no solution. TAD thus embodies the

exploration-exploitation tension in a manner that recalls, but is essentially different from, Bayesian optimization and optimal experimental design.

Carlo Graziani

Argonne National Laboratory
cgraziani@anl.gov

Marieme Ngom
University of Illinois at Chicago
mngom@anl.gov

MS339

What Is Fastmath?

The FASTMath (Frameworks, Algorithms and Scalable Technologies for Mathematics) Institute is a R&D project in applied and computational mathematics and is funded by the U.S. Department of Energy (DOE) through the Scientific Discovery Through Advanced Computing (SciDAC) Program. The objectives of FASTMath are to develop robust mathematical techniques and numerical algorithms to reliably address the challenges of large-scale simulation of complex physical phenomena; deliver highly performant software with strong software engineering to run efficiently and scalably on current and next-generation advanced computer architectures at the DOE Office of Science's major computing facilities; work closely with domain scientists to leverage our mathematical and machine learning expertise and deploy our software in large-scale modeling and simulation codes; and build and support the broader computational mathematics and computational science communities across the DOE complex. In this talk, we briefly describe the SciDAC Program and the role that FASTMath plays, as well as some of the achievements FASTMath has made.

Esmond G. Ng

Lawrence Berkeley National Laboratory
egng@lbl.gov

MS340

Transforming (not Just) DG-FEM Array Expressions (not Just) on GPUs

While domain-specific languages for finite-element-based PDE solvers have enjoyed considerable success, composing them with supporting computations (e.g. chemistry, neural networks, non-variational discretization schemes) has proved challenging. We solve this by developing a lazy-evaluation based array package (Pytato) with numpy-like semantics that also serves as an Intermediate Representation (IR) for operations on multi-dimensional arrays. In our framework, we lower the Pytato IR to an imperative, scalar, polyhedrally-based IR, based on our Loopy tool for loop transformations. Using these IRs we developed a transformation strategy that delivers near-roofline performance for a class of array computation graphs that have the same memory access pattern as that of chained Einstein summation expressions. We claim that the most-common DG-FEM operators have a similar computation structure and support it by applying the transformation to a range of operators seen in real world applications from Combustion Physics, Wave Mechanics and Electrodynamics. We end by showing the relative-roofline performance metrics for a suite of DG-FEM operators and also show a comparison with off-the-shelf array packages and hand-tuned

DG-FEM libraries.

Kaushik Kulkarni

University of Illinois at Urbana-Champaign
kgk2@illinois.edu

Andreas Kloeckner
University of Illinois
andreask@illinois.edu

MS340

Solving Coupled FEM-BEM Problems Using FeniCSx and Bempp

In a range of applications—for example, wave transmission problems or the modelling of ultrasound—solvers that couple the finite element method (FEM) and boundary element method (BEM) can be very powerful. FeniCSx and Bempp are popular open source FEM and BEM libraries. In the past, multiple versions of these libraries have successfully been coupled together to perform small-to medium-sized simulations. More recently, work has been undertaken to move to a more scalable and parallelisable implementation of a coupled FEM-BEM solver. In this talk, we will present progress towards this solver and discuss some of the challenges in its implementation.

Matthew Scroggs

University of Cambridge
matthew.scroggs.14@ucl.ac.uk

Timo Betcke
University College London
t.betcke@ucl.ac.uk

Matthew Archer
University of Cambridge
United Kingdom
ma595@cam.ac.uk

MS340

pyop3: A New Domain-Specific Language for Automating High-Performance Mesh-Based Simulation Codes

With the growing complexity of modern computing hardware, the composition of appropriate domain-specific abstractions is becoming increasingly important to achieving the vaunted 'three Ps' of performance, portability and productivity. Here, we present preliminary work on pyop3, a new domain-specific language embedded in Python for automating mesh-based simulations. The primary contribution of pyop3 is the introduction of a hierarchical data layout language that associates topological information with the locations of the degrees-of-freedom. This enables a generic description for any partial structure in the mesh; for example extrusion or refinement. With such a description, we exploit this structure to avoid extraneous indirect addressing, and the resulting penalty of additional data movement, resulting in possible order-of-magnitude increases in performance.

Connor Ward

Imperial College London
c.ward20@imperial.ac.uk

MS341

Portable Batch Solvers in Structure Preserving Ki-

netic Methods with PETSc

I will introduce our integrated program for both developing structure-preserving (SP) methods for kinetic applications and deploying these methods as high-performance tools in PETSc (Portable, Extensible, Toolkit for Scientific computing). The metriplectic formalism used to develop these methods is introduced. We briefly discuss several methods developed for this code: a mixed Poisson solver for a C^0 electric field; strictly conservative mapping between particle and finite element bases; monotonic entropy time integrators for collisions; and new particle based Landau collision operators. A mature, high-order accurate, finite element based Landau collision operator with adaptive mesh refinement, that has been optimized for accelerator architectures using the portable Kokkos programming language, is presented with results using the NVIDIA A100 and AMD MI250X architectures. We present verification studies using plasma resistivity and compare with Spitzer resistivity. New batch GPU linear solvers have been developed for this work. This work is integrated into the PETSc “solver” framework to provide a fully GPU implicit time advance of the Landau collision operator. We show that collision time advance with many species is practical for $3X + 2V$ models of tokamaks on today's large-scale computers in cylindrical coordinates and that fully $3V$ models should be feasible in the near future.

Mark Adams

Lawrence Berkeley National Laboratory
mfadams@lbl.gov

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Eero Hirvijoki
Aalto University
eero.hirvijoki@gmail.com

Joseph Pusztay
University at Buffalo
Department of Computer Science and Engineering
josephpu@buffalo.edu

Daniel Finn
University at Buffalo
dsfynn@buffalo.edu

MS341

Batched Time Integrators for Chemistry Applications

Current supercomputers are increasingly relying on GPUs to achieve high throughput while maintaining a reasonable power consumption. Consequently, scientific applications are adapting to this new environment and new algorithms are designed to leverage the high concurrency of GPUs. One example of that adaptation is the demand for more batched kernels that concurrently operate on many inputs in a SIMD fashion. In this talk we will present time integration algorithms implemented in the Kokkos Kernels library that are callable within parallel kernels. These algorithms are used by chemistry applications to simulate reacting flows in CFD simulations, in this case the time integration scheme can be called at each integration point of the fluid discretization. Interface, capabilities and imple-

mentation details will be discussed, performance examples on various architectures will be presented.

Luc Berger-Vergiat, Jennifer A. Loe
Sandia National Laboratories
lberge@sandia.gov, jloe@sandia.gov

MS341

Batched Sparse Iterative Solvers on GPU for the XGC Fusion Plasma Application Collision Operator

Batched linear solvers, which are employed to solve many small related but independent problems, are important for several applications. This is particularly the case for accelerators such as graphics processing units (GPUs), which require a substantial amount of work to keep them operating efficiently and solving smaller problems one after another would be highly inefficient. An example use case is found in the XGC gyrokinetic particle-in-cell (PIC) code used for modeling magnetically confined fusion plasma devices. The collision operator employs a LAPACK linear solver, which does not run on a GPU, for current production runs. As these matrices are well-conditioned, batched iterative sparse solvers are an attractive option. A proxy application has been created for facilitating optimizations and porting of the collision operator to GPUs. We describe how Ginkgo's batched solver can be integrated into the collision kernel and accelerate the simulation process. Comparisons for the solve times on NVIDIA V100 and A100 GPUs and AMD MI100 GPUs with a CPU node are presented for matrices from the collision kernel of XGC. The results suggest that Ginkgo's batched sparse iterative solvers are well suited for efficient utilization of the GPU for this problem, and the performance portability of Ginkgo in conjunction with Kokkos (used within XGC as the heterogeneous programming model) allows seamless execution for exascale oriented heterogeneous architectures.

Paul Lin

Lawrence Berkeley National Laboratory
paullin@lbl.gov

Aditya Kashi
Oak Ridge National Laboratory
kashia@ornl.gov

Pratik Nayak
Karlsruhe Institute of Technology
pratik.nayak@kit.edu

Dhruva Kulkarni
Lawrence Berkeley National Laboratory
dkulkarni@lbl.gov

Aaron Scheinberg
Jubilee Development
aaron@jubileedev.com

Hartwig Anzt
University of Tennessee, U.S.
hantz@icl.utk.edu

MS341

Leveraging Batched Kernels in Hypre

The need for batched kernels has emerged to make more efficient use of GPUs for many small and sparse linear al-

gebra problems. This talk will discuss some of our recent efforts in utilizing batched kernels within the hypre linear solver library to improve performance and make effective use of existing and emerging hardware. Numerical results will be presented to demonstrate potential performance improvements when using batched kernels within algebraic multigrid solvers and preconditioners with modern hardware accelerators.

Sarah Osborn

Lawrence Livermore National Lab
osborn9@llnl.gov

Victor A. Paludetto Magri, Ulrike Meier Yang
Lawrence Livermore National Laboratory
paludettomag1@llnl.gov, yang11@llnl.gov

MS342

Asymptotic Models for Evolving Nematic Liquid Crystal Films

Asymptotic methods are used to derive a 4th-order nonlinear partial differential equation governing the evolution of a free surface thin film of nematic liquid crystal (NLC) spreading and/or dewetting on a solid substrate. Within a certain temperature range, NLCs behave as anisotropic liquids, with anisotropy imparted by the tendency of the long, rod-like, molecules to align locally. We consider the evolution of "nanoscale" films, of thickness around 100nm, for which van der Waals-type intermolecular interactions between the film and the substrate are important. By implementing an ADI-type numerical method within a GPU computing environment [Lam et al., Computing dynamics of thin films via large scale GPU-based simulations, J. Comp. Phys. X 2, 100001, 2019], we are able to simulate our model efficiently to describe 3D flows on large computational domains (physically, hundreds of microns across) for long times. Comparison of our results with available experimental data shows very good qualitative agreement. We also use our model to investigate the influence of the so-called "anchoring" conditions at the substrate (the preferred orientation of the NLC molecules), and show that it can affect the instability development quite strongly under certain conditions.

Linda Cummings

New Jersey Institute of Technology
linda.cummings@njit.edu

Lou Kondic

Department of Mathematical Sciences, NJIT
University Heights, Newark, NJ 07102
kondic@njit.edu

Michael Lam

Coastal & Hydraulics Laboratory
US Army Engineer Research & Development Center
michaelangelo.yh.lam@gmail.com

MS342

On Some Novel Instabilities in Viscous Multilayer Shear Flows

Shear flows that are stable (e.g. Couette flow or Poiseuille flow at moderate Reynolds numbers) can become unstable to an interfacial mode when multiple immiscible layers are present. Such instabilities are typically long wave and were discovered about 50 years ago by Yih. The rule of thumb is that in simple shear flows instability arises if the more

viscous layer is also the thinner one, the so-called thin layer effect. Motivated by experiments and molecular dynamics simulations, we consider fluid systems that allow for slip at fluid-fluid interfaces (slip at fluid-solid surfaces is orders of magnitude smaller and is ignored). The basic states are viscous but support a velocity jump across interfaces. We have a viscous analogue of the Kelvin-Helmholtz problem seen in inviscid flows, and the natural question is to analyse such slip models to gain an understanding of their well-posedness and map out their stability characteristics. I will address two aspects: (i) A linear theory of the general problem and a semi-analytical demonstration of short-wave instabilities when the slip is asymptotically large; (ii) the development of a thin-layer long-wave nonlinear model and a discussion of its stability characteristics. It will be shown that for viscosity ratios that render the no-slip flow stable, even a small amount of slip can cause short wave instabilities of the Turing type. Nonlinear computations will be reported.

Demetrios Papageorgiou

Department of Mathematics
Imperial College London, UK
d.papageorgiou@imperial.ac.uk

MS342

Laminar Drag Reduction in Surfactant-Contaminated Superhydrophobic Channels

Although superhydrophobic surfaces (SHSs) show promise for drag reduction (DR) applications, their performance can be compromised by traces of surfactant that accumulate at liquid-gas interfaces, generating Marangoni stresses that increase drag. This question is addressed for a three-dimensional laminar flow in a plane periodic channel with SHSs along both walls, in the presence of soluble surfactant. We consider the regime in which bulk diffusion is sufficiently strong for concentration gradients normal to the SHSs to be small. Seeking solutions that are periodic in the streamwise and spanwise directions, and exploiting a long-wave theory that accounts for rapid transverse Marangoni-driven flow and shear-dispersion effects, we thereby reduce this high-dimensional problem to a one-dimensional model for the surfactant distribution. The system exhibits multiple regimes where asymptotic solutions can be constructed, which compare favourably with numerics. Some are characterised by advection and diffusion-dominated processes, where the liquid-gas interface exhibits shear-free behaviour and the DR is at its maximum. In contrast, others are dominated by Marangoni effects, where the liquid-gas interface exhibits no-slip behaviour and the DR vanishes. This analysis provides a guide for designing surfactant-contaminated SHSs to maximise the DR for applications.

Samuel Tomlinson

University of Manchester
samuel.tomlinson@manchester.ac.uk

Frédéric Gibou, Paolo Luzzatto-Fegiz

University of California, Santa Barbara
fgibou@ucsb.edu, fegiz@engineering.ucsb.edu

Fernando Temprano-Coleto

Princeton University
ftempranocoleto@princeton.edu

Oliver E. Jensen

University of Manchester, United Kingdom
Oliver.Jensen@manchester.ac.uk

Julien Landel
University of Manchester
julien.landel@manchester.ac.uk

MS343

Discarding the Magnitude of Client-Server Messages for Privacy-Preserving Federated Learning

There is a dearth of convergence results for differentially private federated learning (FL) with non-Lipschitz objective functions (i.e., when gradient norms are not bounded). The primary reason for this is that the clipping operation (i.e., projection onto an l_2 ball of a fixed radius called the clipping threshold) for bounding the sensitivity of the average update to each client's update introduces bias depending on the clipping threshold and the number of local steps in FL, and analyzing this is not easy. For Lipschitz functions, the Lipschitz constant serves as a trivial clipping threshold with zero bias. However, Lipschitzness does not hold in many practical settings; moreover, verifying it and computing the Lipschitz constant is hard. Thus, the choice of the clipping threshold is non-trivial and requires a lot of tuning in practice. In this work, we provide the first convergence result for private FL with clipping on smooth convex objectives for a general clipping threshold—without assuming Lipschitzness. We also look at a simpler alternative to clipping (for bounding sensitivity) which is normalization—where we use only a scaled version of the unit vector along the client updates, completely discarding the magnitude information. The resulting normalization-based private FL algorithm is theoretically shown to have better convergence than its clipping-based counterpart on smooth convex functions.

Abolfazl Hashemi
Purdue University
abolfazl@purdue.edu

MS343

Verification of the Integration Process and Aggregation of Health Data from Portable Sensors

Wearable devices are increasingly useful for the personalization of medical treatments. The manufacturers of these devices apply algorithms for aggregation and dimensionality reduction of the raw data collected to show the client a more user-friendly version that is easier to interpret. However, neither the raw data nor the details of the aggregation process are provided by the manufacturer, so in most cases they function as a black box for the user. Given the relevance of the information presented to the user, its application and the consequences of its interpretation (behavior modification, medication administration, etc.), it is essential to verify the algorithms used in the process to guarantee that the information presented really corresponds to the information collected by the sensors. In this work we present a suitable system for the verification of aggregated data from personal activity monitoring sensors. The system includes a parsing algorithm that does the data structure and relates it to the output. The effectiveness of the algorithm has been tested with real data over a period of two years and both for daytime activity and for monitoring sleep quality. The algorithm is perfectly scalable to be used in any device, so the computer system presented can be useful for the future computer audit of this type of process.

Victoria López
CUNEF University

Madrid
victoria.lopez@cunef.edu

Carlos Díez, Pavel Llamocca
CUNEF University
carlos.diez@cunef.edu, pavel.llamocca@cunef.edu

MS343

Application of Privacy Preserving Federated Learning in Biomedical Applications Lessons Learned

AI/ML models are known to be vulnerable to dataset shift and under specification because of the inability of current deep learning methods to learn the casual structure. This problem manifests when a model is deployed in a real test domain, where even simple changes in demographics or image formats could lead to unexpected poor performance, straining credibility. The solution for this is to train deep learning models on as many as real world datasets as possible. But access to biomedical datasets is governed by complex data usage agreements, time-consuming IRBs. One possible solution is to send models to data and develop frameworks to perform secure federated learning with privacy guarantees. In this talk I will present the Argonne Privacy Preserving Federated Learning Framework (APPFL) and adopting the framework to challenges in biomedicine.

Ravi Madduri
Argonne National Laboratory
madduri@anl.gov

MS344

New Stories About SZ Lossy Compression for Scientific Datasets

Today's scientific simulations and advanced instruments are producing extremely large amount of data everyday, introducing significant burdens in data storage and transferring at runtime. SZ (the R&D100 2021 award winner) is an efficient error-bounded lossy compressor developed for scientific datasets. SZ offers different types of error controls such as absolute error bound, relative error bound and peak signal to noise ratio (PSNR). SZ supports diverse execution environments including CPU, openMP and GPU. SZ has been integrated in multiple I/O libraries such as HDF5 and ADIOS, as well as various scientific packages such as cosmology simulation package. SZ is a modular composable framework allowing users to customize different compression pipelines according to specific requirement or data characteristics. In this presentation, I will present the latest research work and software development in the regard of the SZ compression framework, especially about diverse fresh compression methods which can adapt to a wide range of use-cases (e.g., ultra high speed, high compression ratio or high fidelity on specific quality metrics) in practice.

Sheng Di
Argonne National Laboratory
sdi1@anl.gov

Franck Cappello
ANL

cappello@mcs.anl.gov

MS344

QuadConv: Quadrature-Based Convolutions with Applications to Non-Uniform Data Compression

We present a new convolution layer for deep learning architectures which we call QuadConv — an approximation to continuous convolution via quadrature. Our operator is developed explicitly for use on non-uniform data, and accomplishes this by learning a continuous kernel which facilitates evaluation at any point in the domain. When compressing partial differential equation simulation data, we show that a QuadConv-based autoencoder can match the performance of standard discrete convolutions on uniform grid data (i.e., autoencoding CNNs) and maintain this accuracy on non-uniform data.

Kevin Doherty, Cooper Simpson
University of Colorado, Boulder
kevin.doherty@colorado.edu,
cooper.simpson@colorado.edu

Stephen Becker
University of Colorado at Boulder
stephen.becker@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS344

Lossy Scientific Data Compression with SPERR

As the need for data reduction in HPC continues to grow, this paper introduces a new and highly effective compressor to help achieve this goal, which we name as SPERR. SPERR is built on top of an advanced wavelet compression algorithm, and enhances it with abilities to bound user-prescribed point-wise error. Evaluations show that among popular lossy scientific data compressors, SPERR achieves one of the best rate-distortion curves over a wide range of error tolerance levels at the cost of increased computational load.

Samuel Li
National Center for Atmospheric Research
shaomeng@ucar.edu

MS344

Convincing Climate Scientists that Compression is a Good Idea; and Making Sure it is!

Climate models typically produce enormous amounts of output data, which storage capacities have not kept up with. That has forced climate scientists to make hard choices about which variables to save, data output frequency, simulation lengths, etc, all of which can negatively impact science objectives. Therefore, we have been investigating lossy data compression techniques as a means of reducing data storage. As with any data reduction approach, we must exercise extreme care when applying lossy compression to climate output data to avoid introducing artifacts in the data that could affect scientific conclusions. Our focus has been on better understanding the effects of lossy compression on spatio-temporal climate data and on gaining user acceptance via careful analysis and testing.

We will describe the challenges and concerns that we have encountered and will discuss climate-specific metrics and tools to enable scientists to evaluate the effects of lossy compression and facilitate optimizing compression for each variable. In particular, we will present our Large Data Comparison for Python (LDCCPy) package for visualizing and computing statistics on differences between multiple datasets. We will also demonstrate the usefulness of an alternative to the popular SSIM that we developed, called the Data SSIM (DSSIM), that can be applied directly to the floating-point data in the context of evaluating differences due to lossy compression on large volumes of simulation data.

Alexander Pinard, Dorit Hammerling
Colorado School of Mines
apinard@mines.edu, hammerling@mines.edu

Allison H. Baker
National Center for Atmospheric Research
abaker@ucar.edu

MS345

Accelerating Hydrodynamics Simulations with Reduced Order Models

Although many model reduction schemes have been developed to reduce the computational cost of simulations while minimizing the error introduced in the reduction process, there are challenges especially in nonlinear advection-dominated problems such as sharp gradients, moving shock fronts, and turbulence, which hinder those model reduction schemes from being practical. In this talk, we consider the time-dependent Euler equations of compressible gas dynamics in a moving Lagrangian frame, and demonstrate some novel techniques to address the challenges imposed by the advection-dominated solutions. Lagrangian hydrodynamics is formulated as a nonlinear problem, which requires a proper hyper-reduction technique to reduce the complexity due to the nonlinear terms in projection-based reduced order models. We will present various hyper-reduction techniques, including over-sampling DEIM and S-OPT. For reducing the Kolmogorov n -width of the solution manifold, we employ a manifold decomposition approach to build local reduced order models in submanifolds. We also present both a posteriori and a priori error bounds associated with the projection-based reduced order model, and the performance comparison of the reduced order modeling approaches in terms of accuracy and speed-up with respect to the corresponding full order model for several numerical examples, namely Sedov blast, Gresho vortices, Taylor-Green vortices, triple-point problem and Rayleigh-Taylor instability problem.

Siu Wun Cheung
Lawrence Livermore National Laboratory
cheung26@llnl.gov

MS345

Reduced Basis Method for Stabilized Smagorinsky Models

In this work we present the numerical analysis of some Reduced Basis Smagorinsky models. In particular, we present a Local Projection Stabilization (LPS) VMS-Smagorinsky model and the construction of the reduced basis spaces via the Greedy Algorithm. We construct the reduced velocity space by two different strategies, by considering or not the enrichment of the reduced velocity space with the so-called

inner pressure *supremizer*. We present the development of an *a posteriori* error estimator for the snapshot selection through a Greedy algorithm, based on the Brezzi-Rappaz-Raviart (BRR) theory. Moreover, the Empirical Interpolation Method (EIM) is considered for the approximation of the non-linear terms. Finally, we present some numerical tests in which we show an speedup on the computation of the reduced basis problem with the LPS pressure stabilization, with respect to the method using pressure supremizers.

Enrique Delgado
Differential Equations and Numerical Analysis
Universidad de Sevilla
edelgado1@us.es

Tomás Chacón Rebollo
Universidad de Sevilla
Spain
chacon@us.es

Macarena Gómez Mármol
Differential Equations and Numerical Analysis
Universidad de Sevilla
macarena@us.es

MS345

Dynamics of Weather and Climate on Reduced Manifolds

While operational weather forecasting models have been dramatically improving over the past several decades, the prediction of weather extremes, within actionable time scales, remains elusive. In this work, we propose a new framework, based on concept derived from nonlinear dynamical system theory, that can complement and possibly improve the forecasting skills of existing operational weather models.

Chenyu Dong
National University of Singapore
e0974131@u.nus.edu

Davide Faranda
LSCE, CEA Saclay l'Orme des Merisiers, CNRS
davide.faranda@lsce.ipsl.fr

Tommaso Alberti
INAF-Istituto di Astrofisica e Planetologia Spaziali, Italy
tommaso.alberti@inaf.it

Adriano Gualandi
INGV-Istituto Nazionale di Geofisica e Vulcanologia, Italy
adriano.gualandi@ingv.it

Gianmarco Mengaldo
National University of Singapore
College of Design and Engineering
mpegin@nus.edu.sg

MS345

Rom for Large-Scale Modelling of Urban Air Pollution

Urban air pollution is a major global challenge responsible for damages to climate, ecosystems, and health. A framework that combines direct measurements and computational modeling techniques is therefore an important

analytical tool, able to extract various insights from the collected statistics and open new strategies for mitigation policies. Since pollutant dispersion depends on daily weather conditions, CFD models with low time scales, repeated evaluation, and fine mesh discretization must be used. The former requirements translate into huge memory requirements, making it essential to use HPC facilities. However, the problem is suitable for the employment of Reduced Order Models (ROMs) to achieve fast converged solutions with limited loss of accuracy. For this reason, the present work consists of the exploration of a Proper Orthogonal Decomposition (POD) coupled with the Galerkin projection approach for the acceleration of external environmental flow problems. The evolution of the pollutant is described through the transport equation, where the convective field is given by the solution to the RANS equations, while the source term consists of an empirical time series. Here we propose a data-driven approach based on a POD-NN reconstruction of the flux field, which is used to recover in a non-intrusive fashion the reduced-order operators required for the online evaluation. Our framework is validated on different scales and using real traffic measurements.

Moaad Khamlich
SISSA
mkhamlic@sissa.it

MS347

Striking the Right Chord (Part II): Learning to Hear

In this follow-up to the presentation "Striking the Right Chord (Part I): What note is this?" we discuss how a novel data-consistent approach is modified to "hear" the properties associated with the tension of a drum. Particular attention is paid to applying feature extraction techniques commonly associated with machine learning to construct the parameter-to-observable maps required in this data-consistent approach. In the context of the application discussed here, we are identifying the most informative "chords" (a harmonious sound made up of individual "notes") to listen to from the striking of a drum that help us determine how tightly stretched the drum is over its frame. Numerical results demonstrate our ability to reconstruct the distributions of different families of tensions through the hearing of the correct chords.

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

MS347

Striking the Right Chord (Part I): What Note Is This?

How does the sound of the drum depend on its material properties? We model the drum as an elastic membrane and compute the associated variable diffusion eigenproblem, where the diffusion can be interpreted for instance as the tension of the drum. We choose as the shape of a drum a disk with a cut sector. In this case the exact eigenvalues for the constant coefficient problem are known, and the non-symmetric domain guarantees that there are no clusters of eigenvalues at the low end of the spectrum. This is an important property of the problem, since as the diffusion coefficient is perturbed, not only do the eigenvalues change but the order of the modes can change too. This phenomenon is referred to as mixing of modes. The nu-

merical experiments are computed with hp-FEM where it is ensured that the discretisation is such that in the reference case all first 60 eigenpairs are in correct order. Thus, when in the perturbed cases mixing is observed, we can with high confidence say that it is due to the variation of the tension of the drum. Using this setup we can sample the coefficients from given distributions and observe the responses.

Harri H. Hakula
Aalto University
Harri.Hakula@aalto.fi

MS347

Sequentially Optimized Projections in X-Ray Imaging

We review the framework of Bayesian experimental design and consider how to apply it to selecting optimal projection geometries in X-ray tomography assuming the prior and the additive noise are Gaussian. The introduced greedy exhaustive optimization algorithm proceeds sequentially, with the posterior distribution corresponding to the previous projections serving as the prior for determining the design parameters, i.e. the imaging angle and the lateral position of the source-receiver pair, for the next one. Both A- and D-optimality are considered, with emphasis on efficient evaluation of the corresponding objective functions. Two-dimensional numerical experiments demonstrate the functionality of the approach.

Nuutti Hyvönen
Aalto University
nuutti.hyvonen@aalto.fi

MS347

Edge-Promoting Sequential Experimental Design for X-Ray Imaging

We consider sequential edge-promoting Bayesian experimental design in X-ray imaging. The process of computing a total variation type reconstruction of the absorption inside the imaged body via lagged diffusivity iteration is interpreted in the Bayesian framework. Assuming a Gaussian additive noise model, this leads to an approximate Gaussian posterior with a covariance structure that contains information on the location of edges in the posterior mean. The next projection geometry is then chosen through A- or D-optimal Bayesian design. Two- and three-dimensional numerical examples based on simulated data demonstrate the functionality of the introduced approach.

Juha-Pekka Puska
Aalto University
juha-pekka.puska@aalto.fi

MS348

Computing Optimal Designs for the Approximate Synthesis of Four-Bar Mechanisms

Exact synthesis of robotic mechanisms is limited by the number of design positions, which Wampler, Sommese, and Morgan found in 1992 that there are 1,442 distinct coupler curves to exactly trace a maximum of nine design positions. Approximate synthesis has become an alternative where one instead specifies a larger number of positions and finds optimal mechanism designs. Such approximate synthesis problems can be described by solving an unconstrained

optimization problem where the objective is a polynomial. Employing homotopy continuation methods allows one to compute all critical points. This talk will describe the techniques used in computing the total solution set to an approximate synthesis formulation independent of the number of design positions and present an example associated with solving real design solutions. This is joint work with Aravind Baskar, Jonathan Hauenstein, and Mark Plecnik. Authors: Caroline Hills, chills1@nd.edu, University of Notre Dame; Mark Plecnik, plecnikmark@nd.edu, University of Notre Dame; Jonathan Hauenstein, hauenstein@nd.edu, University of Notre Dame; Aravind Baskar, abaskar@nd.edu, University of Notre Dame

Caroline Hills
University of Notre Dame
chills1@nd.edu

MS348

Normalizing Flows based Mutual Information Estimation

Mutual Information (MI), as a measure of mutual dependence on random quantities without specific modelling assumptions, plays an important role in many scientific disciplines. However, estimating mutual information numerically from high-dimensional data remains a difficult problem. Classic non-parametric estimators of entropy and mutual information depend either on density estimation and Monte Carlo integration, or on the computation of k nearest neighbours (kNN). The first type of algorithm suffers from the curse of dimensionality, while kNN methods lack accuracy in the presence of correlation. Non-parametric estimators are usually fast and accurate without further assumptions, but they do not scale well for small-sized samples in high-dimensional spaces. This impediment has motivated the parametric estimators, which take advantage of the expressive power of deep learning architectures in parameterizing the space of functions, to find the variational bounds for mutual information. Some recent research include MINE (Belghazi et al., 2018; Song and Ermon, 2019), Contrastive Predictive Coding (van den Oord et al., 2018), DEMI (Liao et al., 2020) and Normalizing Flows based entropy estimators (Ao and Li, 2021). As shown by Song and Ermon (2019), these models based on discriminative learning can introduce high variance in the estimation, while generative models such as entropy estimators based on normalizing flows need to learn three distributions to compute mutual information. We propose a principled mutual information estimator based on a generalization of normalizing flows, a deep generative model that has gained considerable popularity in recent years. This estimator not only estimates the marginal and joint entropy in a systematic way, but also uses an autoregressive structure to obtain better performance in estimating mutual information. Empirical results demonstrate that our proposed estimator exhibits improved bias-variance trade-offs on standard benchmark tasks. Authors: Haoran Ni, Mathematics Institute, haoran.ni@warwick.ac.uk; Martin Lotz, Mathematics Institute, martin.lotz@warwick.ac.uk

Haoran Ni
University of Warwick, U.K.
Haoran.ni@warwick.ac.uk

MS348

Time-Marching in Physics Informed Neural Net-

works in the Context of Shear Waves in Soft Solids

Millions of people suffer from traumatic brain injury (TBI) every year. Recently, the shear shock was observed in brain tissue which may be the primary cause of TBI [1]. Shear shock wave simulation in complex geometry like the human head can be computationally exhaustive. In 2019, Raissi et al. proposed the concept of Physics Informed Neural Networks (PINNs) [2]. We intend to develop an artificial neural network-based numerical method for the propagation of shear shock waves in the brain. In this work, we developed a fully connected dense multi-layer perceptron to describe the propagation of shear waves in soft solids formulated as a system of first-order PDEs (Partial Differential Equation) based on the work of Tripathi et al. [1]. The PINN (Physics Informed Neural Networks) was trained without any labeled data using a novel time-marching algorithm designed to incorporate the causality of the physical phenomenon, i.e., wave propagation. The non-dimensionalized physical system of equations was formulated to facilitate efficient training of PINN [3]. The developed data-free surrogate neural network model is trained sequentially using a novel training algorithm that enforces the boundary condition and minimizes the summation of boundary loss and physics-based loss. The data-free PINN was constrained by 1) initial and boundary conditions, 2) uniformly distributed internal collocation points, and 3) the non-dimensionalized system of PDEs. PINNs with two input neurons containing spatial and temporal variables were used to approximate PDE (Partial Differential Equation) solutions. The weights and biases of the neural network were initialized using Glorot normal method. The optimal set of hyperparameters (activation function, optimizer, etc.) and the network architecture were chosen through multiple iterations performed systematically. The data-free PINNs solution for the system of linear wave equation was compared with the analytical solution. The PINNs solution for the system of linear shear wave equation has a relative error of 1% as compared with the analytical solution. The results have demonstrated the merit of the proposed data-free PINNs. It motivates the natural extension of the PINN to a nonlinear system of first-order PDEs as formulated for shear waves in soft solids [1]. Authors: Vikrant Pratap, University of Galway, v.pratap1@nuigalway.ie; Dr. Bharat B Tripathi, University of Galway

Vikrant Pratap
University of Galway
v.pratap1@nuigalway.ie

MS348**Invariant Domain Preserving and Exactly Conservative Approximation of the Lagrangian Hydrodynamics Equations**

The Euler equations serve as the cornerstone of computational fluid dynamics and can be described in two different reference frames, either by tracking the motion of a specific fluid parcel or viewing the motion of the fluid in a fixed frame. These perspectives are called the Lagrangian and Eulerian, respectively. In this talk, we will briefly discuss the applications and benefits of each perspective. Then we present a first order finite element method for approximating the compressible Euler equations in Lagrangian coordinates. This method is explicit in time, uses a combination of continuous and discontinuous finite elements in space, is exactly mass conservative (in the Lagrangian sense), and is invariant-domain preserving. Authors: Jean Luc Guermond, Texas A&M University, guermond@tamu.edu; Bojan Popov, Texas A&M University, popov@tamu.edu; Ben-

nett Clayton, Texas A&M University, bgclayto@tamu.edu; Laura Saavedra, ETSIAE-UPM, School of Aeronautics

Madison Sheridan
Texas A&M University
sheridanm@math.tamu.edu

MS348**Developing an Antimatter Gravity Interferometer**

The assumption that the effects of gravity on antimatter and matter are equivalent has permeated throughout almost all of modern physical theory and experiment. However, no direct observation of this effect from gravity has been made on a particle in freefall. A muonium beam diffracting through a series of gratings has proven to be a suitable method for recording such freefall. Simulating this with the best current understanding of diffraction and interferometry is vital in determining this antimatter-gravity relationship, since a physical construction requires a picometer-precise atom interferometer and muonium beam. The development of these simulations has both demonstrated the relative feasibility of experimentation and brought to question the viability in applying certain physical modeling and simulation methods.

Jacob H. Thomas
Illinois Institute of Technology
jthomas22@hawk.iit.edu

MS349**Filtering in Non-Intrusive Data-Driven Reduced-Order Modeling of Large-Scale Systems**

We present a method for enhancing data-driven reduced-order modeling with a preprocessing step in which the training data are filtered prior to training the reduced model. Filtering the data prior to training has a number of benefits for data-driven modeling: it attenuates (or even eliminates) wavenumber or frequency content that would otherwise be difficult or impossible to capture with the reduced model, it smoothens discontinuities in the data that would be difficult to capture in a low-dimensional representation, and it reduces noise in the data. This makes the reduced modeling learning task numerically better conditioned, less sensitive to numerical errors in the training data, and less prone to overfitting when the amount of training data is limited. We first illustrate the effects of filtering in one-dimensional advection and inviscid Burgers' equations. We then consider large-scale rotating detonation rocket engine simulations with millions of spatial degrees of freedom for which only a few hundred down-sampled training snapshots are available. A reduced-order model is derived from these snapshots using operator inference. Our results indicate the potential benefits of filtering to reduce overfitting, which is particularly important for complex physical systems where the amount of training data is limited.

Distribution Statement A: Approved for Public Release; Distribution is Unlimited. PA# AFRL-2022-4312

Ionut-Gabriel Farcas
The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences
ionut.farcas@austin.utexas.edu

Rayomand Gundevia
 Jacobs Engineering Group, Inc.
 Edwards Air Force Base
 rayomand.gundevia.ctr@afml.af.mil

Ramakanth Munipalli
 Edwards Air Force Base
 Air Force Research Laboratory
 ramakanth.munipalli@us.af.mil

Karen E. Willcox
 UT Austin
 kwillcox@oden.utexas.edu

MS349

Extensions and Open-Source Algorithms for Dynamic Mode Decomposition for Data-Driven Modeling

The dynamic mode decomposition (DMD) has grown to become one of the leading methods of data-driven modeling of dynamical systems across a multitude of scientific disciplines. PyDMD is a Python package that provides the tools necessary for executing the DMD regression framework. However, since the package's initial release, several major innovations and extensions to DMD have arisen, such as Optimized DMD, Physics-informed DMD, and Randomized DMD to name a few. Thus in order to enhance users' ability to manage and examine practical, real-world datasets, we have extended the functionality of PyDMD to include several of these recent DMD developments. In this talk, an overview is given for a variety of modern DMD innovations that will debut in this major update to PyDMD. I will discuss the theoretical underpinnings of these methods in addition to providing practical tips regarding DMD usage. I will then demonstrate how one may utilize the PyDMD package in order to model and analyze real-world data.

Sara M. Ichinaga
 Department of Applied Mathematics
 University of Washington
 sarami7@uw.edu

Nicola Demo
 SISSA, Mathematics Area, mathLab
 International School for Advanced Studies
 ndemo@sissa.it

Marco Tezzele
 Oden Institute for Computational Engineering and Sciences
 The University of Texas at Austin
 marco.tezzele@austin.utexas.edu

Gianluigi Rozza
 SISSA, International School for Advanced Studies,
 Trieste, Italy
 grozza@sissa.it

J. Nathan Kutz
 University of Washington, Seattle
 Dept of Applied Mathematics
 kutz@uw.edu

Steven Brunton
 University of Washington

sbrunton@uw.edu

MS349

The Shifted Operator Inference Method for Learning Solar Wind Models

Solar wind conditions are predominantly predicted via three-dimensional numerical magnetohydrodynamic (MHD) models. Despite their ability to produce highly accurate predictions, MHD models require computationally intensive high-dimensional simulations. This renders them inadequate for making time-sensitive predictions and for large-ensemble analysis required in uncertainty quantification. This paper presents a new data-driven reduced-order model (ROM) capability for forecasting heliospheric solar wind speeds. Traditional model reduction methods based on Galerkin projection have difficulties with advection-dominated systems—such as solar winds—since they require a large number of basis functions and can become unstable. A core contribution of this work addresses this challenge by extending the non-intrusive operator inference ROM framework to exploit the translational symmetries present in the solar wind caused by the Sun's rotation. The numerical results show that our method can adequately emulate the MHD simulations and outperforms a reduced-physics surrogate model, the Heliospheric Upwind Extrapolation model.

Boris Kramer
 University of California San Diego
 bmkramer@ucsd.edu

Opal Issan
 University of California, San Diego
 oissan@ucsd.edu

MS349

Model Reduction and Control of Dynamical Systems Based on Frequency Data

In this talk, a new unified framework for model reduction and control based on frequency response data will be presented that is particularly well-suited to deal with structured models. In this talk, our focus will be on port-Hamiltonian systems. In the context of model reduction, it is often desired to obtain a reduced-order model of the same structure to benefit from the same structural properties of the original model such as stability or passivity. In our framework this goal is achieved by parametrizing a reduced-order model of prescribed structure and then optimizing the parameters with respect to a specially designed objective functional that typically results in reduced-order models with small (and often, close-to-optimal) \mathcal{H}_∞ -error. Likewise, in the context of designing \mathcal{H}_∞ -controllers for port-Hamiltonian plant models, one can parametrize a low-order port-Hamiltonian controller that is in negative feedback interconnection with the plant. This construction automatically leads to stable closed-loop systems and the \mathcal{H}_∞ -performance can be optimized using a similar objective functional as in model reduction.

Matthias Voigt
 UniDistance Suisse
 matthias.voigt@fernuni.ch

Paul Schwerdtner
 TU Berlin

schwerdt@math.tu-berlin.de

MS350

Sensitivity Analysis of the Information Gain in Infinite-Dimensional Bayesian Linear Inverse Problems

We consider sensitivity analysis of Bayesian linear inverse problems with respect to modeling uncertainties. To this end, we consider sensitivity analysis of the information gain, as measured by the KullbackLeibler divergence from the posterior to the prior. This choice provides a principled approach that leverages key structures within the Bayesian inverse problem. Also, the information gain admits a closed-form expression in the case of linear Gaussian inverse problems. The derivatives of the information gain are extremely challenging to compute. To address this challenge, we present accurate and efficient methods that combine eigenvalue sensitivities and hyper-differential sensitivity analysis that take advantage of adjoint based gradient and Hessian computation. This results in a computational approach whose cost, in number of PDE solves, does not grow upon mesh refinement. These results are presented in an application-driven model problem, considering a simplified earthquake model to infer fault slip from surface measurements.

Abhijit Chowdhary

North Carolina State University
achowdh2@ncsu.edu

Shanyin Tong

Columbia University
st3503@columbia.edu

Georg Stadler

Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Alen Alexanderian

North Carolina State University, U.S.
alexanderian@ncsu.edu

MS350

Variational Bayesian Optimal Experimental Design with Normalizing Flows

Optimal experimental design (OED) seeks experimental design conditions that lead to the greatest expected information gain (EIG). The EIG is typically approximated through a double-nested Monte Carlo estimators, requiring $O(N^2)$ forward model evaluations. This becomes prohibitive when working with computationally intensive models commonly encountered in studies of complex and physical systems. Recently, Foster et al. proposed a variational EIG approximation technique that requires only $O(N)$ forward model evaluations, by maximizing a stochastic lower bound of the EIG that emerges from replacing the true posterior with a Gaussian variational distribution. However, Gaussian distributions are unable to capture non-Gaussian posterior behavior such as multi-modality and skewness, and it can lead to large bias (loose bound) in estimating the EIG. We improve the variational OED by deploying normalizing flows that transform standard normal latent variables to general non-Gaussian distributions through expressive neural networks, leading to lower bias estimators as well as an ability to produce accurate non-

Gaussian posteriors. We illustrate this method on benchmark problems as well an application to design optimal observing times for inferring parameters of an aphid population model.

Jiayuan Dong, Christian Jacobsen

University of Michigan
jiayuand@umich.edu, csjacobs@umich.edu

Karthik Duraisamy

University of Michigan Ann Arbor
kdur@umich.edu

Xun Huan

University of Michigan
xhuan@umich.edu

MS350

Differential Methods for Assessing Sensitivity in Biological Models

Differential sensitivity analysis is indispensable in fitting parameters, understanding uncertainty, and forecasting the results of models. Although there are many methods currently available for performing differential sensitivity analysis, it can be difficult to determine which method is best suited for a particular model. We explain a variety of differential sensitivity methods and assess their value in some typical biological models. First, we explain the mathematical basis for three numerical methods: adjoint sensitivity analysis, complex perturbation sensitivity analysis, and forward mode sensitivity analysis. We then carry out four instructive case studies. (a) The CARRGO model for tumor-immune interaction highlights the additional information that differential sensitivity analysis provides beyond traditional naive sensitivity methods, (b) the deterministic SIR model demonstrates the value of using second-order sensitivity in refining model predictions, (c) the stochastic SIR model shows how differential sensitivity can be attacked in stochastic modeling, and (d) a discrete birth-death-migration model illustrates how the complex perturbation method of differential sensitivity can be generalized to a broader range of biological models. Finally, we compare speed, accuracy, and ease of use. We find that forward mode automatic differentiation has the quickest computational time, while the complex perturbation method is the simplest to implement and the most generalizable.

Rachel Mester

University of California Los Angeles
rmester@ucla.edu

MS350

A Bayesian Framework for Coupling Optimal Experimental Design and Optimal Control

Optimal experimental design (OED) allows us to determine, a-priori, how to collect the most informative experimental data. Traditional OED approaches focus on optimizing the solution to the Bayesian inverse problem, which is often an intermediate step to the true objectives of making accurate model predictions and determining optimal control policies. In this work, we provide a novel formulation of a Bayesian OED problem that couples the goal-oriented experimental design and optimal control problems, thus minimizing uncertainty in the control objective directly. Through a numerical example, we illustrate how the resulting control-oriented OED approach leads to a better allocation of experimental resources that decreases un-

certainty regarding the optimal control policy in comparison to classical OED strategies.

Rebekah White

Sandia National Laboratories
rebwhit@sandia.gov

Alen Alexanderian

North Carolina State University, U.S.
alexanderian@ncsu.edu

Arvind Saibaba

North Carolina State University
saibab@ncsu.edu

Bart G. Van Bloemen Waanders

Sandia National Laboratories
bartv@sandia.gov

MS351

Properties and Transitions of Mesoscale Convective Organisation During EUREC4A Using Unsupervised Learning

The representation of shallow tradewind cumulus clouds in climate models accounts for the majority of inter-model spread in climate projections, highlighting an urgent need to understand these clouds better. In particular, their spatial organisation appears to cause a strong impact of their radiative properties and dynamical evolution. The precise mechanisms driving different forms of convective organisation which arise both in nature and in simulations are, however, currently unknown. Using unsupervised learning for identifying regimes of convective organisation in the tropical Atlantic, we will show results from analysing: a) what the radiative properties of different forms of organisation are, b) what atmospheric characteristics coincide with different forms of organisation and c) what transitions occur when following air-masses along Lagrangian trajectories. Specifically, we find: a) net radiation changes significantly between different forms of organisation, b) agreement with previous studies on the importance of boundary layer wind-speed and to some degree atmospheric stability, and c) we are able to succinctly capture what transitions occur between regimes.

Leif Denby

University of Leeds
l.c.denby@leeds.ac.uk

Franziska Glassmeier

TU Delft
f.glassmeier@tudelft.nl

MS351

BOSS: A Probabilistic Parameterization Framework to Improve the Representation of Microphysical Process Rates in Weather and Climate Models

Weather and climate models represent cloud microphysics using computationally-efficient bulk schemes that predict only a few variables such as drop number and mass concentrations. Current bulk schemes exhibit deficiencies that are due in part to their simplified representation of a complex natural state, and in part due to a fundamental lack of understanding of microphysical processes. These schemes also employ numerous ad-hoc assumptions, parameter choices, and model structural as-

sumptions. To synthesize information from detailed process models and observations for improving bulk schemes, we have developed a novel probabilistic parameterization framework called the Bayesian Observationally-constrained Statistical-physical Scheme (BOSS). BOSS combines existing, though limited, process level microphysical knowledge with flexible process rate formulations and parameters trained to observations and data from detailed process models through Bayesian inference. The approach is flexible and can be used to examine structural choices in bulk schemes, particularly the functional form of process rates and the number of predicted cloud microphysical variables. This allows for systematic quantification of both parametric and structural uncertainty, which is not possible using traditional schemes. Results using an MCMC sampler with BOSS to constrain process rates and parameters with synthetic "observations" generated by a detailed bin microphysical model will be presented.

Hugh Morrison

National Center for Atmospheric Research
morrison@ucar.edu

Marcus Van Lier-Walqui

Center for Climate Systems Research
Columbia University/NASA GISS
mv2525@columbia.edu

Sean P. Santos

Columbia University
sean.santos@pnml.gov

Karly Reimel

Pennsylvania State University
karlyreimel@gmail.com

Matthew Kumjian

Penn State
kumjian@psu.edu

Olivier Prat

NORTH CAROLINA STATE UNIVERSITY
opprat@ncsu.edu

MS351

When Will It Start to Rain? Producing Accurate Data-Driven Approximations for Microphysical Systems with Transient Instabilities

Collision-coalescence between cloud drops is typically modeled using the quasi-stochastic collection equation (SCE), a nonlinear partial integro-differential equation. The long-term behavior of a system governed by the SCE is predictable, but trajectories can have a significant *temporary* sensitivity to initial conditions. In atmosphere models, this corresponds to uncertainty about when the system starts to produce a significant number of "rain-sized" drops. Such models can use "bin" microphysics schemes, which discretize the drop size distribution, to solve equations such as the SCE numerically. However, most models use cheaper "bulk" microphysics schemes, which describe the population of drops using only a few variables, e.g. total mass and number. These bulk schemes are reduced order models for microphysical systems, but the strong nonlinearity of the SCE makes conventional reduced order modeling techniques difficult to apply. We present a new method, Jacobian Estimation of Functional Error (JEFE), which uses an adjoint model to measure the predictability of bin model outputs from a small set of prognostic variables.

This allows us to identify whether the given variables provide enough information to capture the system's dynamics, a prerequisite for an effective bulk scheme. Having chosen a set of variables, we can then produce a bulk scheme using data-driven methods, such as the Bayesian Observationally constrained Statistical-physical Scheme (BOSS).

Sean P. Santos
Columbia University
sean.santos@pnml.gov

Hugh Morrison
National Center for Atmospheric Research
morrison@ucar.edu

Marcus Van Lier-Walqui
Center for Climate Systems Research
Columbia University/NASA GISS
mv2525@columbia.edu

Adele Igel
University of California - Davis
aigel@ucdavis.edu

MS351

Bayesian Tuning of Earth System Models in Global and Single-Column Simulations, Constrained by Satellite Observations and High Resolution Models

Earth System Models solve the equations of motion for geophysical fluid flow on the globe numerically at discrete points on a 3D grid that spans the Earth. Numerous processes occur on sub-grid scales and cannot be explicitly resolved by grid-scale variables. And yet, these processes strongly affect grid-scale fluid motion as well as thermodynamical and radiative properties. Examples of these sub-grid processes include the microphysics of how cloud and precipitation particles grow and interact. There is necessary approximation in the modeling of sub-grid scale effects on the grid-scale variables; additionally many of the sub-grid processes are uncertain at any scale. These approximations and uncertainties result in forecast errors. Recently, efforts have been proposed to systematically optimize parameters of models of sub-grid processes using observations. One such approach has been used with the NASA Goddard Institute for Space Studies ModelE climate model, informed by satellite observational data. The approach is Bayesian and employs a MCMC sampler in conjunction with machine learning emulators to estimate the probability density over 45 parameters for various sub-grid processes. We present this methodology, which uses a relatively small number (500) of one-year global climate simulations, and also discuss the potential benefit of incorporating computationally inexpensive single-column model simulations that focus on specific regimes of climate system behavior.

Marcus Van Lier-Walqui
Center for Climate Systems Research
Columbia University/NASA GISS
mv2525@columbia.edu

MS352

Error Analysis of Measure Transport Problems in Data Science

Measure transport problems have become prevalent in data science and machine learning, from generative adversarial

networks and normalizing flows to optimal transport. In applications we often solve these transport problems via a variational/optimization problem involving various parameterizations of transport maps. In this talk I will talk about a theoretical framework for connecting the parameterization error of transport to classic results from functional approximation theory in high-dimensions.

Hosseini Bamdad
University of Washington
bamdadh@uw.edu

MS352

Stein VI for Bayesian Neural Network Prediction of Helicopter Rotor Performance

In this work, we develop a novel framework for rapid prediction of helicopter rotorblade performance with quantified uncertainty—under various icing conditions by bringing together concepts in computational fluid dynamics, computational aeroacoustics, and Bayesian neural networks. We begin by using CFD and CAA to compute the acoustic signatures created by rotorblades under a range of iced blade geometries. Simultaneously, we extract aerodynamic performance coefficients, including lift, drag, and moment, of the respective icing geometries from the flow solutions. Using this dataset, we train a BNN to create a mapping from the iced rotor acoustic signatures to predict the corresponding performance coefficients with quantified epistemic uncertainty in the neural network weights. The creation of BNN involves solving a very high dimensional inference problem on all the network parameters. We tackle this challenge by employing the projected Stein variational gradient descent that identifies and leverages the low dimensional innate structures embedded in the highly temporally correlated acoustic signatures. We show these BNNs to accurately predict rotorblade performance coefficients together with uncertainty in the prediction, which is tremendously valuable to safety-critical pilot decision-making. Furthermore, these BNN models only need be trained offline, and can achieve extremely fast online predictions, making them suitable for real-time, in-flight use.

Jeremiah Hauth
University of Michigan
hauthj@umich.edu

Myles Morelli
Polytechnic University of Milan
mylescarlo.morelli@polimi.it

Alberto Guardone
Politecnico di Milano
Department of Aerospace Science and Technology
alberto.guardone@polimi.it

Xun Huan
University of Michigan
xhuan@umich.edu

Beckett Zhou
University of Bristol
beckett.zhou@bristol.ac.uk

MS352

Probabilistic Approaches to Transfer Learning for

Sparse and Noisy Data Environments

Machine learning (ML) models have thus far been applied to tasks and domains that, while impactful, have sufficient volume of data. For predictive tasks of national security relevance, ML models of great capacity are often needed to capture the complex underlying physics. Such models normally require an abundance of training data to exhibit sufficient predictive accuracy, which might not be available due to (1) excessive expense of computer simulations, (2) prohibitive experimental data acquisition cost, or (3) limited access to classified/sensitive data. To alleviate such difficulties, transfer learning (TL) may be used in which similar data from existing datasets or domains is used. We present a novel probabilistic TL framework to enhance the trust in ML models within noisy and sparse data settings. The framework will assess when it is worth applying TL, which ML model to use in TL, and how much knowledge is to be transferred. We rely on extensions of concepts and techniques from the fields of Bayesian inversion, sequential data assimilation, uncertainty quantification, and information theory. We provide insights through an application to polynomial-based surrogate model construction. We investigate that extent to which TL alleviates sparsity in training data that may jeopardize the reliability of such surrogates.

Mohammad Khalil
Sandia National Laboratories
mkhalil@sandia.gov

MS352

Encoding Spatial Priors with VAEs for Small Area Estimation

Gaussian processes (GPs), implemented through multivariate Gaussian distributions for a finite collection of data, are the most popular approach in small-area spatial statistical modelling. In this context, they are used to encode correlation structures over space and can generalize well in interpolation tasks. Despite their flexibility, off-the-shelf GPs present serious computational challenges which limit their scalability and practical usefulness in applied settings. Here, we propose a novel, deep generative modelling approach to tackle this challenge, termed PriorVAE: for a particular spatial setting, we approximate a class of GP priors through prior sampling and subsequent fitting of a variational autoencoder (VAE). Given a trained VAE, the resultant decoder allows spatial inference to become incredibly efficient due to the low dimensional, independently distributed latent Gaussian space representation of the VAE. Once trained, inference using the VAE decoder replaces the GP within a Bayesian sampling framework. This approach provides tractable and easy-to-implement means of approximately encoding spatial priors and facilitates efficient statistical inference. We demonstrate the utility of our VAE two-stage approach on Bayesian, small-area estimation tasks.

Elizaveta Semanova
University of Oxford
elizaveta.p.semanova@gmail.com

MS353

An Nonlinear Acceleration Method That Exploits Symmetry of Hessian

Nonlinear acceleration methods are powerful techniques to speed up fixed-point iterations. However, many accelera-

tion methods require storing a large number of previous iterates and this can become impractical if computational resources are limited. In this work, we propose a nonlinear Truncated Generalized Conjugate Residual method (nTGCR) whose goal is to exploit the symmetry of the Hessian to reduce memory usage. The proposed method can be interpreted as either an inexact Newton or a quasi-Newton method. We show that, with the help of global strategies like residual check techniques, nTGCR can converge globally for general nonlinear problems and that under mild conditions, nTGCR is able to achieve superlinear convergence. We further analyze the convergence of nTGCR in a stochastic setting. Numerical results demonstrate the superiority of nTGCR when compared with several other competitive baseline approaches on a few problems.

Huan He, Shifan Zhao
Emory University
huan_he@hms.harvard.edu, shifan.zhao@emory.edu

ziyuan tang
University of Minnesota
tang0389@umn.edu

Joyce Ho
Emory University
joyce.c.ho@emory.edu

Yousef Saad
University of Minnesota
saad@umn.edu

Yuanzhe Xi
Department of Mathematics
Emory University
yuanzhe.xi@emory.edu

MS353

Numerical Optimization Algorithm for Estimating a Conditioned Symmetric Positive Definite Matrix Under Constraints

We present RCO (Regularized Cholesky Optimization): a numerical algorithm for finding a symmetric Positive Definite (PD) $n \times n$ matrix with a bounded condition number that minimizes an objective function. This task arises when estimating a covariance matrix from noisy data or due to model constraints, which can cause spurious small negative eigenvalues. A special case is the problem of finding the nearest well-conditioned PD matrix to a given matrix. RCO explicitly optimizes the entries of the Cholesky factor. This requires solving a regularized nonlinear optimization problem, for which we apply Newton-CG and exploit the Hessian's sparsity. The regularization parameter is determined via numerical continuation with an accuracy-conditioning trade-off criterion. We apply RCO to our motivating educational measurement application of estimating the covariance matrix of an Empirical Best Linear Prediction (EBLP) of school growth scores. RCO outperformed general-purpose near-PD algorithms (Higham's method; Tanaka-Nakata) in terms of the EBLP estimate bias and accuracy of its mean squared error. We present results for two empirical datasets: a large urban school district dataset and state dataset. For finding the nearest PD matrix, RCO yields similar results to near-PD methods. While all methods admit an $O(n^3)$ runtime complexity, RCO's constant is $100\times$ larger, thus it is best suited for general objective functions as opposed to this particular

task.

Oren E. Livne
Department of Human Genetics
University of Chicago
olivne@ets.org

Katherine Castellano, Dan McCaffrey
Educational Testing Service
kecastellano@ets.org, dmccaffrey@ets.org

MS353

Newton-MR Algorithms with Complexity Guarantees for Non-Convex Optimization

Classically, the conjugate gradient (CG) method has been the dominant solver in most inexact Newton-type methods for unconstrained optimization. In this talk, we consider replacing CG with the minimum residual method (MINRES), which is often used for symmetric but possibly indefinite linear systems. We show that MINRES has an inherent ability to detect negative-curvature directions. Equipped with this advantage, we discuss algorithms, under the general name of Newton-MR, which can be used for optimization of general non-convex objectives, and that come with favorable complexity guarantees. We also give numerical examples demonstrating the performance of these methods for various non-convex problems.

Yang Liu
University of Queensland
yang.liu2@uq.edu.au

Fred Roosta
University of Queensland, Australia
fred.roosta@uq.edu.au

MS353

Competitive Gradient Descent Algorithms

Competitive gradient descent and related method extend gradient descent to multi-agent optimization by solving, at each iteration, for the Nash equilibrium of a local, multilinear approximation of the agents' loss functions. This talk presents extensions and applications of competitive gradient descent, including mirror-descent variants based on ideas from information geometry and competitive physics-informed networks (CPINNs) that achieve order-of-magnitude improvements in accuracy compared to PINNs.

Florian Schaefer
Georgia Institute of Technology
florian.schaefer@cc.gatech.edu

Anima Anandkumar
Amazon and California Institute of Technology
anima@caltech.edu

Spencer H. Bryngelson
University of Illinois
shb@gatech.edu

Yash Kothari
Georgia Institute of Technology
ykothari3@gatech.edu

Houman Owhadi

Applied Mathematics
Caltech
owhadi@caltech.edu

Qi Zeng
Georgia Institute of Technology
qizeng@gatech.edu

Hongkai Zheng
California Institute of Technology
hzzheng@caltech.edu

MS354

Overview of Earth Science Challenges in the Exascale Era

The increase in Earth System Models (ESMs) capabilities is strongly linked to the amount of computing power and data storage capacity available. Exascale ESMs are expected to open up a new range of opportunities, from increasing model resolution and larger ensemble simulations to the integration of new components to increase the accuracy of our predictions. To make this possible, exascale supercomputers will rely on accelerators and specialised hardware. While diversity for hardware gives more flexibility for co-design, it will increase the complexity of our ESMs. The exascale challenges are: leveraging accelerators, coupling models, ensuring accuracy and stability, dealing with increased I/O, enabling big-data workflows, etc. Solutions to these challenges must also guarantee future-proofing codes, pursue performance-portability across pre and exascale architectures, domain-specific languages and code sustainability. Most applications will require some degree of rewriting to expose more parallelism, and many face severe strong-scaling challenges if they are to effectively progress to exascale. Additionally, ESMs are expected to be performant in terms of both computational and energy efficiency. This talk will include a exascale overview for Earth modeling, covering main challenges to achieve the exascale goal for our ESMs and the range of solutions adopted by the community to be ready for the new era.

Mario Acosta
Barcelona Supercomputing Center
Universitat Politècnica de Catalunya
mario.acosta@bsc.es

Sarat Sreepathi
Oak Ridge National Laboratory
sarat@ornl.gov

MS354

Global Cloud-Resolving Simulations with NICAM on the Supercomputer Fugaku

I will briefly introduce our recent activity in relation to the supercomputer Fugaku, the latest flagship machine of Japan, and show some issues in terms of both computational and modeling aspects toward an era of practical use of global cloud-resolving climate model. The Nonhydrostatic ICosahedral Atmospheric Model (NICAM), a global model with an icosahedral grid system, has been developed and used for two decades. The first grand challenge simulation of NICAM was global cloud-resolving (here, 3.5 km mesh) simulation around 2005, and its extension toward climate scale is now in sight with Fugaku. The value of such climate simulation is guaranteed by a series of recent model improvements to better simulate both mean state

and disturbances. We are also revealing the climatological performance of the ocean-coupled version of NICAM through a series of seasonal to decadal simulations. Another important challenge is an $O(10^3)$ ensemble global prediction experiment with 14 km mesh for better forecast of tropical cyclones over a week. As a resolution-intensive approach, we are trying to run NICAM with a mesh size of 200 m globally to open an era of global large-eddy simulation. The above challenges on Fugaku would be harder without an effective co-design activity with the application developers. In terms of post-processing, we recently modified our MPI-based remapping tool and confirm its applicability to a global 870 m mesh data.

Chihiro Kodama

Japan Agency for Marine-Earth Science and Technology (JAMST)
kodamac@jamstec.go.jp

Hisashi Yashiro

National Institute for Environmental Studies
yashiro.hisashi@nies.go.jp

Tamaki Suematsu

RIKEN Center for Computational Science
tamaki.suematsu@riken.jp

Daisuke Takasuka, Yuki Takano, Yohei Yamada, Ryusuke Masunaga

Japan Agency for Marine-Earth Science and Technology
dtakasuka@jamstec.go.jp, takanoy@jamstec.go.jp,
yoheiy@jamstec.go.jp, masunagar@jamstec.go.jp

Shuhei Matsugishi

The University of Tokyo
matsugishi@aori.u-tokyo.ac.jp

Takashi Arakawa

CliMTech
arakawa@climtech.jp

Tatsuya Seiki, Masuo Nakano

Japan Agency for Marine-Earth Science and Technology
tseiki@jamstec.go.jp, masuo@jamstec.go.jp

Hiroaki Miura

The University of Tokyo
h_miura@eps.s.u-tokyo.ac.jp

Masaki Satoh

Atmospheric and Ocean Research Institute
The University of Tokyo
satoh@aori.u-tokyo.ac.jp

Tomoki Ohno

Meteorological Research Institute
t-ohno@mri-jma.go.jp

Akira Noda

Japan Agency for Marine-Earth Science and Technology
a_noda@jamstec.go.jp

Woosub Roh

The University of Tokyo
ws-roh@aori.u-tokyo.ac.jp

Tomoe Nasuno

Japan Agency for Marine-Earth Science and Technology

nasuno@jamstec.go.jp

Tomoki Miyakawa

The University of Tokyo
miyakawa@aori.u-tokyo.ac.jp

MS354

Can Software Be Portable, Performant and Productive? Co-Designing a Weather and Climate Model with a Domain Specific Language

Weather and climate models simulate complex, multi-scale physics. They can deliver substantial value for society utilising their predictive power. However, It can take a decade or more to develop a new model. Moreover, computer architectures and the programming models used to develop the software have evolved on a much shorter timescale and become more complex. In this presentation the Domain Specific Language (DSL) approach being adopted by the Met Office in developing its next generation model, LFRic, will be discussed and how different programming models for different architectures can then be employed without having to re-write the science code will be presented.

Chris Maynard

UK Met Office
University of Reading
christopher.maynard@metoffice.gov.uk

MS354

Numerical Simulation of an Idealized Coupled Ocean-Atmosphere Climate Model

We present numerical simulations for an idealized coupled ocean-atmosphere climate model. Our climate model belongs to the class of intermediate coupled models which are much simpler than the coupled general circulation models of the ocean-atmosphere system but still allow to study the fundamental aspects of ocean-atmosphere interactions. Our model couples an atmosphere system, described by the compressible two-dimensional (2D) Navier-Stokes equations and an advection-diffusion equation for temperature, to an ocean system, given by 2D incompressible Navier-Stokes equations and an advection-diffusion equation for temperature. Finite element method is used to discretize the system of partial differential equations representing the climate model on a 2D periodic domain and the discrete model is solved using Firedrake, which is an efficient automated finite element method library. The accuracy of simulation results of the climate model is ensured by carrying out detailed numerical investigation of its atmosphere and ocean components separately and then testing our codes against some benchmark problems available in the literature. Our final goal is to incorporate stochasticity into the coupled ocean-atmosphere model following the Hasselmann's paradigm [Klaus Hasselmann, Stochastic climate models part I. Theory, 1976] and use the model to study key features of climate phenomena like El-Nino Southern Oscillation. This will be the subject of our future work.

Kamal Sharma

University of Hamburg
kamal.sharma@uni-hamburg.de

Peter Korn

Max Planck Institute for Meteorology

peter.korn@mpimet.mpg.de

MS355

Progress on Faster Butterfly Construction Based on Randomized Matrix-Vector Multiplication

Butterfly factorization is a promising technique to develop fast direct solvers for solving differential equations or integral equations for high-frequency wave equations. In the past, several butterfly-based direct solvers have been developed to demonstrate reduced complexities. The critical component for constructing these solvers is to reconstruct a butterfly representation of a block from fast randomized black-box matrix-vector multiplications. Previously we developed a randomized butterfly algorithm that yields an $O(N^{1.5} \log N)$ complexity for a $N \times N$ block. In this work, we show some recent progress on developing an improved randomized algorithm that detects and prunes the sparsity patterns of a butterfly on-the-fly. Empirical studies show that an $O(N \log^2 N)$ complexity is achievable for several classes of butterfly matrices.

Yang Liu

Lawrence Berkeley National Laboratory
liuyangzhu@lbl.gov

MS355

Random Butterfly Matrices and Growth Factors

Random butterfly matrices were introduced by Parker in 1995 to remove the need for pivoting when using Gaussian elimination. The growing applications of butterfly matrices have often eclipsed the mathematical understanding of how or why butterfly matrices are able to accomplish these given tasks. To help begin to close this gap using theoretical and numerical approaches, we explore the impact on the growth factor of preconditioning a linear system by butterfly matrices. These results are compared to other common methods found in randomized numerical linear algebra. In these experiments, we show preconditioning using butterfly matrices has a more significant dampening impact on large growth factors than other common preconditioners and a smaller increase to minimal growth factor systems. Moreover, we are able to determine the full distribution of the growth factors for a subclass of random butterfly matrices. Previous results by Trefethen and Schreiber relating to the distribution of random growth factors were limited to empirical estimates of the first moment for Ginibre matrices. This is joint work with Tom Trogdon.

John Peca-Medlin

Department of Mathematics
University of Arizona
johnpeca@math.arizona.edu

MS355

Butterfly-Accelerated Manifold Harmonic Transforms

Manifold harmonics are the eigenfunctions of the Laplace-Beltrami operator (LBO) on a Riemannian manifold. A variety of fast special function transforms for orthonormal bases consisting of discretized manifold harmonics for highly symmetric geometries exist. For example, the fast Fourier transform (FFT) is a fast algorithm for applying the discrete Fourier transform, which itself is an orthonormal basis comprised of sampled eigenfunctions of the LBO on a torus. At the same time, the butterfly factorization

provides a linear algebraic framework generalizing the FFT based on factoring an $N \times N$ matrix into a product of $O(\log N)$ sparse factors, leading to fast transforms. It has been used to develop fast transforms for systems of orthogonal polynomials. We present an approach to designing subquadratic fast transforms based on butterfly factorizing the LBO eigenmatrix obtained under a finite element discretization of the eigenproblem. Numerical examples taken from computer graphics will be used for illustration.

Samuel F. Potter

Courant Institute of Mathematical Sciences
New York University
sfp@cims.nyu.edu

MS355

Fast Orthogonal Polynomial Transformations

FastTransforms is a free and open-source software package in C with a Julia wrapper that provides a number of fast algorithms for orthogonal polynomial transforms (<https://github.com/MikaelSlevinsky/FastTransforms>).

For univariate classical orthogonal polynomials, the connection coefficients of a measure-perturbing problem are the matrices of eigenvectors of a triangular-banded generalized eigenvalue problem. We use a divide-and-conquer approach that is accelerated by the fast multipole method (FMM). For bivariate analogues of the classical orthogonal polynomials, including spherical harmonics, Zernike polynomials, and Prorol polynomials on the triangle, a measure-preserving connection is developed to convert expansions into tensor product bases that are suitable for FFTW-based synthesis and analysis. It is empirically well-known that these measure-preserving connection problems (and overall syntheses and analyses) may be compressed by the butterfly factorization. We provide an incremental contribution to the theoretical underpinning by relating the connection coefficients to Fourier integral operators.

Richard M. Slevinsky

Department of Mathematics
University of Manitoba
richard.slevinsky@umanitoba.ca

MS356

A New Strategy in Developing Location Model

The location model (LM) is designed to enable classification when a dataset contains both continuous and categorical variables. Due to the issue of empty cells, a smoothed location model (smoothed LM) is introduced. However, the smoothing process caused changes in the original information of the non-empty cells. It is well known that original information is valuable and important that should be maintained. Thus, a new strategy is proposed by amalgamating of maximum likelihood and smoothing estimations to construct a new LM. Consequently, maximum likelihood estimation will be used if the cell was found to be non-empty, otherwise smoothing estimation will be used instead. The analysis shows that the newly constructed LM can provide optimal classification results and demonstrates better performance compared to the old models, i.e. classical LM and smoothed LM, where the estimation used is based on the cells conditions. The new proposed strategy of parameter estimation could handle all situations; whether the cells are empty or not, limited sample size with many variables

measured mainly the binary.

Hashibah Hamid

School of Quantitative Sciences, Universiti Utara
Malaysia
06010 UUM Sintok, Kedah
hashibah@uum.edu.my

MS356

Neural Sampling Machine with Stochastic Synapse Allows Brain-Like Learning and Inference

Estimating the confidence of neural networks predictions from real-world data is a critical requirement for mission-critical AI technologies. The brain is capable of such uncertainty estimation from data, but the principles underlying this feat and its implementation in a compact, low-power hardware remains a challenge. In this presentation, I will first introduce Neural Sampling Machines (NSM), a stochastic binary neural network inspired by the stochastic behavior of biological neurons that uses to perform probabilistic inference. We demonstrate a novel hardware fabric that can implement NSMs by exploiting the stochasticity in the synaptic connections. We experimentally demonstrate a hybrid stochastic synapse by pairing a ferroelectric field-effect transistor (FeFET)-based analog weight cell with a two-terminal stochastic selector element. We show that the stochastic switching characteristic of the selector between the insulator and the metallic states resembles the multiplicative synaptic noise of the NSM. The stochastic NSM can not only perform highly accurate image classification with 98.25% accuracy on standard MNIST dataset, but also estimate the uncertainty in prediction (measured in terms of the entropy of prediction) when the digits of the MNIST dataset are rotated. Building such a probabilistic hardware platform that can support neuroscience inspired models can enhance the learning and inference capability of the current artificial intelligence (AI).

Emre Neftci

RWTH Aachen
Forschungszentrum Jülich
e.neftci@fz-juelich.de

MS356

Stochastic Computing: From Classical to Morphological Neural Networks

During the last few years, Stochastic Computing (SC) has emerged as a promising design technique for the implementation of energy-efficient Artificial Neural Networks (ANN) hardware. This efficiency is mainly due to the possibility of parallelizing the network through the use of SC and, therefore, minimizing its dependence on a central memory node. In the classical ANN implementation, the SC circuit is adapted to the standard ANN design essentially exploiting the ability to reduce a multiplication within a single logic gate. However, SC has many other advantages such as its ability to implement maximum and minimum functions also with individual gates. Therefore, the possibility of searching for the NN paradigm that best suits the capabilities of SC (and therefore using both products and max-min functions) is opened. In this work we show that NNs known as Morphological Neural Networks (MNN) are ideal to be implemented in hardware using SC. MNNs can be implemented without using activation functions and combine products, sums, maximum and minimum functions. They have been studied in the literature by some researchers in recent years and present the property of universal ap-

proximation. An FPGA implementation of a large SC-based MNN is presented and compared with other hardware methodologies in terms of energy efficiency and performance. The results show that the proposed SC-based MNN is as competitive as binary neural networks.

Josep L. Rossello

Universitat de les Illes Balears
j.rossello@uib.es

Joan Font-Rossello

Universitat de les illes Balears
joan.font@uib.es

Vicens Canals, Alejandro Moran, Christian Frasser, Miquel Roca

Universitat de les Illes Balears
v.canals@uib.es, a.moran@uib.es, christian.franco@uib.es, miquel.roca@uib.es

MS356

Sampling Distributions from Biased Coins

Flips of a fair coin can be exploited to approximate outcomes of distributions in expected ways. Fair coins can even be used to simulate the flips of biased coins, and biased coins can also be exploited to sample distributions. Deterministic computing and PRNGs are able to exploit such simple relations. Stochastic devices that exhibit two state behavior, however, may not be able to be reliably classified as a Bernoulli coin flip with a fixed probability. Devices may behave statistically as such a coin over many, many draws, but they may actually fluctuate between values or behave in some other pathological way. This talk examines the consequences of utilizing a noisy device for probabilistic computation under the assumption it is behaving as a fair or biased Bernoulli coin flip when the underlying dynamics are actually much more complex. We aim to categorize such a TRNG under appropriate metrics analogous to those used for PRNG today. We conclude with a discussion on how rich dynamics of noisy two state devices may be exploited for probabilistic computation.

Darby Smith, William Severa

Sandia National Laboratories
jsmit16@sandia.gov, wmsevera@sandia.gov

MS357

Continuum-Scale Simulation Tools for Understanding the Interaction of Shear and Thermal Gradients in Flows of Dense Suspensions

Suspension flows are notoriously difficult to model due to the variety of fluid-particle and particle-particle interactions that can occur. We propose a new approach based on a two-fluid model (TFM). The TFM can be conveniently implemented in an OpenFOAM solver to allow industrial-scale simulations, including in curvilinear coordinates. Good agreement is found between TFM simulations and previous experiments on shear-induced particle migration. Heat transfer in dense suspension remains a frontier topic. We propose a closure relation for the inter-phase heat transfer coefficient in the TFM, based on calibration against experimental data. This closure allows us to use the OpenFOAM solver developed to simulate non-isothermal dense suspension flow in an annular Couette cell. In this context, we identify a novel ‘thermo-rheological’ migration force, which can oppose (or aid) the well-known shear-induced migration. This work is joint

with Dr. Federico Municchi (Colorado School of Mines) and Purdue PhD student Pranay Nagrani, and it was supported by the American Chemical Society Petroleum Research Fund.

Ivan C. Christov
Purdue University
christov@purdue.edu

MS357

Efficient Neighbor Detection for Polydisperse Gas-Solids Flows

Lagrangian particle tracking methods have become a well established simulation tool for multiphase flows. When employing a 'soft-sphere' methodology, where particle collisions are multi-bodied and enduring, a neighbor list must be constructed. Traditionally, the neighbor list is built with a grid that is 2-3X the largest particle diameter. When considering polydisperse solids, the construction of the neighbor list and computation of collision forces becomes computationally prohibitive since many small particles are contained within a cell. Following [Shire et al., 2021, DOI: 10.1007/s40571-020-00361-2], we implement a multi-grid neighbor search algorithm within AMReX; an exascale-compatible, open-source software for block-structured adaptive mesh refinement [Zhang et al., 2021, DOI: 10.1177/10943420211022811]. Scaling of the multi-grid neighbor search algorithm on CPU and GPU is examined for a variety of gas-solids flows with MFiX-Exa; an open-source Euler-Lagrange solver that is built upon the AMReX library [Musser et al., 2022, DOI: 10.1177/109434202110092].

Aaron M. Lattanzi
Lawrence Berkeley National Laboratory
amlattanzi@lbl.gov

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

MS357

An Unstructured Collocated Finite-Volume Level Set Method for Incompressible Two-Phase Flows with a Geometrical Phase Indicator

The Level Set method for simulating multiphase flows traditionally relies on re-distancing the level-set function. The signed distance is advantageous for the phase indicator computation that relies on algebraic Heaviside function models. Traditionally, the signed-distance property was also motivated by simplified curvature computation and for stabilizing the numerical differentiation. We propose to use a geometrical computation of the zero level-set for the phase indication and the curvature approximation. This approach was disregarded initially for the level set method because of computational costs; however, recent advances in the geometrical computation of volume fractions for the unstructured Volume-of-Fluid method provide high-speed intersection algorithms for non-convex polyhedrons. Our approach omits re-distancing (that amplifies mass loss) while relying on compact-stencil unstructured Finite-Volume discretization of the Level Set equation. The method is sufficiently volume-conservative and accurately advects the phase indicator, while its straightforward equation discretization and efficient geometrical phase indicator deliver a high degree of (parallel) computational

efficiency.

Tomislav Maric
Technische Universität Darmstadt
maric@mma.tu-darmstadt.de

Damir Juric
LISN, CNRS, France
damir.juric@limsi.fr

Jalel Chergui
LISN, CNRS
jalel.chergui@lisn.fr

Seungwon Shin
Hongik University, South Korea
sshin@hongik.ac.kr

MS357

A Set of Machine Learning-Based Filtered Drag Closures for Cohesive Gas-Particle Flows

Coarse-grid simulations of large-scale gas-solid flows require appropriate sub-grid closure models to approximate unresolved physical phenomena. Such sub-grid closures must account for the effects of the inhomogeneous particle distribution. Several closure models are available in the literature for free flowing and non-cohesive gas-solid flows. In our present study we expand this existing knowledge by investigate the effect of cohesion on the drag force closure, and by proposing a drag closure concept based on machine learning (ML). The key novelty is that our concept allows the rapid generation of closures, making it possible to create a set of closure laws rather than a single closure that might have limitations. First, the results of classical unresolved (but unfiltered) CFD-DEM simulations of cohesive gas-particle flow in a pseudo 2D periodic domain are filtered with different filter sizes. Second, corresponding fully 3D simulations were performed, and filtered in an identical manner. Third, closure models were developed from data of 2D and 3D simulations and compared to each other. ML-closures were derived separately with data from the gravitational and non-gravitational (i.e., horizontal) direction, allowing us to demonstrate that filtered drag anisotropy is significant. Our benchmarking of 2D versus 3D closures suggests that only 3D data should be used for closure development in the case of cohesive powders.

Josef Tausendschön, Mohammadsadegh Salehi
Graz University of Technology
josef.tausendschoen@tugraz.at, m.salehi@tugraz.at

Sankaran Sundaresan
Princeton University
sundar@princeton.edu

Stefan Radl
Graz University of Technology
radl@tugraz.at

MS357

Upscaling Particulate Simulations Towards Continuum Modeling of Dense Granular Flows

At high solid volume fractions, the dynamics of dense granular materials are dominated by enduring inter-particle contacts that emerge as intriguing macroscale phenomena such as yield stress, normal stress effects, and geometry-

dependent rheology. Particle-based simulations such as the Discrete Element Method (DEM) are well-suited to resolve particulate features, but ill-suited for modeling at continuum scales. Whereas continuum modeling is ideally suited for macroscale simulations, there is a lack of a general constitutive model that correctly accounts for loading geometry and higher-order rate effects. This talk will describe our computational efforts in upscaling DEM simulations towards continuum modeling by developing tensorial rheological models of dense granular flows. DEM-based numerical methods that can simulate granular flows to infinitely large strains for arbitrary imposed loading geometries will be described. These simulations were used to calibrate a frame-indifferent, tensorial, viscoplastic rheological model that is valid for various deformation geometries beyond shear, such as biaxial and triaxial deformations. This talk will describe our ongoing work on implementing a regularized version of the tensorial rheological model in an incompressible flow solver towards simulating dense granular materials at the continuum.

Ishan Srivastava

Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory
isriva@lbl.gov

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

MS358

A Model of Collective Motion of Self-Propelled Organisms

We present a single-point model for self-propelled microscopic swimmers immersed in a viscous, incompressible fluid. The organism model is based on a specific limit of regularized Stokeslets, resulting in the superposition of a regularized stresslet and a regularized potential dipole at the same location, and an orientation vector. No special treatment of the self-velocity is necessary. Modeling “pusher and “puller organisms is straightforward, and the model can also be extended to flows bounded by a plane wall using a method of images and to organisms with rotating helical flagella by including rotlets. The model derivation will be sketched and its properties will be described through examples that show the motion, interaction and diffusion of these particles as a function of the concentration density.

Ricardo Cortez

Tulane University
Mathematics Department
rcortez@tulane.edu

Alex Hoover
Cleveland State University
a.p.hoover@csuohio.edu

Priya Boindala
Georgia Gwinnett College

pboindal@ggc.edu

MS358

Tubular-Body Theory

Cable-like bodies play a key role in many biological fluid systems but are hard to simulate. Asymptotic theories, called slender-body theories, are effective but apply in specific regimes and can be hard to extend beyond leading order. In this presentation I will develop an exact slender-body-like theory for the surface traction of cable like bodies in viscous flow. This theory expresses the traction as a series of solutions to a well behaved one-dimensional Fredholm integral equation of the second kind. This process can be simply generalised to other systems. We test this theory against known solutions and then use it to look at the swimming of a tightly wound helix exposed to an external torque.

Lyndon Koens

University of Hull
l.m.koens@hull.ac.uk

MS358

Regularized Stokeslets Inside Spherical Geometries

We formulate a general mapping from singular image systems to regularized image systems for use with the method of regularized Stokeslets and apply it to biologically relevant flows within a fluid-filled, spherical cavity. The method of regularized Stokeslets has been heavily applied to micro-scale biofluids problems such as microorganism locomotion. It is based on the boundary integral form of the Stokes equations where the singular fundamental solution (the “Stokeslet”) is replaced with a smooth approximation. The method can be modified such that certain boundary conditions are satisfied automatically via the addition of an appropriate “image system” to the regularized Stokeslet. Here, we focus on the computation of fluid flows due to distributions of forces within a fluid-filled, spherical cavity because no regularized image system has been derived for this specific setting and because of its importance to flows driven by active structures within cells. We present two example problems in this context: the flow due to a translating particle and the flow due to a relaxing flexible filament contained in the cavity.

Sarah D. Olson, Nicholas G. Chisholm

Worcester Polytechnic Institute
sdolson@wpi.edu, nchisholm@wpi.edu

MS358

Undulatory Swimming in Suspensions and Networks of Flexible Filaments

Microorganisms swimming through viscoelastic fluids are a common feature in naturally-occurring fluids. Approaches to modelling this behaviour normally come in two flavours: a macroscale, course-grained approach where the fluid is modelled as a continuum; and a microscale, detailed approach where the anisotropic structure within the background fluid is specifically modelled. The former approach has the benefit of speed, but in this talk I will share my experience of attempting the latter, in 3D, using high-performance computer simulations, and show how these simulations suggest the physical structure of the viscoelastic fluid directly affects the microorganism swimming through it. We see some surprising differences with 2D and

investigate why this might be so.

Adam K. Townsend
Durham University
adam.k.townsend@durham.ac.uk

Eric Keaveny
Imperial College
e.keaveny@imperial.ac.uk

MS359

Topology Optimization of a Rotating Electric Machine by the Topological Derivative

We consider the topology optimization of a rotating electric machine in magnetoquasistatic operation in two space dimensions. This amounts to a topology optimization problem subject to a parabolic PDE constraint on a moving domain which we intend to solve by means of the topological derivative concept. For that purpose, we consider a topological perturbation of the materials in the space-time cylinder along a trajectory given by the rotation of the machine. Using a Lagrangian approach, we derive the topological derivative formula, which depends on the solution of an exterior problem that is bounded in time and unbounded in the space directions. We present an efficient way for the numerical computation of the topological derivative and use this sensitivity information for the design optimization of an electric motor by means of a level set method.

Peter Gangl
RICAM
Johann Radon Institute for Computational and Applied Mathema
peter.gangl@ricam.oeaw.ac.at

Nepomuk Krenn
Johann Radon Institute for Computational and Applied Mathematics
nepomuk.krenn@oeaw.ac.at

MS359

Space-Time Finite Element Methods for Electric Machines

For the numerical solution of parabolic evolution equations we consider a space-time least squares formulation in appropriate function spaces which results in a mixed finite element scheme. To ensure stability, a discrete inf-sup stability condition has to be satisfied. Since the adjoint is zero in the continuous formulation, its Galerkin discretization can be used to define an a posteriori error estimator for the primal unknown to drive an adaptive scheme. This approach is then applied to parabolic-elliptic interface problems in order to model electric machines.

Christian Köthe
TU Graz
Institute of Applied Mathematics
c.koethe@tugraz.at

Olaf Steinbach
TU Graz, Institute of Applied Mathematics
o.steinbach@tugraz.at

MS359

Surrogate Modeling, Uncertainty Quantification,

and Sensitivity Analysis for Thermal Circuit Models of Electric Machines

Thermal circuit models offer a computationally inexpensive alternative for assessing the performance of electric machine designs, as opposed to resource-demanding thermal and multiphysics simulations, e.g., based on finite element or finite volume methods. However, the simplifications that are necessary to establish a thermal circuit model introduce a number of uncertain model parameters. Variations in these parameters may significantly affect model-based predictions and, eventually, machine performance. In this study, we present surrogate modeling methods based on data-driven, regression-based Polynomial Chaos Expansions (PCEs), which allow to perform uncertainty analyses and estimate the impact of said uncertainties on scalar- and vector-valued machine model outputs, both effectively and accurately. Additionally, the sensitivity of model outputs on the uncertain parameters is assessed by means of a PCE-based Sobol sensitivity analysis, which is additionally extended to the vector-valued output case by means of generalized Sobol indices. Sensitivity analysis results are also used to reduce the number of considered uncertainties to a minimum, by neglecting non-influential model parameters, thus reducing the computational demand of subsequent studies.

Dimitrios Loukrezis
Siemens AG
dimitrios.loukrezis@siemens.com

MS359

Space-Time Methods for Parabolic Quasi-Linear Time-Periodic Evolution Problems

We present new space-time finite element methods for the solution of time-periodic parabolic evolution problems, with possibly nonlinear coefficients. Our approach is based on an all-at-once discretization, which results, in the linear case, in a single, huge linear system. The linear system is then solved by means of an algebraic multigrid preconditioned GMRES method. We present numerical experiments based on applications in magneto-quasi statics as well as optimization problems, where we additionally investigate the parallel performance of our method.

Herbert Egger
Johannes Kepler University Linz
Institute for Computational Mathematics
herbert.egger@jku.at

Andreas Schafelner
Johann Radon Institute for Computational and Applied Mathematics
andreas.schafelner@oeaw.ac.at

MS360

Stochastic Sampling Techniques for Muscle Recruitment Distribution with Longitudinal Prior

Movement in the human body is made possible by a coordinated joint effort from several muscles. A very simple task performed by an individual such as raising an arm or lifting a leg can be accomplished with a wide range of muscle activation levels due to the concept of muscle redundancy. Some quantification is provided for the redundancy while presenting feasible solutions (muscle forces at joints). A sample of feasible solutions is generated for each time instance in the time interval for executing a specific locomo-

tion of the body. This is done in the Bayesian framework, employing Markov Chain Monte Carlo methods. These solutions are in conformance with the uncontrolled manifold theory. Furthermore, inspired by the principles surrounding the Feynman-Kac formula, we propose an approach that smoothly connects the muscle activation pattern at each time instance to mimic the smoothness exhibited in normal human movement.

Mercy G. Amankwah
Case Western Reserve University
mercy.amankwah@case.edu

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and
Statistic
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS360

Efficient Coarse Propagators for Parallel Time Stepping in Plaque Growth Simulations

The numerical simulation of atherosclerotic plaque growth is computationally prohibitive, since it involves a complex cardiovascular fluid-structure interaction (FSI) problem with a characteristic time scale of milliseconds to seconds, as well as a plaque growth process governed by reaction-diffusion equations, which takes place over several months. A resolution of the fast (micro) scale over this period can easily require more than a billion time steps, each corresponding to the solution of a computationally expensive FSI problem. To tackle this problem, we combine a temporal homogenisation approach with parallel time-stepping. First, a temporal homogenisation approach separates the problem in an FSI problem on the micro scale and a reaction-diffusion problem on the macro scale. Second, a parallel time-stepping approach based on the parareal algorithm is applied on the macro-scale of the homogenised system. The parallel scalability is, however, dominated by the computational cost of the coarse propagator. We discuss different possibilities to modify the coarse propagator to further reduce the number of expensive micro problems to be solved and test the numerical algorithms in detailed numerical studies.

Stefan Frei
University of Konstanz
stefan.frei@uni-konstanz.de

Alexander Heinlein
Delft University of Technology
Delft Institute of Applied Mathematics
a.heinlein@tudelft.nl

MS360

Reverse Engineering of a Long-Distance Hormone Transport Mechanism in Plants by Computational Modeling and Data Analysis

Indole-3-acetic acid ('auxin') is a plant hormone that plays an important role in various developmental processes. Its transport through the plant has been studied for more than a century, but the detailed mechanisms of transport, in particular the long-range transport from top to roots, are still

not fully clear. In this talk we discuss how mathematical modeling, simulation and experimental research have been integrated to gain insights into these mechanisms in the model plant *Arabidopsis thaliana*. The model consists of a coupled system of reaction-diffusion-advection equations. Parameter estimation procedures and physiologically relevant ranges for parameters require simulation of the model in regimes where the Peclet number is neither large nor small, causing numerical complications. Moreover, experimental circumstances strongly limit the number of data points (of various types) that can be obtained. A practical parameter identifiability analysis needs to be performed to identify which parameter combinations can still be confidently estimated from the data. We introduce such a method and show that the addition of the measurement of a single well-chosen quantity (and single data point) in the experimental design can drastically influence the collection of identifiable parameter sets. Any of these sets can then be used to draw conclusions on the underlying mechanisms and aid in the 'reverse-engineering' of the long-range auxin transport mechanism.

Sander Hille
Mathematical Institute
Universiteit Leiden
shille@math.leidenuniv.nl

MS360

Applying Machine Learning to Enhance the Computational Efficiency in Agent-Based Wound Healing Model

Deep tissue injuries often go with a contraction of the skin. This contraction can become contracture if it causes a loss of mobility of the joint of the patient. Contracture results from the cellular traction forces exerted by the (myo)fibroblasts during the proliferation phase. In Peng and Vermolen (2020), an agent-based model is developed to mimic the wound healing, in which every cell is treated individually as a fixed circle in two dimensions. To model the cellular forces exerted by the (myo)fibroblasts, the cell membrane is first broken into line segments. Then, the point forces are applied at the midpoint of each line segment using the Dirac Delta distribution and then solving the elasticity equation with finite-element methods; this is the most computationally expensive part of the entire model. Since our final goal is to predict the likelihood that such a contracture occurs and to analyze the impact of several treatments on preventing a serious contraction, the model has to be evaluated many times and a computationally efficient model is required. Moreover, while myofibroblasts are known as a dendritic shape with lamellipodium developed from the cell membrane, a large number of line segments would be needed to model such cell membrane shapes. In the talk, the benefits of the machine learning-enhanced model with respect to the variation of the cell shape and the resulting improvements in the computational efficiency are discussed based on numerical results.

Qiyao Peng
Leiden University
Mathematical Institute
q.peng@math.leidenuniv.nl

Alexander Heinlein
Delft University of Technology
Delft Institute of Applied Mathematics

a.heinlein@tudelft.nl

MS360**Model Selection Using Rare Data in Life-Sciences**

Biological processes are the result of many interacting particles on microscopic scales. While some of the particles and interactions are well known, oftentimes not all dynamics are fully understood. Therefore, pure physics-based modeling approaches are not always successful. Apart from that, some outcomes of the considered biological processes are usually observable on a macroscopic scale. Although, the observations mostly do not include time and space resolved measurements of all interacting components but some information at discrete times. These observations serve as quantitative data for the biological applications. The quantitative data is usually not sufficient for pure data-driven approaches. In this talk, ideas on how to use the rare data in a model selection process are presented: Mesoscopic models connect the different scales by abstracting from the microscopic particles and interactions. The few microscopic information on mechanisms and the observation on the macroscopic scale enable together physics-informed machine learning techniques for the mesoscopic model.

Cordula Reisch

TU Braunschweig
 Institut Computational Mathematics
 c.reisch@tu-bs.de

MS361**Structure-Exploiting Learning for Nested Operator Inference**

Projection-based model order reduction methods are well-established for approximating computationally expensive full-order models for use in predictive, real-time, or many-query applications. By exploiting the structure of the governing equations for a low-dimensional subspace, the constructed reduced-order model typically 1) achieves significant computational savings, 2) builds upon established error theory to guarantee approximation accuracy, and 3) remains interpretable in terms of the governing physical equations. However, by construction projection-based methods require access to the full-order operators—a major shortcoming for large-scale applications with legacy codes or commercial solvers. Operator Inference methods circumvent this requirement by learning the intrusive reduced-order model from available full-order data and the structure of the governing equations. The method guarantees exact reconstruction in the limit or by applying an additional re-projection technique if sufficient data is available to stabilize the learning problem. In this talk, we exploit the projective structure of the reduced operators for a nested operator inference approach that continues to guarantee exact reconstruction with re-projection while significantly reducing the data requirement from polynomial scaling to $O(1)$. This in particular makes the nested Operator Inference approach applicable for high-order polynomial operators that were previously infeasible to infer.

Nicole Aretz

The University of Texas at Austin
 nicole.aretz@austin.utexas.edu

Karen E. Willcox
 UT Austin

kwillcox@oden.utexas.edu

MS361**Physics-Informed Neural Networks for Boltzmann-BGK Equations with Absorbing Boundary Layers**

We introduce the application of physics-informed neural networks (PINN) for Boltzmann-BGK equations for nearly incompressible flows. The governing equations are discretized in the velocity space using Hermite polynomials resulting in a first-order conservation law. A perfectly matching layer (PML) surrounding the computational domain is used to dampen the waves leaving the domain. A multi-domain approach is used with PINN to ensure the continuity of the solution. The damping profile of the PML is an output of the PINN offering a suitable constant for absorbing the waves.

Atakan Aygun, Ali Karakus

Middle East Technical University

atakana@metu.edu.tr, akarakus@metu.edu.tr

MS361**Predicting Emergent Behaviors in Coupled Systems Using Localized Reduced Order Models**

Mutual interactions in multi-component parametrized systems often lead to emergent phenomena, i.e., behaviors that cannot be observed by considering the individual components separately. Practical applications often demand accurate and computationally fast methods able to predict such behaviors, which can be achieved through Domain Decomposition (DD) and Model Order Reduction (MOR) techniques. However, simulations of the coupled model may be unfeasible, as it might not be a-priori available or its computational cost is too large. Thus, in this work we propose two MOR techniques that construct local independent surrogate models, whose interaction accurately predicts the solution of a given coupled system. The offline phase is based on a suitable parametrization of the interface data, whereas the online phase combines the reduced models using ideas from DD methods. We show their potential in terms of accuracy and computational performance in different test cases, including nonlinear and multi-physics problems.

Niccolo Discacciati, Jan S. Hesthaven

EPFL

niccolo.discacciati@epfl.ch, jan.hesthaven@epfl.ch

MS361**Beta-Variational Autoencoders for Nonlinear and Orthogonal Reduced-Order Models in Turbulence**

In this work we propose a deep-learning framework for learning a minimal and near-orthogonal set of non-linear modes in the context of turbulent flows. In particular, we focus on a high-fidelity numerical database of a simplified urban environment. The proposed technique relies on beta-variational autoencoders (beta-VAEs) and convolutional neural networks (CNNs), which enable extracting non-linear modes while encouraging the learning of statistically-independent latent variables and penalizing the size of the latent vector. Moreover, we introduce an algorithm for ordering the resulting modes with respect to their contribution to the reconstruction. We demonstrate that by constraining the shape of the latent space, it is possible to motivate orthogonality and extract a set of

parsimonious modes which enable high-quality reconstruction. Our results show the excellent performance of the method in the reconstruction against linear-theory-based decompositions, where the energy percentage captured by the proposed method from five modes is equal to 87.36% against 32.41% from POD. Furthermore, we show the ability of our approach to extract near-orthogonal modes with the determinant of the correlation matrix, which is equal to 0.99, thus enhancing the interpretability of the obtained reduced-order models (ROMs).

Ricardo Vinuesa
KTH Royal Institute of Technology
rvinuesa@mech.kth.se

MS361

Physics-Informed Neural Networks for Learning the Homogenized Coefficients of Multiscale Elliptic Equations

Multiscale elliptic equations with scale separation are often approximated by the corresponding homogenized equations with slowly varying homogenized coefficients (the G-limit). The traditional homogenization techniques typically rely on the periodicity of the multiscale coefficients, thus finding the G-limits often requires sophisticated techniques in more general settings even when multiscale coefficient is known, if possible. Alternatively, we propose a simple approach to estimate the G-limits from (noisy-free or noisy) multiscale solution data, either from the existing forward multiscale solvers or sensor measurements. By casting this problem into an inverse problem, our approach adopts physics-informed neural networks (PINNs) algorithm to estimate the G-limits from the multiscale solution data by leveraging a priori knowledge of the underlying homogenized equations. Unlike the existing approaches, our approach does not rely on the periodicity assumption or the known multiscale coefficient during the learning stage, allowing us to estimate homogenized coefficients in more general settings beyond the periodic setting. We demonstrate that the proposed approach can deliver reasonable and accurate approximations to the G-limits as well as homogenized solutions through several benchmark problems.

Xueyu Zhu
University of Iowa
Department of Mathematics
xueyu-zhu@uiowa.edu

Jun R Park Park
Korea Advanced Institute of Science and Technology
pjss1223@kaist.ac.kr

MS362

Accurate Formulation of Runge-Kutta Nystrm Method for Direct Integration of Second Order Ordinary Differential Equations

In this work, we extend the classical fourth order Runge-Kutta (RK) method to Runge-Kutta Nystrm Type (RKNT) method for direct integration of Special and General third order initial valued ODEs via the idea as those invented by Nystrm. The theory of Nystrm was adopted in the derivation of the method. The formulation yielded an accurate RKNT method with reduced number of function evaluation thereby making it cost effective. The proposed method were tested with Numerical experiment to illustrate its efficiency and the method can be extended

to solve higher order differential equations. The scheme is simple to implement and converges better with the exact solution.

Zamurat A. Adegboye
NIGERIAN DEFENCE ACADEMY
,KADUNA,NIGERIA
zamurat.adegboye@imsp-uac.org

C. A. Anayochukwu
Kaduna Polytechnic, Kaduna, Nigeria
amunews1721@gmail.com

Helen Edogbanya
Federal University Lokoja, Nigeria
helen.edogbanya@fulokoja.edu.ng

MS362

3-Additive Linear Multi-Step Methods for Diffusion-Reaction-Advection Models

Some systems of differential equations that model problems in science and engineering have natural splittings of the right-hand side into the sum of three parts, in particular, diffusion, reaction, and advection. Implicit-explicit (IMEX) methods treat these three terms with two numerical methods, and this may not be desirable. Accordingly, A recent approach that treat diffusion, reaction, and advection with separate methods was introduced in "[Raed Ali Mara'Beh, Raymond J. Spiteri, P. Gonzalez, Jos M. Mantas, "3-additive linear multi-step methods for diffusion-reaction-advection models", Applied Numerical Mathematics, volume 183, January 2023, Pages 15-38. <https://doi.org/10.1016/j.apnum.2022.08.015>]." In this talk, The construction of 3-additive linear multi-step methods for the solution of diffusion-reaction-advection systems is presented. The stability analysis and order of convergence of the new 3-additive methods are considered. A comparison between the new methods is with some popular IMEX methods in terms of stability and performance is presented. In general, It is found that the new 3-additive methods have larger stability regions than the IMEX methods tested in some cases and generally outperform in terms of computational efficiency.

Raed A. Marabeh
Department of Mathematics, Foundation Program
College of General Studies, Qatar University
rmarabeh@qu.edu.qa

MS362

Generalized Exponential Time Differencing Schemes for Stiff Fractional Systems with Nonsmooth Source Term

Many processes in science and engineering are described by fractional systems which may in general be stiff and involve a nonsmooth source term. In this talk, we discuss a framework for the development of exponential time differencing schemes for solving such systems. In particular, we present robust first, second, and third order accurate schemes. Rather than imposing regularity requirements on the solution to account for the singularity caused by the fractional derivative, we only consider regularity requirements on the source term for preserving the optimal order of accuracy of the proposed schemes. Optimal convergence rates are proved for both smooth and nonsmooth source terms using uniform and graded meshes, respectively. For efficient implementation, high-order global

Padé approximations together with their partial fraction decompositions are developed for Mittag-Leffler functions. Numerical experiments involving a typical stiff system, and other large systems obtained by spatial discretization of a sub-diffusion problem are typical examples to be discussed. Demonstration of the efficiency of the rational approximation implementation technique and the constructed high-order schemes are to be illustrated.

Ibrahim O. Sarumi

King Fahd University of Petroleum and Minerals
ibrahim.sarumi@kfupm.edu.sa

Khaled Furati

Department of Mathematics
King Fahd University of Petroleum & Minerals
kmfurati@kfupm.edu.sa

Abdul Khaliq

Middle Tennessee State University
abdul.khaliq@mtsu.edu

Kassem Mustapha

Department of Mathematics and Statistics
King Fahd University of Petroleum and Minerals
kassem@kfupm.edu.sa

MS362

IFDIFF - A Matlab Toolkit for ODEs with State-Dependent Switches

We present the toolkit IFDIFF for integration and sensitivity generation in parameterized implicitly (state-dependent) switched ODEs whose right-hand side is given as Matlab code containing non-differentiable operators (max, abs, etc.) and conditionals (if). Naive implementations using IF-THEN-ELSE branching give unreliable simulation results without warning, as switching events are undetectable by standard integrators. The widespread belief that this can be countered using more stringent integration tolerances is wrong: We give a simple example where the integrators error estimation always delivers zero. Correct treatment of switched systems requires elaborate formulation of switching functions and tailored integrators, placing high mathematical demands on modelers. Even small model changes often imply considerable reformulation effort. Further, n switches generate up to 2^n possible program flows and switching functions, rendering a-priori formulations not feasible already in medium-sized models. IFDIFF programmatically handles switching events, auto-generating only required switching functions. It determines switching times up to machine precision, and ensures accurate simulation and sensitivity results. Transparently extending the Matlab integrators (ode45, ode15s, etc.), IFDIFF is applicable to existing code with state- and parameter-dependent conditionals, thus enabling fast prototyping and relieving modelers of mathematical-technical effort.

Andreas Sommer, Hans Georg Bock, Ekaterina Kostina
Interdisciplinary Center for Scientific Computing (IWR)
Heidelberg University
andreas.sommer@iwr.uni-heidelberg.de, bock@iwr.uni-heidelberg.de, kostina@iwr.uni-heidelberg.de

MS363

Numerical Schemes for Efficient Kernel-Based

Methods on Graphs

Important kernels in classical settings can be formulated explicitly and implemented directly for interpolation or classification purposes. On irregular graph domains this is in general not possible. Kernels that describe diffusion processes or characterize smoothness classes on graphs are not accessible in closed form and have to be calculated numerically. Often, these kernels can be characterized and computed as matrix functions of the graph Laplacian, as for instance the diffusion kernel which is the exponential of the Laplacian. Calculating these matrix functions becomes prohibitive if the size of the graph is large and if the spectral decomposition of the graph Laplacian is required. In this talk, we will present efficient numerical schemes for the calculation of kernels on large graphs that avoid the spectral decomposition of the graph Laplacian. We will focus on two numerical concepts that help to reduce the computational load: the first idea is the usage of Krylov subspace methods tailored to the needs of kernel methods on graphs. These iterative methods approximate relevant parts of the kernels by low-cost matrix polynomials of the graph Laplacian. The second idea exploits local kernel approximation on subdomains of the graph via a partition of unity method that merges the local approximants. We will present some theoretical aspects of these cost-efficient numerical schemes as well as numerical tests demonstrating their benefits.

Wolfgang Erb

Department of Mathematics, University of Padova
erb@math.unipd.it

MS363

Efficient Kernel Methods with Localized Bases

Despite enjoying very good convergence properties, kernel methods often suffer from high computational costs as the size of the underlying problem grows. This talk will serve as an overview of this challenge, and will present one possible remedy: the use of localized bases to accelerate and stabilize kernel methods. This will include a recent history of localization of kernels and conclude with very recent progress and accompanying improvements to a variety of kernel based methods for solving PDEs on manifolds. These include new stable and efficient multigrid methods for Galerkin problems using localized kernel bases, as well as improvements to kernel-FD and asymmetric collocation for time dependent and independent PDEs. The talk is based on recent joint work with Christian Rieger and other researchers who will be acknowledged during the talk.

Thomas C. Hangelbroek

Department of Mathematics
University of Hawaii at Manoa
hangelbr@math.hawaii.edu

MS363

Dimensionality Reduction and Wasserstein Stability for Kernel Regression

In a high-dimensional regression framework, we study consequences of the naive two-step procedure where first the dimension of the input variables is reduced and second, the reduced input variables are used to predict the output variable. More specifically we combine principal component analysis (PCA) with kernel regression. In order to analyze the resulting regression errors, a novel stability result of kernel regression with respect to the Wasserstein

distance is derived. This allows us to bound errors that occur when perturbed input data is used to fit a kernel function. We combine the stability result with known estimates from the literature on both principal component analysis and kernel regression to obtain convergence rates for the two-step procedure. This talk is based on joint work with Stephan Eckstein and Mathias Trabs.

Armin Iske

University of Hamburg
Department of Mathematics
armin.iske@uni-hamburg.de

MS363

Theoretical Foundations of Kernel Methods on Compact Manifolds

The focus of this talk is the study of non-isotropic kernels on compact Riemannian manifolds with approximation properties similar to those of classically studied radial kernels. We study a class of series kernels applicable for interpolation of arbitrary scattered data on M^d by a linear combination of shifts, where M^d is a compact Riemannian manifold. A class of functions for which the resulting interpolation problem is uniquely solvable for any distinct point sets is the class of strictly positive definite kernels. For kernels possessing a certain series expansion in eigenfunctions of the Laplace-Beltrami operator on M^d , we derive a characterisation of this class. We give more specific results for two-point homogeneous manifolds and the sphere. For the special case of S^{d-1} , we study the connection between geometric properties of the kernel and the coefficients in the series expansion. Examples of such geometric properties are axial-symmetry and invariance under parity.

Janin Jäger

University of Giessen
janin.jaeger@math.uni-giessen.de

MS364

The Active Flux Method for Cartesian Grids with Adaptive Mesh Refinement

We present the first implementation of the Active Flux method on Cartesian grids with adaptive mesh refinement. The Active Flux method is a finite volume method for hyperbolic conservation laws, previously introduced by Eyemann, Roe and coauthors, which uses a continuous, piecewise quadratic reconstruction and Simpson's rule to compute numerical fluxes. While classical finite volume methods only use cell average values of the conserved quantities as degrees of freedom, the Active Flux method uses also point values at grid cell interfaces at the current time as well as at later time levels. These point values together with the cell average value are also used to compute the reconstruction. The resulting method is third order accurate, has a compact stencil in space and time and good stability properties. We implemented the adaptive Active Flux method as a new solver in ForestClaw, a software for parallel adaptive mesh refinement based on a quadtree approach. This is joint work with Donna Calhoun (Boise) and Christiane Helzel (Düsseldorf)

Erik Chudzik

Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany

erik.chudzik@hhu.de

MS364

A Relaxation Model for the Approximation of the Non-Isothermal Navier-Stokes-Korteweg Equations Using High-Order Methods

The non-isothermal Navier-Stokes-Korteweg (NSK) equations are a classical diffuse interface model, which governs two-phase flow that include surface tension and phase transition effects. However, if the numerical discretization of the NSK system is aimed, several issues have to be resolved. First, the non-convex bulk Helmholtz free energy induces a non-monotonous pressure function, which in turn results in mixed hyperbolic-elliptic first-order fluxes. Secondly, the gradient dependence of the underlying energy potential generates a third-order derivative in the momentum equation, which prevents the straight forward use of boundary conditions at the three-phase contact line.

In this talk, we present the non-isothermal extension of our recently presented relaxation formulation for the isothermal NSK equations [T. Hitz, J. Keim, C.-D. Munz, C. Rohde, A parabolic relaxation model for the Navier-Stokes-Korteweg equations, *Journal of Computational Physics*, 421 (2020) 109714]. The model is derived from principles of classical irreversible thermodynamics and hence inherently consistent with the Second Law of Thermodynamics. For an appropriate choice of the relaxation parameters, the model converges to the non-isothermal NSK equations. The model is discretized by the use of a high-order discontinuous Galerkin spectral element method. We demonstrate an excellent agreement with the original NSK model for several one-, two- and three-dimensional benchmark problems.

Jens Keim

University of Stuttgart, Institute of Aerodynamics and Gasdy
keim@iag.uni-stuttgart.de

Claus-Dieter Munz

Institut für Aerodynamik und Gasdynamik (IAG)
munz@iag.uni-stuttgart.de

Christian Rohde

Institute of Applied Analysis and Numerical Simulation
University of Stuttgart
christian.rohde@mathemtaik.uni-stuttgart.de

MS364

Central-Upwind Schemes with Reduced Numerical Dissipation

Godunov-type central-upwind schemes are quite popular tools as they can be applied as "black-box" solvers for a wide variety of hyperbolic systems of conservation and balance laws. These, however, suffer from relatively large amount of numerical dissipation, which makes them less accurate than the modern upwind schemes. I will introduce two new ways of reducing the amount of numerical dissipation present in central-upwind schemes. Both approaches lead to significant enhancement of the resolution, especially in the areas of contact waves, which are typically heavily affected by excessive numerical diffusion. The new schemes are implemented in the second-order finite-volume and fifth-order finite difference frameworks and applied to one- and two-dimensional compressible Euler equations of

gas dynamics in both single- and multi-fluid cases.

Alexander Kurganov
Southern University of Science and Technology of China
alexander@sustech.edu.cn

MS364

Coupling of Hyperbolic Pdes via Approximating Half-Riemann Problems

Modelling of gas networks, fluid-structure systems and multiphase flows requires development of numerical coupling techniques. Recent works on coupling of hyperbolic systems based on solving two half Riemann problems can be useful but require expensive iterative procedures at every quadrature point. The iterative coupling method is in the same spirit as the finite-volume Godunov scheme. Thus, the most natural step to improve its performance is to approximate the solutions of the half Riemann problems using techniques established for finite-volume methods. In this work we investigate multiple approximate-state and the Roe Riemann solvers coupling systems of hyperbolic conservation laws.

Aleksey Sikstel
University of Cologne, Germany
a.sikstel@uni-koeln.de

MS364

Quinpi: Constructing Implicit High Order Schemes for Hyperbolic Systems

Many interesting applications of hyperbolic systems of equations are stiff, and require the time step to satisfy restrictive stability conditions. One way to avoid small time steps is to use implicit time integration. Implicit integration is quite straightforward for first order schemes. This talk is concerned with the challenges of writing high order implicit schemes for hyperbolic systems. High order schemes need to control spurious oscillations, which requires limiting in space and time also in the implicit case. We propose a framework to simplify considerably the application of high order non oscillatory schemes through the introduction of a low order implicit predictor, which is used both to set up the nonlinear weights of a standard high order space reconstruction, and to achieve limiting in time. In this talk, we concentrate on the case of a third order scheme, based on DIRK integration in time and CWENO reconstruction in space.

Giuseppe Visconti
Department of Mathematics "G. Castelnuovo"
Sapienza University of Rome
giuseppe.visconti@uniroma1.it

MS365

Electrostatic Solves for Plasma Physics Using PETSc-PIC

We present a PETSc based particle-in-cell (PIC) method for studying electrostatic plasmas. The PETSc-PIC algorithm is a highly scalable method, with multigrid capabilities, for solving the Vlasov-Poisson equations. In the PETSc-PIC algorithm, the Vlasov equation is solved using a particle representation while a continuum representation is employed for the Poisson solve. We implement our PETSc-PIC method on two example systems: two stream instability and Landau damping. These systems are standard tests in plasma physics and provide a baseline to show

the effectiveness and correctness of our methodology. Furthermore, we compare the implementation of the traditional Poisson equation with H^1 finite elements to a mixed form Poisson method with H^{div} finite elements.

Daniel Finn
University at Buffalo
dsfynn@buffalo.edu

MS365

Two-Fluid Simulations of Ultracold Neutral Plasmas in the Absence of External Magnetic Fields

Ultracold neutral plasmas (UCNPs), created by the photoionization of a cold gas, are an excellent platform for studying neutral plasmas in far more complex environments such as plasma in the Sun's atmosphere, white dwarf stars, and inertial-confinement fusion devices. UCNPs are highly versatile systems with ion temperatures below 1 K, tunable electron temperatures from 1-1000 K, and densities ranging from 10^6 - 10^{12} cm^{-3} and have proven to be an ideal testbed for a wide variety of numerical plasma models due to their clean and controllable initial conditions and precise diagnostics. In this talk, I will discuss our recent progress developing a multi-physics, two-fluid (2F) plasma code for the purpose of simulating plasmas spanning the ultracold in a laboratory environment to the aftermath of interchange reconnection in the solar corona. The code is validated in the absence of external magnetic fields by comparing 2F simulations of the evolution of a UCNPs with various initial conditions to experimental data and kinetic theory, including the evolution of a UCNP with a spherically symmetric Gaussian density distribution, the propagation of planar ion holes, and the evolution of ion acoustic waves excited by a periodic density perturbation. Research supported by the NSF through CAREER award AGS-1450230.

Grant Gorman, Gregory Szytko, Stephen Bradshaw
Rice University
gg25@rice.edu, gms12@rice.edu,
stephen.bradshaw@rice.edu

MS365

Subspace Acceleration Method for Simulation of Plasma in Tokamak Boundary

Large-scale simulations of the plasma boundary are key to the understanding and improvement of tokamak fusion experiments. A collaboration with the Swiss Plasma Center institute led to the development of a new method in order to accelerate the solution of consecutive linear systems which can arise from such simulations. A system of differential algebraic equations is obtained from the 3D drift-reduced Branginskii equations that model plasma dynamics in the edge layer of the tokamak. The explicit time integration of such a system yields the need for solving a sequence of consecutive linear systems, for which both the matrix and the right-hand-side are time-dependent. As the time step is constrained to remain small, a very large number of systems need to be solved, dominating the computational time. In order to mitigate this, we leverage the history of previous time solutions in order to reduce, sometimes drastically, the number of iterations needed by the linear systems solver. Taking inspiration by existing methods for implicitly discretized nonlinear evolution problems, our approach is based on combining the solution of a reduced order model with randomization. The effectiveness of our method is demonstrated for academic problems as well as

large-scale plasma simulations.

Margherita Guido

École polytechnique fédérale de Lausanne (EPFL)
margherita.guido@epfl.ch

MS365

Modeling Interchange Reconnection in the Solar Corona

Interchange Reconnection (IR) is a magnetic reconnection scenario which occurs at the boundary between open- and closed-field regions of the Sun's atmosphere. When IR connects a hot, dense Active Region to the cooler plasma surrounding it, the resulting discontinuity can lead to field-aligned shock formation. In this talk, we detail our efforts to simulate this aftermath of IR using our generalizable fluid plasma code. Specifically, we consider the effect of non-neutrality on post-IR shock formation and propagation through a multi-fluid magnetohydrodynamic approach. We also discuss the relevance of kinetic-scale behavior to such reconnection scenarios. This work was supported by the NSF through CAREER award AGS-1450230.

Gregory Szytko, Stephen Bradshaw, Grant Gorman
Rice University
gms12@rice.edu, stephen.bradshaw@rice.edu,
gg25@rice.edu

MS366

Electron Transport in Bilayer Graphene

In the last decades graphene, thanks to its peculiar characteristics, has attracted the attention of researchers and engineers and opened new ways for the realization of a new generation of electronic devices. Nevertheless, being a gapless conductor in its pristine form, it cannot be used for modeling Field Effect Transistors (FET), consequently other solutions have to be investigated to overcome the problem. A solution is the use of graphene nano-ribbons which exhibit a band gap inversely proportional to the width of the strip, another solution is offered by the use of bilayer graphene where a tunable band gap can be induced by applying an electric field. Here we propose a preliminary study for a FET, with graphene bilayer in the active area, by using a Monte Carlo simulation.

Vito Dario Camiola
University of Catania, Italy
dario.camiola@unict.it

Giorgia Vitanza
University of Catania
giorgia.vitanza@phd.unict.it

MS366

Modelling the Effect of Deep and Shallow Trap States on Polarization and Depolarization of High Voltage Polymeric Insulation

We propose a model to describe polarization and depolarization currents observed in thin samples of polymeric materials under the application of an electric field of magnitude comparable to the working conditions of HVDC power transmission cables. The model requires a reduced number of parameters with respect to commonly used transport models. A single carrier with negative charge is consid-

ered. The polarization curves are empirically modeled by a power law to account for transport assisted by shallow traps. A simplified model for deep traps with Gaussian distribution of the density of states is proposed to describe the depolarization curves. Numerical methods for solving the model equations and for fitting its parameters to measurement data are discussed. The validity of the model is confirmed by excellent agreement with experiments.

Carlo De Falco
MOX Modeling and Scientific Computing
Politecnico di Milano
carlo.defalco@polimi.it

Pasquale Cambareri, Luca Di Rienzo
DEIB Politecnico di Milano
pasquale.cambareri@polimi.it, luca.dirienzo@polimi.it

Gian Carlo Montanari
Center for Advanced Power Systems
Florida State University
gmontanari@fsu.edu

Paolo Seri
Dip. Ingegneria dell'Energia Elettrica e dell'Informazione
Università di Bologna
paolo.seri2@unibo.it

MS366

Ripples in Free Standing Graphene

Suspended graphene sheets are covered with ripples of size ranging from 5 to 10 nm and about 1 nm height believed to be generated by coupling between the bending modes of the sheet and electronic density. We present an atomistic model to calculate the free energy that combines a quantum tight-binding Hamiltonian of electrons in a honeycomb lattice coupled to a classical nonlinear elasticity theory of von Krmn plates at finite temperature. The strength g of the electron-phonon coupling controls a transition from a stable flat configuration of the sheet to a stable rippled phase. We use a scaling procedure that allows us to effectively reach larger sheet sizes. By numerically solving the model for different size scaling, k , we find the critical value of the parameter, g_c , and the stable configuration of graphene. For any value of k , renormalized elastic constants produce a critical coupling, $g_c(T, k)$, which is a decreasing function of temperature. As the sheet size measured by k tends to infinity, $g_c(T, k)$ decreases to a limiting value, $g_c(T)$, approximately as $k^{0.5}$. In this limit, rippling is a second order phase transition with a susceptibility that scales as $|g_c(T, k) - g_c(T)|^{-\gamma}$, with critical exponent γ close to 0.25.

Antonio Fernández Caballero
Department of Mathematics; Universidad Carlos III de Madrid
antonife@math.uc3m.es

Luis Bonilla
Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience & IndMath
bonilla@ing.uc3m.es

MS366

Charge Transport in Graphene and Application to

GFETs

Device engineers devote considerable effort for developing transistor designs in which short-channel effects are suppressed and series resistances are minimized. Scaling theory predicts that a field effect transistor (FET) with a thin barrier and a thin gate-controlled region will be robust against short-channel effects down to very short gate lengths. The possibility of having channels that are just one atomic layer thick is perhaps the most attractive feature of graphene for its use in transistors. Main drawbacks of a large-area monolayer graphene are the zero gap and, for graphene on substrate, the degradation of the mobility [G. Nastasi, V. Romano, A full coupled drift-diffusion-Poisson simulation of a GFET, *Commun. Nonlinear Sci. Numer. Simulat.*, Vol. 87 (2020), 105300]. In [G. Nastasi and V. Romano, An efficient GFET structure, *IEEE Transaction on Electron Devices* 68 (9) 2021, 4729] the authors propose a new geometry which seems very promising to get an efficient Graphene Field Effect Transistor (GFET). The simulations are based on a degenerate drift-diffusion model. Here the same GFET is simulated by directly solving the semiclassical Boltzmann equation by using a Discontinuous Galerkin approach.

Vittorio Romano

University of Catania, Italy
romano@dmi.unict.it

Giovanni Nastasi

University of Catania
giovanni.nastasi@unict.it

MS366**Micro-Macro Markov Chain Monte Carlo Using Reaction Coordinate Proposals for Molecular Dynamics**

In many problems and applications in molecular dynamics, one is typically interested in sampling from the time-invariant probability distributions of the system. Within many of these molecular systems, there is a natural time-scale separation present between the fast (microscopic) dynamics, and the slowly changing global conformation of the molecule determined by a few (macroscopic) variables. Markov chain Monte Carlo (MCMC) is a general sampling algorithm that has been designed for sampling from probability distributions determined by a potential energy function. The objective is to construct a stochastic process whose time-invariant distribution is the invariant distribution of the molecular system. By ergodicity, we can record samples from a single path of the process, and then these samples are consistent with invariant distribution. However, when the underlying system has a medium to large time-scale separation, the MCMC method may remain stuck in one of the local minima of the potential energy function, prohibiting a swift exploration of the complete state space of the molecular system. In my talk, I will present a new micro-macro MCMC method (mM-MCMC) in which we first sample from the macroscopic variables, before reconstructing a new molecular instance that is consistent with the macroscopic value. I will give a detailed explanation of the algorithm, and show its efficiency on two molecular examples if time permits.

Hannes Vandecasteele

University of Leuven
hannes.vandecasteele@cs.kuleuven.be

Giovanni Samaey

KU Leuven

giovanni.samaey@kuleuven.be

MS367**A Goal-Oriented Adaptive Hierarchical Method for Random Elliptic PDEs**

This talk presents a goal-oriented adaptive hierarchical method for approximating deterministic, real-valued, bounded linear functionals that depend on the solution of a linear elliptic PDE with a lognormal coefficient field and geometric singularities in bounded domains of \mathbb{R}^d . Our method adopts a hierarchy with its levels defined by a geometrically decreasing sequence of error tolerances; in contrast to standard multilevel methods, where a hierarchy is chosen based on mesh-element or time-stepping sizes. We use isoparametric d-linear quadrilateral finite element approximations. Consider a fixed error tolerance: for each realization of the random coefficient, we select an associated mesh from a sequence of deterministic, non-uniform auxiliary meshes determined by a goal-oriented adaptive algorithm based on a dual-weighted residual error representation. We demonstrate our approach for a random elliptic PDE problem with a geometric singularity similar to that at the tip of a slit modeling a crack. The source of randomness is a lognormal coefficient field, based on a Fourier expansion with random coefficients. Our adaptive approach performs well for the lognormal case treated here, despite the lack of uniform coercivity that cause functional outputs to vary over orders of magnitude between outcomes.

Joakim Beck

King Abdullah University of Science and Technology (KAUST)
joakim.beck@kaust.edu.sa

Yang Liu

King Abdullah University of Science and Technology
yang.liu.3@kaust.edu.sa

Erik von Schwerin

King Abdullah University of Science and Technology
Department of Applied Mathematics and Computational Science
erik.vonschwerin@kaust.edu.sa

Raul F. Tempone

Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

MS367**Boosting Variance Reduction with Meta Multifidelity Estimators**

This work introduces meta estimators that combine multiple multi-fidelity techniques based on control variates, importance sampling, and information reuse, to yield a quasi-multiplicative amount of variance reduction. The proposed meta estimators are particularly efficient within outer-loop applications, such as reliability-based design and shape optimization, when the input distribution of the uncertainties changes with the outer-loop iteration. We derive asymptotic formulae for the variance reduction of the meta estimators in the limit of convergence of the outer-loop applications. We demonstrate the meta estimators, using data-driven surrogates and biasing densities, on an uncertainty propagation problem from nuclear fusion, namely the esti-

mation of energetic particle confinement statistics during a stellarator coil design optimization loop. The numerically measured speedups agree well with the asymptotic variance formulas. The meta estimators outperform all of their constituent variance reduction techniques alone, ultimately yielding two orders of magnitude speedup compared to standard Monte Carlo estimation at the same computational budget.

Frederick Law

Courant Institute of Mathematical Sciences
New York University
fl1284@nyu.edu

Antoine Cerfon
Courant Institute NYU
cerfon@cims.nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

Florian Wechsung
Courant Institute of Mathematical Science
wechsung@nyu.edu

MS367

Multifidelity Methods to Actively Learn from Diverse Model Ensembles

Optimization problems characterize a variety of decision making processes in engineering, whenever a choice e.g. about a design configuration or an operational action is expected as an outcome. Simulation-based optimization methodologies have been proven tremendously powerful to support these processes for applications across all engineering domains. Nevertheless, the computational cost associated with search and assessment of the best final choice can be very high and, in turn, limit the use of these methods when problem complexity increases. With focus on the general case of black-box optimization, this talk will discuss methods to actively learn from diverse model ensembles to inform the decision processes. Different computational strategies will be considered ranging from established frameworks for multifidelity Bayesian optimization to novel original formulations targeting more efficient use of both the available models and the computational resources. Methods will be demonstrated over a set of established analytical benchmarks; applications to engineering design problems will also be presented to provide discussion elements for the adoption of these methods to address real-world optimization challenges.

Laura Mainini

Politecnico di Torino
laura.mainini@polito.it

MS367

UM-Bridge: Coupling Models and Methods for HPC-Scale Uncertainty Quantification

We present UM-Bridge, a new, universal, language-agnostic interface for coupling uncertainty quantification (UQ) and model software. At the example of a parallelized multilevel Markov chain Monte Carlo (MLMCMC) method inferring the source of a tsunami, we illustrate the numerous challenges to methods and software arising in UQ in high-performance computing (HPC) applications. Our

new interface addresses the technical challenges, making it easy to couple any UQ software to any model software. In addition, it massively simplifies both the UQ software for HPC-scale problems as well as the required parallel scaling of numerical model software by optionally leveraging modern cloud environments. Finally, we present a new library of UQ models and benchmarks based on UM-Bridge. It is easily accessible from virtually any programming language and provides portability and reproducible results through containerization.

Linus Seelinger

Heidelberg University
Institute for Applied Mathematics
mail@linusseelinger.de

Anne Reinarz
Durham University
anne.k.reinarz@durham.ac.uk

Robert Scheichl
Universitaet Heidelberg
r.scheichl@uni-heidelberg.de

MS367

Multifidelity Methods for Uncertainty Quantification of Pollutant Dispersion Models Using Control Variates with Estimated Means

Air quality is a pressing public health issue facing modern society. The World Health Organisation (WHO) estimate 4.2 million deaths are linked annually to ambient air pollution. This issue having a huge impact on society leading to public measures of frustration such as current attempts by a group of citizens sue the German government over its lack of progress in meeting WHO recommendations on air quality levels presented in 2021. Understanding the key factors driving air pollution the combination of accurate pollution measurement techniques as well as detailed computational models of the underlying physical processes. The uncertainty in the model input leads to uncertainty in the model output. To quantify this uncertainty we use a Monte-Carlo (MC) approach based on random samples from the input parameter PDFs. To improve the computational efficiency of the MC method we consider a range of multifidelity computational models and use a control variate approach to minimise the mean squared error of estimate of the mean of the model output. The classical control variate approach assumes known values for the means of the low fidelity models. In this talk we present results of techniques used to estimate these means and the optimal strategy of combining the multifidelity models.

Stephen Smith

University of Strathclyde
stephen.smith.100@strath.ac.uk

MS368

Handling Corruption in CG and S-Approximate Conjugate Directions

Solves are a critical operation used in applications ranging from machine learning to simulations. As computational architectures continue to diversify, particularly with increased parallelization, highly asynchronous solvers are required. Furthermore, a need arises for fault tolerant methods that can natively handle unreliability present in these new environments. However, for linear systems, few asynchronous methods have been the focus of Algorithm Based

Fault Tolerance including a robust Push-Sum algorithm, a robust Alternating Directions Method of Multipliers, and Asynchronous Jacobi. In this presentation, we motivate the need for handling corruption in the Conjugate Gradient Method (CG) and an asynchronous method based on CG known as s-Approximate Conjugate Directions (s-ACD). S-ACD aims to leverage the power of CG through local orthogonality constraints while using Conjugate Directions globally to reduce synchronization points and the s-step framework to further reduce communication. This presentation investigates various corruption schemes, including corruption during communication. We demonstrate how the analogs of the orthogonality constraints evolve during the solves and corruption. Using these metrics, we propose a detection scheme, and motivate correction schemes. Finally, we demonstrate the effectiveness of these methods through numerical experiments. Prepared by LLNL under Contract DE-AC52-07NA27344. Funded by LLNL LDRD project 22-ERD-045. LLNL-ABS-840243

Lucas Erlandson, Alyson Fox
Lawrence Livermore National Laboratory
erlandson3@llnl.gov, fox33@llnl.gov

Colin Ponce
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
ponce11@llnl.gov

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

Agnieszka Miedlar
Virginia Tech
amiedlar@vt.edu

MS368

Optimal Size of the Block in Block GMRES on GPUs: Computational Model and Experiments

The block version of GMRES (BGMRES) is most advantageous over the single right hand side (RHS) counterpart when the cost of communication is high while the cost of floating point operations is not. This is the case on modern Graphics Processing Units (GPUs), while it is generally not the case on traditional Central Processing Units (CPUs). In this talk, experiments on both GPUs and CPUs are shown that compare the performance of BGMRES against GMRES as the number of RHS increases. The experiments indicate that there are many cases in which BGMRES is slower than GMRES on CPUs, but faster on GPUs. Furthermore, when varying the number of RHS on the GPU, there is an optimal number of RHS where BGMRES is clearly most advantageous over GMRES. A computational model is developed using hardware specific parameters, showing qualitatively where this optimal number of RHS is, and this model also helps explain the phenomena observed in the experiments.

Andrew J. Higgins
Temple University, U.S.
andrew.higgins@temple.edu

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Daniel B. Szyld
Temple University
szyld@temple.edu

MS368

Asynchronous Chebyshev Methods

Iterative methods typically contain many synchronization points which can scale poorly on massively parallel computers. Additionally, these synchronization costs can be further amplified when some cores progress more slowly than others, e.g., when using heterogeneous computers or when faults occur. While asynchronous methods have recently gained interest, asynchronous versions of certain fast-converging iterative solvers have yet to be developed. We present the first asynchronous Chebyshev method which combines an asynchronous formulation of Chebyshev with an asynchronous version of the BPX multigrid method. Our initial experiments show that using asynchronous Jacobi within asynchronous Chebyshev can result in divergence in many situations, including for large problem sizes. This indicates the need for a preconditioner that provides grid-size independent convergence, which is why we use a multigrid preconditioner. We present experimental results from an OpenMP implementation of asynchronous Chebyshev.

Jordi Wolfson-Pou
Lawrence Berkeley National Laboratory
jwolfsonp@lbl.gov

MS369

Adaptive Nonlinear Domain Decomposition Preconditioning Strategies for Robust and Efficient Simulation of Field-Scale Subsurface Flow

Macro-scale models of subsurface flow are usually simulated using implicit methods, which lead to large, stiff, and strongly coupled systems of discrete nonlinear equations. Newton-type methods work well as long as the coupling/nonlinearity is balanced in space, but may end up wasting a lot of computational effort when the dominant nonlinearity is localized. A better alternative is then to use a nonlinear domain-decomposition preconditioning strategy (e.g., ASPIN) that first resolves challenging nonlinearities in a local stage before accounting for long-range iterations in a global, Newton-like stage. This improves robustness to time-step sizes and initial guess but also introduces higher computational cost per global iteration compared with standard Newton. In this talk, we will discuss methods to hybridize the two approaches, so that standard Newton is used when this method converges well, with an adaptive switch to ASPIN when faced with severe local nonlinearities. We demonstrate the efficacy of the hybrid method on examples carefully designed to stress-test the nonlinear solver as well as on real, industrial simulation models.

Øystein Klemetsdal
SINTEF Digital
oystein.klemetsdal@sintef.no

Olav Møyner
SINTEF ICT
olav.moyner@sintef.no

Knut-Andreas Lie
SINTEF ICT, Dept. Applied Mathematics

Knut-Andreas.Lie@sintef.no

MS369

A Nonlinear Preconditioning Strategy for Solving Phase-Field Fracture Problems

One of the state-of-the-art methodologies to model crack propagation is the phase-field approach. In this approach, a new variable called phase-field is introduced which characterizes the state of the material from intact to fully broken. Modeling such problems is computationally challenging due to the highly nonlinear, non-convex, and non-smooth nature of the underlying energy function. In this work, we aim to solve such problems, from the perspective of the optimization methods. We propose to solve the arising nonlinear problems efficiently using the domain decomposition methods. For the domain decomposition approach, we take advantage of the field-split method to decouple the problem into two sets that are related to the displacement and the damage variable. This split allows us to take advantage of the additive and multiplicative variants of the Schwarz preconditioned inexact Newton (ASPIN/MSPIN) method to solve the nonlinear system. In this work, we will demonstrate the convergence behavior of the proposed solution strategies using several benchmark problems, and also discuss an approach to handle the variational inequalities. We will also compare these methods with the widely-used alternate minimization method and show the benefits of our solution schemes.

Hardik Kothari

Università della Svizzera italiana
hardik.kothari@usi.ch

Alena Kopanicakova
Università della Svizzera italiana
Switzerland
alena.kopanicakova@brown.edu

Rolf Krause
Università della Svizzera italiana, Switzerland
rolf.krause@usi.ch

MS369

Adaptive Nonlinear Domain Decomposition Methods

Highly scalable and robust Newton-Krylov domain decomposition approaches are widely used for the solution of nonlinear implicit problems. In these methods, the nonlinear problem is first linearized and afterwards decomposed into subdomains. By changing this order, that is, by first decomposing the nonlinear problem, many nonlinear domain decomposition methods have been designed in the last decades. These methods often show a higher robustness and a faster nonlinear convergence compared with classical Newton-Krylov variants. In this talk, Nonlinear-FETI-DP (Finite Element Tearing and Interconnecting - Dual Primal) will be discussed. Here, the choice of the nonlinear elimination set, that is, the set of degrees of freedom which are eliminated in a nonlinear fashion, and the choice of the coarse space have a huge impact on the nonlinear and linear convergence behavior. In this talk, we will show new results combining recently developed approaches for the adaptive choice of the nonlinear elimination set with adaptive coarse spaces.

Martin Lanser

Universitaet zu Koeln

Mathematisches Institut
martin.lanser@uni-koeln.de

Axel Klawonn
Universität zu Köln
axel.klawonn@uni-koeln.de

MS369

Cost Efficient Preconditioners for Vector Field Problems

When designing preconditioners based on domain decomposition methods, the coarse space plays a key role. In order to keep the scalability, the coarse space of low computational complexity is essential. We introduce a new coarse space of reduced dimension for vector field problems. Numerical results for the problems in three dimensions are also presented.

Duk-Soon Oh

Chungnam National University
duksoon@cnu.ac.kr

MS370

Analysis of Latent States Modeling of Biological Systems Using Hybrid Differential Equation Recurrent Neural Network Models

Physiological dynamic models are often represented using a system of ordinary differential equations (ODEs). The system of ODEs modeling a dynamic system frequently relies on measurements from multiple nodes. In realistic biological systems, it might not be possible to measure these nodes for a variety of reasons. Incorporating latent states in the model helps to compensate for the lack of measurements in such cases. In this talk, we utilize cubature Kalman filters and recursive Bayesian state estimation with a hybrid ODE-Neural Network model to approximate the output of missing ODEs and measurements. We apply this approach to two models of biological systems: a model of compartmental pressures in the human retina and the Hodgkin-Huxley model for neuron action potentials. Both of these models include states which cannot be readily measured, and we demonstrate how our approach can produce accurate estimates in spite of this hurdle.

Elie Alhajjar

USMA
USMA
eliealhajjar@gmail.com

MS370

Transfer and Multi-Task Learning in Physics-Based Applications with Deep Neural Operators

Traditional machine learning algorithms are learned in isolation, i.e., a predictive model is trained for a single task. In cases where multiple tasks are considered, this learning approach can be computationally prohibitive. Furthermore, when only few and insufficient labeled data are available for a given task, training a model from scratch might lead to overfitting. Transfer learning allows us to leverage information from a model trained on a source domain with sufficient labeled data and transfer it to a different but closely related target domain for which only a small number of data is available. We propose a new TL framework for task-specific learning of partial differential equations (PDEs) under multiple domains that are heterogeneous

but subtly correlated, based on the deep operator network (DeepONet). After the training of a source operator regression model, additional given tasks are learned via the fine-tuning of task-specific layers based on a hybrid loss function that allows for the matching of individual target samples while also preserving the global properties of the conditional distribution of target data. The task-specific training is based on the conditional embedding operator theory where conditional target distributions are embedded onto a reproducing kernel Hilbert space. We demonstrate the advantages of our approach for various TL scenarios involving nonlinear PDEs under diverse conditions due to shift in the geometric domain and model dynamics.

Katiana Kontolati
Johns Hopkins University
kontolati@jhu.edu

Somdatta Goswami
Brown University
somedatta_goswami@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

Michael D. Shields
Johns Hopkins University
michael.shields@jhu.edu

MS370

RaISE: A Framework to Characterize Surrogate Models in Scientific Machine Learning

We present a novel and unifying framework of characterizing surrogate models for the emerging field of scientific machine learning (SciML). In this context, the computational cost of data collection dominates the cost of surrogate construction or evaluation. Accordingly, we define *robustness*, *scalability*, and *efficiency* in terms of *accuracy*. These three concepts then collectively inform a user-specific notion of *interpretability*. We apply this framework to methods of both function and operator approximation in order to (i) create a comprehensive and intuitively accessible catalogue for SciML and (ii) demonstrate the utility of our framework in practice.

Zachary Morrow, Bart G. Van Bloemen Waanders
Sandia National Laboratories
zmmorro@sandia.gov, bartv@sandia.gov

John D. Jakeman
Sandia National Labs
jddjakem@sandia.gov

MS370

Learning High-Dimensional Parametric Derivatives with Neural Operators

In this talk we will present efficient strategies for learning high-dimensional derivative information via neural operators. By exploiting low-dimensional information of a high-dimensional map, if it exists, one can both generate and learn high dimensional derivative information where the dominant computational costs can be made independent of the discretization dimensions. Numerical results demonstrate that this additional derivative information improves

the function approximation, and additionally neural operators that are not trained on derivative information are unlikely to produce reliable derivatives with respect to high-dimensional parameters. Numerical examples will demonstrate how these derivative informed neural operators can be used to accelerate the solutions of stochastic optimization problems and high-dimensional inference problems.

Thomas O’Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Dingcheng Luo
Oden Institute for Computational Engineering and Sciences,
The University of Texas at Austin, Austin, TX, USA
dc.luo@utexas.edu

Peng Chen
The University of Texas at Austin
peng@cc.gatech.edu

Umberto Villa
Oden Institute for Computational Engineering & Sciences
uvilla@oden.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS370

Expression Rate Bounds for Neural and GPC Operator Surrogates

Approximation rates are analyzed for deep surrogates of maps between infinite-dimensional function spaces, which arise e.g. as data-to-solution maps for linear and nonlinear partial differential equations. Such surrogates may be used to speed up computations in parameter estimation problems in engineering. We study in particular deep neural surrogate operators for holomorphic maps between separable Hilbert spaces, where the operator inputs are parametrized by stable, affine representation systems such as frames. Additionally, we discuss an interpolation based alternative to this framework, that allows for a deterministic construction and therefore does not require training of the network weights. Algebraic and dimension independent convergence rates are established in both cases.

Lukas Herrmann
RICAM, Austrian Academy of Sciences
lukas.herrmann@oeaw.ac.at

Christoph Schwab
ETH-Zurich
schwab@sam.math.ethz.ch

Jakob Zech
Heidelberg University
jakob.zech@uni-heidelberg.de

MS371

Robust Control Methods for Mean-Field Models with Uncertainty

We present a class of numerical methods for the control

of large-scale interacting agent systems, and their mean-field limit in presence of uncertainties. In particular, we focus on designing robust controls that allow bounding the variance of the controlled system over time, and promote consensus emergence. First, in the spirit of non-linear model-predictive control, we will discuss the adaptive synthesis of feedback controls based on dynamic information on moments of linear mean-field dynamics. Second, we will quantify the robustness of controls deriving bounds in \mathcal{H}_∞ setting for stabilizing linear controllers independent of the number of agents. Numerical realization of the high-dimensional uncertain dynamics is assessed by means of asymptotic mean-field Monte-Carlo methods combined with generalized polynomial chaos expansion. The performance and robustness of the proposed methodologies are assessed through different examples of interacting agent systems.

Giacomo Albi
Università degli Studi di Verona
giacomo.albi@univr.it

MS371

A Multifidelity Approach to Model-Based Bayesian Optimal Experimental Design

In many engineering and scientific applications, experiments are vital but costly and time-demanding. As such, it is often important to find experimental design conditions that maximize the value of these experiments. Model-based Bayesian optimal experimental design (OED) provides a rigorous framework for identifying the ideal design by leveraging a mathematical model that simulates the experimental outcomes. We achieve this by quantifying the value of an experiment using the expected information gain (EIG). This quantity is typically evaluated numerically using a double-nested Monte Carlo (DNMC) estimator. However, DNMC can be prohibitively expensive when using a complex and computationally intensive model. We propose a multi-fidelity approach to accelerate the OED process, in which an ensemble of EIG estimators of varying accuracy and cost (e.g. from lower-fidelity models that use simplified physics, coarsened meshes, or partial convergence) are combined into a single estimator via the approximate control variate method. By leveraging the accuracy-cost tradeoffs of the available models, the multi-fidelity estimator achieves lower variance in estimating the EIG criterion. This approach is demonstrated on a nonlinear benchmark and in designing turbulent flow experiments for inferring the Reynolds-averaged NavierStokes (RANS) turbulence closure model parameters.

Thomas Coons, Xun Huan
University of Michigan
tcoons@umich.edu, xhuan@umich.edu

MS371

Reinforcement Learning Framework for Bayesian Sequential Optimal Experimental Design

Experiments are indispensable for learning and developing models in engineering and science. When experiments are expensive, a careful design of these limited data-acquisition opportunities can be immensely beneficial. Optimal experimental design, while leveraging the predictive capabilities of a simulation model, provides a rigorous framework to systematically quantify and maximize the value of experiments. We focus on designing a finite sequence of experiments, seeking fully optimal design policies (strategies)

that can (a) adapt to newly collected data during the sequence (i.e. feedback) and (b) anticipate future changes (i.e. lookahead). We cast this sequential decision-making problem in a Bayesian setting with information-based utilities, and solve it numerically via policy gradient methods from reinforcement learning. In particular, we directly parameterize the policies and value functions by neural networks thus adopting an actor-critic approach and improve them using gradient estimates produced from simulated design and observation sequences. The overall method is demonstrated on an algebraic benchmark and a sensor movement application for source inversion. The results provide intuitive insights on the benefits of feedback and lookahead, and indicate substantial computational advantages compared to previous numerical approaches based on approximate dynamical programming.

Wanggang Shen, Xun Huan
University of Michigan
wgshen@umich.edu, xhuan@umich.edu

MS371

Bayesian Optimal Experimental Design for High-Dimensional Combustion Models

Many scientific and engineering experiments are developed to study specific questions of interest. Unfortunately, time and budget constraints make operating these experiments at a wide range of conditions intractable, thus limiting the amount of data collected. In this talk, we consider an application of Bayesian optimal experimental design to identify conditions that are expected to be most informative, measured by the expected information gain. We focus on inferring specific parameters of a chemically reacting system subject to both parametric and model uncertainty. A physics-based model was developed to simulate the gas-phase reactions occurring between highly reactive intermediate species in a high-pressure photolysis reactor coupled to a mass spectrometer. The measured time-of-flight mass spectrum evolves in both time and energy producing a high-dimensional output at each design. We discuss how accurate low-dimensional representations of the high-dimensional mass spectrum are necessary for computing the expected information gain. Additionally, we employ Bayesian optimization to efficiently explore over the constrained design space. We present our results, discuss trade-offs during the optimization, and outline a workflow for efficiently computing optimal experimental designs for high-dimensional models.

James Oreluk
Sandia National Laboratories
joreluk@sandia.gov

Leonid Sheps
Sandia National Labs
lsheps@sandia.gov

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

MS371

Machine Learning Strategies for Scaling Up Optimal Experimental Design

Objective-based uncertainty quantification via MOCU (mean objective cost of uncertainty) enables one to mea-

sure the impact of model uncertainty on the operational objective to be achieved based on an uncertain model. Various real-world science and engineering applications involve modeling complex systems based on insufficient data, where the resulting uncertain model may be used for subsequent tasks such as prediction or control. The ability to accurately estimate the impact of the uncertainty present in the current model on the expected performance of these tasks is crucial for optimal experimental design (OED), as it allows one to prioritize experiments that can effectively reduce the model uncertainty that practically matters. While the efficacy of MOCU-based OED has been demonstrated in several problems, the high computational cost for MOCU estimation has been a significant obstacle to its application. In this talk, we present machine learning (ML) strategies that can remarkably accelerate MOCU computation, thereby substantially scale-up MOCU-based OED capabilities.

Hyun-Myung Woo
Brookhaven National Laboratory
hwoo@bnl.gov
larwind@gmail.com

Youngjoon Hong
Sungkyunkwan University
hongyj@skku.edu

Bongsuk Kwon
Department of Mathematical Sciences
Ulsan National Institute of Science and Technology
bkwon@unist.ac.kr

Byung-Jun Yoon
Brookhaven National Laboratory
bjyoon@tamu.edu

MS372

Stochastic Methods for Machine Learning

We develop a backward stochastic differential equation based probabilistic machine learning method, which formulates a class of stochastic neural networks as a stochastic optimal control problem. An efficient stochastic gradient descent algorithm is introduced with the gradient computed through a backward stochastic differential equation. Convergence analysis for stochastic gradient descent optimization and numerical experiments for applications of stochastic neural networks are carried out to validate our methodology in both theory and performance.

Feng Bao
Florida State University
bao@math.fsu.edu

Richard Archibald
Computational Mathematics Group
Oak Ridge National Laboratory
archibaldrk@ornl.gov

Yanzhao Cao
Department of Mathematics & Statistics
Auburn University
yzc0009@auburn.edu

MS372

Reduced-Order Neural Operator Surrogates for

Dynamical System Models

A primary challenge in achieving high accuracy in simulations of physical systems is the expense of resolving the high dimensional and multiscale nature of dynamical interactions. In order to make many-query predictive simulation of such systems tractable, as required for e.g. design optimization and uncertainty quantification efforts, accurate surrogate models are typically required. In this work we employ neural operator surrogates using the deep operator network (DeepONet) architecture to construct surrogates for the full temporal solution of chemical dynamical systems models arising in plasma and reacting flow applications. In particular we explore the opportunity for constructing such surrogates on reduced dimensional representations of the chemical species state, propose network augmentations to overcome expense associated with temporal symmetries arising in the training data, and investigate user-defined temporal modes for accelerating network training. Results are presented for surrogates applied to spatially homogeneous chemistry models relevant to CFD solvers for reacting flows and low-temperature fluid plasmas.

Tiernan Casey
Sandia National Laboratories
tcasey@sandia.gov

Simone Venturi
Sandia National Laboratories
sventur@sandia.gov

MS372

Representing Subgrid-scale Models with Neural Ordinary Differential Equations

We explore a new approach to learning the subgrid-scale model effects when simulating partial differential equations (PDEs) solved by the method of lines and their representation in chaotic ordinary differential equations based on neural ordinary differential equations (NODEs). Solving systems with fine temporal and spatial grid scales is computationally challenging, and closure models are generally difficult to tune. We propose a machine-learning strategy to represent the coarse- to fine-grid map, which can be viewed as subgrid-scale parameterization. Our approach is based on NODEs and partial knowledge of the coarse system; therefore, we learn the source dynamics operator. Our method inherits the advantages of NODEs and can be used to parameterize subgrid scales, approximate coupling operators, and improve the efficiency of low-order solvers. Numerical results are used to illustrate this approach using the two-scale Lorenz 96 equation, the convection-diffusion equation, and the Navier–Stokes equations.

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

Shinhoo Kang
Department of Aerospace Engineering
The University of Texas at Austin
shinhoo.kang@anl.gov

MS372

Sann: Stiffness-Aware Neural Network for Learn-

ing Hamiltonian Systems

In this talk, I will present stiffness-aware neural network (SANN), a new method for learning Hamiltonian dynamical systems from data. SANN identifies and splits the training data into stiff and nonstiff portions based on a stiffness-aware index, a simple, yet effective metric that can be used to quantify the stiffness of the dynamical system. This classification along with a resampling technique allows one to apply different time integration strategies such as step size adaptation to better capture the dynamical characteristics of the Hamiltonian vector fields. SANN has been evaluated using complex physical systems including a three-body problem and billiard model. The results show that SANN is more stable and can better preserve energy when compared with the state-of-the-art methods, leading to significant improvement in accuracy.

Hong Zhang
Argonne National Laboratory
hongzhang@anl.gov

MS373

Leveraging Exascale Computing Resources for Particle-In-Cell on GPU (PIConGPU)

This talk will discuss challenges and success stories while migrating the plasma application, PIConGPU to ORNL Frontier exascale system. PIConGPU is one of the 8 CAAR (Center for Accelerated Application Readiness) codes chosen by ORNL to stress test the hardware and software of Frontier. The talk will also cover recent results and software tools used to analyze performance while preparing PIConGPU for the Frontier system.

Sunita Chandrasekaran
University of Delaware
schandra@udel.edu

MS373

Scalable, Performance Portable Particle-in-Cell Modeling with WarpX

WarpX is an electromagnetic and electrostatic particle-in-cell code that is being developed as part of the Exascale Computing Project. Originally designed for particle accelerator modeling, WarpX has also been used for a variety of other scientific applications, including laser-plasma interaction and astrophysical plasmas. WarpX includes several advanced algorithmic options, including pseudo-spectral analytical time-domain (PSATD) solvers and the ability to operate in Lorentz-boosted reference frame. Powered by the AMReX framework, WarpX can run on a variety of CPU and GPU-based platforms and has been scaled up to some of the most powerful supercomputers in the world, such as Frontier. In this talk, I will give an overview of the WarpX project and discuss some of the optimizations that allowed WarpX to take advantage of these machines.

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

MS373

The Cabana Particle Library and Cabana-Based PIC Applications

We first describe the Cabana library, developed as part of the Exascale Computing Project (ECP) by the Co-design

center for Particle Applications (CoPA). Cabana extends Kokkos, a library for performance portability across parallel architectures (of particular interest, Summit and Frontier), for particle applications. Cabana includes particle data structures, particle algorithms, and particle communication (MPI) capabilities, as well as structured grids with corresponding grid algorithms, grid communication, and particle-grid algorithms. Next, we detail the development of Picasso, a library built on Cabana for PIC algorithms, and PicassoMPM, a new MPM application for additive manufacturing built in collaboration with the ECP ExaAM center. In addition, use of Cabana within the XGC plasma physics application is discussed. Finally, the use of Cabana for PIC algorithmic exploration is described, including sparse grid and collision kernel examples.

Sam Reeve
Oak Ridge National Laboratory
reevest@ornl.gov

MS373

Dynamic Parallelization of Multi-Dimensional Lagrangian Random Walk, Mass-Transfer Particle Tracking Schemes

Lagrangian particle tracking schemes allow a wide range of flow and transport processes to be simulated accurately, but a major challenge is numerically implementing the inter-particle interactions in an efficient manner. Our work improves upon a multi-dimensional, parallelized domain decomposition (DDC) strategy for mass-transfer particle tracking (MTPPT) methods. We show that this can be efficiently parallelized by employing large numbers of CPU cores to accelerate run times. Current, static DDC methods fix each processor's area of responsibility for the duration of the simulation and provide ample speedup when particle density is approximately constant throughout the domain. However, this type of DDC does not guarantee proper load balance in simulations with, say, a highly heterogeneous velocity field. Our algorithm relies on particles that move within each time step, and their corresponding information must be manually moved between processors when necessary. In this work, we investigate and compare two different procedures for dynamically decomposing the domain to address work imbalance amongst processors—one being a moving geometric DDC and the other being a tree-based neighbor search. In theory, each method presents advantages over the other, which prompted this study for the method we should use moving forward. Both new techniques provide significant speedup over their serial versions, and we observe further advantages than their parallelized predecessors.

Lucas Schauer
Colorado School of Mines
Department of Applied Mathematics and Statistics
lschauer@mines.edu

Michael J. Schmidt
Sandia National Laboratories
mjschm@sandia.gov

Nicholas Engdahl
Washington State University
Department of Civil and Environmental Engineering
nick.engdahl@wsu.edu

Stephen Pankavich, David A. Benson
Colorado School of Mines

pankavic@mines.edu, dbenson@mines.edu

Diogo Bolster
University of Notre-Dame
dbolster@nd.edu

MS373

Hipace++: Progress on the Quasi-Static Particle-in-Cell Method on GPU

Simulations of plasma-based accelerators with the particle-in-cell (PIC) method are notoriously expensive, as they require a nanometer-scale resolution over a meter-scale propagation distance. Due to the Courant condition, this scale discrepancy results in a very large (many millions) number of time steps for a production simulation, which can be unworkable. The quasi-static PIC method is a reduced model not subject to the same Courant condition, thus accelerating these simulations by orders of magnitude. Developed in the last few years, HiPACE++ is the first quasi-static PIC code fully running on GPU. In this algorithm, a 3D problem is decomposed into many 2D problems (one per longitudinal cell). HiPACE++ exploits this property to reduce the volume of communications by orders of magnitude and harness the effectiveness of single-GPU FFTs and multi-grid solves. We will discuss the main algorithm tweaks needed to fully port the quasi-static PIC method on GPU (Nvidia or AMD) based on the AMReX library.

Maxence Thevenet
Deutsches Elektronen Synchrotron (DESY), Germany
maxence.thevenet@desy.de

MS374

Quantum Algorithms for Solving Partial Differential Equations

Recently we have developed quantum algorithms for solving the Navier-Stokes equations governing the motion of classical fluids, as well as arbitrary systems of partial differential equations. In this talk we briefly describe the construction of these quantum algorithms and the regimes where they provide a quantum speed-up. We then discuss ongoing work extending these algorithms in multiple directions. First, we describe quantum circuits that implement the quantum oracles appearing in these algorithms so that they are now expressed entirely as quantum circuits. We then discuss how the original spatial discretization based on finite-difference methods can be replaced by: (i) finite-element; (ii) finite-volume; and (iii) spectral methods. Lastly, we discuss the extension of the quantum Navier-Stokes algorithm to problems in magnetohydrodynamics and plasma dynamics. We emphasize that these quantum algorithms open up a large new application area for quantum computing quantum simulation of classical nonlinear continuum systems and fields which have substantial economic and scientific impact.

Frank Gaitan
Laboratory for Physical Sciences
University of Maryland, College Park
fgaitan@lps.umd.edu

MS374

Efficient Quantum Lattice Boltzmann Method Implemented in Qiskit

The emergent field of quantum computing is of increas-

ing interest to the computational scientist and engineer. Specifically the quantum computers capability to work in a space that is exponentially large with the respect of quantum bits available, can lead to some exciting speed ups when properly exploited. In the field of computational fluid dynamics the prospect of working with an exponentially large storage space is particularly interesting, but currently little research is being done on this topic as it cannot yet be realised on today's Noisy Intermediate Scale Quantum (NISQ) computers. In this presentation we introduce a novel start-to-end quantum algorithms for the efficient simulation of the collisionless lattice Boltzmann equation on universal quantum computers. Our approach supersedes earlier work in this direction both in its computational efficiency and its ability to handle all types of flow configurations. We verify the correct functioning of our approach and its true quantum-readiness by providing a reference start-to-end implementation in Qiskit and presenting numerical results for difference benchmark computations. Next to this verification by example we give correctness proofs for the novel building blocks of our algorithm.

Merel A. Schalkers
Delft University of Technology
m.a.schalkers@tudelft.nl

Matthias Moller
Delft Institute of Applied Mathematics (DIAM)
Delft University of Technology
M.Moller@tudelft.nl

MS374

Hybrid Quantum-Classical Reservoir Computing for Thermally Driven Fluid Flows

Quantum computing and machine learning have changed our ways to process data fundamentally in the last years. A combination of both can be a promising way to accelerate data processing in the near future. Here, we discuss the porting of a classical recurrent neural network algorithm in the form of a reservoir computing model to a quantum computer. We simulate the nonlinear chaotic dynamics of Lorenz-type models for a classical two-dimensional thermal convection flow by a hybrid quantum-classical reservoir computing model. The high-dimensional quantum reservoir dynamics are established by universal quantum gates that rotate and entangle the individual qubits of the tensor product quantum state. A comparison of the quantum reservoir computing model with its classical counterpart shows that the same prediction and reconstruction capabilities of classical reservoirs with thousands of perceptrons can be obtained by a few strongly entangled qubits. This strong entanglement is shown to be essential for the performance of the quantum reservoir computing algorithm. Furthermore, the quantum reservoir computing model is implemented on a real noisy IBM quantum computer and compared with ideal quantum simulators. Our work opens the door to model the dynamics of classical complex systems in a high-dimensional phase space effectively with an algorithm that requires a small number of qubits.

Joerg Schumacher
Technische Universitaet Ilmenau
Institute of Thermodynamics and Fluid Mechanics
joerg.schumacher@tu-ilmenau.de

Philipp Pfeffer, Florian Heyder
Technische Universitaet Ilmenau
philipp.pfeffer@tu-ilmenau.de, florian.heyder@tu-

ilmenau.de

MS375

Higher Order Generalized Nystrom Approximation

The growing abundance of data that is naturally represented in a multi-way format, such as high-dimensional arrays of numbers, has become a common occurrence in various fields such as computer vision, natural language processing, and bioinformatics. These types of data often poses a computational challenge as they can be too large to store in memory. To effectively analyze and extract insights from such data, efficient low-rank tensor approximation methods are needed. In this paper, we introduce a higher-order Generalized Nystrom approximation method for low-rank Tucker approximations of tensors. We provide theoretical guarantees for the expected approximation error and computational complexity of our approach, establishing it as promising solution for working with large, multi-dimensional datasets.

Haim Avron

Tel Aviv University
haimav@tauex.tau.ac.il

Uria Mor

Tel Aviv University
Weizmann Institute of Science
uriamor@mail.tau.ac.il

MS375

Block Subsampled Randomized Hadamard Transform for Low-Rank Approximation on Distributed Architectures

In this talk we discuss a novel structured random matrix composed blockwise from subsampled randomized Hadamard transforms (SRHTs). The block SRHT is expected to outperform well-known dimension reduction maps, including SRHT and Gaussian matrices, on distributed architectures with not too many cores compared to the dimension. We discuss the oblivious subspace embedding property of block SRHT transform and outline that the required number of rows is similar to that of the standard SRHT. We then present different experimental results showing the accuracy and the performance on parallel architectures obtained when using this transform.

Oleg Balabanov, Matthias Beaupere

INRIA
oleg.balabanov@inria.fr, matthias.beaupere@inria.fr

Laura Grigori

INRIA
France
laura.grigori@inria.fr

Victor Lederer

INRIA
victor.lederer@inria.fr

MS375

A Fast Randomized Algorithm for Computing the Null Space

In this talk, we examine the effect of sketching on the right singular vectors corresponding to the smallest singular val-

ues of a tall-skinny matrix. We describe a randomized algorithm for approximating the trailing right singular vectors and examine the quality of the solution using multiplicative perturbation theory. For an $m \times n$ ($m \gg n$) matrix, the algorithm runs with complexity $\mathcal{O}(mn \log n + n^3)$ which is faster than the standard $\mathcal{O}(mn^2)$ methods.

Taejun Park, Yuji Nakatsukasa

University of Oxford
taejun.park@worc.ox.ac.uk, nakatsukasa@maths.ox.ac.uk

MS375

Exploring the Potential of Sketched Rayleigh-Ritz in Unrestarted Lanczos and Restarted Davidson Methods

Sketching methods have been increasingly popular for solving large scale least-squares problems. Recently a sketched Rayleigh-Ritz method proposed by Nakatsukasa and Tropp promises to avoid orthogonalization in Krylov methods. This is important for eigenvalue problems where ghost eigenvalues rise quite fast unless some form of reorthogonalization is used in the Lanczos method. In this work, we explore the computational tradeoffs of sketched Krylov methods by providing a high performance implementation within the PRIMME software package and comparing with traditional methods.

Andreas Stathopoulos, Heather Switzer

College of William & Mary
Department of Computer Science
andreas@cs.wm.edu, hmswitzer@wm.edu

MS376

High Resolution Level-Set Methods in Fluids and Materials

Free boundary problems are ubiquitous in science and engineering but their numerical solutions are challenged by the fact that (1) smaller scales influence larger one in a non-trivial manner, which demands efficient adaptive gridding; (2) boundary conditions, including sharp jump conditions, must be imposed on a moving and irregular boundary and (3) nonlinearity requires advanced numerical schemes. I will present a framework that seeks to address these challenges and I will present some of their applications in the field of fluids and materials.

Frederic Gibou

UCSB
fgibou@ucsb.edu

MS376

A Numerically Consistent Unstructured Volume-of-Fluid Discretization for the Two-Phase Momentum Convection with High-Density Ratios

We extend the numerically consistent discretization of the single-field two-phase momentum convection from the ρ LENT method an unstructured Level Set / Front Tracking method to the isoAdvection VOF method. The numerical consistency of the momentum convection term in the discretized single-field Navier-Stokes equations is crucial for the numerical stability when simulating two-phase flows with high density ratios. We use an auxiliary mass transport equation to ensure the numerical consistency. A careful derivation of the unstructured VOF discretization indicates that an exact solution of the volume fraction equation

is equivalent to solving this auxiliary mass equation. However, equation discretization is unavoidable, and it introduces the inequality between the numerical integration of the mass flux $\rho_f F_f$ and the geometrical integration of the phase-specific volume by the VOF method. Two numerically consistent discretizations are possible with the isoAdvection VOF method. First, an auxiliary mass equation can be solved, while carefully dealing with the face-centered density. Second, an appropriate discretization scheme chosen for the momentum convection term stabilizes the solution. Multiple droplet studies with the density-ratio range $\rho^-/\rho^+ \in [1, 10000]$ are used to verify the discretization.

Jun Liu

Technical University of Darmstadt
liu@mma.tu-darmstadt.de

Henning Scheufler

DLR German Aerospace Center, Institute of Space
Systems
henning.scheufler@web.de

Johan Rønby

Mathematics and Physics (IMFUFA), Roskilde University
johan@ruc.dk

Davide Zuzio, Jean-Luc Estivalezes

The French Aerospace Lab ONERA
davide.zuzio@onera.fr, jean-luc.estivalezes@onera.fr

Tomislav Maric

Technische Universität Darmstadt
maric@mma.tu-darmstadt.de

MS376

Challenges and Opportunities in Cut-Cell Grid Generation: An Applied Perspective

Adaptive cut-cells grids are exceptionally useful for simulations of low-temperature filamentary plasma, with applications in high-voltage technology, aerodynamics, and plasma medicine. A common theme in these applications is that they often combine complex geometries with the need for deep mesh refinement. Additional challenges for cut-cell grid generation arise due to: 1) Inefficient representation of implicit functions (e.g. for surface tessellations), and 2) load imbalance. In this talk we discuss how to efficiently represent complex geometries by reconstructing signed distance functions using bounding volume hierarchies. We also discuss approaches for robustly load-balancing the cut-cell generation, and provide simple strategies for dealing with the inherent load imbalance associated with the evaluation of the resulting implicit function. Finally, we provide algorithms for accelerating implicit function evaluations for composite geometries.

Robert Marskar

SINTEF
robert.marskar@sintef.no

MS376

A High-Order Cartesian Grid, Finite Volume Method for Elliptic Interface Problems

In this work we solve the elliptic PDE $-\beta u + \nabla \cdot (\eta \mathbf{F}(u)) = f$ where β , η and f may be discontinuous across an interface. Our scheme builds upon the finite volume, Cartesian grid embedded boundary method developed by Devendran et.

al (2017), which utilizes weighted least-squares interpolation to construct stencils in complex geometries. We extend this technique to 1) enforce jump conditions at an interface and 2) solve elliptic PDEs with variable and discontinuous coefficients. We consider the flux $\mathbf{F}(u) = \nabla(u)$ and demonstrate high-order convergence for various geometries, source terms, and coefficients.

Will Thacher

Lawrence Berkeley Lab
University of California, Berkeley
wthacher@lbl.gov

Hans Johansen

Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

MS376

A 3D High-Order, Multiresolution Immersed Geometry Method for Multiphysics Problems

We present an overview of our work on a 3D high-order, finite-difference based immersed interface method for multiphysics problems, implemented within a multiresolution adaptive grid software framework. In this talk we specifically highlight on three features of our approach. First, we discuss our method to achieve a high-order (third or more) spatial discretization all the way up to the boundary for three-dimensional geometries with both convex and concave features. Second, we will present our two-sided interface treatment which enables the solution of separate physical problems on either side of the interface, which may be coupled through the boundary conditions. Third, we will discuss the implementation of this methodology in our in-house multiresolution adaptive grid solver MURPHY, and the scalability of this implementation on massively parallel compute architectures.

Wim van Rees, James Gabbard, Thomas Gillis

MIT
wvanrees@mit.edu, jgabbard@mit.edu, tgillis@mit.edu

MS377

Data-Driven Reduced Order Models for Nonlinear Optimal Control

Modeling and control of high-dimensional, nonlinear robotic systems remains a challenging task. While various model- and learning-based approaches have been proposed to address these challenges, they broadly lack generalizability to different control tasks, rarely preserve the structure of the dynamics, and are unable to tractably bridge the gap between having accurate, but low-dimensional models. In this work, we propose a new, data-driven approach for extracting low-dimensional models from data using Spectral Submanifold Reduction (SSMR). In contrast to other data-driven methods which fit dynamical models to training trajectories, we identify the dynamics on generic, low-dimensional attractors embedded in the full phase space of the robotic system. This allows us to obtain computationally-tractable models for control which preserve the system's dominant dynamics and better track trajectories radically different from the training data. We demonstrate the superior performance and generalizability of SSMR in dynamic trajectory tracking tasks vis-à-vis the state of the art, including Koopman operator-based

approaches.

John Alora

Autonomous Systems Laboratory
Stanford University
jjalora@stanford.edu

Mattia Cenedese
ETH Zurich
Switzerland
mattiacenedese15@gmail.com

Edward Schmerling
Autonomous Systems Laboratory
Stanford University
schmrlng@stanford.edu

George Haller
ETH, Zurich
Switzerland
georgehaller@ethz.ch

Marco Pavone
Autonomous Systems Laboratory
Stanford University
pavone@stanford.edu

MS377

Piecewise Quadratic Approximation Manifold for Nonlinear Projection-Based Model Order Reduction

A piecewise quadratic approximation manifold is proposed for mitigating the Kolmogorov barrier to projection-based model order reduction (PMOR). After the computed solution snapshots are clustered, each local model is trained in two steps using an approach that builds on the data-driven procedure underlying the construction of traditional (affine) projection-based reduced order models (PROMs). The overall approach is application-independent and linearization-free. Consequently, it is more robust than alternative approaches for the solution of highly nonlinear problems. If for a convection-dominated turbulent flow problem a PMOR method equipped with the traditional piecewise-affine approximation delivers a certain level of accuracy using a local PROM of average dimension n , the proposed PMOR approach produces a piecewise quadratic PROM that delivers a similar level of accuracy using however an average local dimension of roughly $\sqrt{2n}$. It is shown that when implemented within the least-squares Petrov-Galerkin (LSPG) PMOR method and equipped with the energy conserving sampling and weighting (ECSW) hyperreduction method, the proposed piecewise quadratic approximation manifold accelerates the solution of a benchmark turbulent flow problem by more than four orders of magnitude. It also accelerates the solution of such a problem using a traditional PROM by two orders of magnitude.

Joshua Barnett, Charbel Farhat
Stanford University
jb0@stanford, CFarhat@stanford.edu

MS377

Dynamics-Based Model Reduction of Nonlinear Finite Element Problems

The prediction and continuation of steady-state response

poses a serious computational challenge for full-scale nonlinear finite element models of realistic engineering systems. Model reduction techniques based in rigorous dynamical systems theory aim to reduce such nonlinear systems over low-dimensional, attracting invariant manifolds. The low-dimensional dynamics of these invariant manifolds acts as a mathematically exact reduced-order model for the full, high-dimensional system. In particular, the recent theory of Spectral Submanifolds (SSM) allows for such a rigorous model reduction and leads to reliable steady-state response predictions within feasible computation times. Further developments have made the direct computation of such invariant manifolds and their reduced dynamics scalable to realistic, nonlinear finite-element models. In this talk, we review the basics of SSM theory and computation. We demonstrate applications of SSM-based reduced-order models towards the extraction of nonlinear forced response curves in realistic finite-element models of mechanical systems.

Shobhit Jain

TU Delft
shobhit.jain@tudelft.nl

MS377

Sensitivity Analysis of Frequency Response Curves Using a Parametric Reduced Order Model and Harmonic Balance

In MEMS industry, designers often resort the large FE models to characterize new devices. This process is computationally cumbersome, even more so when dealing with nonlinearities due to large deformations. On top of this, as MEMS production processes lead to considerable geometric imperfections, the performance spread of each device must be assessed during the design phase. A popular way to do so is the Monte Carlo method, but it requires thousands of simulations. To alleviate the computational burden of these repeated nonlinear analyses, not be feasible in practice, we recently developed a parametric ROM for shape defects in the geometry. Using the Harmonic Balance method (HB), we were able to compute Frequency Response Curves (FRC) in the matter of minutes. However, even if way more efficient than running the full order model, repeating these simulations thousands of times is still unacceptable. To overcome these difficulties, we propose a HB Sensitivity Analysis, which approximates the nominal solution (i.e. without defects) with a Taylor series. This way, after the nominal solution and the sensitivities are computed, the FRC for each new parameter realization will be given simply through an update. Different strategies to achieve this goal are discussed on a simple nonlinear oscillator, with special care for the closure equation required for the well-posedness of the problem. Finally, the computational benefits of the method are assessed on a FE model of a MEMS device.

Alexander Saccani, Jacopo Marconi
Politecnico di Milano
alexander.saccani@mail.polimi.it,
jacopo.marconi@polimi.it

Paolo Tiso
ETHZ
ptiso@ethz.ch

MS377

Parametric Reduced-Order Modelling of Contact Mechanics for the Efficient Dynamic Simulation of

Bolted Structures with Changes in Pretension

Bolted joints play an important role in mechanical system assemblies to connect different components and materials. The varying contact conditions at the large frictional contact area, which can highly depend on the pretension, inevitably change the properties of joints. Consequently, the dynamic behavior of the system is influenced and it is desirable to predict and assess these influences from a design point of view. However, commonly used finite element models can result in large sizes for complex systems while the contact mechanism requires complicated time-domain iterative solvers, making the computation time-consuming. Therefore, a parametric model order reduction strategy is proposed for the efficient simulation of bolted structures with changes in pretension. This strategy uses a Craig-Bampton rationale, in which the second-order Krylov subspace is applied for reduction and parameterization for pretension. Due to the large dimension of contact surfaces, the contact surface reduction is applied based on the combination of characteristic constraint modes and joint interface modes. To compute the time-domain dynamics of the reduced model, the dual loop contact algorithm for the augmented Lagrange multiplier method in conjunction with Newmark method is employed. The efficiency and accuracy of the proposed reduction method is illustrated by the time domain displacements and detailed contact conditions of a proof-of-concept bolted structure under different pretensions.

Shuyang Zhang, Hendrik Devriendt
 KU Leuven
 shuyang.zhang@kuleuven.be,
 hendrik.devriendt@kuleuven.be

Lucas Van Belle
 KU Leuven, Department of Mechanical Engineering,
 DMMS Core lab, Flanders Make.
 lucas.vanbelle@kuleuven.be

Wim Desmet
 KU Leuven, Department of Mechanical Engineering
 wim.desmet@kuleuven.be

MS378

Massive Graph Analytics in Arkouda

Due to the emergence of massive real-world graphs, whose sizes may extend to terabytes, new tools must be developed to enable data scientists to handle such graphs efficiently. These graphs may include social networks, computer networks, and genomes. In this talk, we present our novel graph package, Arachne, to make large-scale graph analytics more effortless and efficient based on the open-source Arkouda framework. Arkouda has been developed to allow users to perform massively parallel computations on distributed data with an interface similar to NumPy. In this package, we developed a fundamental sparse graph data structure and then built several useful graph algorithms around our data structure to form a basic algorithmic library. Benchmarks and tools were also developed to evaluate and demonstrate the use of our graph algorithms. The graph algorithms we have implemented thus far include breadth-first search (BFS), connected components (CC), k-Truss (KT), Jaccard coefficients (JC), triangle counting (TC), and triangle centrality (TCE). Their corresponding experimental results based on realworld and synthetic graphs are presented. Arachne is organized as an Arkouda extension package and is publicly available on

GitHub (<https://github.com/Bears-R-Us/arkouda-njit>).

David A. Bader
 New Jersey Institute of Technology
 bader@njit.edu

Zhihui Du, Oliver Alvarado Rodriguez
 NJIT
 zhihui.du@njit.edu, oaa9@njit.edu

MS378

GBTLX: High Performance Code Generation of Graph Applications

We propose GBTLX, a code generation framework that translates a linear algebraic graph program written with a high level library into a program that would be written by a domain expert. In the front end, GBTLX captures the set of graph library calls within the program. This is then sent to the code generation system, SPIRAL, for semantic analysis and optimization. If a known algorithm is found, SPIRAL will generate a domain expert algorithm in place of those library calls. The code generation process is built with the same building blocks that makes linear algebraic graph processing so appealing, and uses a rule based system to transform those building blocks into known algorithms. The process also supports various graph representations and hardware architectures. The end result is a single source framework that provides domain knowledge at both the algorithmic and architectural level, while keeping a user-friendly medium, a high level library. We show that the generated code for GBTLX achieves performance comparable to others in the larger community.

Sanil Rao
 CMU
 sanilr@andrew.cmu.edu

Scott McMillan
 SEI/CMU
 smcmillan@sei.cmu.edu

Tze Meng Low, Franz Franchetti
 Carnegie Mellon University
 lowt@andrew.cmu.edu, franzf@andrew.cmu.edu

MS378

HPC Graph Analytics on the OneGraph Model

Current graph databases generally support one of the two graph models: W3Cs Resource Description Framework (RDF) or Labeled Property Graphs (LPG). Each comes with its own modeling features and disjoint query languages. At Amazon Neptune, we are developing a single, unified graph data model that embraces both RDF and LPG, that we named OneGraph (1G). We are also building a scalable computational infrastructure that will support both OLTP and OLAP workloads on 1G model. In this talk, I will present challenges of supporting HPC Graph analytics on this new computational framework.

Umit V. Catalyurek
 Georgia Institute of Technology
 School of Computational Science and Engineering
 umit@gatech.edu

Brad Bebee
 Amazon Neptune Team
 Amazon Web Services

beeb@amazon.com

Dave Bechberger, Willem Broekema, Hal Cooper, Olaf Hartig, Ankesh Khandelwal, Ora Lassila, Kelvin Lawrence, Carlos Manuel Lopez Enriquez, Michael Schmidt, Ronak Sharda, Lefteris Sidorouros, Gabriel Tanase, Bryan Thompson, Gregory Williams

Amazon Web Services
 dbechbe@amazon.com, metawilm@amazon.de,
 halcoope@amazon.com, ohartig@amazon.com,
 ankeshk@amazon.com, ora@amazon.com,
 lakelvin@amazon.com, lemanolo@amazon.de,
 schmdtm@amazon.com, shardars@amazon.com,
 sidirol@amazon.com, igtanase@amazon.com,
 bryant@amazon.com, ngregwil@amazon.com

MS378

Writing Custom GPU Linear Algebra Kernels for Fun and Peak Performance

GPU programs that achieve peak performance are notoriously difficult to write, as they require the author to not only reason about concurrent execution, but also have an appreciable understanding of the underlying microarchitecture. Although vendor libraries provide basic kernels such as the BLAS suite, with modern workloads, it is becoming increasingly critical to accelerate their custom computations that have been codesigned with the application algorithm in mind and are less likely to be supported by vendors out of the box due to their bespoke nature. In this talk, we motivate the need for custom GPU linear algebra kernels by taking our efforts to run all pairs shortest path on Summit and Frontier supercomputers as a case study, and highlight the features of our library, CUDA Algebra for Semirings (cuASR) that emerged from it. Highlighting this as a user story, we go over how it influences the design philosophy for open source vendor libraries that are expected to be extended by users, and we conclude with the new features in NVIDIA CUTLASS and how one can extend its abstractions to author custom kernels that achieve peak performance on modern CUDA hardware.

Vijay Thakkar
 Georgia Tech
 thakkarv@gatech.edu

Richard Vuduc
 Georgia Institute of Technology
 richie@cc.gatech.edu

MS379

Estimating Rare Event Probabilities with Stein Variational Gradient Descent

Stein variational gradient descent (SVGD) is a recently proposed approach to sampling from Bayesian posterior distributions. The method – unlike many variational inference approaches – avoids explicitly parametrizing posterior approximations. Instead, approximations are formulated implicitly as a sequence of measure transports starting from the prior distribution. In this work, we aim at repurposing SVGD for estimating rare event probabilities in the order of 10^{-3} – 10^{-10} . We use importance sampling (IS) and employ SVGD to approximate the optimal importance sampling density. The latter is straight-forward based on a well-known analogy between sampling from Bayesian posteriors and optimal IS distributions: the indicator function of the rare event is treated as the likelihood and estimating the rare event probability amounts to estimating the

marginal likelihood or model evidence. In order to avoid non-smooth (indicator) likelihoods, we employ a smooth approximation of the rare event indicator function. We discuss the choice of the smooth indicator approximation and adaptively choosing the smoothing parameter appearing therein, which yields a rare event pendant to annealed SVGD approaches for Bayesian inference. Several examples are used to benchmark the efficacy of our approach against state-of-the-art methods for estimating rare event probabilities.

Max Ehre
 TU München
 max.ehre@tum.de

MS379

Power System Event Location via DEIM

Real-time event analysis is critical for the successful security and operation of electrical power systems. Gone unchecked, local disturbances can cascade into wide area failures (see e.g., the 2003 Northeastern United States blackout). Data-driven analysis techniques based on low-rank matrix factorizations of system measurements have gained popularity in recent years due to the increased deployment of phasor measurement units (PMUs). In this talk, we introduce a selection procedure for locating the source power system events from real-time data based on the Discrete Empirical Interpolation Method (DEIM). The DEIM selection algorithm parses the dominant singular vectors of the collected measurements in order to select a set of representative row indices, each corresponding to measurements at a particular bus. These selected indices serve as a proxy for identifying the buses most affected by an event, and locate the source of the initial fault with high accuracy. Numerical results involving synthetically generated PMU data illustrate the effectiveness of the proposed approach.

Sean Reiter
 Department of Mathematics, Virginia Tech
 seanr7@vt.edu

Mark Embree
 Department of Mathematics
 Virginia Tech
 embree@vt.edu

Serkan Gugercin
 Virginia Tech
 Department of Mathematics
 gugercin@vt.edu

MS379

Camera: A Method for Cost-Aware, Adaptive, Multifidelity, Efficient Reliability Analysis

Estimating probability of failure in aerospace systems is a critical requirement for flight certification and qualification. We propose a method to use models of multiple fidelities that trade accuracy for computational efficiency. Specifically, we propose the use of multifidelity Gaussian process models to efficiently fuse models at multiple fidelity, thereby offering a cheap surrogate model that emulates the original model at all fidelities. Furthermore, we propose a novel sequential *acquisition function*-based experiment design framework that can automatically select samples from appropriate fidelity models to make predictions about quantities of interest in the highest fidelity.

We use our proposed approach in an importance sampling setting and demonstrate our method on the failure level set estimation and probability estimation on synthetic test functions as well as two real-world applications, namely, the reliability analysis of a gas turbine engine blade using a finite element method and a transonic aerodynamic wing test case using Reynolds-averaged Navier–Stokes equations. We demonstrate that our adaptively constructed multifidelity surrogate model is guaranteed to predict the true failure level set with high probability.

Ashwin Renganathan
University of Utah
ashwin.renganathan@utah.edu

MS379

Adaptive Failure Probability Estimation with Gaussian Processes

Efficiently approximating the probability of system failure has gained increasing importance as expensive simulations begin to play a larger role in reliability quantification tasks in areas such as structural design, power grid testing, and safety certification among others. This work develops an adaptive approximation scheme for efficiently quantifying the probability a simulation exceeds a given failure threshold subject to a prior distribution on system parameters. Our method iteratively updates a Gaussian process by IID sampling from a density proportional to the pointwise expected-weighted-misclassification-rate. Probabilistic error bounds are derived to enable an adaptive stopping criterion. We show this method achieves the same error as standard Monte Carlo techniques using substantially fewer simulations.

Aleksei Sorokin
Illinois Institute of Technology
asorokin@hawk.iit.edu

MS379

Identification of Network Vulnerabilities to Cascading Failures

Measures of power grid vulnerability are often assessed by the amount of direct and instantaneous damage an adversary can exact on the network. The cascading impact of such attacks is often overlooked, even though it is well known that cascades are one of the primary causes of large-scale blackouts. This paper explores modifications of the power network's transmission line flow thresholds as candidates for adversarial attacks, which are undetectable as long as they don't affect line flows at equilibrium. This forms the basis of a black-box function in a Bayesian optimization procedure, where the goal is to find line threshold settings that maximize cascade damage. Notably, the proposed method is agnostic to the choice of underlying cascade simulator. Extensive experiments conducted using a recently developed quasi-steady state cascade simulator reveal that, contrary to what one might expect, the most restrictive thresholds do not necessarily produce the most damaging cascades. More surprisingly, even in a limited resource setting where attacks are resource-constrained, one can find attacks that produce cascades comparable in severity to those found when there are no resource constraints.

Anirudh Subramanyam
Penn State University
subramanyam@psu.edu

Albert Lam
Argonne National Laboratory
albert3033@gmail.com

MS380

Bayesian Sparse Dictionary Learning

Dictionary learning can be seen as a possible data-driven alternative to solve inverse problems by identifying the data with possible outputs that are either generated numerically using a forward model or the results of earlier observations of controlled experiments. Sparse dictionary learning is particularly interesting when the underlying signal is known to be representable in terms of a few vectors in a given basis. In this talk we propose to use hierarchical Bayesian models for sparse dictionary learning that can capture features of the underlying signals, e.g., sparse representation and nonnegativity. The same framework can be employed to reduce the dimensionality of an annotated dictionary through feature extraction, thus reducing the computational complexity of the learning task. Computed examples where our algorithms are applied to hyperspectral imaging and classification of electrocardiogram (ECG) will be presented.

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and Statistic
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

Nathan Waniorek
waniorek@uchicago.edu
nathan

MS380

Sparse Bayesian Inference with Regularized Gaussians

In this talk, we will present a method for Bayesian inference that, unlike many existing Bayesian methods, results in posterior distributions that assign positive probability to sparse vectors. We combine Gaussian distributions with the deterministic effects of sparsity-inducing regularization like l_1 norms, total variation and/or constraints. The resulting posterior distributions assign positive probability to various low-dimensional subspaces and therefore promote sparsity. Samples from this distribution can be generated by solving regularized linear least-squares problems with properly chosen data perturbations. We will discuss some properties of the underlying prior and use this methodology to derive an efficient algorithm for sampling from a Bayesian hierarchical model with sparsity structure.

Jasper M. Everink
Technical University of Denmark
jmev@dtu.dk

Yiqiu Dong, Martin S. Andersen
Technical University of Denmark
Department of Applied Mathematics and Computer Science

yido@dtu.dk, mskan@dtu.dk

MS380

Certified Coordinate Selection for Linear Bayesian Inversion with Laplace Prior

In this talk we are presenting a new and comprehensive method to tackle large-dimensional linear Bayesian inverse problems with Gaussian likelihood and Laplace prior. The unknown parameter is represented on an ad hoc basis the dimension of which can be very large. The inverse problem is then formulated w.r.t. the basis coefficients (coordinates), and we show how to select a possibly small subset of them which is most informative w.r.t. the data relative to the heavy-tailed prior. To this end, we propose a novel technique for the diagnosis of the contribution of each basis coefficient that is based on the MAP of the exact posterior density. In the new framework, we perform likelihood-informed dimension reduction by constructing efficiently a low-dimensional approximation to the likelihood function which accounts for the informed coefficients only. In addition, we provide a tractable upper bound on the resulting approximation error measured with the Hellinger distance. We show how MCMC sampling from the approximated but also exact posterior can be accelerated by using pseudo marginal MCMC sampling that is performed only on the selected set of coordinates. In the end of this talk, we present numerical results from examples in computed tomography and image deblurring that affirm the versatility of our method.

Rafael Flock

Technical University of Denmark
raff@dtu.dk

MS380

Computational Strategies for Bayesian Inversion with Conditionally Gaussian Sparsity Priors

In many practical Bayesian linear inverse problems, important parameters such as the noise covariance and the ideal strength of regularization are unknown a priori. Furthermore, adopting a sparse Bayesian learning (SBL) approach we may wish to employ a sparsity-promoting prior that strongly promotes sparsity in a linear transform of the unknown. It has been shown that both concerns may be addressed using hierarchical Bayesian models that employ conditionally Gaussian priors with various choices of hyper-priors. However, the resulting posterior densities are typically no longer log-concave, which creates convergence concerns when using either optimization or sampling to characterize the posterior distribution. With this in mind, here we study iterative algorithms for posterior point estimation, particularly with regards to more general measurement and regularization settings. In particular, we consider the scenario that the noise covariance of one or several data sources are unknown and made part of the inference procedure, and when sparsity is promoted in multiple transforms of the unknown. We then apply recent data-augmentation techniques to accelerate the methods for large-scale problems such as image-deblurring and synthetic aperture radar (SAR) despeckling. We demonstrate the performance of our methods on a suite of numerical examples.

Jonathan Lindbloom

Dartmouth College

jonathan.t.lindbloom.gr@dartmouth.edu

MS380

Varying Map Estimators with Varying Assumptions in Sparsity Promoting Gaussian Hierarchical Models

Maximum a posteriori (MAP) estimation, like all Bayesian methods, depends on prior assumptions. These assumptions are often chosen to promote specific features in the recovered estimate like sparsity. The form of the chosen prior determines the shape of the posterior distribution, thus the behavior of the estimator, and the complexity of the associated optimization problem. Here, we consider a family of Gaussian hierarchical models with generalized gamma hyperpriors designed to promote sparsity in linear inverse problems. By varying the hyperparameters we can move continuously between priors that act as smoothed ℓ_p penalties with flexible p , smoothing, and scale. We then introduce methods for tracking MAP solution paths along paths through hyper parameter space. Path following allows a user to explore the space of possible MAP solutions under varying assumptions and to test the robustness and sensitivity of solutions to changes in the prior assumptions. By tracing paths from a convex region to a non-convex region, the user can find local minimizers in strongly sparsity promoting regimes that are consistent with a convex relaxation of the same problem derived using a consistent family of prior assumptions. We show experimentally that these solutions are less error prone than direct optimization of the non-convex problem.

Alexander Strang, Zilai Si, Yucong Liu

University of Chicago
alexstrang@uchicago.edu, zilai@uchicago.edu, yucong-liu@uchicago.edu

MS381

Modeling and Design of Single-Atom Alloy Catalysts

In this contribution, we apply data-driven modeling - i.e., molecular simulations and machine learning (ML) - for the study of CO₂ hydrogenation processes (e.g., reverse water-gas shift) on single-atom alloy (SAA) catalysts, i.e., diluted bimetallic materials able to break the scaling relationships that limit conventional catalysts [RT Hannagan, et al., Chem Rev 120, 12044-12088, 2020]. We target a wide combinatorial space of elements of the periodic table, which makes a direct study with density-functional theory (DFT) computationally prohibitive. Therefore, we produce a database of DFT-calculated energies on a limited number of SAAs and we apply physics-inspired ML techniques for the extrapolation to a wide range of materials. We use a graph-based Gaussian Process Regression ML model (i.e., WWL-GPR) to calculate adsorption energies [W Xu, et al., Nat Comput Sci 2, 443-450, 2022], and simpler models (e.g., multivariate regressions) to estimate the activation energies of the new materials. We employ microkinetic modeling to simulate the reaction kinetics; then, we apply Sensitivity Analysis and Uncertainty Quantification to identify the parameters that can improve the model predictions, and we refine them with additional DFT calculations. The application of the framework to CO₂ hydrogenation allows us to rationalize how reaction mechanisms and catalytic activity change with the catalyst composition, paving the way toward the design and

nano-engineering of SAA catalysts.

Raffaele Cheula, Mie Andersen
Aarhus University
raffaele.cheula@phys.au.dk, mie@phys.au.dk

MS381

Solving Kinetic Master Equations with Tensor Trains

The performance of a catalytic surface depends on physical processes taking place on multiple time scales. On the time scale of molecular motion, chemical surface reactions are rare transitions between metastable basins in free-energy landscapes. The interplay of these elementary reactions making up the catalytic cycle can be modeled as a Markov jump process on the lattice of adsorption sites. The corresponding Markovian master equation describing the stochastic reaction kinetics, however, is usually too high-dimensional to be solved with standard numerical techniques and one has to rely on sampling approaches such as the kinetic Monte Carlo method. In this talk, we show how to mitigate the curse of dimensionality for the direct solution of the kinetic master equation by exploiting tensor decompositions. We use the tensor-train format to construct probability distributions as well as transition-rate matrices in form of networks of so-called tensor cores representing the different adsorption sites. The performance of the approach is demonstrated on a reduced model for the CO oxidation on the RuO₂(110) surface. We investigate the complexity for increasing system size as well as for various reaction conditions and illustrate the advantage over stochastic simulation approaches by a sequence of problems with increasing stiffness.

Patrick Gelß
FU Berlin
gelss@zib.de

MS381

From Transition States to Dynamics Over Long Time Scales

For many chemical reactions and diffusion mechanisms in materials, transition states can be characterized by saddle points on the potential energy surface. Saddle point finding methods can be classified into two categories. First, there are methods including the nudged elastic band and string methods which require knowledge of both reactant and product states to determine a minimum energy path connecting them. These so-called double-ended methods are widely used with density functional theory to determine rates of anticipated reaction mechanisms. A second class of single-ended methods, generally based upon minimum-mode following techniques, require only the initial state to determine possible reaction mechanisms leading to possibly unknown product states. These single-ended methods require more sampling but they have the power to map out reactions on a potential surface with minimal user bias. I will discuss application of these saddle point finding methods to accelerating molecular dynamics simulations, such as through the hyperdynamics approach. Also, upon discovering a set of reaction mechanisms from an initial state, minimum mode following saddle searches can be coupled to kinetic Monte Carlo to model the evolution of atomic scale systems over the time scales of the reactions. This adaptive kinetic Monte Carlo method can, in favorable cases, even model the evolution of atomic scale systems to experimental time scales of seconds or minutes – orders of

magnitude longer time scales than can be reached with molecular dynamics. A frank assessment of the limitations of the methods will also be given. Examples will include molecular reactions on metal surfaces and surface segregation in bimetallic nanoparticles.

Graeme Henkelman
Assistant professor at the University
of Texas at Austin
henkelman@mail.utexas.edu

MS381

Conquering the Mesoscale in the Simulation of Surfaces and Interfaces via Distributed Lattice-Based Kinetic Monte Carlo

Kinetic Monte Carlo (KMC) simulations are instrumental in the multi-scale modelling of surfaces and interfaces. As an example, such simulations are frequently used to elucidate the complex reaction dynamics exhibited by solid catalysts and predict macroscopic performance metrics, such as activity and selectivity. However, the length- and time-scales accessible by sequential KMC implementations are limited, and handling lattices with millions of sites has been prohibitive due to large memory requirements and long simulation times. Domain decomposition approaches could address these limitations, but they require sophisticated algorithms for conflict resolution at the boundaries between subdomains. Jeffersons Time-Warp algorithm addresses this challenge via rollbacks and re-simulations which correct any and all causality violations arising transiently during simulation. Thus, the exact dynamics of the chemical master equation are finally reproduced. In this work, we have coupled the Time-Warp algorithm with the Graph-Theoretical KMC framework enabling the handling of complex adsorbate lateral interactions and reaction events within large lattices. This approach has been implemented in our general-purpose KMC software Zacros, and has been validated and benchmarked for efficiency in model systems as well as realistic chemistries. This work makes Zacros the first-of-its-kind general-purpose KMC code with distributed parallelisation capability to study heterogeneous catalysts.

Michail Stamatakis
University College London
m.stamatakis@ucl.ac.uk

MS382

Tile-Centric Approximations of The Maximum Likelihood Estimation for Large-Scale Spatial Statistics Modeling

Maximum Likelihood Estimation (MLE) is one of the most popular approaches to model spatial and spatio-temporal data in Geoscience. It relies on generating an $npt \times npt$ covariance matrix to describe the underlying geospatial random field, where n represents the number of spatial locations, p is the number of variables, and t is the number of time slots. When the npt value is large, the MLE calculation becomes prohibitive with $O((npt)^3)$ time complexity. In the literature, studies have adopted two ways to mitigate the computational cost of the MLE operation: high-performance computing and algebraic approximations. In this work, we couple parallel processing with matrix approximation techniques to speed up the MLE operation on large-scale systems. Moreover, we combine two approximation techniques, i.e., mixed precisions and low-rank matrix approximations, to deliver higher compression rate and

faster time-to-solution, while maintaining applications' accuracy. The results show that our approximation technique can provide accurate estimates for the statistical parameters with a considerable performance improvement. Our experiments were conducted on three different systems, KAUST Shaheen-II (an Intel Haswell system), HLRS HAWK (an AMD Epyc Rome system), and Riken Fugaku (a Fujitsu A64FX system).

Sameh Abdulah

Extreme Computing Research Center (ECRC)
King Abdullah University of Science and Technology
sameh.abdulah@kaust.edu.sa

Qinglei Cao

University of Tennessee Knoxville
qcao3@vols.utk.edu

Rabab Alomairy, Pratik Nag

KAUST
rabab.omairy@kaust.edu.sa, pratik.nag@kaust.edu.sa

Yu Pei

University of Tennessee
ypei2@vols.utk.edu

George Bosilca

University of Tennessee Knoxville
bosilca@icl.utk.edu

Jack J. Dongarra

University of Tennessee, Oak Ridge National Laboratory,
USA
dongarra@icl.utk.edu

Marc G. Genton

King Abdullah University of Science and Technology
Statistics Program
marc.genton@kaust.edu.sa

David E. Keyes

King Abdullah University of Science and Technology
(KAUST)
david.keyes@kaust.edu.sa

Hatem Ltaief

KAUST, Saudi Arabia
hatem.ltaief@kaust.edu.sa

Ying Sun

KAUST
Saudi Arabia
ying.sun@kaust.edu.sa

MS382

Algorithms, Applications, and Software for \mathcal{H}^2 Hierarchical Matrices

The \mathcal{H}^2 hierarchical matrix representation makes it possible to store and apply certain dense kernel matrices with linear (optimal) complexity. This makes it promising to use \mathcal{H}^2 hierarchical matrices for large-scale scientific and machine learning applications where these kernel matrices arise. This talk will discuss recent developments in the efficient and accurate low-rank or interpolative decomposition needed in \mathcal{H}^2 hierarchical matrix construction for different types of kernel matrices from different types of applications. Recently developed software for these applications

will also be presented.

Edmond Chow

School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

MS382

Approximate Multifrontal Solver with GPU Accelerated Block Low Rank Compression

We present an approximate sparse multifrontal solver with block low rank compression, used as a preconditioner for GMRES or BiCGStab. The solver runs on modern GPU architectures, relying on non-uniform batched dense linear algebra kernels from the MAGMA and KBLAS libraries. Various low rank compression kernels are compared, including methods based on adaptive randomized sampling, in terms of accuracy and performance on the GPU. The multifrontal solver can furthermore combine multiple rank structured formats, such as block low rank and hierarchically off-diagonal butterfly, where the choice of approximation method depends on the size of the frontal matrix being compressed. This combination reduces both the asymptotic complexity as well as the runtime for medium scale problem sizes. We consider a range of industrial and academic applications: high frequency Helmholtz and Maxwell, incompressible Navier-Stokes, etc.

Pieter Ghysels

Lawrence Berkeley National Laboratory
Computational Research Division
pghysels@lbl.gov

Lisa Claus

Lawrence Berkeley National Laboratory
lclaus@lbl.gov

Wajih Boukaram

Lawrence Berkeley National Lab
wajih.boukaram@lbl.gov

Xiaoye S. Li

Computational Research Division
Lawrence Berkeley National Laboratory
xsli@lbl.gov

MS383

Learning Dynamics from Images Using Lagrangian/Hamiltonian Structure

In many real-world settings, image observations of physical systems, such as satellites, may be available when low-dimensional measurements are not. However, the high-dimensionality of image data precludes the use of classical estimation techniques to learn the dynamics, and a lack of interpretability reduces the usefulness of standard deep learning methods. In this talk, I will discuss our work on leveraging Lagrangian and Hamiltonian formalisms in neural network design for physically plausible neural network based video prediction and generation. In our prediction pipeline we explicitly construct the equations of motion from learned representations of the underlying physical quantities, and in our generative model we implicitly discover the structure of the configuration space.

Christine Allen-Blanchette

Princeton University

ca15@princeton.edu

MS383

Structured Barycentric Forms for Data-Driven Modeling of Structured Dynamics

Data-driven rational approximation is a key tool in learning dynamical systems from data. In this talk, we will first briefly review how the barycentric form of a rational approximant is employed in learning linear (unstructured) dynamics from data. Then, we will extend this analysis and methodology to structured linear systems, such as mechanical systems described second-order dynamics, by developing new barycentric-like forms for these classes of dynamical systems. Then, we will further extend the analysis to special classes of structured nonlinear systems where the nonlinearity appears as a quadratic output.

Serkan Gugercin
Virginia Tech

Department of Mathematics
gugercin@vt.edu

Ion Victor Gosea
Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

Steffen W. R. Werner
Courant Institute of Mathematical Sciences
New York University
steffen.werner@nyu.edu

MS383

Preserving Lagrangian Structure in Data-Driven Reduced-Order Modeling of Large-Scale Mechanical Systems

We present a nonintrusive physics-preserving method to learn reduced-order models (ROMs) of Lagrangian mechanical systems. Existing intrusive projection-based model reduction approaches construct structure-preserving Lagrangian ROMs by projecting the Euler-Lagrange equations of the full-order model (FOM) onto a linear subspace. This Galerkin projection step requires complete knowledge about the Lagrangian operators in the FOM and full access to manipulate the computer code. In contrast, the proposed Lagrangian operator inference approach embeds the mechanics into the operator inference framework to develop a data-driven model reduction method that preserves the underlying Lagrangian structure. The method does not require access to FOM operators or computer code. The numerical results demonstrate Lagrangian operator inference on an Euler-Bernoulli beam model and a large-scale discretization of a soft robot fishtail with 779,232 degrees of freedom. Accurate long-time predictions of the learned Lagrangian ROMs far outside the training time interval illustrate their generalizability.

Harsh Sharma, Boris Kramer
University of California San Diego
hasharma@ucsd.edu, bmramer@ucsd.edu

MS383

Data-Driven Structure-Preserving Model Reduction for Stochastic Hamiltonian Systems

In this work we demonstrate that SVD-based model reduction techniques known for ordinary differential equa-

tions, such as the proper orthogonal decomposition, can be extended to stochastic differential equations in order to reduce the computational cost arising from both the high dimension of the considered stochastic system and the large number of independent Monte Carlo runs. We also extend the proper symplectic decomposition method to stochastic Hamiltonian systems, both with and without external forcing, and argue that preserving the underlying symplectic or variational structures results in more accurate and stable solutions that conserve energy better than when the non-geometric approach is used. We validate our proposed techniques with numerical experiments for a semi-discretization of the stochastic nonlinear Schrödinger equation and the Kubo oscillator.

Tomasz Tyranowski
Max Planck Institute for Plasma Physics
tomasz.tyranowski@ipp.mpg.de

MS384

Unlimited Sampling Meets Super-Resolution

Almost all physical sensors are limited by their dynamic range. This fundamental bottleneck also limits the practical utility of the Shannon-Nyquist sampling theorem which is at the heart of Digital Revolution. This is because amplitudes exceeding the dynamic range of an analog-to-digital converter (ADC) result in permanent loss of information due to clipping or saturation. To reconcile this gap between theory and practice, we introduce a computational sensing approach the Unlimited Sensing framework (USF) that is based on a co-design of hardware and algorithms. On the hardware front, our work is based on a radically different design, which allows for the ADCs to produce modulo or folded samples. On the algorithms front, we develop new, mathematically guaranteed recovery strategies. We show that recovering signals from modulo folded measurements is related to the super-resolution and exponential fitting problems. Thereon, we also consider the reconstruction of sparse signals from modulo samples which leads to a new twist in that one ends up into a double super-resolution problem. We show that our approach is empirically robust to noise and allows for recovery of signals upto 25 times modulo threshold, when working with modulo ADC hardware.

Ayush Bhandari
Imperial College London
Department of Electrical and Electronic Engineering
a.bhandari@imperial.ac.uk

MS384

Reconstruction of Unresolved Small-Scale Structures in Isotropic Turbulence Using Deep Learning

Recently, deep learning has been extensively utilized for modeling and analysis of fluid turbulence. One such application is the use of super-resolution (SR) algorithms to reconstruct small-scale structures from their large-scale counterparts for turbulent flows. To date, all SR algorithms are either supervised, or require unpaired reference data at high-resolution for training. This renders the model inapplicable to practical fluid flow scenarios, where the generation of high-resolution ground truth by resolving all scales down to the Kolmogorov scales becomes prohibitive. Hence, it is imperative to develop physics-guided models exploiting the multi-scale nature of turbulence. Casting SR as an inverse problem, we present a self-supervised workflow based on generative adversarial networks to re-

construct small-scale structures relevant to homogeneous isotropic turbulence. In addition to visual similarity, we assess the quality of the obtained reconstruction using spectra, structure functions, and probability density functions of velocity components and a passive scalar. From the analysis, we infer that the outputs of the workflow are in statistical agreement with the ground truth, which is excluded from the training. Additionally, we observe that the trained network generalizes well across Reynolds numbers. The results from this work lay the foundation to reconstruct small-scale structures from large-eddy simulation data.

Priyabrat Dash
Indian Institute of Science, Bangalore
priyabratd@iisc.ac.in

Konduri Aditya
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
konduriadi@iisc.ac.in

MS384

Espira for Reconstruction of Exponential Sums and Its Application

In this talk we present our recent results regarding the ESPIRA (Estimation of Signal Parameters by Iterative Rational Approximation) method for reconstruction of exponential sums. As signal model we consider the complex exponential sums which strongly generalize usual Fourier sums. As input information the ESPIRA method employs a set of equidistant signal values and is based on rational approximation of DFT data. The rational approximation is done with the use of AAA algorithm developed by Nakatsukasa, Ste and Trefethen (2018). Some applications of ESPIRA will be presented during the talk.

Nadiia Derevianko
Institute for Numerical and Applied Mathematics
University of Göttingen
n.derevianko@math.uni-goettingen.de

Gerlind Plonka
Institute for Numerical and Applied Mathematics
University of Goettingen
plonka@math.uni-goettingen.de

Markus Petz
Institute for Numerical and Applied Mathematics
University of Göttingen
m.petz@math.uni-goettingen.de

MS384

Optimal Algorithm for Super-Resolving Close Point Sources from Bandlimited and Noisy Data

We consider the problem of recovering a linear combination of Dirac masses from noisy Fourier samples, also known as the problem of super-resolution. Following recent derivation of min-max bounds for this problem when two or more sources collide, we develop an optimal algorithm which provably achieves these bounds in such a challenging scenario.

Rami Katz, Nuha Diab
Tel-Aviv University
ramikatz@mail.tau.ac.il, nuhadiab@tauex.tau.ac.il

Dmitry Batenkov
Tel Aviv University
dbatenkov@tauex.tau.ac.il

MS384

On the Stability of Sparse Super-Resolution

We consider super resolution (SR) as the nonlinear mapping of the low frequency Fourier coefficients of a discrete measure on $[0, 1]^d$ to its support and weights. We focus on weak SR assuming a condition on the separation of the involved measures similar to the Rayleigh criterion and prove that the reconstruction map satisfies a local Lipschitz property giving explicit estimates for the Lipschitz constant depending on the dimension d and the sampling effort while improving a recent bound on the assumed separation from Chen/Moitra for $d = 2$. Technically, we heavily rely on the construction of localising functions (extremal functions, minorants) as introduced by Beurling/Selberg in the 20th century (being less magic than their famous lattice-interpolating counterparts by Cohn/Kumar/Miller/Radchenko/Viazovska). We finally conclude that weak SR is globally Lipschitz continuous and hence well conditioned when the space of measures is equipped with the Wasserstein distance as metric.

Stefan Kunis
University of Osnabrueck
stefan.kunis@uos.de

MS385

Adaptive Choice of Near-Optimal Expansion Points for Interpolation-Based Reduced Order Modeling

Interpolation-based methods are well-established and effective approaches for the efficient generation of accurate reduced-order surrogate models. Common challenges for such methods are the automatic choice of good or even optimal interpolation points as well as a suitable size for the reduced-order model. An approach that tackles the first problem for linear, unstructured systems is the Iterative Rational Krylov Algorithm (IRKA), which computes optimal interpolation points by iterative updates solving linear eigenvalue problems. However, in the case of structured systems, the eigenvalue problems resulting from the corresponding transfer functions are nonlinear such that the number of potential new interpolation points exceeds the order of the reduced model. In our work, we propose an IRKA-inspired iterative interpolation method for structured systems to compute near-optimal interpolation points as well as a suitable size for the reduced-order model. Thereby, the iterative updates of the expansion points are chosen to yield accurate approximations in specified frequency ranges of interest.

Quirin Aumann
Max Planck Institute for Dynamics
of Complex Technical Systems
aumann@mpi-magdeburg.mpg.de

Steffen W. R. Werner
Courant Institute of Mathematical Sciences
New York University
steffen.werner@nyu.edu

MS385

Optimization-Based Model Order Reduction for

Stochastic Systems

In this talk, we will connect ideas from model order reduction for stochastic linear systems and \mathcal{H}_2 -optimal model order reduction for deterministic systems. In particular, we supplement and complete the theory of error bounds for model order reduction of stochastic differential equations. With these error bounds, we establish a link between the output error for stochastic systems (with additive and multiplicative noise) and modified versions of the \mathcal{H}_2 -norm for both linear and bilinear deterministic systems. When deriving the respective optimality conditions for minimizing the error bounds, we will see that model order reduction techniques related to iterative rational Krylov algorithms (IRKA) are very natural and effective methods for reducing the dimension of large-scale stochastic systems with additive and/or multiplicative noise. We apply modified versions of (linear and bilinear) IRKA to stochastic linear systems and show their efficiency in numerical experiments.

Melina Freitag
University of Potsdam
melina.freitag@uni-potsdam.de

Martin Redmann
Martin Luther University of Halle Wittenberg
martin.redmann@mathematik.uni-halle.de

MS385

Certified Optimization with Reduced Order Models

This talk presents new approaches for the optimization of expensive large-scale optimization problems using reduced order models (ROMs). The fundamental idea is to approximate computationally expensive evaluations of objective and constraint functions and derivatives by inexpensive ROMs. The challenge is that ROMs need to be adjusted during the optimization process, and that only asymptotic ROM error bounds are available. New trust-region and line search based optimization approaches and their integration with efficient ROM constructions are presented that are implementable when only asymptotic ROM error bounds are available and have proven convergence properties.

Matthias Heinkenschloss
Department of Computational and Applied Mathematics
Rice University
heinken@rice.edu

MS385

Optimization-Based Model Order Reduction with Applications in Finance

Solving optimal stopping problems by backward induction in high dimensions is often very complex since the computation of conditional expectations is required. Typically, such computations are based on regression, a method that suffers from the curse of dimensionality. Therefore, the objective of this talk is to establish dimension reduction schemes for large-scale asset price models and to solve related optimal stopping problems (e.g. Bermudan option pricing) in the reduced setting, where regression is feasible. The proposed algorithm is based on an error measure between linear stochastic differential equations. We establish optimality conditions for this error measure with respect to the reduce system coefficients and propose a particular method that satisfies these conditions up to a small devi-

ation. We illustrate the benefit of our approach in several numerical experiments, in which Bermudan option prices are determined.

Martin Redmann
Martin Luther University of Halle Wittenberg
martin.redmann@mathematik.uni-halle.de

MS385

Damping Optimization of the Excited Mechanical System Using Dimension Reduction

In this work a mechanical system excited by a periodic external force is considered. The main problem is to determine the best damping matrix to be able to minimize the system average displacement amplitude. Damping optimization usually includes optimization of damping positions and corresponding damping viscosities. Since the objective function is non-convex, a standard optimization approach typically requires a large number of objective function evaluations. We first propose a dimension reduction approach that calculates approximation of the average displacement amplitude and additionally we efficiently use a low rank update structure that appears in the external damping matrix. Moreover, an error bound which allows determination of appropriate approximation orders is derived and incorporated within the optimization method. We also present a theoretical error bound that allows determination of effective damping positions. The methodology proposed here provides a significant acceleration of the optimization process. The gain in efficiency is illustrated in numerical experiments.

Zoran Tomljanovic
University of Osijek
ztomljan@mathos.hr

MS386

Data-Driven Uncertainty Quantification for Linear Systems

Uncertainty propagation and uncertainty quantification is at the core of many control and optimization problems under uncertainty. In this talk, we discuss the prospect of uncertainty propagation for finite-dimensional linear systems in discrete time without explicit knowledge of the actual system model. We combine the framework of polynomial chaos expansions with data-driven system representations stemming from the fundamental lemma. The latter lemma states that under suitable controllability assumptions, any finite horizon input-output trajectory of a linear system is contained in the column space of suitable Hankel matrices constructed from persistently exciting input data and the corresponding output trajectories. We show that under suitable assumptions on the available input-output measurement data, this enables to propagate Gaussian and non-Gaussian uncertainties through linear systems without explicit model knowledge. We draw upon examples from energy system networks and from other domains to illustrate our findings.

Timm Faulwasser, Guanru Pan, Ruchuan Ou
TU Dortmund University
timm.faulwasser@tu-dortmund.de, gunaru.pan@tu-dortmund.de, ruchuan.ou@tu-dortmund.de

MS386

Uncertainty Quantification in Hierarchical Vehicu-

lar Traffic Models

It is well known that evolution of vehicular traffic is exposed to the presence of uncertainties. In this talk, starting from the established hierarchy between microscopic, kinetic and macroscopic scales, we will investigate the propagation of uncertainties through the models, applying the intrusive stochastic Galerkin approach. Connections between the scales will be presented in the stochastic scenario and numerical simulations will be performed.

Elisa Iacomini, Michael Herty
RWTH Aachen University
iacomini@igpm.rwth-aachen.de, herty@igpm.rwth-aachen.de

MS386

Uncertainty Aware Optimal Allocation for Gas Transmission Networks Using Finite Element Methods

We develop a finite volume representation of uncertainty in solutions of hyperbolic partial differential equation systems on graph-connected domains with nodal coupling and boundary conditions. The representation is used to state the physical constraints in stochastic optimization problems subject to uncertain parameters. The method is based on the Stochastic Finite Volume (SFV) approach, and can be applied for uncertainty management of fluid flow over actuated transport networks. The method is examined for steady-state optimization subject to probabilistic constraints where the SFV-based representation of the uncertainty is directly incorporated as constraints into a non-linear programming formulation. The computational advantages of using higher order quadrature for evaluation of the probabilistic constraints is explored and demonstrated using numerical examples.

Sidhant Misra
Los Alamos National Laboratory
sidhant@lanl.gov

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

Anatoly Zlotnik
Los Alamos National Laboratory
azlotnik@lanl.gov

MS386

Stochastic Finite Volume Method for Uncertainty Quantification of Transient Flow in Gas Pipeline Networks

We develop a weakly intrusive framework to simulate the propagation of uncertainty in solutions of generic hyperbolic partial differential equation systems on graph-connected domains with nodal coupling and boundary conditions. The method is based on the Stochastic Finite Volume (SFV) approach, and can be applied for uncertainty quantification (UQ) of the dynamical state of fluid flow over actuated transport networks. The numerical scheme has specific advantages for modeling intertemporal uncertainty in time-varying boundary parameters, which cannot be characterized by strict upper and lower (interval) bounds. We describe the scheme for a single pipe, and then formulate the controlled junction Riemann problem

(JRP) that enables the extension to general network structures. We demonstrate the method's capabilities and performance characteristics using a standard benchmark test network.

Anatoly Zlotnik
Los Alamos National Laboratory
azlotnik@lanl.gov

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

Vitaliy Gyrya
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
vitaliy_gyrya@lanl.gov

MS387

Posterior Predictive Variational Inference for Uncertainty Quantification in Machine Learning

Machine Learning (ML) is increasingly becoming an integral part of the computational science and engineering (CSE) toolkit, yet the rigor of uncertainty quantification (UQ) for ML has not kept pace with its utilization. Bayesian inference provides a principled framework for UQ, yet ML models like Deep Neural Networks (DNNs) challenge many traditional inference algorithms. This has motivated the use of approximate Bayesian inference methods like Variational Inference (VI) within ML. While VI has significant computational advantages over traditional algorithms like Markov Chain Monte Carlo, many challenges remain with VI for Bayesian UQ in ML such as the robustness of UQ to often ad hoc assumptions of variational family and prior distribution. We seek to mitigate some of these deficiencies by taking a goal-oriented approach. Our approach, Posterior Predictive Variational Inference (PPVI) seeks to optimize VI to make specific predictions. This means that uncertainty is optimally represented to answer specific questions at hand. This makes UQ more accurate for these specified quantities of interest and makes it easier to understand the robustness of these methods to assumptions. We will present theory, algorithms, and illustrative examples for PPVI to describe its utility to CSE.

Thomas A. Catanach
Sandia National Laboratories
tacatan@sandia.gov

MS387

A Study of Bias-Variance in Variational Inferencing Using Delta Method

Variational inference (VI) is a method from machine learning that approximates probability densities through optimization. VI has been used in many applications and tends to be faster than classical methods, such as Markov chain Monte Carlo sampling. Instead of KL divergence, the evidence lower bound as a cost function is used in optimizing the parameters of a variational distribution, especially for high dimensional problems. We explore the technicalities in using KL divergence in such a scenario, synthesizing estimators, and leveraging upon the delta method to study the bias-variance trade-off. Our objective is to study the gradient space of the KL divergence with respect to the variational parameter and to understand how we can re-

duce its variance using importance sampling. With lower variance in the gradient information, we can achieve larger optimization steps, lower the number of samples needed, which we compare against the ELBO based optimization.

Niladri Das
Sandia National Laboratory
ndas@sandia.gov

Thomas A. Catanach
Sandia National Laboratories
tacatan@sandia.gov

MS387

Deep Ensembles of Variational Approximations

log-likelihood is estimated using the standard evidence lower bound (ELBO), or improved versions as the importance weighted ELBO (IWELBO). We propose the multiple importance sampling ELBO (MISELBO), a versatile yet simple framework. MISELBO is applicable in both amortized and classical VI, and it uses ensembles, e.g., deep ensembles, of independently inferred variational approximations. As far as we are aware, the concept of deep ensembles in amortized VI has not previously been established. We prove that MISELBO provides a tighter bound than the average of standard ELBOs, and demonstrate empirically that it gives tighter bounds than the average of IWELBOs. MISELBO is evaluated in density-estimation experiments that include MNIST and several real-data phylogenetic tree inference problems. First, on the MNIST dataset, MISELBO boosts the density-estimation performances of a state-of-the-art model, nouveau VAE. Second, in the phylogenetic tree inference setting, our framework enhances a state-of-the-art VI algorithm that uses normalizing flows. On top of the technical benefits of MISELBO, it allows to unveil connections between VI and recent advances in the importance sampling literature, paving the way for further methodological advances.

Oskar Kviman
KTH Royal Institute of Technology
okviman@kth.se

MS388

Shape-Changing Trust-Region Methods Using Multipoint Symmetric Secant Matrices

We consider methods for large-scale and nonconvex unconstrained optimization. We propose a new trust-region method whose subproblem is defined using a so-called shape-changing norm together with densely-initialized multipoint symmetric secant (MSS) matrices to approximate the Hessian. Shape-changing norms and dense initializations have been successfully used in the context of traditional quasi-Newton methods, but have yet to be explored in the case of MSS methods. Numerical results suggest that trust-region methods that use densely-initialized MSS matrices together with shape-changing norms outperform MSS with other trust-region methods.

Roummel F. Marcia
University of California, Merced
rmarcia@ucmerced.edu

Johannes J. Brust
University of California, San Diego
Mathematics Department

jjbrust@ucsd.edu

Jennifer Erway
Wake Forest University
erwayjb@wfu.edu

MS388

Distributed Stochastic Inertial-Accelerated Methods with Delayed Derivatives for Nonconvex Problems

Stochastic gradient methods (SGMs) are predominant approaches for stochastic optimization. On smooth nonconvex problems, acceleration techniques have been applied to improve the convergence of SGMs. However, little exploration has been made on applying acceleration to a stochastic subgradient method (SsGM) for nonsmooth nonconvex problems. Also, few efforts have been made to analyze an (accelerated) SsGM with delayed derivatives. The information delay naturally happens in a distributed system. In this talk, I will present an inertial proximal SsGM for nonsmooth nonconvex stochastic optimization. Our method has guaranteed convergence even with delayed derivatives in a distributed environment. Convergence rate results are established to three problem classes: weakly-convex nonsmooth problems with a convex regularizer, composite nonconvex problems with a nonsmooth convex regularizer, and smooth nonconvex problems. In a distributed environment, the convergence rate of the proposed method will be slowed down by the information delay. Nevertheless, the slow-down effect will decay with the number of iterations for the latter two problem classes. We test the proposed method on three applications. The numerical results clearly demonstrate the advantages of using the inertial-based acceleration. Furthermore, we observe higher parallelization speed-up in asynchronous updates over the synchronous counterpart, though the former uses delayed derivatives.

Yangyang Xu
Rensselaer Polytechnic Institute
xuy21@rpi.edu

MS389

On Machines and CPMIP Metrics Underlying CMIP Simulations

Since 1995, the Coupled Model Intercomparison Project (CMIP) has coordinated climate model experiments involving multiple international modelling teams. This has led to a better understanding of the Earth's climate, and CMIP model experiments have routinely been the basis for future climate change assessments made by the Intergovernmental Panel on Climate Change (IPCC). Along with all other aspects of the workflow for the sixth phase, CMIP6, the Earth System Documentation (ES-DOC) project manages the creation, analysis and dissemination of documentation about the computing platform used to run the various ensembles of simulations and their representative performance for the given hardware and CMIP6 model and experiment, with a dedicated eco-system of tools. The computational performance evaluation is characterised by the CPMIP Metrics, a set of measures specifically designed to capture the actual performance of Earth System Models (ESMs). We present an overview of the results of all present submissions by modelling groups towards CMIP6 machine and performance documentation coordinated by ES-DOC, whilst describing the challenges of collecting and communicating such information. We briefly look to the

past and future with regards to the realities and importance of this task, for example thinking ahead to CMIP7.

Sadie L. Bartholomew

National Centre for Atmospheric Science (NCAS)
Department of Meteorology, University of Reading
sadie.bartholomew@ncas.ac.uk

MS389

Deep Learning of Seismograms

Seismology is a data-rich and data-driven science in which the application of machine learning (ML) techniques has been growing rapidly. Among the various ML techniques, seismologists quickly recognized and realized the potential of deep neural networks (DNNs) to address a broad array of seismological applications. Much of this is due to the availability of large-scale datasets and the suitability of deep-learning techniques for seismic data processing; however, some aspects of applying AI to seismology are likely to prove instructive for geosciences more broadly. Deep learning is a powerful approach, but its application has subtleties and nuances. I will present a systematic overview of trends, challenges, and opportunities in applications of deep-learning methods in seismology.

Mostafa Mousavi

stanford university
smousavi05@gmail.com

MS389

Global Ocean Modeling on GPUs with Julia

In this talk, we examine the performance of the open-source Oceananigans.jl software in modeling low to high-resolution large-scale ocean simulations. The software is written entirely in the Julia programming language and uses a rich system of abstractions not easily implemented in the traditional languages of climate science, such as Fortran and C. Additionally, by leveraging the Julia infrastructure, Oceananigans allows switching the computation between CPU and GPU architectures seamlessly, making it the perfect tool to prototype, test, and run large-scale ocean simulations. We demonstrate the usability and the performance of Oceananigans by integrating two relevant global setups, an eddy-permitting simulation at a 0.25-degree resolution and an eddy-resolving one at a 0.1-degree resolution. While the former fits entirely on one gaming-grade GPU, the latter requires parallelizing the computation on a minimum of three scientific programming GPUs. We examine the software's scalability and accuracy along with the performance for both test cases, showing optimal results compared to state-of-the-art ocean models. We argue that it is primarily the use of Julia that allows us to implement a high-level model that is easy and engaging but also attains good scaling and absolute performance on large-scale climate problems.

Simone Silvestri

MIT
silvestri.simone0@gmail.com

Gregory Wagner

Massachusetts Institute of Technology
glwagner@mit.edu

Chris Hill

Earth, Atmospheric and Physical Sciences
M.I.T.

cnh@mit.edu

MS389

DestinE - Digital Twins of the Earth System

As part of the European Commission's Destination Earth programme, global, coupled storm-resolving simulations can contribute to building digital replicas of Earth, thanks to recent advances in Earth system modelling, supercomputing and the adaptation of weather and climate codes for novel computing architectures. Combined with novel data-driven deep learning advances, this offers a window into the future, with a promise to significantly increase the timeliness and realism of Earth system information. Despite the significant compute and data challenges, there is a real prospect to better support global to local climate change mitigation and adaptation efforts, and complement today's existing information. Digital Twins of Earth thus encapsulate both the latest science as well as technology advances to provide near-real time information on Extremes and climate change in a wider digital environment. Here users can interact, modify and ultimately create or complement their own tailored information. Partnering with ESA and EUMETSAT, this is facilitated through complex workflows akin to a game engine environment provided by ECMWF's digital twin engine that closely connects EuroHPC resources for the production of and access to digital twin data. The underlying system architecture and design choices will be described and justified.

Nils Wedi, Peter Bauer, Irina Sandu, Tiago Quintino,

Simon Smart, James Hawkes

European Centre for Medium-Range Weather Forecasts
(ECMWF)

nils.wedi@ecmwf.int, peter.bauer@ecmwf.int,
irina.sandu@ecmwf.int, tiago.quintino@ecmwf.int,
simon.smart@ecmwf.int, james.hawkes@ecmwf.int

Balthasar Reuter

ECMWF

balthasar.reuter@ecmwf.int

Andrew Bennett, Vasileios Baousis, Thomas Geenen

European Centre for Medium-Range Weather Forecasts
(ECMWF)

andrew.bennett@ecmwf.int, vasilis.baousis@ecmwf.int,
thomas.geenen@ecmwf.int

MS390

A Descent Algorithm for the Optimal Control of Relu Neural Network Informed PDEs Based on Approximate Directional Derivatives

We introduce a numerical solver for a class of optimal control problems with learning-informed semilinear partial differential equations (PDEs). The latter have constituents that are in principle unknown and are approximated by nonsmooth ReLU neural networks. We argue that a direct smoothing of the ReLU network with the aim to make use of classical solvers can have significant disadvantages. This motivates us to devise a numerical algorithm that treats directly the nonsmooth optimal control problem, by employing a descent algorithm inspired by a bundle-free method. Several numerical examples are provided and the efficiency of the algorithm is shown.

Michael Hintermüller

Weierstrass Institute for Applied Analysis and Stochastics
Humboldt University Berlin

michael.hintermueller@wias-berlin.de

MS390

Trust-Region Method for PDE-Constrained Optimization with Total Variation Regularization of Distributed Integer-Valued Parameters

We are interested in solving mixed-integer PDE-constrained optimization problems with integer-valued distributed parameters. The latter are regularized with a total variation term in the objective. The regularization allows us to derive first-order optimality conditions of the optimization problem and trust-region subproblems by means of local variations of the level sets of the distributed parameter functions. The latter also allow us to prove that the iterates of a function space trust-region algorithm (solving the aforementioned subproblems) converge to first-order optimal points. We also show how inexactness may be introduced into the trust-region subproblems and one may obtain a convergence of the method when discretized subproblems are solved.

Paul Manns

Mathematik

TU Dortmund

paul.manns@tu-dortmund.de

MS390

Optimization Problems over the Space of Functions of Bounded Variation with Distributional Gradient Constraints

We consider non-diffusive variational problems with mixed boundary conditions and with gradient distributional constraints. These problems arise as first order optimality conditions of optimization ones over BV (or Sobolev) spaces with appropriate gradient constraints, and are useful in the modelling of the growth of non-homogeneous material piles. In this setting, the upper bound in the constraint is either a function or a Borel measure which allows the pile growth to observe discontinuities. We rigorously identify the Fenchel pre-dual problem, and consider primal-dual type algorithms for the approximation of solutions. Further, we investigate the perturbation of solutions to the problem by means of perturbation of the upper bound of the gradient constraint and establish stability results.

Carlos Rautenberg

George Mason University

crautenb@gmu.edu

MS390

Risk-Averse Optimization of Random Elliptic Variational Inequalities

Mathematical programs with equilibrium constraints (MPEC) have long since been a challenging topic in nonsmooth optimization. The core difficulties in the derivation of meaningful optimality conditions and development of provably convergent numerical optimization methods lie in the inherent degeneracy of the constraint set, which can often be rewritten as a mixed complementarity problem. We consider a class of MPECs in which the equilibrium constraint is modeled by an elliptic variational inequality subject to random inputs. Due to the resulting presence of uncertainty in the objective function \mathcal{J} , we model our risk preference by considering composite objective functions of the type $\mathcal{R}(\mathcal{J}(z)) + \rho(z)$, where \mathcal{R} is a numerical surrogate

for risk and $\rho(z)$ is the cost of the decision z . We propose to derive optimality conditions using an adaptive penalty technique, which can later be linked to a function-space based optimization method. In doing so, we highlight a number of unforeseen challenges, not unique to the infinite-dimensional setting, which give rise several open questions. Finally, we conclude with some initial numerical studies, in which we make use of a path-following stochastic approximation algorithm with a variance reduction step that periodically makes use of the full gradient.

Thomas M. Surowiec

Department of Mathematics and Computer Science

Philipps University of Marburg

thomasms@simula.no

MS390

Application of Composite Dfo Solver Limols to Bending Machine Modelling

Complex physical simulations often don't allow for automatic differentiation and numerical differentiation is costly and fidgety when the simulation output is not very smooth. Optimization of simulator parameters therefore has a need for derivative-free optimizers. As the simulations are expensive, sample efficiency is critical. Simulators often have many detailed outputs which get lumped into a single objective value. Avoiding all that loss of structure (e.g. by retaining a least squares structure) can greatly reduce the number of evaluations till convergence. Sioux Technologies has developed the composite derivative-free solver LIMOLS which has a clean interface for exploiting a plethora of composite structures and is very sample efficient as a result. We illustrate 10x faster convergence than Nelder-Mead on an application of modelling a bending machine and matching simulated shapes with measured shapes through a Hausdorff metric / Iterative closest point computation.

Lense Swaenen, Guus Bollen

Sioux Technologies

lense.swaenen@sioux.eu, guus.bollen@sioux.eu

MS391

Massively Parallel Probabilistic Computing with P-Bits

The slowing down of Moore's Law has led to a crisis as the computing workloads of Artificial Intelligence (AI) algorithms continue skyrocketing. There is an urgent need for scalable and energy-efficient hardware catering to the unique requirements of AI algorithms and applications. In this environment, probabilistic computing with p-bits [Camsari et al., 2017] has emerged as a scalable, domain-specific and energy-efficient computing paradigm, particularly useful for probabilistic applications and algorithms. In this talk, I will describe two general applications of p-computing: combinatorial optimization and probabilistic sampling as relevant problems to Machine Learning and AI. I will discuss recent and representative [Borders et al. 2019] experiments illustrating how both problems can be efficiently addressed by a suitably modified magnetoresistive random access memory (MRAM) technology. I will also show standard CMOS-based implementations of p-computing applied to practical optimization problems in large scale [Aadit et al., 2022] to stress why nanodevice-based implementations of p-computing may be a crucially needed ingredient.

Kerem Camsari

University of California, Santa Barbara
camsari@ucsb.edu

MS391

Spiking Neural Network Representation of Partial Differential Equation Evolution Maps

We present a novel method based on Spiking Neural Networks for the construction of numerical integrators of ordinary (ODE) and partial differential equations (PDEs). For illustration purposes, we provide applications to a host of ODEs and PDEs of variable complexity.

Adar Kahana, Qian Zhang
Brown University
adar_kahana@brown.edu, qian_zhang1@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george.karniadakis@brown.edu

Panos Stinis

Pacific Northwest National Laboratory
panos.stinis@pnnl.gov

MS391

Neuromorphic Circuits for Generating Useful Correlations from Random Bit Streams

Hardware approaches to random number generation from specific distributions require the ability to tune covariances between random bits. Correlated variability is also observed to be tightly controlled between spike trains from biological neurons. Computational neuroscientists have investigated the properties of neuronal correlations by devising algorithms to generate spike trains with prescribed correlations, and by studying plasticity rules that relate learned weights to neuronal correlations. I will present neuromorphic circuits that use these principles to generate bit streams with controlled correlations, and to solve discrete optimization problems such as graph MAXCUT. SAND2022-13164 A

Bradley H. Theilman
Sandia National Laboratories
bhtheil@sandia.gov

MS392

Stellarator Optimization with Automatic Differentiation in DESC

Next-generation fusion reactor designs require expensive numerical optimization to ensure good confinement and achievable engineering tolerances. DESC is a newly developed code that computes 3D Magnetohydrodynamic (MHD) equilibria and performs PDE constrained optimization of these solutions. DESC solves for equilibria by solving the force balance equation $J \times B = \nabla p$ using a pseudo-spectral method. In tests, it has been shown to be significantly more accurate than legacy codes such as VMEC, particularly near the coordinate singularity at the magnetic axis. DESC is also the first 3D MHD equilibrium code built with automatic differentiation, which allows accurate calculation of the Jacobian matrices at negligible additional cost. These derivatives can then be used for optimizing the solutions to obtain desirable physics and engineering properties. This reduces the computation time by three orders

of magnitude in tests compared to the legacy code STELLOPT, and enables exploration over a higher-dimensional parameter space. Examples such as quasisymmetry, metrics for particle confinement, stability, and turbulent fluxes will be shown. DESC also allows novel boundary conditions, which yields understanding on the existence of non-axisymmetric equilibria with nested flux surfaces and reveals their connection to axisymmetric solutions. Examples of these computational advances are demonstrated along with a discussion of the new physics insights they provide.

Rory Conlin, Daniel Dudt, Dario Panici
Princeton University
wconlin@princeton.edu, ddudt@princeton.edu, dpanici@princeton.edu

Patrick Kim
University of Maryland
pkim1236@terpmail.umd.edu

Kaya Unalmis, Egemen Kolemen
Princeton University
kunalmis@princeton.edu, ekolemen@princeton.edu

MS392

Multi-Objective Non-Convex Stellarator Optimization in Julia

Designing new stellarator configurations with improved plasma confinement presents a challenging numerical optimization problem. A characteristic feature of stellarator optimization is the need to optimize for different competing physics objectives that span a range of disparate time and length scales. Further complicating matters, the standard method for computing stellarator equilibria by specifying the plasma boundary is an ill-posed problem. When combined, these properties produce a complex optimization landscape that is both non-convex and not simply connected. A final barrier to stellarator optimization traditionally has been in the numerical software itself. In general, the existing software for stellarator optimization is very heterogenous in complexity/resource requirements and implementation language, presenting a large barrier to leveraging new developments in both optimization techniques and physics objective models. To address these needs, a new suite of modular optimization packages has been created using Julia that provides a common framework for interfacing equilibrium solvers and physics objectives with robust optimization routines. Built with high-performance computing and usability in mind, new physics objective functions can be seamlessly incorporated into optimizations that performantly scale to thousands of CPUs. This framework has been employed to optimize for turbulent transport simultaneously with good fast particle confinement and quasisymmetry.

Benjamin Faber, Aaron Bader
University of Wisconsin-Madison
bfaber@wisc.edu, aabder@wisc.edu

Joseph Duff
University of Wisconsin-Madison
Department of Physics
jduff2@wisc.edu

John Schmitt
Auburn University
jcs0114@auburn.edu

Chris Hegna, David Anderson
University of Wisconsin-Madison
cchegna@wisc.edu, dtanders@wisc.edu

MS392

Quantifying and Reducing Uncertainty in Inertial Confinement Fusion Experiments Using Optimal Experimental Design

High Energy Density (HED) science is the study of the behavior of material under extreme conditions of temperature and pressure. Understanding the growth and properties of hydrodynamic instabilities and the transition into turbulence is important in many HED processes and also yields insights into other areas where hydrodynamic instabilities occur. In contrast to classical fluids experiments, examining hydrodynamic instabilities in the HED regime poses novel challenges. HED experiments are expensive and often performed at oversubscribed facilities. Additionally, there are many limitations for the available diagnostics and experiments are typically multi-physics in nature. Such complexity means that modeling can become prohibitively expensive. Improving the models from limited experimental and high-fidelity simulation data is therefore of great importance. We will discuss the use of statistical techniques on a large suite of 2D xRAGE simulations of multimode Richtmyer-Meshkov and Rayleigh-Taylor experiments performed on Omega-EP to better understand turbulence in the HED regime. We will preform Bayesian inference to quantify and reduce uncertainty on our model parameters using experimental data. We will also explore experimental design techniques to identify the most informative experiments. These computations will involve many repeated forward model runs, which can be enabled through building inexpensive surrogate models that approximate the high-fidelity XRAGE.

Codie Fiedler Kawaguchi
University of Michigan
cfkawa@umich.edu

Kirk A. Flippo, Elizabeth C. Merritt, Alexander M. Rasmus
Los Alamos National Laboratory
kflippo@lanl.gov, emerritt@lanl.gov, rasmus@lanl.gov

Xun Huan, Eric Johnsen
University of Michigan
xhuan@umich.edu, ejohnsen@umich.edu

MS392

Optimizing Stellarators to Preserve Magnetic Surfaces at High Plasma Pressure

Stellarators are a promising concept for the development of controlled fusion power plants. They rely on the confinement of a hot plasma in a topologically toroidal vessel by using magnetic fields that are entirely generated externally. Since the charged particles in the plasma tend to follow magnetic field lines, a magnetic field that foliates nested surfaces is desirable. These magnetic surfaces can act as heat and particle transport barriers and in particular this is guaranteed if certain symmetries are ensured. The presence of the plasma, however, implies the existence of electric currents inside the confinement volume that scale with the plasma pressure. These currents perturb the magnetic field and, beyond a certain pressure limit, can break the magnetic surfaces. The breaking of the surfaces, which often leads to magnetic field line chaos, allows magnetic

field lines to undergo large excursions across the device, which can be detrimental for confinement. In this talk I will show how these pressure limits can be predicted and how they can be increased by optimizing the external magnetic field or by externally perturbing the plasma current. This involves solving the magneto-hydrodynamic equilibrium problem in three dimensions, a fantastic numerical challenge that requires treating magnetic field discontinuities as well as chaotic magnetic field lines.

Joaquim Loizu, Antoine Bailod
Ecole Polytechnique Fédérale de Lausanne (EPFL)
joaquim.loizu@epfl.ch, antoine.bailod@epfl.ch

Zhisong Qu
Nanyang Technological University
zhisong.qu@ntu.edu.sg

Stuart Hudson
Princeton University
shudson@pppl.gov

MS392

Level Set Learning for Poincaré Plots of Symplectic Maps

Many important qualities of plasma confinement devices can be determined via a Poincaré plot of the symplectic magnetic field return map. These qualities include the locations of chaotic regions of phase space and the locations of magnetic islands, which are commonly used for ejecting unwanted particles in the plasma. However, every evaluation of the magnetic return map requires an ODE solve, meaning a detailed Poincaré plot can be expensive to create. In recent years, a variety of symplecticity-preserving neural network based learning methods have been proposed that could be used as proxies to speed up return map evaluations. Unfortunately, these methods still require long training times, and are too expensive to be used for stellarator optimization objectives. In this talk, we propose a kernel-based method where we instead learn a single labeling function that is approximately invariant under the symplectic map. From the labeling function, we show that we can recover important qualities of the Poincaré plot, including the location of stable orbits, islands and chaos. Additionally, the labeling function can be found with relatively few evaluations of the underlying symplectic map, while being robust to error in the symplectic map.

Maximilian Ruth, David Bindel
Cornell University
mer335@cornell.edu, bindel@cornell.edu

MS393

Propulsion of Flexible Rodlike Magnetic Micro-robots

Inspired by bacterial swimming we design and develop rodlike soft robots that can be propelled in viscous Newtonian fluids. The soft robots consist of tapered, hollow tubes constructed of a polyacrylamide-based hydrogel polymer, with magnets embedded at both ends. The magnets allow actuation via torque applied by a uniform rotating magnetic field. We demonstrate the propulsion of the soft-robots in viscous Newtonian fluids, and connect it to the deformations of the soft robot under magnetic actuation. Our results suggest that the geometrical asymmetry of tapering enables propulsion, and suggest ways to further refine the

design of such soft robots.

Henry Fu
University of Utah
henry.fu@utah.edu

MS393

Image-Based Flow Simulation of Platelet Aggregates under Variable Shear Flow

As the initial step of hemostasis and thrombosis, platelet aggregates significantly affect the formation and stability of thrombi. The mechanical properties of such aggregates are influenced by their microstructure and their interaction with the flow during formation. However, due to technical limitations, it is difficult to study the intrathrombus conditions of platelet aggregates only by experimental or computational approaches. In this work, an image-based computational model is proposed based on two sets of microscopy images of three different platelet aggregates formed under different external flow rates to study the blood flow behavior within such aggregates. These image sets are captured from whole blood perfused microfluidic chambers coated with collagen. One set of images captured the shape of the aggregate outline, while the other employed platelet labeling to infer the internal density. The platelet aggregates are considered as porous mediums. The computational model is subsequently applied to study the blood flow behavior and the interaction between the platelet aggregates and the blood flow under different shear rates. The penetration depth, represented by the boundary of convective and diffusive transport, is also investigated. The results show that external flow velocity has a significant influence on the agonist transport in the outer layers of the aggregate, while the internal core stably remains in the diffusion dominated regime.

Yue Hao, Gabor Zavodszky
University of Amsterdam
y.hao@uva.nl, g.zavodszky@uva.nl

Claudia Tersteeg
KU Leuven
claudia.tersteeg@kuleuven.be

Alfons Hoekstra
University of Amsterdam
a.g.hoekstra@uva.nl

MS393

Understanding Bacterial Swimming in *E. Coli*

We present an in silico microswimmer motivated by peritrichous bacteria such as *E. coli*, which can run and tumble by spinning their flagellar motors counterclockwise (CCW) or clockwise (CW). Runs are the directed movement driven by a flagellar bundle and tumbles are reorientations of cells caused by some motor reversals from CCW to CW. Our simulations reveal that physical properties of the hook and the counterrotation of the cell body are important factors for efficient flagellar bundling, and that longer hooks in mutant cell models create an instability and disrupt the bundling process, resulting in a limited range of movement. Moreover, we demonstrate that cells can explore environment near surface by making variations of tumble events as they swim close to the surface.

Sookkyung Lim
University of Cincinnati

sookkyung.lim@uc.edu

Wanho Lee
National Institute for Mathematical Sciences
lwh3958@gmail.com

Yongsam Kim
Chung-Ang University
kimy@cau.ac.kr

Charles S. Peskin
Courant Institute of Mathematical Sciences
New York University
peskin@cims.nyu.edu

MS393

Effects of Prey Capture on the Swimming and Feeding Performance of Choanoflagellates

Locomoting organisms often carry loads such as captured prey or young. Load-carrying effects on high Reynolds number flight have been studied, but the fluid dynamics of load-carrying by low Reynolds number microorganisms have not. We studied low Reynolds number load-carrying using unicellular choanoflagellates, which wave a flagellum to swim and create a water current transporting bacterial prey to a food-capturing collar of microvilli. A regularized Stokeslet framework was used to model the hydrodynamics of a swimming choanoflagellate with bacterial prey on its collar. Both the model and microvideography of choanoflagellates showed that swimming speed decreases as number of prey being carried increases. Flux of water into the capture zone is reduced by bacteria on the collar, which redirect the water flow and occlude parts of the collar. Feeding efficiency (prey captured per work to produce the feeding current) is decreased more by large prey, prey in the plane of flagellar beating, and prey near microvillar tips than by prey in other locations. Some choanoflagellates can attach themselves to surfaces. We found that the reduction in flux due to bacterial prey on the collars of these attached thecate cells was similar to the reduction in flux for swimmers.

Hoa Nguyen, Emma C. Ross
Trinity University
hnguyen5@trinity.edu, eross1@trinity.edu

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

Lisa J. Fauci
Tulane University
Department of Mathematics
fauci@tulane.edu

M.A.R. Koehl
University of California Berkeley
Integrative Biology
cnidaria@berkeley.edu

MS393

Results on Classical Elastohydrodynamics for a Swimming Filament

We consider two models of an immersed inextensible filament undergoing planar motion in \mathbb{R}^3 : (1) the classical

elastohydrodynamic model using resistive force theory coupled with Euler-Bernoulli beam theory, and (2) a novel curve evolution formulation incorporating the effects of linear viscoelasticity. We mention our recent PDE results on these models and highlight how this analysis can help to better understand undulatory swimming at low Reynolds number. This includes the development of a novel numerical method to simulate inextensible swimmers in Newtonian and viscoelastic media.

Laurel Ohm

Princeton University
laurel.ohm@princeton.edu

MS394

On a Dynamic Variant of the Iteratively Regularized Gauss-Newton Method with Sequential Data

In this talk we will present a variant of the iterative regularized Gauss-Newton method, which is a well-known methodology for solving inverse problems. This talk, and work, is motivated from the question of whether one can improve the computation based on sequential data rather than using one single instance of the data. As a result, we present a dynamic version which considers sequential data at every iteration. We introduce two new algorithm, which we analyze both numerically and computationally. We provide various results such as convergence, well-posedness and error bounds, while testing our new algorithms on a range of PDEs.

Neil Chada

King Abdullah University of Science and Technology
neilchada123@gmail.com

Marco Iglesias

University of Nottingham
marco.iglesias@nottingham.ac.uk

Shuai Lu

School of Mathematics
Fudan University
slu@fudan.edu.cn

Frank Werner

Julius-Maximilians-Universität of Würzburg
frank.werner@mathematik.uni-wuerzburg.de

MS394

Projected Particle Filtering

Data assimilation of high-dimensional nonlinear models is subject to curse of dimensionality. It is when an ensemble of small size is unable to reduce an error of the estimate. Typical remedy to the curse of dimensionality is distance-based localization. Distance-based localization reduces the model state dimension by taking into account only a few numerical cells of the model state near each observation. Even though distance-based localization reduces the error substantially for both linear data-assimilation methods such as ensemble Kalman filter and nonlinear data-assimilation methods such as particle filtering, linear data-assimilation methods still considerably outperform nonlinear data-assimilation methods in linear and quasi-linear regimes. We propose a further dimension reduction based on projection. We analyze the proposed projected ensemble Kalman filter and the projected particle filter in terms of error propagation. The numerical results show considerable error decrease when used with small ensemble sizes.

Svetlana Dubinkina

VU Amsterdam
s.b.dubinkina@vu.nl

Jana de Wiljes

University of Potsdam
wiljes@uni-potsdam.de

MS394

Optimal Regularized Hypothesis Testing in Statistical Inverse Problems

In many inverse problems, one is not primarily interested in the whole solution $u^\dagger \in \mathcal{X}$, but in specific features of it that can be described by a family of linear functionals of u^\dagger . We perform statistical inference for such features by means of hypothesis testing. This problem has recently been treated by multiscale methods based upon unbiased estimates of those functionals. Constructing hypothesis tests using unbiased estimators, however, has two severe drawbacks: Firstly, unbiased estimators only exist for sufficiently smooth linear functionals, and secondly, they suffer from a huge variance due to the ill-posedness of the problem, so that the corresponding tests have bad detection properties. We overcome both of these issues by considering hypothesis tests with maximal power among all tests based upon linear estimators that have a given level of significance. While the construction of such optimal tests requires knowledge of the true solution u^\dagger , we present a way to compute approximately optimal hypothesis tests adaptively. We study this approach both analytically and numerically for linear inverse problems and compare it with unregularized hypothesis testing.

Remo Kretschmann

University of Würzburg
remo.kretschmann@mathematik.uni-wuerzburg.de

Daniel Wachsmuth

Institute of Mathematics
University of Würzburg
daniel.wachsmuth@mathematik.uni-wuerzburg.de

Frank Werner

Julius-Maximilians-Universität of Würzburg
frank.werner@mathematik.uni-wuerzburg.de

MS394

Inhomogeneous Regularization for Ensemble Kalman Inversion

Regularization stabilizes ill-conditioned inverse problems by imposing prior information/characteristics of the unknown signal to recover. l_1 regularization, for example, imposes sparsity of the unknown signal, while l_2 regularization imposes smoothness of the signal. For a signal with both sparsity and smoothness, it is natural to use regularization that changes over different signal locations to account for corresponding characteristics, which we call inhomogeneous regularization. In this work, we use Ensemble Kalman Inversion for inhomogeneous regularization. The ensemble provides prior information for classifying characteristics to determine the power p for l_p regularization. Once the characteristics are determined, Ensemble Kalman Inversion uses the transformation-based approach for l_p regularization as an efficient solver for inhomogeneous regularization problems. The work validates the effectiveness

and robustness of the inhomogeneous regularization problem using Ensemble Kalman inversion through a suite of stringent 1D and 2D test problems, including sea ice image recovery that shows both smooth heterogeneous variations and local cracks.

Yoonsang Lee
Dartmouth College
Department of Mathematics
yoonsang.lee@dartmouth.edu

MS394

MCMC Sampling For Sparsity Promoting Bayesian Hypermodels

In this talk, we review some conditionally Gaussian hypermodels that are shown to favor sparse solutions to the Maximum A Posteriori (MAP) estimation problem. A natural question is whether the whole posterior distribution is concentrated on sparse solutions. In this presentation, the question is addressed by MCMC sampling.

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS395

Projection-Based Model Order Reduction for Large-Scale Nonlinear Problems in Structural Mechanics Using An Industrial Code

We present our work on parametric model order reduction (pROM) for a generic class of parametric mechanical problems with internal variables in a nonlinear quasi-static framework. Our methodology is integrated in the qualified industrial grade finite element solver for structural mechanics studies `code_aster`. As part of engineering design, engineers may perform repeated simulations for slightly differing configurations with different physical parameters or geometrical parameters (parametric study). Nevertheless, successive evaluations for nonlinear mechanical problems can entail prohibitive computational costs. We develop an adaptive algorithm based on a POD-Greedy strategy. Since the operator is nonlinear, the computational complexity of the operator assembly scales with the size of the high-fidelity model. We develop an hyper-reduction strategy using an element-wise empirical quadrature in order to speed-up online assembly costs. We introduce a cost-efficient error indicator which relies on the reconstruction of the stress field by a Gappy-POD strategy, which can be used to drive the Greedy procedure. As an illustration and validation of our methodology, we provide namely numerical results on an elasto-plastic holed plate under traction, with physical variable parameters (strain hardening or elasticity coefficients).

Eki Agouzal
Inria, Team MEMPHIS
Institut de Mathématiques de Bordeaux (IMB)
eki.agouzal@inria.fr

Jean-Philippe Argaud
EDF R&D
jean-philippe.argaud@edf.fr

Michel Bergmann
Inria Bordeaux Sud Ouest
michel.bergmann@inria.fr

Guilhem Ferté
EDF R&D
guilhem.ferte@edf.fr

Tommaso Taddei
INRIA Bordeaux South-West
tommaso.taddei@inria.fr

MS395

Non-Intrusive Machine Learning Models of PDEs with Differentiable Programming: An Interpretability Study

Surrogate models of partial differential equations (PDEs) are an important area of research for applications where rapid, accurate predictions are desired with low computational costs. Deep learning is a popular approach, but they typically lack the strong physical constraints intrinsic in PDEs. Differentiable programming is an emerging paradigm that aims to enable the expressivity of neural networks inside PDEs, such that the learned model is intimately connected to the physics of the problem by construction. Recent efforts in differentiable programming have shown promise in learning accurate parameterizations from simulation data with known numerical properties. However, several earth/climate applications have incomplete or partially known PDEs that need non-intrusive parameterization from observational training data. This leads to a significantly challenging learning problem, where the strengths and weaknesses of differentiable programming are less known. Furthermore, PDEs often exhibit chaotic non-local dynamics, which are harder to model than the local dynamics seen in several canonical PDEs. In this work, we systematically study differentiable programming-based strategies to learn such non-local dynamics by training neural networks embedded directly inside a non-local generalized Burgers system. In particular, we represent the nonlinear and non-local terms as neural networks and use backpropagation to train the equations while simultaneously solving the surrogate PDE in the forward pass. Additionally, we investigate the properties of the learned surrogate PDEs, including their sensitivity to system noise, external forcing, impact on prediction accuracy, and comment on potential applications. Our results show that differentiable programming as a paradigm can accurately model PDEs while surpassing vanilla neural networks. Interestingly, it succeeds even when strong assumptions are made about the missing physics while requiring lesser data and computational cost. However, we also discover that the problem specification and numerical methods employed have a non-trivial impact on the quality and stability of the learned surrogate model, with significant implications for various applications.

Arvind T. Mohan
Los Alamos National Lab
arvindm@lanl.gov

Dustin Nguyen
Ohio State University
nguyen1971@osu.edu

Divya Sri Praturi
Max Planck Institute for Meteorology
divya-sri.praturi@mpimet.mpg.de

MS395

Discovery of Interpretable Structural Model Er-

rors by Combining Bayesian Sparse Regression and Data Assimilation

We introduce MEDIA (Model Error Discovery with Interpretability and Data-Assimilation), a framework designed to identify structural errors in models of nonlinear dynamical systems. While traditional calibration methods can reduce parametric uncertainties, it remains a challenge to deduce structural uncertainties from observations. We demonstrate the performance of our framework on a canonical prototype of geophysical turbulence, the twolevel quasigeostrophic system. A Bayesian sparsity-promoting regression framework is used to select the relevant terms from a library of interpretable kernels. As calculating the library from noisy and sparse data (e.g., from observations) using conventional techniques leads to interpolation and numerical errors, here we propose using a coordinate-based multilayer embedding to impute the sparse observations. We demonstrate the importance of alleviating spectral bias, especially with the multiscale nature of turbulent flows, and show a random Fourier feature layer can sufficiently increase accuracy of the kernel terms to enable an accurate discovery. Our framework successfully identifies structural model errors due to linear and nonlinear processes (e.g., radiation, surface friction, advection), as well as misrepresented orography.

Rambod Mojjani, Ashesh K. Chattopadhyay
Rice University
rambod.mojgani@rice.edu, akc6@rice.edu

Pedram Hassanzadeh
Mechanical Engineering and Earth Science
Rice University
pedram@rice.edu

MS396

A PCA Based Nonlinear Preconditioner and Applications

We discuss a novel nonlinearly preconditioned inexact Newton method for solving nonlinear system of algebraic equations from the discretization of highly nonlinear PDEs. The preconditioner is constructed by a decomposition of the nonlinear residual space into two subspaces; one corresponds to the low frequency subspace and the other corresponds to the high frequency subspace. Such a decomposition is obtained by an unsupervised learning algorithm. The new method features a low computational cost and is capable of balancing the overall nonlinearity effectively.

Xiao-Chuan Cai
University of Colorado, Boulder
Dept. of Computer Science
xccai@um.edu.mo

Li Luo
Faculty of Science and Technology
University of Macau
liluo@um.edu.mo

MS396

Physics-Informed Neural Networks to General Domains

The application of machine learning to different engineering fields shows an increasing trend in recent decades. One sub-field in particular, deep learning, emerges from the rest. Deep-learning algorithms are more complex and ab-

stract than traditional machine learning algorithms, which allows them to automate some pieces of the training process and eliminate some of the human intervention. A mesmerizing recent application of deep-learning is the physics-informed neural network, a mesh-free method that can be easily implemented using automatic differentiation from popular deep-learning frameworks. This application is nothing else but a neural network mapping multiple independent variables to a physics field and whose predictions conform to the physics laws described by differential equations. This is achieved by guiding its training with a loss function that includes the equations residuals and the constraint of initial and boundary conditions. Although previous studies solved a wide variety of differential equations, most applications are still limited to simple one- or two-dimensional domains. However, real-life engineering problems involve more complex geometries. This work demonstrates the application of physical-informed neural networks to such geometries. The demonstration cases solve the steady-state heat conduction equation with both Dirichlet and Neumann conditions in two-dimensional geometries with non-orthogonal boundaries to the Cartesian coordinates.

Roberto Fairhurst-Agosta, Tomasz Kozłowski
University of Illinois at Urbana-Champaign
ref3@illinois.edu, txk@illinois.edu

MS396

Hybrid Data- and Knowledge-Driven Multiscale and Multiphase Simulation of the Human Liver

The liver is the central organ for metabolic processes with a complex function-perfusion relationship. This coupled relation can be investigated by numerical multiscale simulations of the total organ, liver lobules and liver cells. To enable robust, efficient, and patient-specific diagnosis and prediction of hepatic diseases through simulations, the purely knowledge-based models are complemented by data-based approaches. To better understand the relation between hepatic perfusion, metabolism and tissue, we have developed a multicomponent, poroelastic multiphase and multiscale function-perfusion model, using a mixture theory based on the Theory of Porous Media. To increase the robustness of the model and allow for near real-time patient-specific simulations, computationally critical parts are enhanced by data-driven methods through the integration of experimental, clinical and in silico data from cooperation partners. In addition to purely data-driven methods, we will increase the model robustness and the overall accuracy of the model predictions by exploring hybrid knowledge- and data-driven approaches. The overarching goal is the definition and development of an accurate decision-supporting framework for clinical applications.

Lena Lambers, Steffen Gerhäuser, Luis Mandl, André Mielke, Tim Ricken
University of Stuttgart
lena.lambers@isd.uni-stuttgart.de,
steffen.gerhaeuser@isd.uni-stuttgart.de,
luis.mandl@isd.uni-stuttgart.de, andre.mielke@isd.uni-stuttgart.de, tim.ricken@isd.uni-stuttgart.de

MS396

Understanding the Relation Between Microstructure and Macroscopic Mechanical Properties of Arteries by Deep Learning

Biological tissues exhibit substantial interindividual differ-

ences and also change their properties during aging. Therefore, modeling the constitutive behavior of biological tissues one cannot simply rely on a library of once measured material parameters as it is typically possible for engineering materials. Rather one has to apply constitutive models tailor-made for a specific tissue of a specific patient at a specific age. To tackle the related challenges in constitutive modeling, it is key to understand the relation between the microstructure of biological tissues and their macroscopic mechanical properties. Here, we introduce a novel machine learning framework that can bridge the gap between the microscale and macroscale in arteries. Using data from mechanical tests, histological analyses and advanced imaging, it can predict the nonlinear macroscopic mechanical properties of arterial tissue. The incorporation of substantial prior knowledge from continuum mechanics and materials theory enables our framework to do so even on the basis of a relatively small amount of data ($10^1 - 10^2$ samples) and yet with high accuracy ($R^2 > 0.9$). Moreover, using so-called relevance propagation, our framework can quantify the role of different microstructural features for the macroscopic mechanical properties in an automatic, systematic and largely unbiased way opening up a host of new insights into the foundations of soft tissue biomechanics.

Kevin Linka

Technical University of Hamburg
kevin.linka@tuhh.de

Cristina Cavinato, Jay Humphrey
Yale University
cristina.cavinat@yale.com, jay.humphrey@yale.com

Christian Cyron
Technical University of Hamburg
christian.cyron@tuhh.de

MS396

Neural Networks Accelerated Simulations for Burn Contraction

Patients with severe burn injuries often suffer from hypertrophic scars, which impair their aesthetic appearance. Next to hypertrophy, skin contraction is a common long-term side effect. If contraction has reached the extent that patients suffer from joint immobility, then the contraction is referred to as a contracture. In order to devise therapies that prevent the formation of contractures, it is important to be acquainted with the underlying biological mechanisms. Since clinical observations are represented in terms of patterns and numbers, mathematical quantification in order to validate biological theory is indispensable. Since the model formulations consist of a set of coupled nonlinear partial differential equations, the finite element strategies result in expensive simulations. In particular, the assessment of uncertainty quantification used to estimate the probability distribution of medical scenarios makes the computations very expensive. In order to have feasible access to the simulations in clinical settings, the simulations are accelerated using artificial intelligence by means of feed-forward neural networks. During the presentation, the mathematical model and the applied artificial intelligence model will be introduced.

Fred Vermolen

Hasselt University
Computational Mathematics Group
Fred.Vermolen@uhasselt.be

Ginger Egberts
Delft University of Technology
g.egberts@tudelft.nl

MS397

Stable Discrete Closure Models Through Embedded Learning

For turbulent fluid flows, neural network closure models can learn to predict sub-grid scale stresses to high accuracy, while still producing instabilities once included in a solver [M. Kurz and A. Beck, Investigating model-data inconsistency in data-informed turbulence closure terms, 14th WCCM-ECCOMAS Congress, CIMNE, 2021]. By minimizing the a-posteriori error of the same models (embedded learning), stability can be enforced [Y. Guan et al., Stable a posteriori LES of 2D turbulence using convolutional neural networks, Journal of Computational Physics, 458:111090, 2022]. This requires a careful choice of hyper-parameters however, e.g. number of time steps unrolled [B. List et al., Learned turbulence modelling with differentiable fluid solvers, arxiv:2202.06988, 2022] and choice of initial parameters [S. D. Agdestein and B. Sanderse, Learning filtered discretization operators, arxiv:2208.09363, 2022]. Care must be taken to avoid local minima and exploding gradients. We investigate the accuracy and stability of closure models learned for PDEs subject to explicit discrete filters using embedded learning. The discrete setting allows for computing exact commutator errors and loss functions. We show that the learned closure models produce stable and accurate trajectories for the solution and energy in time upon extrapolation. We also compare the stability of the models trained while embedded in a solver to the one of those trained to predict the discrete commutator error only.

Syver Døving Agdestein
Centrum Wiskunde & Informatica
sda@cwi.nl

Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cwi.nl

MS397

Model Reduction and Uncertainty Quantification Using Deep Orthogonal Decomposition

In the context of parametrized PDEs, Reduced Order Models (ROMs) allow for an efficient approximation of the parameter-to-solution map, thus enabling expensive many-query routines such as those typical of Uncertainty Quantification (UQ). Motivated by the limitations of classical approaches such as the Reduced Basis method, many authors have now been considering the use of Deep Learning techniques for building nonintrusive ROMs. Here, we propose a novel approach, Deep Orthogonal Decomposition (DOD), where a deep neural network yields a representation of the solution manifold in terms of an adaptive local basis. In principle, designing and training these models can be a challenging task because of the high dimensions involved. As a remedy, we exploit Mesh-Informed Neural Networks (MINNs), a novel class of architectures designed to handle mesh-based data. MINNs embed hidden layers into discrete functional spaces of increasing complexity, obtained through a sequence of meshes defined over the underlying PDE domain. This results in sparse models that are computationally less demanding and thus bet-

ter suited for the DOD. We assess the proposed approach over a broad set of numerical experiments, including complicated domains and high dimensional parameter spaces, and compare it with other state-of-the-art methods, such as principal orthogonal decomposition and autoencoders. Finally, we present some applications to (inverse) UQ, showcasing the advantages entailed by the DOD approach.

Nicola Rares Franco, Andrea Manzoni
Politecnico di Milano
nicolarares.franco@polimi.it, andrea1.manzoni@polimi.it

Paolo Zunino
Politecnico di Milano
Department of Mathematics
paolo.zunino@polimi.it

MS397

Anomaly Detection for Dynamical Systems Using Bayesian Online Change-point Detection

Data-driven analyses of complex fluid flows are frequently limited by the lack of labeled data. In applications that relate to anomaly or regime transition detection, such a scenario precludes the use of supervised learning techniques. In this study, we deploy unsupervised data-driven transition detection for Kolmogorov flow, as a prototype for real-world applications, where anomaly or regime change detection is desirable without historical data. In particular, we propose the use of a deep variational autoencoder, which encodes high dimensional turbulent flow into disentangled low dimensional representations. Benefiting from extremely weak correlation between latent dimensions, we implement latent Bayesian online change-point detection algorithm to identify regime transition. We deploy this state-of-art two-stage data-driven approach on the Kolmogorov flow dataset, revealing an accurate correlation between the detected change-point and a regime transition.

Sen Lin
University of Houston
slin41@cougarnet.uh.edu

Romit Maulik
Argonne National Laboratory
rmaulik@anl.gov

MS397

Embedded Learning Strategies for Neural Closure Models

Discovering physics models is an ongoing, fundamental challenge in computational science. In fluid flow problems, this problem is usually known as the closure problem, and the art is to discover a closure model that represents the effect of the small scales on the large scales. Well-known examples appear in large eddy simulations (LES) and in reduced-order models (ROMs). Recently, it appeared that highly accurate closure models can be constructed by using neural networks. However, integrating the neural network into the physics models (neural closure models) is typically prone to numerical instabilities as the training environment does not match the prediction environment. Instead, we investigate learning closure models *while they are embedded in a discretized PDE solver* by using differentiable programming software. This "trajectory fitting" or "embedded learning" approach is a more difficult learning problem than the more classical approach of operator fitting, but promises better accuracy and stabil-

ity. We compare two strategies to solve the trajectory fitting problem: discretise-then-optimise and optimise-then-discretise. Furthermore, we present a new neural closure model form which allows us to preserve structure, namely kinetic energy conservation, and therefore non-linear stability bounds.

Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cw.nl

Hugo Melchers
Centrum Wiskunde & Informatica (CWI)
ham@cw.nl

Daan Crommelin
CWI Amsterdam and University of Amsterdam
Daan.Crommelin@cw.nl

Barry Koren
Department of Mathematics and Computer Science
Eindhoven University of Technology
b.koren@tue.nl

MS397

Non-Intrusive Reduced-order Modeling for Blood Flow Simulations with Surface Registration

In cardiovascular science and engineering, high-fidelity blood flow simulations are computationally expensive. Snapshot-based reduced-order modeling is one of the common remedies to the problem. However, the anatomical geometries used for simulations are either segmented from clinical image data or generated indirectly by computational models, which are typically unparameterized and lead to distinct spatial discretisations. In this work, a data-driven surrogate model based on non-intrusive reduced-order modelling with surface registration is proposed. The surrogate model approximates the fluid dynamics of blood flow within vessels of distinct but similar shapes. Different from existing approaches, the proposed surrogate model performs a surface registration to approximate the shape of the target domain from the reference shape. The diffeomorphism constructed during surface registration provides the mapping between the reference domain and the target domain, as well as the parametrisation of the shape. The simulations are subsequently performed on the reference domain to ensure the spatial compatibility of the snapshots. Proper orthogonal decomposition is applied to construct reduced bases and the projection coefficients are computed via a radial basis function interpolation method. Two numerical examples based on stenosis and bifurcation vessels are presented and discussed to demonstrate the efficiency and accuracy of the surrogate model.

Dongwei Ye, Valeria Krzhizhanovskaya, Alfons Hoekstra
University of Amsterdam
d.ye@uva.nl, v.krzhizhanovskaya@uva.nl,
a.g.hoekstra@uva.nl

MS398

Spectral Deferred Correction Methods for Second-Order Problems

Spectral deferred corrections (SDC) is an iterative method for the numerical solution of ordinary differential equations. SDC can be understood as applying an appropri-

ate preconditioner to a Picard iteration to obtain faster and more robust convergence. SDC has been studied primarily for first-order problems, using low-order methods as preconditioners (e.g., explicit, implicit, or implicit-explicit Euler). It has been shown that SDC can achieve arbitrarily high order of accuracy and possesses good stability properties. In this talk, we will present a theoretical analysis of the convergence and stability properties of SDC when applied to second-order ODEs and using velocity-Verlet as a preconditioner. While there are a few numerical studies of the performance of SDC applied to the second-order Lorentz equations (so-called Boris-SDC), there exists so far no theoretical analysis of the properties of SDC for general second-order problems. We will present a theorem that shows that the order of convergence depends on whether the force on the right-hand side of the system depends on the velocity (like in the Lorentz equation) or not (like in the undamped harmonic oscillator). We also show a theorem that guarantees convergence of SDC under certain conditions. We compare the stability domain of SDC with the Picard iteration and validate our theoretical analysis in numerical examples.

Ikrom Akramov, Daniel Ruprecht
Institute of Mathematics
Hamburg University of Technology
ikrom.akramov@tuhh.de, ruprecht@tuhh.de

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich GmbH
r.speck@fz-juelich.de

Michael Minion
Lawrence Berkeley National Lab
mlminion@lbl.gov

Sebastian Götschel
Hamburg University of Technology
sebastian.goetschel@tuhh.de

MS398

Time-Adaptive Spectral Deferred Corrections

The complexity of the dynamics of most computational problems is not uniform throughout time. This means that fixed resolution schemes will either over-resolve less complex areas or under-resolve the more complex parts. In principle, adaptively adjusting resolution provides a remedy. However, it can be computationally expensive since it requires the estimation of local errors and often costly bookkeeping of the discretization data structure. We present here temporal adaptivity based on an embedded version of spectral deferred corrections (SDC), which estimates the local error similarly to embedded Runge-Kutta methods. We are able to estimate the local error without additional right-hand-side evaluations and provide dynamic temporal resolution within the pySDC software framework. In addition, we show how adaptivity enables other improvements like fault-tolerant time-stepping.

Thomas Baumann
Forschungszentrum Juelich
Technische Universitaet Hamburg Harburg
t.baumann@fz-juelich.de

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich GmbH

r.speck@fz-juelich.de

Daniel Ruprecht
Institute of Mathematics
Hamburg University of Technology
ruprecht@tuhh.de

Thibaut Lunet, Sebastian Götschel
Hamburg University of Technology
thibaut.lunet@tuhh.de, sebastian.goetschel@tuhh.de

MS398

Rosenbrock-Wanner and W-Methods for Stochastic Galerkin Systems

Rosenbrock-Wanner methods represent linear-implicit techniques to solve initial value problems of stiff ordinary differential equations (ODEs). In W-methods, the exact Jacobian matrix of the right-hand side is replaced by an approximation to reduce the computation work. We consider a system of ODEs including random variables, which is used in uncertainty quantification, for example. The stochastic Galerkin approach changes the random-dependent system into a larger deterministic system of ODEs. We apply a W-method to the stochastic Galerkin system, where the approximation of the Jacobian matrix has a block-diagonal form. Each block is the Jacobian matrix of the right-hand side in the original random-dependent ODEs inserting the expected value. We investigate the stability of the W-method for stiff linear systems of ODEs. The concept of A-stability is extended using a scalar linear ODE with a single negative random variable. It follows that the stability of the W-method depends on the eigenvalues of the resulting system matrix, the expected value of the random variable, and the coefficients of the numerical scheme. We discuss the stability conditions for W-methods up to four stages with maximum order of convergence. The construction of a W-method satisfying the stability condition is possible for each probability distribution. We present results of numerical computations for a test example.

Roland Pulch
University of Greifswald
pulchr@uni-greifswald.de

Adrian Sandu
Virginia Polytechnic Institute and State University
Computational Science Laboratory
sandu@cs.vt.edu

MS398

Efficient Implicit Multiderivative Time Integration for Compressible Flows

In this talk, we show how to use an implicit two-derivative time integration process efficiently with a discontinuous Galerkin solver. The area of application is the compressible Navier-Stokes equations. We focus on a stable implementation as well as on the possibilities of time-parallelism through a pipelining procedure; numerical results demonstrating the efficiency will be shown.

Jochen Schuetz
Universiteit Hasselt
jochen.schuetz@uhasselt.be

David C. Seal

United States Naval Academy
seal@usna.edu

Jonas Zeifang
Universiteit Hasselt
jonas.zeifang@uhasselt.be

MS398

Interleaving of SDC with DoF Subset Selection for Improved Efficiency in Electrophysiology Simulations

The electrical excitation of cardiac tissue is described by PDE models of reaction-diffusion type coupled to pointwise ODEs, ranging from the simplest monodomain equation to EMI models with cellular resolution. Due to fast reactions and small spatial extent of activation fronts, small time steps and mesh widths are required for a sufficiently accurate discretization, incurring a substantial computational effort. The highly localized dynamics of these models call for adaptive simulation methods. Unfortunately, the overhead incurred by classical mesh adaptivity turns out to outweigh the performance improvements achieved by reducing the problem size. Here, we explore a different approach to adaptivity based on progressive subset selection of algebraic degrees of freedom during time stepping with spectral deferred correction methods. This realizes a kind of multirate higher order integration but allows for a natural and efficient definition of artificial boundary conditions justified by a priori error estimates based on the maximum principle. Numerical experience indicates a significant performance increase compared to simulations on uniform grids.

Martin Weiser
Zuse Institute (ZIB)
Berlin, Germany
weiser@zib.de

Fatemeh Chegini
Zuse Institute Berlin
chegini@zib.de

MS399

An Efficient Kernel-Based Method to Solve Fractional Differential Equations

We describe a reproducing kernel Hilbert space framework that gives rise to a naturally linked efficient Galerkin-based numerical solver for fractional differential equations. Some theory as well as numerical examples will be presented. This is joint work with Hamed Mohebalizadeh and Hoja-tollah Adibi (Amirkabir University of Technology, Tehran).

Greg Fasshauer
Colorado School of Mines
fasshauer@mines.edu

MS399

PDE-Greedy Kernel Methods for Adaptive Approximation of Partial Differential Equations

Kernel greedy methods have been shown to be provably and empirically efficient for general function approximation tasks. These methods can be extended to meshless approximation of linear PDEs on domains or manifolds by symmetric collocation. We present different collocation point selection strategies, especially the PDE-f-greedy and

PDE-P-greedy procedure. We provide analytical results on convergence rates and demonstrate the applicability by computational experiments for elliptic PDEs.

Bernard Haasdonk, Tizian Wenzel, Daniel Winkle
University of Stuttgart
haasdonk@mathematik.uni-stuttgart.de,
tizian.wenzel@ians.uni-stuttgart.de,
daniel.winkle@mathematik.uni-stuttgart.de

Gabriele Santin
Center for Information and Communication Technology
Bruno Kessler Foundation, Trento, Italy
gsantin@fbk.eu

MS399

Hp-Adaptive RBF-FD Method

In general, the desired accuracy of the numerical solution varies over the domain of interest. A typical example of such a situation is a contact between two bodies with concentrated stress under the contact region, which requires a more accurate simulation than in the rest of the domain. Similarly, in flow simulations we often deal with geometries with narrow passages or other geometric complications for which a more refined solution is desired. In fact, most real-world problems we wish to simulate exhibit such behaviour at some point. This can be addressed with two conceptually different adaptive approaches, namely p-adaptivity or h-/r-adaptivity. In p-adaptivity, the solution accuracy is varied by changing the order of approximation, while in h-/r-adaptivity the resolution of the spatial discretisation is adjusted for the same purpose. Ultimately, h- and p-adaptivities can be combined into hp-adaptivity for maximum performance. Since the regions where higher accuracy is required are usually not known a priori, and to avoid human intervention in the solution procedure, a measure of the quality of the solution, commonly referred to as an a posteriori error indicator, is a necessary additional step in an adaptive solution procedure. In this talk we will present our work on the hp-adaptive RBF-FD method for solving problems from linear elasticity and fluid flow.

Gregor Kosec
Parallel And Distributed Systems Laboratory
Joef Stefan Institute
gkosec@ijs.si

Mitja Jancic
Institute Jozef Stefan
mitja.jancic@ijs.si

MS399

Improved Error Estimates for RBF Interpolation

The usual error estimates for interpolation with radial basis functions (RBF) provide convergence rates if the error is measured in a low smoothness norm. Those rates deteriorate if the error is measured in stronger norms. Numerically, often higher convergence rates than the predicted ones are observed if smooth functions are approximated. One explanation is the so-called doubling trick. In this talk we will extend those doubling results to the case, where the error is measured in norms which are stronger than the so-called native space norm. This is based on joint work with Thomas Hangelbroek

Christian Rieger
Philipps-Universität Marburg

riegerc@mathematik.uni-marburg.de

MS399

A New Fractional Differentiation Formula for Caputo Fractional Derivatives and Its Application

In this talk, a new fractional differentiation formula is proposed for the Caputo fractional derivative of order $\alpha \in (1, 2)$. We have applied Cubic Hermite Interpolation polynomial in the subinterval $[t_0, t_2]$ and cubic Lagrange's Interpolation polynomial in the subinterval $[t_{j-1}, t_j]$, for $j \geq 3$. This new differentiation formula is of second order. A difference scheme is proposed with the help of this approximation formula to solve the time-fractional wave equation (TFWE). The stability and convergence of the numerical scheme are also analyzed. To show the effectiveness and accuracy of the proposed numerical scheme, two numerical examples are also provided. The scheme is observed to be fast, easy, and provide highly accurate results.

Nikhil Srivastava

Indian Institute of Technology (Banaras Hindu University)
Varanasi, U.P., INDIA
nikhils.iitbhu@gmail.com

MS400

Using Both Particle and Continuum Discretizations

PIC discretizations require a coordination between discrete fields supported on meshes, and discrete fields supported on particles. We will discuss the infrastructure for this coordination in PETSc, including conservative projections between these spaces, accurate integration in RBF overlaps, adaptivity in both domains, and integration with scalable solvers.

Matthew G. Knepley, Joseph Puszta
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com, josephpu@buffalo.edu

Daniel Finn
University at Buffalo
dsfenn@buffalo.edu

Mark Adams
Lawrence Berkeley National Laboratory
mfadams@lbl.gov

MS400

The Pseudo-Inverse Mapping in the Global Gyrokinetic Magnetic Fusion Code XGC

The X-point gyrokinetic code (XGC) is a 5D total-f gyrokinetic particle-in-cell (PIC) code being used to study multiscale kinetic turbulence and transport phenomena in magnetic fusion plasmas [<https://xgc.pppl.gov>]. XGC is a US ECP (Exascale Computing Project) and SciDAC code. XGC evaluates dissipative operations such as Fokker-Planck Coulomb collisions and heat sources/sinks on a 5D (3D configuration and 2D velocity space) grid at each time step. This requires a mapping back-and-forth between the marker particles and the 5D phase-space grid. To enable exact conservation properties (of particle number, momentum and energy) and to avoid significant error accumulation over time in plasmas with steep density and tempera-

ture gradients, as appears in a tokamak edge, the code has recently been equipped with a novel mapping in 2D velocity space based on the calculation of a pseudo-inverse. This mapping has been implemented in XGC using the PETSc library and generalized to use a discretization space of any order. A complication of the pseudo-inverse mapping is that the calculation is only well-posed if the particle coverage is good enough in the simulations. Here we explore the performance and the robustness of the mapping in various scenarios using different numerical solvers in PETSc and possible solutions to obtain sufficient particle coverage. We also discuss the possibility of using the mapping as a coarse-graining technique to control particle noise.

Albert Mollen

Princeton Plasma Physics Laboratory
amollen@pppl.gov

Mark Adams
Lawrence Berkeley National Laboratory
mfadams@lbl.gov

Varis Carey
Theodon Consulting
varis.carey@theodonconsulting.com

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Robert Hager
Princeton Plasma Physics Laboratory
rhager@pppl.gov

Aaron Scheinberg
Jubilee Development
aaron@jubileedev.com

Choong-Seock Chang
Princeton Plasma Physics Laboratory
cschang@pppl.gov

MS400

Structure Preserving Hybrid Models for Plasma Physics

We present STRUPHY (STRUcture-Preserving HYbrid codes) - a Python package for energetic particles (EPs) in plasma. The package is easy-to-use (PyPI and Docker installations) and features a collection of PDE solvers for fluid-kinetic hybrid models such as MHD-kinetic current and pressure coupling schemes, among others. Such models represent a decent compromise between physical content and numerical cost, as the bulk plasma is treated with a relatively "cheap" fluid model and only the EPs are described kinetically. The discretization is based on the GEMPIC framework [Kraus et al., GEMPIC: Geometric ElectroMagnetic Particle-In-Cell Methods, J. Plas. Phys. 2017] and leads to provable stability due to exact conservation laws. We will discuss some of the implemented hybrid models, their discretization and some numerical results in the context of magnetized plasmas. STRUPHY kernels are accelerated to Fortran speed using the open source accelerator Pyccel [<https://github.com/pyccel/pyccel>]; an MPI/OpenMP hybrid parallelization allows for HPC application of STRUPHY. The package is available under

<https://gitlab.mpcdf.mpg.de/struphy/struphy>.

Stefan Possanner

Max-Planck-Institut für Plasmaphysik
stefan.possanner@ipp.mpg.de

MS400

Particle Basis Landau Collisions in PETSc

In recent years, theoretical advances have demonstrated the conservative properties of a particle basis discretization of the Landau collision integral. Additional time stepping schemes have been proposed to maintain the conservative properties of this discretization. I present an implementation of the conservative particle basis Landau collision operator in the Portable Extensible Toolkit for Scientific Computing (PETSc) for use in plasma codes without the need for basis projection at the collision step.

Joseph Pusztay

University at Buffalo
Department of Computer Science and Engineering
josephpu@buffalo.edu

MS401

Simulation of Reverse Osmosis Membrane Compaction Using Material Point Method

Reverse Osmosis (RO) is a promising technology to address the impending water-crisis in the upcoming decades. RO at high salinities and high pressures is challenging due to membrane compaction that changes its porosity and permeability. In this study, we present a simulation methodology for membrane structural mechanics to understand pore size distribution and permeability variations under high pressure. We use the material point method (MPM), that solves the solid mechanics equations in a Lagrangian framework. MPM provides many features that make it well-suited for simulating mesoporous membranes. The Lagrangian framework allows for large deformations, easy integration of constitutive models and direct import of complex geometries as particles. The spatial discretization in our MPM solver is achieved using linear or cubic-spline shape functions while the time integration is carried out using the explicit Euler method. A series of images containing detailed pore structures obtained from X-ray tomography is first converted to a collection of material points. Compressive loads are applied to the top layer of the membrane to simulate the application of pressure. The membrane deformation and pore size distribution before and after load application are reported and compared with the experimental measurements. The presentation discusses the numerical methods used, the performance of the solver on high-performance computing machines, and the results of membrane compact in detail.

Sreejith Appukuttan

National Renewable Energy Laboratory, Colorado
sreejith.nadakkalappukuttan@nrel.gov

Hariswaran Sitaraman

National Renewable Energy Laboratory
hariswaran.sitaraman@nrel.gov

Marc Day

National Renewable Energy Laboratory, Colorado
marcus.day@nrel.gov

Yara Suleiman, Sina Shahbazmohamadi

University of Connecticut, CT

yara.suleiman@uconn.edu,
sina.shahbazmohamadi@uconn.edu

Jishan Wu, Eric Hoek

University of California Los Angeles, CA
wujishan123@g.ucla.edu, emvhoek@ucla.edu

MS401

A Transport Model for Cell Motion

Two phase flow models have been shown to be useful to understand cell motility. Simple 1D, two-phase, poroviscous, reactive flow model capture behaviors relevant to cell crawling. These models assume constant cell length. We release that assumption and study the behavior of the resulting transport equations. We devise adequate numerical schemes and simulate travelling solutions capturing new features of cell motion.

Ana Carpio

Universidad Complutense de Madrid
Departamento de Matematica Aplicada
carpio@mat.ucm.es

Rafael González-Albaladejo

Universidad Complutense de Madrid, Spain
rafael09@ucm.es

Falko Ziebert

Heidelberg University, Institute for Theoretical Physics
f.ziebert@thphys.uni-heidelberg.de

MS401

Transport Equation from a Stochastic Model of Angiogenesis

Angiogenesis is a multiscale process of blood vessel growth from a primary vessel. It is of paramount importance in healthy processes such as organ growth and repair and wound healing, and also in pathologies such as cancer or retinopathies. It is triggered by the emission of vessel growth substances from a region of tissue lacking oxygen. On mesoscopic scales large compared to cell size but small compared to organ size, it can be modeled using stochastic differential equations and birth/death processes for branching of blood vessels and for their merging (anastomosis). From the stochastic process, we have derived a transport equation for the density of active tips of blood vessels, which is parabolic and contains source terms for branching and anastomosis that include a memory term. It is coupled to a reaction-diffusion equation for the growth factor concentration. Numerical simulation of the transport equation and the underlying stochastic process show that the density evolves towards a moving soliton-like wave whose velocity and shape change slowly according to some ordinary differential equations for collective coordinates. We have also developed a theory of fluctuations about the density by adding multiplicative noises to the transport equation. The moments of the density turn out to be given by powers of the soliton solution with appropriate exponents.

Manuel Carretero

Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience & IndMath
manuel.carretero@uc3m.es

Bjorn Birnir
University of California Santa Barbara
Department of Mathematics
birnir@math.ucsb.edu

Luis Bonilla
Universidad Carlos III de Madrid
G. Millan Institute of Fluid Dynamics, Nanoscience
& IndMath
bonilla@ing.uc3m.es

Filippo Terragni
Universidad Carlos III de Madrid, Spain
fterragn@ing.uc3m.es

MS401

Electrothermal Monte Carlo Simulation of a Resonant Tunneling Diode

Monte Carlo technique for the solution of the Wigner transport equation has been developed, based on the generation and annihilation of signed particles. A stochastic algorithm without time discretization error has been recently introduced (Muscato et al. *Kin.Relat. Models*, 12(1), 59-77, 2019), based on the theory of piecewise deterministic Markov processes. Numerical experiments have been performed in the case of a GaAs Resonant Tunneling Diode (RTD), improving the CPU consumption (Muscato, *J. Comput. Electron.*, vol 20., pagg. 2062-2069, 2021). In such nano-devices the presence of very high and rapidly varying electric fields is the cause of thermal heating of the carriers and the crystal lattice. In fact, the external electric field transfers energy to the electrons and in turn to the lattice through the scattering mechanism. This self-heating process can influence significantly the electrical behaviour because the dissipated electrical energy causes a temperature rise in the device resulting in increased power dissipation. The mechanism through which lattice self-heating occurs is that of electron scattering with phonons, and therefore a model which deliberately incorporates all scattering events will also capture such energy dissipation details. Following the guidelines developed in Muscato et al. *Comp. Math. with Applications*, **65**, 520-527, (2013), the heating phenomenon in a RTD will be tackled.

Orazio Muscato
Dipartimento di Matematica e Informatica,
Universita' di Catania
orazio.muscato@unict.it

MS402

Performance-Portable Solvers for Nonlinear Mechanics

Finite element analysis in solid mechanics typically uses low order elements and assembled sparse matrices. We demonstrate that matrix-free methods based on libCEED are faster for all finite element orders, and that increasing the order of accuracy further reduces the cost per degree of freedom, leading to order of magnitude reduction in simulation cost to reach engineering tolerances. We discuss the design of a matrix-free p-multigrid method and explore solver robustness and efficiency in nonlinear analysis of lattice structures and bonded particulate materials with billions of degrees of freedom, with special attention to recent GPU-based architectures. We also demonstrate how the same principles enable transformative capability

in simulation of compressible turbulence.

Jed Brown, Valeria Barra
University of Colorado Boulder
jed@jedbrown.org, valeria@caltech.edu

Natalie N. Beams
University of Tennessee
nbeams@icl.utk.edu

Leila Ghaffari
University of Colorado Boulder
leila.ghaffari@colorado.edu

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Rezgar Shakeri, Karen Stengel, Jeremy Thompson, James Wright
University of Colorado Boulder
rezgar.shakeri@colorado.edu, karen.stengel@colorado.edu,
jeremy.thompson@colorado.edu,
james.wrightiii@colorado.edu

Junchao Zhang
Argonne National Laboratory
jc Zhang@mcs.anl.gov

MS402

Applications of a Tiled Monte Carlo Algorithm to the Computation of Matrix Functions

We extend our prior work on Monte Carlo algorithms for the solution of large linear equations to compute other matrix functions such as exponential and logarithm. One of the only available tools for huge matrices is the Monte Carlo algorithm, which also has the advantage of being able to compute matrix-free, i.e., computing matrix elements as needed, which is essential for truly huge problems for which it is not even possible to store a full vector. To facilitate this latter scenario, we compute the projection of the various matrix functions onto the vector we are interested in with the assumption that we can compute the vector elements also on demand. However, it has been shown that this one-element-at-a-time approach cannot guarantee convergence of the statistical error. On the other hand, our recent algorithm that computes with matrix tiles can guarantee convergence for sufficiently large tiles. We compute matrix functions by summing a polynomial approximation (e.g., Taylor, Chebyshev, etc.). We first investigate the convergence conditions for each matrix and provide an expression for the statistical error. Then, we optimize the algorithm in terms of operation time and accuracy by adjusting the parameters such as the tile size and the number of samples for given target error. We also apply variance reduction techniques to reduce the statistical error with the same computational cost. Finally, we show numerical examples from quantum chemistry.

Hyeji Choi
SUNY Stony Brook
Institute for Advanced Computational Science

hyeji.choi.1@stonybrook.edu

MS402

Global-Coarsening Multigrid for hp -adaptive Finite Element Computations

hp -adaptive finite element methods allow to choose both the mesh size h and the polynomial degree p locally on every cell. Despite their excellent convergence properties, they are not widely used, possibly, due to their challenging implementation. Recent comprehensive support for hp -FEM in the deal.II library attempts to close this gap, yet reveals new challenges: We observed that common off-the-shelf matrix-based solvers are inadequate for hp -FEM applications both from a memory-consumption and a performance point of view. As a compelling alternative solver, we present herein a hybrid multigrid preconditioner that globally coarsens both p and h and can be evaluated in a state-of-the-art matrix-free fashion on parallel distributed systems. We compare its performance and scalability to algebraic multigrid methods, using realistic computations on up to many thousands of MPI processes as test cases.

Marc Fehling
Colorado State University
Department of Mathematics
marc.fehling@colostate.edu

Peter Munch
University of Augsburg
High-Performance Scientific Computing
peter.muench@uni-a.de

Wolfgang Bangerth
Colorado State University
Department of Mathematics
bangerth@colostate.edu

MS402

Application and Tailoring of Matrix-Free Algorithms to Many-Particle Solid-State-Sintering Phase-Field Implementations

Biodegradable magnesium implants bear a great potential since they degrade after a certain period, during which they temporarily support the healing process. Solid-state-sintering processes are able to produce patient-specific implants. In order to understand and improve the manufacturing process, we perform experiments and develop numerical tools, e.g., based on phase-field (PF) methods, at the Helmholtz-Zentrum Hereon. To be able to simulate problems at late simulation stages when a spherical approximation of particles is not applicable with many particles (1000), we have reimplemented parts of our PF code, leveraging state-of-the-art matrix-free finite-element building blocks. In this presentation, we summarize our efforts in porting a sintering application to matrix-free FEM and discuss challenges we faced in solving such a highly nonlinear and dynamic problem: in particular, 1) the challenge of dealing with a high number of scalar components (up to 15), which dynamically changes, when particles move or merge; and 2) the development of a block preconditioner for the solution of the Jacobian matrix, which we evaluate in a matrix-free fashion. Our solver employs deal.II for spacial discretization and adaptive mesh refinement as well as Trilinos non-linear and linear solver capabilities. Results will also contain performance analyses conducted

on SuperMUC-NG supercomputer.

Peter Munch
University of Augsburg
High-Performance Scientific Computing
peter.muench@uni-a.de

Vladimir Ivannikov
Helmholtz-Zentrum Hereon
vladimir.ivannikov@hereon.de

Martin Kronbichler
University of Augsburg
martin.kronbichler@uni-a.de

MS402

Matrix-Free Low-Order-Refined Preconditioners for High-Order $\mathbf{H}(\text{curl})$ and $\mathbf{H}(\text{div})$ Problems with GPU Acceleration

While efficient algorithms are well-known for the application of high-order finite element operators, the development of effective solvers and preconditioners for the resulting large-scale linear systems remains a more challenging problem. One key challenge is that the associated matrices are generally too dense to compute and store, necessitating the development of matrix-free preconditioners. In this talk, I will describe the construction of a class of *low-order-refined* preconditioners for problems posed in all spaces of the high-order finite element de Rham complex. This construction is facilitated by the use of an *interpolation-histopolation basis* that guarantees spectral equivalence of the high-order and low-order-refined operators, independent of the mesh size h and polynomial degree p . Specialized algebraic multigrid methods can then be used to obtain scalable, high-performance, GPU-accelerated, matrix-free solvers for the high-order problems. Additionally, I will describe the development of new matrix-free solvers for grad-div problems and mixed finite element problems in $\mathbf{H}(\text{div})$. These solvers are based on solving the saddle-point system, and also use properties of the interpolation-histopolation basis to achieve enhanced sparsity.

Will Pazner
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
pazner@pdx.edu

MS403

Approximate Bayesian Computation Parameter Estimation for a Model of Self-Generated Cell Chemotaxis

Approximate Bayesian computation (ABC) has gained considerable interest in recent years as it allows for approximate inference for models with intractable or expensive to calculate likelihoods. ABC replaces the likelihood computation by simulating from the underlying generative model and using the parameters corresponding to the generated observations which are “close enough” to the true data to approximate the posterior distribution. As a realistic test we present results of the use of an ABC algorithm to perform parameter estimation for a complex model of collective cell migration. We consider a hybrid discrete-continuum model for the migration of a population of cells which interact with a background chemical chemoattractant field. Each individual cell acts as moving sink term reflecting the degradation of the chemoattractant through the action of membrane bound enzymes. We perform a

sensitivity analysis of the model using a Morris screening approach to indicate which parameters in the model are identifiable. Our results suggest that regardless of the algorithm used, the ABC-posterior distributions depend crucially on the value of the measurement time interval T and that capturing the true parameter uncertainty of the SDE model parameters might therefore be a challenging task for ABC.

John A. Mackenzie
Strathclyde University
j.a.mackenzie@strath.ac.uk

MS403

Quantifying Model Bias with Information Inequalities for Probability Divergences and Concentration Inequalities

We discuss the foundations of model-form uncertainty by constructing suitable neighborhoods of models around a baseline model. These neighborhoods are based on KL-divergence and other new related divergences, the so-called (f, Γ) divergences, which are used in Machine Learning algorithms (for example GANs). We derive corresponding information inequalities which allows us to estimate the size of possible error on suitable expectation of observables provided the model belong to the neighborhood. The information inequalities are made practical and computable by using concentration inequalities.

Luc Rey-Bellet
Department of Mathematics & Statistics
University of Massachusetts
luc.reybellet@gmail.com

MS403

Model Error Estimation for Active Learning of Interatomic Potential Models

Solutions to inverse problems often assume the computational model can replicate the true mechanism behind data generation. However, physical models carry misspecification due to different parameterizations and assumptions. Ignoring such model errors can lead to overconfident calibrations and poor predictive capability, even when high-quality data are used. As a result, outer-loop tasks, such as active learning or optimal design lead to biased results with poorly calibrated uncertainties. This work will present a Bayesian inference framework for representing, quantifying, and propagating uncertainties due to model structural errors by embedding stochastic correction terms in the model. The physical input parameters and model-error parameters are then simultaneously optimized in an inverse modeling context. We will demonstrate the methodology on example problems developing machine-learned interatomic potential (MLIAP) models. The resulting predictive uncertainties capture model error and will be employed in an active learning loop to enable efficient construction of uncertainty-augmented MLIAPs.

Khachik Sargsyan, Logan Williams
Sandia National Laboratories
ksargsy@sandia.gov, logwill@sandia.gov

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA

hnnajm@sandia.gov

MS403

Bayesian Inference under Model Misspecification and Optimal Transport

Central to any inverse problem is the choice of a misfit function that best represents the discrepancy between the observations (data) and model predictions. In the context of model error, different choices of misfit measures can play a crucial role in mitigating or magnifying the impact of the misspecification. This is particularly important in inverse problems, such as those in full waveform inversion, where the misfit should measure data features that are most relevant for learning the unknown parameters. For various statistical models appearing in supervised learning and inverse problems, we show that specific choices of optimal-transport-based misfit functions can reduce the gap between model space and the observed data. For both deterministic and Bayesian inference problems, we compare the recovered solution to traditional discrepancy measures based on the ℓ_2 norm. A variety of independent scoring rules that account for different predictive goals are used to quantify the error of the recovered solution in both state and data space for these inverse problems.

Andrea Scarinci
Massachusetts Institute of Technology
scarinci@mit.edu

Ricardo Baptista
MIT
rsb@caltech.edu

Youssef M. Marzouk, Michael Fehler
Massachusetts Institute of Technology
ymarz@mit.edu, fehler@mit.edu

MS403

Optimal Experimental Design for Prediction for Data-Consistent Inversion

One of the ultimate goals in science is to develop a symbiotic relationship between two of the pillars, experimentation and simulation, where the data from the experiments informs the computational models and the computational models are used to guide the optimal acquisition of new data. In many situations, the collection of experimental data can be costly and time consuming. Thus, we may only be able to afford to perform a limited number of experiments, so we must choose the experiments that are likely to produce informative data. Moreover, the optimality of the experiment must be chosen with respect to the ultimate objective. We present a mathematically rigorous approach for optimal experimental design for data-consistent inversion. Data-consistent inversion is a relatively new approach for solving a particular class of stochastic inverse problems where a probability measure is sought on model inputs such that the corresponding push-forward measure matches a given target measure. We demonstrate that the optimal experimental design often depends on whether our objective is to characterize the uncertainty in model input parameters or in the prediction of quantities of interest that cannot be observed directly. Numerical results will be presented to illustrate both types of optimal experimental design and to highlight the differences.

Tim Wildey
Sandia National Laboratories

tmwilde@sandia.gov

Troy Butler
University of Colorado Denver
butler.troy.d@gmail.com

MS404

Benchmarking Multi-Fidelity Optimization Methods: Recent Experiences and Lesson Learned

The talk overviews the efforts of the AVT-331 technical team under the NATO STO Applied Vehicle Technology Panel, which is studying multi-fidelity methods through application to vehicle design. The objectives of the team are to understand the potential benefits of multi-fidelity methods in vehicle design and to document the relative strengths and weaknesses of different multi-fidelity methods using a common benchmark suite developed by the team. The benchmark suite has multiple levels of complexity, beginning with analytic functions with well-understood properties and culminating with representative air and sea vehicle benchmarks. The suite also includes intermediate complexity benchmarks of relevance to air, sea, and space vehicles. Benchmark features are presented, including descriptions of benchmark objectives, and enabling software needed to reproduce vehicle-level results, as well as experimental setup for analytic benchmarks. The principal multi-fidelity methods utilized across the team are also summarized. Lastly, the process is described by which the relative strengths and weaknesses of different multi-fidelity methods is assessed over eight assessment categories.

Matteo Diez, Andrea Serani
Consiglio Nazionale delle Ricerche
Istituto di Ingegneria del Mare (CNR-INM)
matteo.diez@cnr.it, andrea.serani@cnr.it

Philip Beran, Dean Bryson
Air Force Research Laboratory
United States
philip.beran@us.af.mil, dean.bryson@us.af.mil

Andrew Thelen
University of Dayton Research Institute
United States
andrew.thelen@udri.udayton.edu

Laura Mainini
Politecnico di Torino
laura.mainini@polito.it

MS404

Multi-Output Approximate Control Variates for Trajectory Optimization

Monte Carlo estimators with many Quantities of Interest (QoIs) can often lead to an extremely high computational burden. Variance reduction approaches are critically needed to reduce the cost of analysis in computationally expensive models. By using multiple models of varying fidelity, a control variate estimator can significantly reduce the variance of an estimator for one QoI using the correlation between the models. However, with many QoIs, a vast majority of approaches use multiple independent control variate estimators for each QoI. In this talk, we develop a multi-output approximate control variate estimator and show that it outperforms a set of multiple single-output control variate estimators. Overall, it significantly reduces

the cost of analysis by additionally extracting correlations amongst multiple model outputs. We demonstrate our technique on applications in flight trajectory estimation in simulations of entry, descent, and landing. Also, using highly correlated QoIs, examples will be shown that significantly surpass current state-of-the-art techniques.

Thomas Dixon, Alex Gorodetsky
University of Michigan
tdixono@umich.edu, goroda@umich.edu

James Warner
Nasa Langley Research Center
james.e.warner@nasa.gov

Geoffrey Bomarito
NASA Langley Research Center
geoffrey.f.bomarito@nasa.gov

MS404

A Combination Technique for Pde Constrained Optimal Control Problems under Uncertainty

We present a combination technique based on mixed differences of both spatial approximations and quadratures formulae for the stochastic variables to solve efficiently Optimal Control Problems (OCPs) under uncertainty. The method requires solving the OCP for several low-fidelity spatial grids and quadrature formulae for the objective functional. All the computed solutions are then linearly combined to get a final approximation which, under suitable regularity assumptions, preserves the same accuracy of fine tensor product approximations, while drastically reducing the computational cost. The combination technique involves only tensor product quadrature formulae, thus the discretized OCPs preserve the convexity of the continuous OCP. Hence, the combination technique avoids the inconveniences of Multilevel Monte Carlo and/or sparse grids approaches, but remains suitable for high dimensional problems. We will present an a-priori procedure to choose the most important mixed differences and an asymptotic convergence analysis, which states that the asymptotic convergence rate is exclusively determined by the spatial solver. Numerical experiments validate the results.

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Tommaso Vanzan
École Polytechnique Fédérale de Lausanne
tommaso.vanzan@epfl.ch

MS404

Towards Multilevel Slice Sampling

For sampling approximately a partially known distribution, e.g., posterior distributions in Bayesian inverse problems, the slice sampling methodology provides a machinery for the design and simulation of a suitable Markov chain without the necessity to tune any parameters as in many Metropolis-Hastings algorithms. In the machine learning community slice sampling is a frequently used approach, which appears not only there as standard sampling tool. In particular, the elliptical slice sampler attracted in the last decade considerable attention as a dimension-robust algorithm. However, from a theoretical point of view it is not well understood. In this talk, we present theoretical results regarding the geometric ergodicity of Markov chains

generated by elliptical slice sampling as well as simple slice sampling. Moreover, we outline how multilevel or multifidelity ideas can be combined with (elliptical) slice sampling in order to reduce the (usually) significant computational cost of this promising Markov chain Monte Carlo method.

Björn Sprungk

TU Freiberg

bjoern.sprungk@math.tu-freiberg.de

MS405

Towards a Dominant SVD Computation via ADMM-Like Decentralized Consensus Optimization with Projection Splitting

The Singular Value Decomposition (SVD) of a matrix provides an optimal low-rank approximation of the data, which is useful across science and engineering. Existing methods for computing this approximation when data is held privately on multiple machines require a central parameter server or the transmission of local gradients; otherwise, they scale poorly with the amount of data. To develop a communication-efficient and scale-invariant decentralized algorithm, we frame the computation as a decentralized consensus optimization problem and introduce reasonable approximations to simplify the Alternating Direction Method of Multipliers (ADMM) solution. Limited numerical experiments show that our algorithm can outperform its centralized version and is competitive with decentralized stochastic gradient descent. While theoretical guarantees of convergence are absent, this work furthers the development of reliable decentralized solutions to problems posed in non-traditional computing environments. Prepared by LLNL under Contract DE-AC52-07NA27344. Funded by LLNL LDRD project 22-ERD-045. LLNL-ABS-840146

Cory A. Gillette

University of Notre Dame

cgillet2@nd.edu

MS405

Resilient Workflow Management for Interfacing Among Black Box Algorithms

Computational scientific discovery is increasingly driven by carefully linking independently developed research codes and deploying them on heterogeneous large scale high performance computing resources. We present Merlin, a workflow management tool, to aid in resolving the subtleties of interfacing complex scientific codes and, in particular, enable robust, machine learning-friendly ensembles of simulations. We will discuss some of Merlins features that allow for multi-tiered fault recovery and low per-task overhead. Our showcase example - currently in development - is to design inertial confinement fusion experiments via a combination of simulation, optimization, and machine learning algorithms on the exascale machine El Capitan, which is projected to be the worlds fastest supercomputer when fully deployed in 2023.

Yamen Mubarka

Lawrence Livermore National Laboratory

mubarka1@llnl.gov

Nai-Yuan Chiang

Lawrence Livermore National Laboratory

7000 East Avenue, Livermore, CA 94550

chiang7@llnl.gov

Andrew Gillette

Lawrence Livermore National Laboratory

Center for Applied Scientific Computing

gillette7@llnl.gov

Irene Kim

Lawrence Livermore National Laboratory

kim88@llnl.gov

Joseph M. Koning

University of California at Davis

ISCR/LLNL

koning1@llnl.gov

Eugene Kur, Ryan Lee, Luc Peterson, Ryan Tran

Lawrence Livermore National Laboratory

kur1@llnl.gov, lee1098@llnl.gov, peterson76@llnl.gov,

tran67@llnl.gov

Jingyi Wang

Lawrence Livermore National Laboratory

7000 East Avenue, Livermore, CA 94550

wang125@llnl.gov

MS405

Random Sketching for Improving the Performance of Gram-Schmidt Algorithm on Modern Architectures

Randomized Gram-Schmidt was introduced in 2022 in a paper authored by Oleg Balabanov and Laura Grigori. The authors showed how to use their randomized scheme in GMRES and demonstrated the results featuring both: a speedup of their Matlab GMRES code and stability comparable to GMRES with regular modified Gram-Schmidt. We extended this work by connecting it to low-synchronization variants of Gram-Schmidt. We also created a GPU implementation to test the practicality of the proposed algorithm. As it turns out, using random sketching in conjunction with low-synchronization GMRES results in a significant reduction of computational time and is ultimately an appropriate and highly performant GPU strategy.

Katarzyna Swirydowicz

Pacific Northwest National Laboratory

kasia.swirydowicz@pnnl.gov

MS406

Neural Operators: Learn to Predict in Computational Sciences

Traditional deep neural networks are maps between finite dimension spaces, and hence, are not suitable for modeling phenomena such as those arising from the solution of partial differential equations (PDE). We introduce neural operators that can learn operators, which are maps between infinite dimension spaces. By framing neural operators as non-linear compositions of kernel integrations, we establish that they are universal approximators of operators. They are independent of the resolution or grid of training data and allow for zero-shot generalization to higher resolution evaluations. We find that neural operators can solve turbulent fluid flow, seismic wave equation, CO₂ storage, and many more hard problems with 100000x speedup compared to numerical solvers. I will outline several applications where neural operators have shown order

of magnitude speedup.

Kamyar Azizzadenesheli
Nvidia
kaazizzad@gmail.com

MS406

Operator Learning in Vascular Flows Modeling: Online Biomarker Prediction and Uncertainty Quantification from Non-Invasive 4D Flow MRI Data

Unlike diseases with strong genetic predisposition, such as cancer, there exist no accurate personalized diagnostic tools for disorders such as Hypertensive Pregnancy Disorders (HPD) and its mechanism remain understudied. For the purpose of performing diagnostics, we require disease indicators (biomarkers) that are strongly correlated with the underlying condition, i.e. absolute vascular pressure for hypertension. Measuring absolute pressure requires an invasive procedure and it is impossible for pregnant subjects, but an informed estimate is possible via the use of Computational Fluid Dynamics. Unfortunately, not only the computational cost of CFD is high, but also patient specific information required in modeling is missing in the clinic. The goal of this talk is to synthesize AI, CFD and medical imaging (US/MRI) techniques to provide trustworthy predictions of strong biomarkers for early diagnosis of HPD in a real-time clinical setting. We will present a three-step procedure based on a reduced order Navier-Stokes and an ensemble of Operator Learning models with a nonlinear manifold decoder and show how it can be employed to provide accurate biomarker predictions for different topologies together with a quantification of the associated uncertainty.

George Kissas, Jacob Seidman, Paris Perdikaris, George Pappas
University of Pennsylvania
gkissas@seas.upenn.edu, seidj@sas.upenn.edu,
pgp@seas.upenn.edu, pappasg@seas.upenn.edu

MS406

Continuous Generative Neural Network for Inverse Problems

We study Continuous Generative Neural Networks (CGNNs), namely, generative models in the continuous setting. The architecture is inspired by DCGAN, with one fully connected layer, several convolutional layers and nonlinear activation functions. In the continuous L^2 setting, the dimensions of the spaces of each layer are replaced by the scales of a multiresolution analysis of a compactly supported wavelet. We present conditions on the convolutional filters and on the nonlinearity that guarantee that a CGNN is injective. This theory finds applications to inverse problems, and allows for deriving Lipschitz stability estimates for (possibly nonlinear) infinite-dimensional inverse problems with unknowns belonging to the manifold generated by a CGNN.

Silvia Sciuotto
University of Geneva
silvia.sciuotto@edu.unige.it

Matteo Santacesaria
MaLGa Center, DIMA
University of Genoa
matteo.santacesaria@unige.it

Giovanni S. Alberti
University of Genoa
giovanni.alberti@unige.it

MS406

Multi-Model Communication and Data Assimilation for Improving Model Dynamics

Developing a single model to capture all the dynamics of the Earth's atmosphere is an immense task. Each model will have a deficiency in at least a small number of areas, such as the Madden-Julian Oscillation or the El Niño-Southern Oscillation. Improving these aspects of atmosphere models often comes at the cost of introducing or worsening deficiencies in other areas. In contrast, it is possible to create specialized, low-dimensional models to capture these phenomena by simplifying many aspects of the full system. Here, we propose a strategy to allow different models communicate throughout a simulation. The communication lets us leverage the strengths of each model without needing to change their dynamics. The strategy is based on common data assimilation techniques used to combine models and real-world data. We implement this strategy in a test case consisting of a deficient tropical atmosphere model and a model of the Madden-Julian Oscillation. The models have different state spaces and different dynamics. We show that our strategy enriches the dynamics of both models beyond the information that is communicated between them.

Jason Torchinsky
University of Wisconsin-Madison
jason.torchinsky@wisc.edu

Sam Stechmann
University of Wisconsin-Madison
stechmann@wisc.edu

MS406

Morphological Parameters Estimation of Neurons Using a Machine Learning Algorithm on Diffusion MRI Data

Diffusion magnetic resonance imaging is a promising non-invasive modality that can be used to probe tissue microstructure. Existing methods that estimate brain microstructure parameters often use biophysical models that are a) based on analytical expressions of the diffusion MRI signal, b) arising from simplified geometrical models of neuronal tissue. We take a different approach by a) numerically solving a gold standard model for the diffusion MRI signals, given by the Bloch-Torrey partial differential equation that governs the transverse magnetization, b) while using realistic neuron geometries. In previous work, we created accurate finite element meshes for over 1000 realistic neurons whose surface descriptions have been uploaded to the database NeuroMorpho.Org, and solved the Bloch-Torrey equation to obtain the simulated diffusion MRI signals of these neurons in 64 diffusion directions and at 31 diffusion-encoding gradient strengths. In this work, this set of simulated data was used to train a support vector machine to estimate morphological parameters such as soma volume, dendrite volume, and dendrite total length. The results reveal that the direction-averaged dMRI signals are significantly related to certain morphological parameters, such as soma volume, soma volume fraction, etc. This study is part of the effort to show that machine learning algorithms can serve as an efficient estimator for certain

neuroanatomical parameters in diffusion MRI experiments.

Zheyi Yang
Ensta Paris
zheyi.yang@polytechnique.edu

Chengran Fang, Jing-Rebecca Li
INRIA Saclay
chengran.fang@inria.fr, jingrebecca.li@inria.fr

MS407

Quadrature by Parity Asymptotic eXpansions (QPAX) for Light Scattering by High Aspect Ratio Plasmonic Particle

The study of scattering by a high aspect ratio particle has important applications in sensing and plasmonic imaging. To illustrate the effect of particle's narrowness (that can be related to parity properties) and the need for adapted methods (in the context of boundary integral methods), we consider the scattering by a penetrable, high aspect ratio ellipse. This problem highlights the main challenge and provides valuable insights to tackle general high aspect ratio particles. We find that boundary integral operators are nearly singular due to the collapsing geometry to a line segment. We show that these nearly singular behaviours lead to qualitatively different asymptotic behaviours for solutions with different parities. Without explicitly taking this into account, computed solutions incur large errors. We introduce the Quadrature by Parity Asymptotic eXpansions (QPAX) that effectively and efficiently addresses these issues. We demonstrate the effectiveness of QPAX through several numerical examples (potentially including high-order QPAX results).

Camille Carvalho
UC Merced
ccarvalho3@ucmerced.edu

MS407

Multitask Optimization of Laser-Plasma Accelerators Using Simulation Codes with Different Fidelities

Laser-plasma accelerators can generate GeV electron beams in an ultra-compact, cm-scale setup, but have yet to demonstrate sufficient beam quality and stability for demanding applications. To overcome this challenge, broad optimization of the accelerator design with numerical simulations is essential. However, due to the high computational cost of general particle-in-cell simulations, optimization over a large parameter space is prohibitively expensive. Here, we show that cheaper simulations based on reduced physical models can be effectively incorporated into an optimization to strongly reduce the need for fully-detailed simulations. This is enabled by a Bayesian optimization algorithm using a multitask Gaussian process, where two tasks corresponding to the output of each simulation code are defined. The algorithm identifies correlations between both tasks, allowing it to learn from inexpensive evaluations and perform only a few, well-targeted simulations of high fidelity. By means of a proof-of-principle study combining the FBPIC and Wake-T codes, we demonstrate an order-of-magnitude reduction in computing time and cost, paving the way towards cost-effective optimization over wide parameter spaces.

Angel Ferran Pousa
Deutsches Elektronen-Synchrotron (DESY)

angel.ferran.pousa@desy.de

Soeren J alas
Hamburg University
soeren.jalas@desy.de

Manuel Kirchen, Alberto Martinez de la Ossa
Deutsches Elektronen-Synchrotron (DESY)
manuel.kirchen@desy.de,
alberto.martinez.de.la.ossa@desy.de

Maxence Thevenet
Deutsches Elektronen Synchrotron (DESY), Germany
maxence.thevenet@desy.de

Stephen Hudson, Jeffrey Larson
Argonne National Laboratory
shudson@anl.gov, jmlarson@anl.gov

Axel Huebl
Lawrence Berkeley National Laboratory (LBNL), USA
axelhuebl@lbl.gov

Axel Huebl
Lawrence Berkeley National Laboratory (LBNL), USA
Helmholtz-Zentrum Dresden-Rossendorf (HZDR),
Germany
axelhuebl@lbl.gov

Jean-Luc Vay
Lawrence Berkeley National Laboratory
Berkeley, CA
JLVay@lbl.gov

Rémi Lehe
Lawrence Berkeley National Laboratory
ATAP
rlehe@lbl.gov

MS407

Numerical Optimization of Secondary Beamlines at the Argonne Tandem Linac Accelerator System

Physicists at the Argonne National Laboratory's ATLAS In-Flight Facility seek automated optimization processes for the production and delivery of secondary beamlines. Secondary beam production relies on a primary beam provided by ATLAS to impinge on a production target, generating numerous new and interesting radioactive ions. A selection of the isotope of interest is first made, and then optimized (for yield, purity, and beam quality) from the production point to the experimental target. In this talk, we highlight approaches for utilizing numerous outputs from beamlines in order to decrease the number of experimental setups required during an optimization campaign.

Jeffrey Larson
Argonne National Laboratory
jmlarson@anl.gov

MS407

Optimization and Machine Learning to Reconstruct Multiple Compton Scattering Events for GRETINA Spectroscopy

Rare isotopes with unstable nuclear structures created by ion-beam and target interactions are probed by sci-

entists at Argonne National Laboratory's ATLAS facility and elsewhere. In order to investigate these rare structures, the gamma-rays emanating from the target must be reconstructed from recorded spectra. The partially complete spherical shells of high-purity germanium (HPGe), GRETA (GRETINA) and AGATA, represent the state-of-the-art in gamma-ray spectroscopy. GRETA is able to capture individual gamma-ray-matter interactions and their energies allowing the reconstruction of the scattered paths of individual gamma-rays that have undergone Compton scattering. However, current methods for reconstruction recover gamma-rays with an efficiency much lower than predicted and lower than older detectors (i.e., Gammasphere). Data-driven optimization using modern machine learning (ML) methods allows for more reliable separation of multiple gamma-rays and ordering of their interactions. Additionally, ML methods are used to improve existing metrics to better differentiate between high and low quality data and employ more computationally intensive methods only when needed. These methods, in conjunction a mixed integer non-linear program formulation, allow for increased efficiency of the ion-beam approaching the predicted efficiency for AGATA and GRETA as well as enabling further scientific investigations into rare isotopes.

Thomas F. Lynn
Northwestern University
tlynn@anl.gov

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

Torben Lauritsen
Argonne National Laboratory
torben@anl.gov

Amel Korichi
National institute of nuclear and particle physics
amel.korichi@ijclab.in2p3.fr

MS407

Using Machine Learning to Optimize the Performance of the Superconducting ECR Ion Source Venus

The fully-superconducting electron cyclotron resonance (ECR) ion source VENUS is key to providing both the high-intensity, medium-mass ions needed for the superheavy element science program and the very-highly-charged ions for the chip testing program at the 88 Inch Cyclotron at Lawrence Berkeley National Laboratory (LBNL). The success of this source has led to the construction of a VENUS replica as primary injector source at the Facility for Rare Isotope Beams at Michigan State University: a facility that will serve at the forefront of Nuclear Structure research for decades to come. At LBNL a Python library has been developed to communicate with the VENUS programmable logic controller (PLC) to monitor and fully control the ion source programmatically. This permits a variety of data analytics and machine learning techniques to be applied to enhance and facilitate the operation of the ion source. In this talk, the physics, mechanics and control of the ion source will be discussed. Some initial results demonstrating how techniques such as bayesian optimization and neural networks enable a better understanding, automatization and output stabilization will be

presented.

Victor Watson
Lawrence Berkley National Lab
vwatson@lbl.gov

Wenhan Sun, Yubin Hu, Alex Kireeff
LBL / University of California Berkeley
wenhan0112@lbl.gov, yubihu@lbl.gov, akireeff@lbl.gov

Marco Salathe, Damon Todd, Heather L. Crawford
Lawrence Berkeley National Lab
msalathe@lbl.gov, dstodd@lbl.gov, hlcrawford@lbl.gov

MS408

Use of Multiple Functional Approximations to Infer Uncertainty in Density Functional Theory

Kohn Sham Density functional theory (KS-DFT) is a first principles electronic structures method widely used in chemistry, physics, and materials science, but predictions rely on an exchange correlation (XC) functional, the form of which is unknown. Many density functional approximations (DFAs) have been developed in attempt to capture the XC functional, but it is difficult to determine which approximation is best suited to a given problem, and the choice of approximation is a major source of predictive uncertainty. This talk will present an approach to infer probability distributions over predicted quantities by leveraging a data set of DFT predictions computed using an ensemble of different DFAs for a range of molecules. To perform inference, we require some model of the relationship between molecules as well as DFAs. We use a multitask Gaussian Process regression framework to relate regression problems corresponding to different DFAs and define kernel functions to relate descriptors of individual molecules. Tests informed by predictions of the ionization potential of small organic molecules indicate that incorporating multiple DFAs through the multitask framework can improve predictive accuracy compared to a single Gaussian Process regression model informed by one DFA. The challenge of producing robust uncertainty predictions from the distributions predicted by GP regression is ongoing work.

Katharine Fisher
Massachusetts Institute of Technology
kefisher@mit.edu

Michael F. Herbst
RWTH Aachen University
Germany
herbst@acom.rwth-aachen.de

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS408

Leveraging Model Assumptions for Model Form Uncertainty Quantification

Model-form uncertainty arises from physical assumptions made in constructing models either to reduce the physical complexity or to model physical processes that are not well understood. In Reynolds-Averaged Navier-Stokes (RANS) turbulence modeling and turbulent combustion modeling, some of these assumptions are known and can be leveraged as a priori knowledge in quantifying the uncertainty. Three

such approaches are presented in this talk that take advantage of the physics inherently neglected by a model in order to characterize the structural uncertainties and model-form error. The peer models approach compares the model output between two equally plausible models with divergent assumptions, such that model variability is indicative of model error. The hierarchical models approach identifies a hierarchy of models of increasing fidelity and functionalizes the error in a lower fidelity model based on the principles of a higher fidelity model. These first two approaches are applied to turbulent reacting flows, where the model-form error is propagated through Large Eddy Simulations (LES) using stochastic collocation. The final implied models approach analyzes the model error through the derivation of a transport equation for the model error that is implied by the model for the quantity of interest. The implied models approach is applied to RANS turbulence model for multiple turbulent flow configurations in order to better understand the sources of model-form error including the role of error cancellation and error superposition.

Michael E. Mueller
Princeton University
muellerm@princeton.edu

Kerry Klemmer
Massachusetts Institute of Technology
kklemmer@mit.edu

MS408

Quantifying Model Prediction Sensitivity to Model-Form Uncertainty

In the process of model development, limited knowledge and computational resources necessitate simplifying assumptions that are not valid in all cases, leading to model-form uncertainty (MFU). When model predictions must be used to inform high-consequence engineering-design and policy decisions, it is critical to understand how they are impacted by uncertainty in these assumptions. We present a novel method to quantify this impact by combining parameterized modifications to modeling assumptions, representing MFU, with grouped variance-based sensitivity analysis. We demonstrate the approach on a subsurface contaminant transport model, and we present theoretical and numerical results on the robustness of the approach to the mathematical form of the MFU parameterization. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525

Teresa Portone, Joseph Hart, Rebekah White
Sandia National Laboratories
tporton@sandia.gov, joshart@sandia.gov, rebwhit@sandia.gov

MS408

Uncertainty Propagation in Pore Water Chemical Composition Calculation Using Surrogate Models

Performance assessment in deep geological nuclear waste repository systems necessitates an extended knowledge of the pore water chemical conditions prevailing in host-rock formations. In the last two decades, important progress has been made in the experimental characterization and thermodynamic modeling of pore water speciation, but the influence of experimental artifacts and uncertainties of thermodynamic input parameters are seldom evaluated. In this respect, we conducted an uncertainty propagation study in a reference geochemical model describing the pore

water chemistry of the Callovian-Oxfordian clay formation. Nineteen model input parameters were perturbed, including those associated to experimental characterization (leached anions, exchanged cations, cation exchange selectivity coefficients) and those associated to generic thermodynamic databases (solubilities). A set of 13 quantities of interest were studied by the use of polynomial chaos expansions built non-intrusively with a least-squares forward stepwise regression approach. Training and validation sets of simulations were carried out using the geochemical speciation code PHREEQC. The statistical results explored the marginal distribution of each quantity of interest, their bivariate correlations as well as their global sensitivity indices. The influence of the assumed distributions for input parameters uncertainties was evaluated by considering two parametric domain sizes.

Pierre Sochala
CEA (Atomic Energy and Alternative Energies
Commission)
pierre.sochala@cea.fr

Christophe Chiaberge, Francis Claret
BRGM/WATER
BP6009 - 45060 ORLEANS CEDEX 2
c.chiaberge@brgm.fr, f.claret@brgm.fr

Christophe Tournassat
ISTO, Université d'Orléans-CNRS-BRGM, Orléans,
France.
Lawrence Berkeley National Laboratory, Berkeley CA,
USA
christophe.tournassat@univ-orleans.fr

MS409

Predicting Composite Structures Failure with the Help of Quantum Neural Networks

Composite materials failure entails many physical phenomena, which make its prediction challenging from a computational standpoint. Especially in a design context, where material properties, geometry and stacking sequence are varied, the computing time required by the finite element analyses becomes prohibitive. This has motivated the idea of data-driven composite failure prediction, where a data-based model learns from relatively few numerical simulations and it can later generalize to unseen input (e.g. loading) conditions. Even though most machine learning models are trained on classical computers, quantum computers can also learn from data. In fact, operations on quantum bits are continuous and can be tuned in a hybrid quantum-classical optimization loop. Since the model is a quantum circuit, it is evaluated on quantum hardware, while the parameters are updated classically. Thanks to the properties of quantum mechanics, quantum models proved able to learn functions that classical machine learning models cannot learn, especially in classification tasks. However, the datasets used were 'fabricated' such to be easily classified by the quantum circuit. As one of the first practical applications, this work explored the power of quantum machine learning models in the practical scenario of learning composite structures failure and compared it to that of classical data-driven models.

Giorgio Balducci, Boyang Chen
Delft University of Technology
G.B.L.TostiBalducci@tudelft.nl, b.chen-2@tudelft.nl

Matthias Moller

Delft Institute of Applied Mathematics (DIAM)
Delft University of Technology
M.Moller@tudelft.nl

Roeland de Breuker
Delft University of Technology
r.debreuker@tudelft.nl

MS409

Early Applications of Quantum Computing to Subsurface Imaging

We show an application of quantum annealing to solving an NP-hard combinatorial optimization problem encountered on land, where strength and timing of scattered waves is used to generate the subsurface image. Geological anomalies such as underground rivers, caves, or boulders, have different wave propagation velocity than their surroundings. This results in abrupt local misalignment in arrival times of detected waves. This can blur the image and misrepresent the geological information. Optimizing the relative alignment of the measurements is described by a multi-modal objective function. Getting stuck in local optima is common, and often a second best optimum, far away in parameter space, gives a false sense of certainty about the subsurface. Solving this conceptually simple, yet mathematically challenging, problem has far reaching implications to the economics of civil engineering, hydrology, gas storage and hydrocarbon exploration. We convert the classical optimization to a Discrete Quadratic Model, and through various discrete (classical) variable to qubit encodings to a Binary Quadratic Optimization and its constraints. We discuss how a custom hybrid classical-quantum workflow enables us to solve problems nearing business-relevance and show how the encodings influence the output generated by the 5000-qubit quantum annealer. We discuss some features of this application which could help establishing other use cases in similar disciplines.

Marcin Dukalski

Aramco Overseas Company B.V.
Delft Global Research Centre
marcin.dukalski@aramcooverseas.com

MS409

Quantum Algorithms for Scientific Machine Learning and Solving Differential Equations

In the talk, I will discuss the advances in the field of quantum scientific computing. This is based on the series of works for solving differential equations for applications in fluid dynamics and generative modelling. First, I will show how to differentiate quantum circuits with feature maps and embed derivatives of multidimensional functions. This approach relies on automatic differentiation to represent derivatives in an analytical form, thus avoiding inaccurate finite difference procedures. We refer to the underlying circuits as derivative quantum circuits (DQCs). I will describe the proposed hybrid quantum-classical workflow where DQCs are trained to satisfy nonlinear differential equations and specified boundary conditions. As an example, I will apply the algorithm to solve a problem from computational fluid dynamics. This corresponds to the fluid flow in a convergent-divergent nozzle (described by Navier-Stokes equations), where we predict density, temperature and velocity profiles for the stiff system of equations that can challenge classical solvers. Second, I will also show that DQCs can provide an advantage for generative modelling from stochastic differential equations, giving ac-

cess to sampling from financially relevant models. Finally, I will discuss the protocols that may be scaled to treat large industrial problems in the future.

Oleksandr Kyriienko, Annie Paine
Quantum Dynamics, Optics, and Simulation
University of Exeter
o.kyriienko@exeter.ac.uk, ap877@exeter.ac.uk

Vincent Elfving
Pasqal
vincent.elfving@pasqal.com

MS409

Finite Element Computations on NISQ Quantum Devices using FEqa

While quantum devices are under development they offer a fresh and exciting perspective for scientific and engineering computation. FEqa is a recently published state-of-the-art technique capable of solving practical problems using NISQ devices using quantum annealing on quantum annealers and the quantum adiabatic optimization algorithm (QAOA) on gate-based quantum computers. This mini-tutorial will provide a quick introduction to the two major quantum computing models, quantum annealing and gate-based quantum computing, give insight into the theoretical foundations of FEqa and its algorithms, formulate problems with results, and present recent advances since publication. The remainder of the tutorial will be focused on open avenues in quantum computing for scientific and engineering computation.

Osama Muhammad Raisuddin, Suvranu De
Mechanical, Aerospace, and Nuclear Engineering
Rensselaer Polytechnic Institute
raisuo@rpi.edu, des@rpi.edu

MS410

Acceleration of Randomized Tensor Decompositions

In this talk, we discuss the application of acceleration methods to randomized tensor low rank approximations (LRAs). Randomized tensor LRA algorithms have been proposed to address the need for efficient decompositions algorithms for large dense tensors; such algorithms rely on stochastic first-order or second-order methods. Acceleration methods have been designed and applied for general first-order optimization problems. We theoretically and empirically investigate the advantages of combining these two lines of work.

Alex Gittens

Rensselaer Polytechnic Institute
Computer Science
gittea@rpi.edu

Dong Hu
Rensselaer Polytechnic Institute
hud3@rpi.edu

MS410

Approximate Euclidean Lengths and Distances Beyond Johnson-Lindenstrauss

In this presentation we will dive in the core of Randomized Numerical Linear Algebra and present out recent re-

sults (Sobczyk and Luisier, NeurIPS 2022) related to the Johnson-Lindenstrauss (JL) lemma and its applications. We revisit a key problem: how to approximate the Euclidean lengths (or distances) of a set of high-dimensional vectors. Can we do better than JL? Based on techniques which are inspired by the recently published Hutch++ algorithm for trace estimation, we proposed algorithms which can (provably) improve the standard JL bounds, both in theory and in practice. We will describe all these algorithms and sketch the proof techniques. These improvements can ultimately lead to more practical and efficient randomized algorithms for scientific computing, especially when high accuracy is required, which is a known shortcoming of JL-based random projections. We will also discuss how these results can be applied to other problems, namely to compute Euclidean distances between high-dimensional data points, to estimate the statistical leverage scores of a matrix, and to approximate the charge densities in a quantum mechanical system.

Aleksandros Sobczyk
IBM Research Europe
ETH Zürich
obc@zurich.ibm.com

Mathieu Luisier
Integrated Systems Laboratory
ETH Zurich
mluisier@iis.ee.ethz.ch

MS410 Faster Accurate Sketching for Tensor Networks

Linear sketches with tensor product structure provide an effective tool for accelerating decomposition of tensors. We describe a new approach for inexact optimization of Tucker decomposition via linear sketching (arXiv:2104.01101). Then, we propose efficient a general scheme for constructing computationally-efficient sketches of tensor network inputs with arbitrary structure (arXiv:2205.13163). We provide bounds on the complexity and accuracy of both schemes and show that the general scheme is optimal or near-optimal among possible tensor networks consisting of Gaussian tensors.

Edgar Solomonik
University of Illinois at Urbana-Champaign
solomon2@illinois.edu

Linjian Ma
University of Illinois at Urbana Champaign
lma16@illinois.edu

MS410 Capacity Analysis of Vector Symbolic Architectures

Hyperdimensional computing (HDC) is a biologically-inspired framework that uses high-dimensional vectors and various vector operations to represent and manipulate symbols. The ensemble of a particular vector space and two choice vector operations (one addition-like for "bundling" and one outer-product-like for "binding") form what we call a "vector symbolic architecture" (VSA). While VSAs have been employed in numerous applications and studied empirically, many theoretical questions about VSAs remain open. We provide theoretical analyses for the representation capacities of two popular VSAs: MAP-I and MAP-B. Representation capacity refers to lower bounds on the

dimensions of the VSA vectors required to perform certain symbolic tasks (such as testing for set membership and estimating set intersection sizes) to a sufficient degree of accuracy. Our analysis of MAP-I also demonstrates a connection between VSAs and sketching (dimensionality reduction) algorithms such as the Johnson-Lindenstrauss transform. Time permitting, we will also describe a relationship between the MAP-I VSA to Hopfield networks, which are simple models for associative memory.

Elizabeth Yang
University of California at Berkeley
elizabeth_yang@berkeley.edu

Kenneth Clarkson
IBM
klclarks@us.ibm.com

Shashanka Ubaru
IBM T. J. Watson Research Center
shashanka.ubaru@ibm.com

MS411 High-Order Method-of-Lines Time Integration for Sharp Immersed Methods with Moving Boundaries

Sharp immersed methods are designed to accurately capture moving boundaries without the computational cost and algorithmic challenges of mesh motion or remeshing. However, these methods face a unique challenge in "fresh" or "emerging" computational elements that change sides of an interface during a given time step. Most existing solutions reconstruct an extended flow field in the vicinity of the interface at each time step with up to second-order accuracy. Here we demonstrate that this methodology can be successfully combined with higher-order spatial discretizations and explicit Runge-Kutta time integrators in a method-of-lines approach, and describe the nature of the associated temporal errors. We show an application of this approach to a high-order 3D immersed interface discretization of the advection diffusion equation with moving interfaces, an important intermediate step on the path to a fully high-order sharp interface Navier-Stokes solver.

James Gabbard, Wim van Rees
MIT
jgabbard@mit.edu, wvanrees@mit.edu

MS411 Optimizations and Algorithm Tradeoffs in Embedded Boundary Discretizations

We present a novel higher-order accurate Embedded Boundary (EB) discretization, implemented in a flexible software framework with good GPU performance. By separating concerns for geometry generation, finite volume stencil generation and solvers, we are able to apply multiple optimizations independently. For example, mesh refinement and subdivision algorithms can be used to achieve very fast, high precision representations of moderately complex surfaces. EB and tensor stencils can be represented as mixed-precision hyper-sparse matrix operations on block-structured grid data, greatly reducing memory storage and bandwidth. And many nested loop structures and solver operations can be recast as optimized matrix operations that better utilize GPU floating point capabilities. When combined with fast algorithms and calls to optimized GPU libraries, we are able to demonstrate significant speedups over traditional CPU implementations and

solvers for real applications.

Daniel T. Graves

Lawrence Berkeley Laboratory
dtgraves@lbl.gov

Hans Johansen

Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

Oscar Antepara

Lawrence Berkeley National Laboratory
oantepara@lbl.gov

Nate Overton-Katz

Lawrence Berkeley National Lab
novertokatz@lbl.gov

MS411

The Hermite-Taylor Correction Function Method for Embedded Boundary and Maxwells Interface Problems

The Hermite-Taylor method is a high-order method for hyperbolic problems that relies on a Hermite interpolation procedure in space and a Taylor method in time. The stability condition of this method depends only on the largest wave-speed, independent of the order. However, the treatment of general boundary and interface conditions is still a challenge since the Hermite-Taylor method requires not only to know the electromagnetic fields on the boundary and the interface but also their derivatives through order m in each Cartesian coordinate to achieve a $(2m+1)$ -order accurate method. In this talk, we investigate the correction function method to estimate all the needed information at the nodes near either the boundary or the interface in the Hermite-Taylor setting. The correction function method relies on the minimization of a functional based on Maxwell's equations. This minimization problem has to be solved only for some nodes in the vicinity of the boundary or the interface at each time step. Numerical examples in 1-D and 2-D are performed and the expected convergence order is observed for reasonable values of m .

Yann-Meing Law, Daniel Appelö

Michigan State University
lawkamci@msu.edu, appeloda@msu.edu

Thomas M. Hagstrom

Southern Methodist University
Department of Mathematics
thagstrom@smu.edu

MS411

Higher-Order Finite-Volume Embedded Boundary Methods on Moving Domains

This works presents a high-order finite-volume embedded boundary method capable of representing moving boundaries. The algorithm automates mesh generation around complex geometries by embedding them in a Cartesian grid with cut-cells on the boundaries, and is capable of providing high-order geometric information as necessary. In space, the scheme achieves fourth-order accurate discretizations that are stable in the presence of small cut-cells. We devise weighted least-squares reconstructions to achieve this without approaches such as mesh modifica-

tion, cell merging, or redistribution. In time, we employ a second-order implicit-explicit Runge-Kutta scheme which treats viscous and diffusive terms implicitly, while treating the advection and boundary motion terms explicitly. We apply this algorithm to modeling the formation of small bubbles of liquid water produced on catalyst sites inside hydrogen fuel-cells. Modeling this system requires coupled multi-phase physics with reactant gasses governed by advection-diffusion equations, produced liquid water governed by the unsteady Stokes equations, and chemical reactions occurring on solid boundaries. Embedded boundaries that evolve in time and depend on the fluid behavior represent the boundaries between the gas, liquid, and solid materials. We demonstrate the formal accuracy of the algorithm using grid refinement studies, and show simulations of practical application.

Nate Overton-Katz

Lawrence Berkeley National Lab
novertokatz@lbl.gov

Daniel T. Graves

Lawrence Berkeley Laboratory
dtgraves@lbl.gov

Hans Johansen

Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

MS411

Robust Cut-Cell Finite Elements in Two and Three Dimensions

This presentation concerns a geometrically simple and computationally robust cut-cell finite element method for Poisson's equation on arbitrarily shaped two- and three-dimensional domains. The equation is solved on a Cartesian axis-aligned grid of 4-node square elements in two dimensions and 8-node cube elements. These intersect the boundary of the domain in a smooth but arbitrary manner. Dirichlet boundary conditions are imposed strongly by means of a projection method, while Neumann boundary conditions require integration over a locally discretized boundary region. Representative numerical experiments demonstrate that the proposed method is stable and attains the asymptotic convergence rates expected of the corresponding unstructured body-fitted finite element method. Limited proofs of stability (in the sense of the inf-sup condition) are also available and will be discussed.

Panayiotis Papadopoulos

Department of Mechanical Engineering
University of California, Berkeley
panos@me.berkeley.edu

Donghao Li

University of California, Berkeley
donghaoli@berkeley.edu

MS412

Fast Data-Driven Model Reduction for Non-Linearizable Dynamics

An area of intense study in nonlinear dynamics is automated model identification of nonlinear systems based entirely on data. An ideal model is simple enough for system interpretability and control, but rich enough to capture

essential nonlinearities. Efficient reduction of very high-dimensional datasets is another challenge. In our approach, we use spectral submanifolds (SSMs) as reduced spaces. Specifically, an SSM is the unique smoothest nonlinear continuation of a spectral subspace from a steady state. After identification of the SSM, we compute its dynamics in the Poincaré normal form. The recently presented SSMLearn algorithm identifies SSMs from data via nonlinear optimization. While broadly applicable, this method is too computationally costly to handle higher-dimensional data. Here, we introduce a simplified method, fastSSM. Our main assumption is to estimate the tangent space of the SSM via singular value decomposition. Furthermore, the normal form dynamics are computed analytically rather than via implicit optimization. This entails a speedup of several orders of magnitude, allowing analysis of higher-dimensional datasets. We demonstrate this by training on a full surface profile signal from tank sloshing experiments. Furthermore, we show that our algorithms can capture the loss of stability and transition between buckling states with data from a FE beam simulation.

Joar Axås

Chair in Nonlinear Dynamics
ETH Zurich
jgoeransson@ethz.ch

Mattia Cenedese
ETH Zurich
Switzerland
mattiacenedese15@gmail.com

George Haller
ETH, Zurich
Switzerland
georgehaller@ethz.ch

MS412

Projection-Based Reduction of Adaptively-Refined Nonlinear Models

Adaptive Mesh Refinement (AMR) is often required and fairly practiced in the context of high-dimensional, mesh-based numerical simulations. However, it is in its infancy in that of low-dimensional, generalized-coordinate-based simulations such as those performed using projection-based reduced-order models (PROMs). This talk presents a complete framework for executing PROMs in the presence of AMR that builds on elements of readily available methods while contributing several critical new ones. In particular, it introduces an efficient algorithm for computing a pseudo-meshless inner product between AMR snapshots for the purpose of projection and clustering. The proposed framework allows for hyperreduction, which is often necessary for parametric and/or nonlinear problems: for this purpose, it incorporates the energy-conserving sampling and weighting (ECSW) method. It relies on the concept of local reduced-order bases to achieve computational tractability in the presence of AMR. All new ideas described in the talk are illustrated with the solution of three-dimensional, nonlinear, convection-dominated turbulent flow problems. Their impact is highlighted by the significant speedups that they deliver for realistic problems.

Clayton Little, Charbel Farhat
Stanford University

cr11@stanford.edu, CFarhat@stanford.edu

MS412

Structural Condition Monitoring Using Parametric Reduced Order Models and Limited Output Measurements

Structural and mechanical systems are oftentimes subjected to damage, which impels the vibration response beyond the normal operational regime. Apart from challenging the system's serviceability and integrity, these events pose a limit to the versatility and applicability of structural health monitoring strategies. The present works address the challenges associated with the response prediction and parameter identification of systems subjected to such events in the context of vibration-based structural health monitoring. The proposed approach relies on projection-based Reduced Order Models for the representation of system dynamics, which are tailored to an inverse framework that aims to estimate characteristic traits of the system's state and features of the induced damage. To assess the damage evolution and capture the changes in system parameters over time, the ROMs are used as predictors in the context of a sequential Bayesian inference problem, whereby the physics is fused with vibration response measurements. This approach allows recursive state or damage-related parameter estimation through a bank of filters and an evolution strategy that captures the parametrized states' dynamics. By continuously monitoring damage-related features and tracking the system's traits, the proposed framework aims to provide an indirect evaluation of the system's condition, while quantifying the induced damage.

Konstantinos Vlachas

Chair of Structural Mechanics & Monitoring
ETH Zurich
vlachas@ibk.baug.ethz.ch

Konstantinos Tatsis
ETH
tatsis@ibk.baug.ethz.ch

Anthony Garland, Carianne Martinez
Sandia National Laboratories
agarlan@sandia.gov, cmarti5@sandia.gov

Eleni Chatzi
Institute of Structural Engineering
ETH Zurich
chatzi@ibk.baug.ethz.ch

MS413

Multi-Scale and Rotation-Equivariant Graph Neural Networks for the Simulation of Fluid Dynamics

The numerical simulation of fluid dynamics is an essential tool in many areas of science and engineering. However, its high computational cost can limit application in practice. Recent deep-learning approaches have demonstrated the potential to yield surrogate models for fluid dynamics simulation. While such models exhibit lower accuracy in comparison, their low runtime makes them appealing for design-space exploration and real-time simulation. We introduce two novel graph neural networks (GNNs) for extrapolating the time evolution of a fluid in an unstructured discretization of the domain. In both models, previous states are processed through multiple coarsenings of

the graph, which enables faster information propagation through the fluid domain and improves the capture and forecast of the system state. One of the models is architecturally equivariant to rotations, which allows the network to learn the underlying physics more efficiently. We analyse these models using two canonical fluid models: advection and incompressible fluid dynamics. Our results show that the proposed models can generalise from uniform advection fields to high-gradient fields on complex domains. The multi-scale graph architecture allows for inference of incompressible Navier-Stokes solutions, within a range of Reynolds numbers and design parameters, more effectively than a single-scale GNN. Simulations are between 10^2 to 10^4 times faster than the numerical solutions on which they were trained.

Mario Lino
Imperial College London
mall1218@ic.ac.uk

MS413

Exploiting Parallelism in Extreme Multi-Label Graph Classification Problems

We are developing graph-based techniques for extreme multilabel text classification problems that arise in e-commerce applications. Our techniques are relatively easy to interpret, use intuitive parameters, and are competitive with fast neural network-based classification methods. This talk will introduce the motivating applications and present data structures that make the training and inference algorithms amenable to exploiting shared-memory multicore parallelism. We will also comment on extending the implementations to GPUs and distributed-memory systems.

Ashirbad Mishra
Pennsylvania State University
ashirbadm@gmail.com

Shad Kirmani
LinkedIn
skirmani@linkedin.com

Kamesh Madduri
Pennsylvania State University
madduri@psu.edu

MS413

Graph Partitioning for Communication-Intensive and High-Concurrency Graph Analytics

Graph partitioning is a critical pre-processing step to ensure load balance and low communication overheads in distributed graph analytic computations. With modern graph data at massive scales and modern high performance computation platforms offering significant concurrency both on-node and across a networked system, graph partitioning algorithms and approaches designed for regular scientific problems are becoming increasingly non-optimal. We consider graph partitioning in the context of irregular computations on irregular large-scale graph datasets on highly parallel systems. We propose several novel optimization criteria to balance work and communication loads for these data-driven computations, and we implement partitioning approaches for our new criteria in a GPU framework optimized for our the target datasets and systems.

Michael Mandulak
RPI

mandum@rpi.edu

George M. Slota
Rensselaer Polytechnic Institute
slotag@rpi.edu

MS413

Graman: Graph Network Based Simulator for Forecasting Molecular Polarizabilities

Realistic, high-quality simulations of complex physical systems often require substantial computational resources making them challenging to scale up to larger systems. A promising approach to overcome these limitations is to use machine learning models trained on experimental and simulation data. One of the many applications that can benefit from such an approach is simulating Raman spectra from ab initio molecular dynamics (AIMD) simulations. Recent works have shown that Graph Network-based simulators can be trained to simulate complex dynamical systems. In this talk, we will present our Graph Network-based simulator to accelerate the simulation of Raman spectra. The proposed simulator constitutes of three components: (1) encoder which constructs the graph given the atomic positions and embeds the atoms/edges into a latent space; (2) Graph-Network processor which learns the spatial interactions among atoms and updates the latent embeddings; and (3) decoder which learns the dynamics and predicts the next state of the simulation. Currently, our simulator is trained to simulate molecule trajectories. We trained our simulator on synthetic simulation data obtained from NWChem. We observed satisfactory performance on single-step prediction. We plan to develop novel methodologies to improve multi-step prediction performance. Finally, we dive into future research opportunities on generalizing this approach to other physical systems.

Vishal Dey
Ohio State University
dey.78@buckeyemail.osu.edu

Marco Minutoli, Mahantesh Halappanavar
Pacific Northwest National Lab
marco.minutoli@pnnl.gov,
mahantesh.halappanavar@pnnl.gov

Edoardo Apra, Patrick El Khoury, Niri Govind
Pacific Northwest National Laboratory
edoardo.apra@pnnl.gov, patrick.elkhoury@pnnl.gov,
niri.govind@pnnl.gov

MS413

Reducing Communication in Graph Neural Network Training

Graph Neural Networks (GNNs) are powerful and flexible neural networks that use the naturally sparse connectivity information of the data. GNNs represent this connectivity as sparse matrices, which have lower arithmetic intensity and thus higher communication costs compared to dense matrices, making GNNs harder to scale to high concurrencies than convolutional or fully-connected neural networks. We show that communication-avoiding matrix multiplication algorithms can accelerate GNN training compared to existing training methods with two families of algorithms. The first asymptotically reduces communication in full-batch training by leveraging 1D, 1.5D, 2D, and

3D SpMM algorithms. The second applies communication-avoiding algorithms for SpGEMM to parallelize GNN mini-batch sampling algorithms (e.g. GraphSAGE, LADIES, GraphSAINT). Altogether, these algorithms optimize communication across the full GNN training pipeline. We experiment on graph datasets with billions of edges on over a hundred GPUs to show the performance benefits of these algorithms.

Alok Tripathy
University of California, Berkeley
alokt@berkeley.edu

MS414

Geometric Effects in Wall Shear Stress Uncertainty Quantification in the Cardiovascular System

The vascular system delivers nutrients and oxygen to the body through blood vessels whose walls are subjected to a constant hemodynamic stress due to blood flow. This wall shear stress (WSS) has proven to be a biomarker related to the thickness, ulceration and rupture of atherosclerotic plaques. In practice, the WSS is not measured directly, but computed from blood flow measurements obtained with phase-contrast 4d flow magnetic resonance imaging. This computation depends explicitly on the flow velocity gradients and the geometry of the vessel. Therefore, the statistics of the WSS measurement noise are different from those of the flow velocity measurement noise. In this work we study the geometric effects that characterize the uncertainty propagation from the flow velocity measurements into the WSS measurements, focusing on velocity measurements obtained with MRI. To achieve this, we propose an additive noise model based on Gaussian processes that allow us to explicitly determine these geometric effects. In particular, we study the spatial correlations in the WSS noise. In addition to this theoretical model, we study a noise model arising from phase-contrast 4d flow magnetic resonance imaging using numerical simulations. We perform additional numerical simulations on synthetic geometries, and on human anatomic models. Our numerical results confirm our theoretical findings and provide evidence of spatial correlations in the WSS noise.

Constanza M. Gainza, Carlos Sing Long
Institute for Mathematical and Computational Engineering
Pontificia Universidad Catolica de Chile
cmgainza@uc.cl, casinglo@uc.cl

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

MS414

Stochastic Chaos and Predictability in Laboratory Earthquakes

Laboratory earthquakes exhibit characteristics of a low-dimensional random attractor with average dimension similar to that of natural slow earthquakes (≈ 5). Instantaneous dimensions reach high values (≈ 10), indicating that some regions of the phase space need a high number of degrees of freedom (dofs). We present methods to explore these dofs via a stochastic perturbation of the underlying low-dimensional deterministic dynamics. We study the dynamics for the spectrum of fault slip modes from stable sliding to stick-slip events and find that a model of stochastic differential equations based on rate- and state-

dependent friction explains the observations. In the fully unstable regime we do not observe truly periodic behavior. Aperiodicities can be explained by small perturbations (≈ 0.5) in the stress state. Frictions nonlinear nature amplifies small scale perturbations, reducing the predictability of the otherwise periodic macroscopic dynamics. As applied to tectonic faults, our results imply that even small stress field fluctuations (≈ 150 kPa) can induce coefficient of variations in earthquake repeat time of a few percent. Moreover, these perturbations can drive an otherwise fast-slipping fault, close to the critical stability condition, into a mixed behavior involving slow and fast ruptures.

Adriano Gualandi
INGV-Istituto Nazionale di Geofisica e Vulcanologia
tommaso.alberti@inaf.it

Davide Faranda
LSCE, CEA Saclay l'Orme des Merisiers, CNRS
davide.faranda@lscce.ipsl.fr

Chris Marone
Pennsylvania State University
cjm@geosc.psu.edu

Massimo Cocco
Istituto Nazionale di Geofisica e Vulcanologia
massimo.cocco@ingv.it

Gianmarco Mengaldo
National University of Singapore
College of Design and Engineering
mpegrim@nus.edu.sg

MS414

Constraint Preserving Particle Flow Filters

Data assimilation can be thought of as a Bayesian inverse problem that combines noisy observations of reality with uncertain model dynamics. Flow based assimilation methods evolve the samples from an uninformed prior to posterior that is informed by the observation via a defined stochastic differential equation. But what happens when the model states are constrained to be on a physical manifold? Think of a pendulum with a constant cord length. Standard filters have a minimal ability to strongly enforce these physical constraints. To this end, we formulate a flow based filter that treats the assimilation as an evolution of a stochastic differential algebraic equation (SDAE). We also specifically discuss the time evolution of the said SDAE.

Amit Subrahmanya
Virginia Tech
amitns@vt.edu

MS414

Bayesian Inference with Gaussian Process Models for Multiscale Data and Model Integration

TBD

Haruko M. Wainwright
Massachusetts Institute of Technology
hmwainw@mit.edu

MS415

Numerical Algorithms for Special Functions in the

Santander-Amsterdam Project

We give a brief overview of recent published algorithms by our group and of current activities and future plans. In particular we discuss in some detail two current lines of research: the computation and inversion of cumulative distributions and the evaluation of confluent hypergeometric functions. This is joint work with Javier Segura and Nico Temme.

Amparo Gil

Depto. de Matematica Aplicada y CC de la Comput
University of Cantabria
amparo.gil@unican.es

MS415**On the Numerical Approximation of Functions with Singularities**

Several classes of special functions have singularities, a particular example being Greens functions of partial differential equations. In that case both the nature and the location of the singularity are known in advance. In this talk we discuss numerical methods for the approximation of functions with such singularities. The methodology is based on enriching an approximation space with singular functions. This leads to redundancy in the space, which in turn leads to ill-conditioning of the approximation problem. Yet, we can prove numerical stability of an algorithm based on the theory of frames, we can devise an efficient implementation with near linear complexity in many cases (called the AZ algorithm), and we make it available as part of a Julia toolbox called FrameFun.

Daan Huybrechs

Department of Computer Science
KU Leuven
daan.huybrechs@cs.kuleuven.be

MS415**Providing Library-Grade Software for Special Functions**

Providing accurate values to a special function over its full domain in an efficient manner can be challenging. This will often involve several different algorithms for different parts of the domain and return values that can be on very different scales of magnitude. We use the example of the Gauss Hypergeometric function as implemented in the NAG Library to illustrate how these challenges can be overcome.

Lawrence Mulholland

NAG
lawrence.mulholland@nag.co.uk

MS415**Input-Output Finite Time Stabilization of Sampled Data Control for Semi Markov T-S Fuzzy System with Affine Matched Membership Function**

This article aims to examine the problem of Input-output finite-time stabilization of sampled-data control for semi-Markov fuzzy systems with time delay and disturbance. In most general cases, mode-dependent non-linear systems are difficult to study. In order to deal with this issue, nonlinear systems can be linearised by fuzzy with membership functions. To be specific, affine matched membership function for system and controller is considered. To stabilize and

reduce the communication burden, sampled data feedback control is utilized for the considered system. By endowing Lyapunov stability theory, a set of sufficient conditions is developed in the frame of linear matrix inequalities to guarantee that the fuzzy system is Input Output Finite time stable. Finally, two numerical examples that are useful in the field of engineering are provided to demonstrate the effectiveness of the proposed method.

Keerthana Nanjundan

Research Scholar
Anna University Regional Campus Coimbatore
keerthanananjundan99@gmail.com

Marshal Anthoni Selvaraj

Anna University
smarshalanthoni@gmail.com

MS415**Nist Digital Library of Mathematical Functions: Updates and Related Work**

In 2010, the National Institute of Standards and Technology (NIST) launched the Digital Library of Mathematical Functions (DLMF), a free online compendium of definitions, recurrence relations, differential equations, and other crucial information about functions useful to applied researchers in the mathematical and physical sciences. Although the DLMF was developed to expand and replace the widely cited National Bureau of Standards (NBS) Handbook of Mathematical Functions commonly known as Abramowitz and Stegun (A&S), it is far beyond a book on the web, incorporating web tools and technologies for accessing, rendering, and searching math and graphics content. This talk will spotlight completed and ongoing modifications to DLMF content and technology, and look at related follow-on work.

Bonita V. Saunders

National Institute of Standards and Technology, USA
bonita.saunders@nist.gov

MS416**Thermodynamically-Consistent Coupling of Fluctuating Hydrodynamics and Kinetic Monte Carlo for Gas-Solid Interfaces**

We consider a hybrid simulation method for a gas-solid interfacial system by combining the fluctuating hydrodynamics (FHD) method for the gas phase and the kinetic Monte Carlo (KMC) method for reactive surface. Since both simulation methods incorporate thermal fluctuations, a correct coupling of the two methods through a computational interface requires careful stochastic analysis. In addition, the fact that the two methods are based on disparate types of descriptions (that is, FHD is a continuum description whereas KMC is particle dynamics on a lattice) makes it a nontrivial task to develop a thermodynamically-consistent coupling between the two phases. By assuming that additional interactions between the two phases that are induced by the inclusion of surface chemistry occur only through adsorption and desorption processes, we develop a physically-inspired flux-exchange model where mass and energy fluxes across the interface are computed based on adsorption/desorption counts. By performing stochastic analysis, we justify the proposed mass and energy updates in FHD cells. We validate the overall hybrid simulation method by performing simulations for gas-solid interfaces undergoing reversible adsorption. We also discuss high-

performance computing implementation strategies.

Changho Kim
University of California, Merced
Applied Mathematics
ckim103@ucmerced.edu

Andy J. Nonaka
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
ajnonaka@lbl.gov

Ishan Srivastava
Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory
isriva@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Alejandro Garcia
San Jose State University
alejgarcia987@gmail.com

MS416

Structure-Dependent Multiscale Analysis in Heterogeneous Catalysis

Multiscale analyses based on structure-dependent microkinetic modeling are acknowledged to be an essential key tool to achieve a detailed mechanistic understanding of the catalyst functionality as they bridge the micro-scale environment of the active site and the macroscale reaction environment. Here, we present a methodology to perform a structure-dependent microkinetic analysis of a catalytic process. The methodology makes it possible to unveil the nature and identity of the active site in a self-consistent manner. The methodology is based on the combination of microkinetic modeling and ab initio thermodynamics with Wulff constructions and Boltzmann statistics at given conditions of chemical potential in the reactor. In doing so, the structure and the coordination geometry of each atom at the catalyst surface are determined. Then, microkinetic analyses are performed on the distribution of the sites in order to unravel the identity of the active site during the reaction. Selected examples in the context of water-gas-shift, reverse-water-gas shift, and CO₂ methanation on metal catalysts will be employed as showcases. As a whole, this methodology makes it possible a concomitant description of the nature of the catalyst material in reaction conditions and of its catalytic consequences in terms of reactivity. As such, it paves the way towards the use of first-principles methods for the interpretation of the experimental evidence in terms of structure-activity relationships.

Matteo Maestri
Politecnico di Milano
matteo.maestri@polimi.it

MS416

First-Principles Multi-Scale Modeling of Electrochemical CO₂ Reduction

Electrochemical CO₂ reduction is at the heart of the world's effort towards a sustainable energy future. Despite this, devices have not yet been optimized to the per-

formance level needed to overthrow conventional chemical processes. Here, we propose a new multi-scale model for electrochemical CO₂ reduction which is based on first-principles quantum chemical calculations. We connected a kinetic model built up from Density Functional Theory calculations with a mass transport model accurately depicting the design of a state-of-the-art electrolyzer. From that, we analyze local concentration variations with macroscopic design parameters, but also possible liquid intermediates and mechanistic effects. We conclude by giving general design principles to the hands of experimentalists to derive optimized electrolyzer cells.

Stefan Ringe
Korea University
sringe@korea.ac.kr

MS416

A Hybrid Kinetic-Continuum Computational Model for Simulations of Laser-Induced Plasma Plumes

A combined kinetic-continuum computational model for simulations of the shielding effect in strongly non-equilibrium plasma plumes. The model accounts for various modes of non-equilibrium, which are characteristic for plasma plumes induced by short- and ultra-short pulses. It includes a two-temperature model of the irradiated target, a kinetic model for flow of atoms and ions, and a continuum model for electrons in gaseous plasma. The model of plasma plume is implemented as a part of the direct simulation Monte Carlo (DSMC) method. The classical DSMC method is combined with the collision-radiation (CR) model using an approach, when a simulate particle represents both ions and atoms in multiple excited states. In the CR model, non-equilibrium distribution of ions, elastic collisions between electrons and ions, multiphoton ionization, inverse Bremsstrahlung, electron impact ionization, three-body recombination, as well as plasma emission are accounted for. The developed model is used to study the efficiency of burst laser ablation under conditions of strong plume accumulation effect and plasma shielding. This work was supported by the MKS Instruments, Inc. Work was also performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Release number LLNL-ABS-840136. A.N.V. also acknowledges support from NSF through RII-Track-1 Future Technologies and Enabling Plasma Processes project (award OIA-2148653).

Alexey N. Volkov, Michael Stokes
University of Alabama
avolkov1@ua.edu, mastokes2@crimson.ua.edu

Zhibin Lin
MKS Instruments, Inc.
zhibin.lin@mksinst.com

Saad Khairallah, Alexander Rubenchik
Lawrence Livermore National Laboratory
khairallah1@llnl.gov, rubenchik1@llnl.gov

MS417

Arithmetic Operations with H-matrices

\mathcal{H}^2 -matrices can be used to represent solution operators both for partial differential and boundary integral equations, but computing these (approximate) representations

still presents a challenge, particularly if we want to guarantee a given accuracy. This talk presents a technique for approximating the product of two \mathcal{H}^2 -matrices efficiently. The cluster bases used for compressing the result are computed adaptively in order to guarantee a prescribed accuracy of the approximated matrix. Using the efficient matrix multiplication algorithm, more complex tasks like computing the inverse or constructing triangular factorizations can be approached.

Steffen Boerm
Christian-Albrechts-Universität zu Kiel
boerm@math.uni-kiel.de

MS417

Reducing Communications and Memory Costs of Parallel Block Low-Rank Solvers

Block Low-Rank (BLR) compression is successful at reducing the high computational cost and memory footprint of sparse direct linear solvers. With the growing sizes of problems and increasing relative costs of data movements on modern architectures, it has become crucial to further optimize the memory and communication costs of parallel BLR solvers. In order to do so, we introduce several new approaches targeting both the LU factorization and solution phases of the solver. The key ideas include using mixed precision representations, designing data access patterns better suited to BLR, and compressing intermediate data, even at the cost of increasing the number of operations. We evaluate the potential of these approaches on a range of matrices coming from real-life problems in industrial and academic applications.

Patrick R. Amestoy
Mumps Technologies
patrick.amestoy@mumps-tech.com

Olivier Boiteau
EDF R&D
olivier.boiteau@edf.fr

Alfredo Buttari
CNRS-IRIT-Université de Toulouse, France
alfredo.buttari@irit.fr

Matthieu Gerest
EDF R&D
matthieu.gerest@gmail.com

Fabienne Jézéquel
Sorbonne Université and Université Panthéon-Assas
fabienne.jezequel@lip6.fr

Jean-Yves L'Excellent
Mumps Technologies
jean-yves.l.excellent@mumps-tech.com

Theo Mary
Sorbonne Université, CNRS, LIP6
theo.mary@lip6.fr

MS417

O(N) Factorization of Dense Matrices on GPUs Without Trailing Submatrix Dependencies

In the previous work, we demonstrated pre-computing the fill-ins and integrate them into the shared basis, allows us

to remove the dependency on trailing-matrices even for \mathcal{H}^2 -matrices. Comparisons with a block low-rank factorization code LORAPO already showed a maximum speed up of 4,700x for a 3-D problem with complex geometry. In this talk, we are aiming to give a performance extension of our previous work using batched LAPACK/BLAS API as well as GPUs. The adoption of a method that is free of data dependencies, has great potential of utilizing parallel computing hardware with optimal efficiency.

Qiangxian Ma, Rio Yokota
Tokyo Institute of Technology
ma@rio.gsic.titech.ac.jp, rioyokota@gsic.titech.ac.jp

MS417

Distributed Multi-GPU algorithms for H2 Matrices

We describe high-performance, distributed-memory, and GPU-accelerated algorithms for matrix-vector multiplication and matrix recompression of hierarchical matrices in the \mathcal{H}^2 format. Results show near-ideal scalability up to 1024 NVIDIA V100 GPUs, with performance exceeding 2.3 Tflop/s/GPU for the matrix-vector multiplication, and 670 Gflops/s/GPU for matrix compression, which involves batched QR and SVD operations. The algorithms are a new module of H2Opus, a performance-oriented package that supports a broad variety of \mathcal{H}^2 matrix operations on CPUs and GPUs.

George M. Turkiyyah
KAUST
George.Turkiyyah@kaust.edu.sa

Wajih Boukaram
Lawrence Berkeley National Lab
wajih.boukaram@lbl.gov

Stefano Zampini
King Abdullah University of Science and Technology
stefano.zampini@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS418

Structure-Preserving Model Order Reduction on Manifolds

Classical model order reduction (MOR) computes a reduced solution in a linear subspace. Thus, its accuracy is limited by the Kolmogorov n -widths, which is known to be especially problematic for advection-dominated problems. This is why recent approaches like “MOR on manifolds” actively investigate techniques which rely on nonlinear embeddings [K. Lee and K. Carlberg, “Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders”, 2020]. However, most of these approaches do not respect special structures in the equations. This concerns e.g. the structures present in Lagrangian or Hamiltonian equations. In our work, we use differential geometry to formulate a rather general theoretical framework which is able to preserve such structures during the reduction. We show that this formalism generalizes structure-preserving MOR techniques with linear embeddings to the case of nonlinear embeddings. In a numerical experiment, we investigate for the one-dimensional wave equation how

techniques from machine learning can be used to derive a nonlinear embedding for structure-preserving MOR of Hamiltonian systems on nonlinear manifolds.

Patrick Buchfink

Universitaet Stuttgart

patrick.buchfink@ians.uni-stuttgart.de

Silke M. Glas

University of Twente

s.m.glas@utwente.nl

Bernard Haasdonk

University of Stuttgart

haasdonk@mathematik.uni-stuttgart.de

MS418

Machine Learning Structure Preserving Brackets for Forecasting Irreversible Processes

Forecasting of time-series data requires imposition of inductive biases to obtain predictive extrapolation, and recent works have imposed Hamiltonian/Lagrangian form to preserve structure for systems with reversible dynamics. In this work we present a novel parameterization of dissipative brackets from metriplectic dynamical systems appropriate for learning irreversible dynamics with unknown a priori model form. The process learns generalized Casimirs for energy and entropy guaranteed to be conserved and nondecreasing, respectively. Furthermore, for the case of added thermal noise, we guarantee exact preservation of a fluctuation-dissipation theorem, ensuring thermodynamic consistency. We provide benchmarks for dissipative systems demonstrating learned dynamics are more robust and generalize better than either “black-box” or penalty-based approaches.

Kookjin Lee

Arizona State University

klee263@asu.edu

Nathaniel Trask

Sandia National Laboratories

natrask@sandia.gov

MS418

Structure-Preserving Model Reduction for Dissipative Differential Equations

In this talk, we consider numerical computation of dissipative gradient systems, such as Hamiltonian systems with friction, and systems that arise as discretizations of certain dissipative partial differential equations. Such systems are often equipped with some energy functional that decays with time, which is important in that it controls asymptotic behavior of solutions. The energy decay property could be preserved by “structure-preserving numerical methods.” They are preferable for better qualitative behaviors, but in many cases schemes become fully-implicit, and prohibitive as is for large systems. It encourages us to introduce some model reduction techniques (for example, the proper orthogonal decomposition [Moore, 1981]), but it in turn generally destroys the important mathematical structure, the gradient structure, which makes the overall strategy very problematic. In this talk, we show a new structure-preserving model reduction technique for dissipative gradient systems, which successfully preserves the gradient structure. The technique is partially inspired by an existing technique for Hamiltonian systems (without friction)

[Peng–Mohseni, 2016], but provides more general guiding principle for structure-preservation. With some numerical examples we illustrate that in fact such a structure-preserving model reduction provides us efficient, stable numerical computations.

Takayasu Matsuo

Graduate School of Information Science and Technology,

the University of Tokyo

matsuo@mist.i.u-tokyo.ac.jp

Masaya Inaba

the University of Tokyo

masaya.inaba0221@gmail.com

MS418

Structured Neural Networks and Their Relevance for Mechanical Systems

Neural networks have gained much interest because they can effectively approximate functions defined in high-dimensional spaces. Moreover, when suitably designed, they can also be employed to approximate functions while reproducing an expected structure or property. An example is working with symplectic neural networks to approximate Hamiltonian systems’ flow maps while preserving the model’s interpretability. Thus, neural networks can be relevant tools for physical problems as well. In this talk, we present a formalism for designing neural networks that satisfy specific desired properties. We then demonstrate their effectiveness in the problem of discovering the dynamics of a mechanical system while maintaining some expected features.

Elena Celledoni, Andrea Leone, Davide Murari

Norwegian University of Science and Technology

elena.celledoni@ntnu.no, andrea.leone@ntnu.no,

davide.murari@ntnu.no

Brynjulf Owren

NTNU

brynjulf.owren@math.ntnu.no

Carola-Bibiane Schönlieb, Ferdia Sherry

University of Cambridge

cbs31@cam.ac.uk, fs436@cam.ac.uk

MS419

Conservative Iterative Solvers for Systems of Conservation Laws

We consider implicit discretizations for systems of nonlinear hyperbolic conservation laws. A successful design principle for such discretizations has been to respect conservation. Iterative methods are typically not designed based on such ideas. The question thus arises if iterative methods are locally conservative. As it turns out, many commonly used methods preserve the local conservation of an underlying implicit scheme. This includes pseudo-time iterations, Newton’s method and Krylov subspace methods. However, there are prominent exceptions, in particular the Jacobi and Gauss-Seidel iterations. We present extensions of the Lax-Wendroff theorem for a fixed, finite number of iterations each time step. The iterative method defines a numerical flux that is inconsistent in general. We can describe the specific inconsistency as a form of slowed down time for pseudo-time iterations, Krylov methods, and thereby also Newton-Krylov methods. A simple technique based on the explicit Euler method can alleviate flux-inconsistency

and can additionally be used to generate improved initial guesses for iterative methods. Numerical experiments illustrate the theory.

Philipp Birken, Viktor Linders
Lund University

philipp.birken@math.lu.se, viktor.linders@math.lu.se

MS419

Multiple Invariants-Preserving Relaxation Runge Kutta Methods for Conservative Dynamical Systems

Relaxation Runge-Kutta methods, which are a slight modification of the RK methods, have been introduced to preserve invariants of initial-value problems. So far, this approach has been applied to preserve only one nonlinear functional in the numerical solution of a problem. In this talk, I will present the generalization of the relaxation approach for RK methods to preserve multiple nonlinear invariants of a dynamical system. The significance of preserving multiple invariants and its impact on long-term error growth will be illustrated via several numerical examples.

Abhijit Biswas

King Abdullah University of Science and Technology
abhijit.biswas@kaust.edu.sa

MS419

Walking into the Complex Plane to "order" Better Integrators

Most numerical methods for time integration use real time steps. Complex time steps provide an additional degree of freedom, as we can select the magnitude of the step in both the real and imaginary directions. By time stepping along specific paths in the complex plane, integrators can gain higher orders of accuracy or achieve expanded stability regions. Complex time stepping also allows us to break the Runge-Kutta order barrier, enabling 5th order accuracy using only five function evaluations. We show how to derive these paths for explicit and implicit methods, discuss computational costs and demonstrate clear advantages for complex-valued systems like the Schrodinger equation.

Jithin D. George

Northwestern University
jithingeorge2023@u.northwestern.edu

MS419

Efficient Implementation of Implicit RungeKutta Methods with Downwind-Biased Operators

Strong stability preserving (SSP) time integrators have been developed to preserve certain nonlinear stability properties of the numerical solution in arbitrary norms, when coupled with suitable spatial discretizations. The existing general linear methods either attain small time steps for strong stability preservation or are only first-order accurate. One way to relax the time-step restrictions is to consider time integrators that contain both upwind- and downwind-biased operators. In this talk, we review downwind SSP Runge-Kutta (DWRK) methods in the framework of perturbations of classical Runge-Kutta methods. We show how downwinding improves the SSP properties of time-stepping methods and breaks some stability barriers. In particular, we focus on a novel one-parameter family of implicit DWRK methods for which the SSP coefficient

varies with respect to the method's coefficients and the CFL-like restriction can be arbitrarily large. We analyze the stability and accuracy of such methods for large CFL numbers. Finally, we discuss the computational complexity of implicit DWRK methods, and we propose a block factorization of the Jacobian that enhances the solution of Newton's method.

Yiannis Hadjimichael

Weierstrass Institute for Applied Analysis and Stochastics
hadjimichael@wias-berlin.de

Sidafa Conde

Mythic Inc.
sconde89@gmail.com

Edward G. Phillips

Sandia National Laboratories
egphill@sandia.gov

MS419

Efficient Numerical Solution of Advection-Diffusion-Reaction Equations by Modified PatankarRungeKutta Methods

In this talk we numerically solve advection-diffusion-reaction equations using a Strang splitting. Thereby, the NPZD model—a stiff, positive and conservative system of ordinary differential equations (ODE) describing the ecosystem of phytoplankton—arises as a sub-problem. For solving the NPZD model we compare standard ODE solvers and an unconditionally positive and conservative modified Patankar-Runge-Kutta method (MPRK) with respect to their efficiency. Thereby, we confirm that only the MPRK scheme is positive for all time step sizes while the remaining schemes face severe time step restrictions in order to maintain positive approximations. As a result, MPRK schemes may choose the maximal time step size coming from an explicit time integration of the advection-diffusion part also for the reaction part, whereas the standard ODE solvers must choose an even smaller time step size leading to an overall higher computational effort.

Thomas Izgin

University of Kassel
izgin@mathematik.uni-kassel.de

Stefan Kopecz

Institute for Mathematics
University of Kassel
kopecz@mathematik.uni-kassel.de

Andreas Meister

University of Kassel
Institute of Mathematics
meister@mathematik.uni-kassel.de

MS420

Multiscale Matrix Pencil Method for Superresolution

We present some recent multiscale factorizations of highly structured matrices encountered in exponential analysis, which is a sparse digital signal processing method. These factorizations give rise to matrix pencil methods for the analysis of the digital signal at different granularity levels of the signal samples. Here the granularity is allowed to not adhere to the Nyquist sampling constraint. We numerically

illustrate how coarsely collected signal samples allow to detect fine and weak signal details, meaning the presence of significantly higher signal frequencies than dictated by the Nyquist bandwidth condition. In addition, the high frequency components can be entirely buried in the signal noise.

Annie Cuyt
University of Antwerp
annie.cuyt@uantwerpen.be

MS420 Graph Signal Processing and Intralogistics

This presentation deals with graphs, graph signal processing (GSP) and possible applications to (intra-)logistics. In particular we focus on modelling a cooperating fleet of forklifters (FL's) as an N -vertex-graph (V, E, w) . The vertex set V splits in stationary (e.g. charging stations, pick/drop-places) and moving vertices (FL's), the edge-set E represents cooperation or source/sink-relations of the connected vertices, the weight function w quantifies these relations. As an example we study the task of optimal placement of charging stations for battery-driven FL's based on recorded motion and task data of the fleet. One idea is to connect stationary (charging stations) with moving vertices (FL's) with a proper edge set E and weight function w and to measure the 'importance' of the charging stations at their individual position by a properly modified pagerank-like algorithm. Alternatively in this graph we may consider the state of charge (SOC) of the individual FL's as a graph signal f and perform a sensitivity analysis to its graph Fourier transform \hat{f} , when the positions of the charging stations are varied. Apart from these examples we also show simulations of GSP-like models of FL-fleets. These simulations rely on CARLA; an open source tool developed mainly in the context of autonomous driving. The reported work comes from a cooperation of the Aschaffenburg University of Applied Sciences with a local manufacturer of fork lifters.

Mustafa Jelibaghu, Hans-Georg Stark
Aschaffenburg University of Applied Sciences
mustafa.jelibaghu@th-ab.de, hans-georg.stark@th-ab.de

MS420 Robust and Tractable Multidimensional Exponential Analysis

Given samples a multidimensional signal of the form

$$f(\mathbf{x}) = \sum_{k=1}^K a_k \exp(-\langle \mathbf{x}, \omega_k \rangle), \quad \mathbf{x}, \omega_1, \dots, \omega_k \in \mathbb{R}^q,$$

we develop a method based on localized trigonometric polynomial kernels to determine the values of the number K of components, and the parameters a_k and ω_k 's. Our method is easy to implement, and is shown to be stable under a very low SNR.

Sippanon Kitimoon
Institute of Mathematical Sciences
Claremont Graduate University
sippanon.kitimoon@cgu.edu

Hrushikesh Mhaskar
California State University, Los Angeles
Dept. of Mathematics

hmhaska@gmail.com

MS420 Towards the Future Removal of Bandwidth Constraints

The Nyquist condition which imposes a condition on the sampling step used when collecting data from a digital signal, in terms of the frequencies present in the signal, is at the basis of many digital signal processing algorithms. The Nyquist constraint can be circumvented by using mutually prime multiples of the sampling step, as in coprime arrays. The performance of the latter can be improved to using mutually prime scale and shift parameters instead of two mutually prime scales. We show that the bandwidth condition can be further stretched by cleverly choosing the multiples of the sampling step. Can we remove the bandwidth condition entirely?

Wen-shin Lee
University of Stirling
wen-shin.lee@stir.ac.uk

MS421 Accelerating PDE-Constrained Optimization Using Adaptive Certified Trust-Region Localized Reduced Basis Methods

PDE-constrained parameter optimization problems, where the involved equality constraint is a parameterized partial differential equation (pPDE), are of interest for many applications, such as parameter identification, decision-making, and inverse problems. In many cases, the underlying pPDE admits a large- or multi-scale structure which easily exceeds computational resources if the optimization method relies on standard approximation methods for the forward problem. In this talk, we are concerned with accelerating the overall computational cost of the optimization routine by utilizing adaptive model order reduction (MOR) methods for the pPDE. The adaptive procedure is based on a trust-region (TR) method that relies on the a posteriori error estimation of the reduction approach. While going beyond a classical offline/online splitting, it can be fully certified and shown to be convergent also for problems with a slow convergence of the Kolmogorov n -width. In addition, the method can reliably be relaxed in the initial phase of the iterative algorithm. We focus on several aspects of the algorithm in both the optimization loop and the full- and reduced-order models. We show results that use the classical reduced basis (RB) method for an elliptic forward problem. As a second step, we extend the TR-RB algorithm with adaptive and certified localized model order reduction techniques to fully replace the high-dimensional model.

Tim Keil
University of Münster, Mathematics Münster
Einsteinstraße 62, 48149 Münster, Germany
tim.keil@uni-muenster.de

Luca Mechelli
University of Konstanz
luca-michelli@uni-konstanz.de

Mario Ohlberger, Stephan Rave
University of Münster
mario.ohlberger@wwu.de, stephan.rave@wwu.de

Felix Schindler

University of Muenster
Applied Mathematics Muenster
felix.schindler@wwu.de

Stefan Volkwein
Universität Konstanz
Department of Mathematics and Statistics
stefan.volkwein@uni-konstanz.de

MS421

Optimization of Energy Consumption in District Heating Networks via Space-Time-Model Adaptive Methods

We propose an adaptive optimization algorithm for operating district heating networks in a stationary regime. The behavior of hot water flow in the pipe network is modeled using the incompressible Euler equations and a suitably chosen energy equation. By applying different simplifications to these equations, we derive a catalog of models. Our algorithm is based on this catalog and adaptively controls where in the network which model is used. Moreover, the granularity of the applied discretization is controlled in a similar adaptive manner. By doing so, we are able to obtain optimal solutions at low computational costs that satisfy a prescribed tolerance w.r.t. the most accurate modeling level. To adaptively control the switching between different levels and the adaptation of the discretization grids, we derive error formulas and a posteriori error estimators. Under reasonable assumptions we prove that the adaptive algorithm terminates after finitely many iterations. Our numerical results show that the algorithm is able to produce solutions for problem instances that have not been solvable before.

Hannes Dänschel, Volker Mehrmann
Technische Universität Berlin
daenschel@math.tu-berlin.de,
mehrmann@math.tu-berlin.de

Jacques-Marius Roland, Martin Schmidt
Universität Trier
roland@uni-trier.de, martin.schmidt@uni-trier.de

MS421

Adaptive Reduced Basis Methods for PDE-Constrained Optimization and Optimal Input Design

We propose an algorithm for the bi-level optimal input design involving a parameter-dependent evolution problem. In the inner cycle a control is fixed and the parameter is optimized in order to minimize a cost function that measures the discrepancy from some data. For this optimization problem a trust-region reduced basis approximation of the model with creation and enrichment of the reduced basis on-the-fly is used. In the outer cycle the found parameter is fixed and the control is now optimized in order to minimize a suitable measure of uncertainty of the parameters. Such measure of uncertainty is defined using the Fishers information matrix, that quantifies how informative an input is for a fixed parameter. Numerical examples illustrate the efficiency of the proposed approach.

Andrea Petrocchi
University of Konstanz
andrea.petrocchi@uni-konstanz.de

Matthias K. Scharrer

Virtual Vehicle Research GmbH
matthias.scharrer@v2c2.at

Stefan Volkwein
Universität Konstanz
Department of Mathematics and Statistics
stefan.volkwein@uni-konstanz.de

MS421

An Adaptive Scheme for the Optimization of Damping Positions by Decoupling Controllability Spaces in Vibrational Systems

In this work, the problem of optimizing damper positions in vibrational systems is investigated. The objective is to determine the positions of external dampers in such a way that the influence of the input on the output is minimized. The energy response serves as an optimization criterion, whose computation involves solving Lyapunov equations. Hence, in order to find the best positions, many of these equations need to be solved, and so the minimization process can have a high computational cost. To accelerate the process of finding the optimal positions, we propose a new reduction method. Our algorithm generates a basis spanning an approximation to the solution space of the Lyapunov equations for all possible positions of the dampers. We derive an adaptive scheme that generates the reduced solution space by adding the subspaces of interest, and then we define the corresponding reduced optimization problem that is solvable in reasonable amount of time. We decouple the solution spaces of the problem to obtain a space that corresponds to the system without external dampers and serves as a starting point for the reduction of the optimization problem. In addition, we derive spaces corresponding to the different damper positions that are used to expand the reduced basis if needed. Our new technique produces a reduced optimization problem of significantly smaller dimension that is faster to solve than the original problem, which we illustrate with some numerical examples.

Jennifer Przybilla
Max Planck Institute for Dynamics of Complex Technical Systems
przybilla@mpi-magdeburg.mpg.de

Matea Ugrica
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg
ugrica@mpi-magdeburg.mpg.de

Ninoslav Truhar
Josip Juraj Strossmayer University of Osijek, Croatia
ntruhar@mathos.hr

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS421

A Relaxation-Based Probabilistic Approach for Pde-Constrained Optimization under Uncertainty with Pointwise State Constraints

We consider a class of convex risk-neutral PDE-constrained optimization problems subject to pointwise control and state constraints. Due to the many challenges associated with almost sure constraints on pointwise evaluations of

the state, we suggest a relaxation via a smooth functional bound with similar properties to well-known probability constraints. First, we introduce and analyze the relaxed problem, discuss its asymptotic properties, and derive formulae for the gradient the adjoint calculus. We then build on the theoretical results by extending a recently published online convex optimization algorithm (OSA) to the infinite-dimensional setting. Similar to the regret-based analysis of time-varying stochastic optimization problems, we enhance the method further by allowing for periodic restarts at pre-defined epochs. Not only does this allow for larger step sizes, it also proves to be an essential factor in obtaining high-quality solutions in practice. The behavior of the algorithm is demonstrated in a numerical example involving a linear advection-diffusion equation with random inputs. In order to judge the quality of the solution, the results are compared to those arising from a sample average approximation (SAA). In addition, we conduct statistical tests to further analyze the behavior of the online algorithm and the quality of its solutions.

Mathias Staudigl
Maastricht University
m.staudigl@maastrichtuniversity.nl

Thomas M. Surowiec
Department of Mathematics and Computer Science
Philipps University of Marburg
thomasms@simula.no

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS422

Mathematical Analysis of a Discrete Covid-19 Epidemic Model

In this talk, we will study the global dynamics of a discrete Covid-19 epidemic model. A unique positive solution for the proposed model with the positive initial conditions is obtained. The stability analysis of the disease-free equilibrium and endemic equilibrium have been investigated. It has been proved that the DFE is globally asymptotically stable when the basic reproduction number $\mathcal{R}_0 \leq 1$. The proposed model has a unique endemic equilibrium that is globally asymptotically stable whenever $\mathcal{R}_0 > 1$. The theoretical results are illustrated by a numerical simulation.

Mahmoud H. DarAssi
Department of Basic Sciences
Princess Sumaya University for Technology
m.assi@psut.edu.jo

MS422

Dynamic Models of the Infection of Lassa Fever Epidemics Incorporating Detected and Undetected Class

In this study, a dynamical model of the infection of Lassa fever which incorporates detected and undetected classes was proposed to investigate the effect of vaccine on the spread of the Lassa disease. The stability analysis of the model was conducted to determine free and endemic equilibrium. The exact solutions of the models were obtained using analytic technique. The basic reproduction number was determined using the next generation matrix approach.

The local and global stability were carried out using Routh-Hurwitz conditions and the Castillo-chavez criterion. The numerical computations of the exact solutions were also done using Mathematica software. The result obtained shows that, the disease free equilibrium is locally asymptotically stable when the basic reproduction otherwise its unstable. These indicate that, the Lassa fever disease will be completely wiped out if the secondary infection is kept below unity.

Hamisu Idi
Federal Polytechnic Bauchi
hidi@fptb.edu.ng

MS422

Computational Modelling of Protein Conformational Changes - Application to the Opening Sars-CoV-2 Spike

We present a new approach to compute and analyze the dynamical electro-geometric properties of proteins undergoing conformational changes. The molecular trajectory is obtained from Markov state models, and the electrostatic potential is calculated using the continuum Poisson-Boltzmann equation. The numerical electric potential is constructed using a parallel sharp numerical solver implemented on adaptive Octree grids. We introduce novel a posteriori error estimates to quantify the solution's accuracy on the molecular surface. To illustrate the approach, we consider the opening of the SARS-CoV-2 spike protein using the recent molecular trajectory simulated through the Folding@home initiative. We analyze our results, focusing on the characteristics of the receptor-binding domain and its vicinity. This work lays the foundation for a new class of hybrid computational approaches, producing high-fidelity dynamical computational measurements serving as a basis for protein bio-mechanism investigations.

Anna Kucherova
University of California - Merced
University of California - Merced
akucherova@ucmerced.edu

Selma Strango
UC Merced
sstrango@ucmerced.edu

Maxime Theillard
University of California, Merced
mtheillard@ucmerced.edu

Shahar Sukenik
UC Merced
ssukenik@ucmerced.edu

MS422

A Deep Learning Approach for the Detection of COVID-19, Abnormal and Normal Lungs from Computed Tomography Slices using Different CNN Architectures

The outbreak of COVID-19 has had a devastating impact on the lives of many people. Due to the rampant epidemic of this pandemic, there have been many evident cases of other lung infections that are a threat to human life misdiagnosed as COVID-19 infection. Hence, an efficient system that distinguishes lungs affected by COVID-19 and other infections is highly essential for clinicians in these times. Therefore, a transfer learning-based convolution

neural network (CNN) is designed for accurate detection that classifies as covid-19 affected, abnormal and healthy from computed tomography slices. Competitive results of different CNN architectures namely, VGG16, VGG19, ResNet and Inception v3 networks are compared. Related works in the literature include [Sahinbas, K. and Catak, F.O., 2021. Transfer learning-based convolutional neural network for COVID-19 detection with X-ray images. In Data Science for COVID-19 (pp. 451-466). Academic Press] and [Isaac A., Nehemiah H.K., Isaac A. and Kannan A., 2020. Computer-Aided Diagnosis system for diagnosis of pulmonary emphysema using bio-inspired algorithms. Computers in Biology and Medicine, 124, p.103940].

Rebecca Mercy Victoria Manoharan, Khanna Nehemiah
H
Anna University, Chennai, India
reberia@gmail.com, nehemiah@annauniv.edu

MT2

How to Give Good Talks

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Katerina Kaouri
School of Mathematics
Cardiff University, UK
kaourik@cardiff.ac.uk

MP1

Matrix Free Multigrid Preconditioner for Dg-Sem for Compressible Flow

High fidelity fluid simulations have many important applications in science and engineering, with examples including numerical weather prediction and simulation aided design. Discontinuous Galerkin (DG) methods are high order discretizations, promising for simulation of unsteady compressible fluid flow in three dimensions. Systems arising from such discretizations of turbulent fluid motion are often stiff, and require implicit time integration. This requires fast, parallel, low-memory solvers for the algebraic equation systems that arise. For (low order) finite volume (FV) methods, multigrid (MG) methods have been successfully applied for this purpose. But for high order DG such solver are currently lacking. This deficiency inhibits wider adoption of DG methods, and motivates our research to construct a matrix free preconditioner for high order DG discretizations. The preconditioner is based on a multigrid method constructed for a low order finite volume discretization defined on a subgrid of the DG mesh. Numerical experiments on atmospheric flow problems show the benefit of this approach.

Johannes Kasimir
Lund University
Lund University
johannes.kasimir@math.lu.se

Robert Klöfkor
International Research Institute of Stavanger
robertk@math.lu.se

Philipp Birken
Lund University
philipp.birken@math.lu.se

Andreas Dedner
University of Warwick
a.s.dedner@warwick.ac.uk

MP1

An Overview on Recent Development in Dune and Dune-Fem

For about two decades the Distributed and Unified Numerics Environment (DUNE:https://dune-project.org) has been an active part in the scientific development of computational software and technology and its C++ routines are the basis for several other well established open source projects, for example, DuMux (https://dumux.org) or the Open Porous Media Initiative (https://opm-project.org). Although the C++ interfaces of DUNE are highly flexible and customizable, a solid knowledge of C++ is necessary to make use of this powerful tool. In this Minisymposium an overview on recent development towards a Python interface for DUNE and in particular DUNE-FEM, a module which provides highly efficient implementations of hp-adaptive Discontinuous Galerkin (DG) methods for solving a wide range of non linear partial differential equations. Providing easier user interfaces based on Python and the Unified Form Language (UFL) opens DUNE-FEM to a broader audience.

Robert Klöfkor
International Research Institute of Stavanger
robertk@math.lu.se

Andreas Dedner
University of Warwick
a.s.dedner@warwick.ac.uk

MP1

Reactive Transport Models for Ice-Ocean Interfaces

Heat and mass exchange processes at the boundary between sea ice and ocean saltwater play a vital role in understanding the coupling between earth's ocean and atmosphere. Ice-ocean interface models are furthermore used to constrain thickness and composition of the kilometer-thick shell of icy moons such as the Jovian moon Europa. Mesoscale models of the ice-saltwater interface treat the transition region between ice and seawater as a continuous mush containing a mix of solid and high-salinity brine [Buffo et al., Dynamics of a Solidifying Icy Satellite Shell, 2021]. This is in analogy to established enthalpy methods for the treatment of alloy solidification [Swaminathan and Voller, On the enthalpy method, 1993]. Governing equations for incompressible flow are coupled with thermal and solute transport through buoyancy and advective terms. In this contribution we focus on the numerical treatment of the mixture solidification problem. Assuming thermodynamic equilibrium, conservation of energy and solute transport are solved with discontinuous Galerkin finite elements. Our prototype uses the DUNE Python API [Dedner and Nolte, The Dune Python Module, 2018] and leverages several recent works on multiphase reactive transport. In particular, we generalize a variable switching technique [Krabbenhoft et al., An implicit mixed enthalpy-temperature method for phase-change problems, 2007] to systems where the liquid fraction is a multivariable function.

Benjamin Terschanski
RWTH Aachen
terschanski@mbd.rwth-aachen.de

Julia Kowalski
RWTH Aachen University
kowalski@mbd.rwth-aachen.de

Mohammad Ghanem
RWTH Aachen
mohammad.ghanem1@rwth-aachen.de

PP1

Modelling and Simulation of Charge Transport in Perovskite Solar Cells

Due to its flexibility, perovskite materials are a promising candidate for many semiconductor devices. For example, Perovskite Solar Cells (PSCs) have become recently one of the fastest growing photovoltaic technologies. However, which exact physical operation mechanisms play a fundamental role within such devices is not fully understood yet. Experiments indicate that besides the movement of electric carriers, ion movement within the perovskite needs to be taken into account. For this reason it is paramount to understand the electronic-ionic charge transport within PSCs better via improved modelling and simulation. In this work, we take volume exclusion effects into account by formulating two different current densities either treating the mobility or the diffusion as density dependent while the other quantity remains constant. Finally, we compare both fluxes within drift-diffusion simulations performed by two different open-source tools.

Dilara Abdel
WIAS Berlin
abdel@wias-berlin.de

Nicola E. Courtier
University of Oxford
nicola.courtier@eng.ox.ac.uk

Patricio Farrell
WIAS Berlin
patricio.farrell@wias-berlin.de

PP1

Stochastic Inversion Strategies for Model Calibration and Optimization under Uncertainty : Applications in Materials Design

While significant advances have been made in forward and backward modeling of complex physical systems for capturing nonlinear and multi-scale processes, fewer efforts have been directed towards quantifying uncertainties. A fully integrated approach to the calibration and optimization of the physical systems necessitates the inclusion of uncertainties. To address this stochastic inversion problem, we propose a novel probabilistic framework to formulate the problem of optimization/design in the presence of uncertainty. We apply the proposed framework to materials design problems, in order to perform stochastic inversion of the material Process-Properties-Performance linkages. We advocate the use of probabilistic, data-driven surrogates for the process-property and property-performance links, that are capable of learning under Small Data, that incorporate physical constraints and are able to quantify their predictive uncertainty. We demonstrate the efficacy of the proposed scheme for concrete which exhibits highly nonlinear and multi scale behavior.

Atul Agrawal, Phaedon-Stelios Koutsourelakis
Technical University of Munich

atul.agrawal@tum.de, p.s.koutsourelakis@tum.de

PP1

An Incremental Tensor Train Decomposition for High-Dimensional Data Streams

We propose a tensor network-based approach to discover latent structures in large-scale tensor-structured data (i.e. solutions of high dimensional PDEs, video sequences, etc.). Moreover, we provide an algorithm to sequentially compress large-scale data as they arise through solvers or data gathering. Our approach is developed for computing the compression in the tensor-train format and we discuss and demonstrate advantages compared to the other incremental algorithms existing in the literature. We also demonstrate that our approach can be used for efficient large-scale data assimilation and Bayesian inverse problems to provide accurate estimations in the limited observation scenarios.

Doruk Aksoy
Department of Aerospace Engineering
University of Michigan
doruk@umich.edu

Alex Gorodetsky
University of Michigan
goroda@umich.edu

PP1

Local Compatibility Boundary Conditions for High-Order Accurate Finite-Difference Approximations of the Wave Equation on Curvilinear and Overlapping Grids

We describe a high-order accurate numerical boundary conditions method for finite-difference (FD) approximations of PDEs. The novel Local Compatibility Boundary Conditions (LCBC) method uses boundary conditions, derivatives of the governing equations, interior and boundary grid values to construct a local interpolating polynomial centered at each boundary point. Such polynomial gives a discrete formula for the solution at ghost points near the boundary. Ghost values are then used in the implementation of high-order accurate centered FD schemes. We develop the LCBC method for the wave equation initial boundary value problems over 2D and 3D Cartesian and curvilinear domains utilizing overset grids for geometrically complex domains.

Nour G. Al Hassanieh, Jeffrey W. Banks, William Henshaw
Rensselaer Polytechnic Institute
alhasn@rpi.edu, banksj3@rpi.edu, henshw@rpi.edu

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

PP1

PapStatistical Inference for Nonergodic Weighted Fractional Vasicek Models Er Here

A problem of drift parameter estimation is studied for a nonergodic weighted fractional Vasicek model defined as $dX_t = \theta(\mu + X_t)dt + dB_t^{\alpha, b}$, $t \geq 0$, with unknown parameters $\theta > 0$, $\mu \in R$ and $\alpha := \theta\mu$, whereas $B_t^{\alpha, b} := B_t^{\alpha, b}$, $t \geq 0$ is a weighted fractional Brownian motion with parame-

ters $a > -1, |b| < 1, |b| < a + 1$. Least square-type estimators $\tilde{\theta}_T, \tilde{\mu}_T$ and $\tilde{\theta}_T, \tilde{\alpha}_T$ are provided, respectively, for (θ, μ) and (θ, α) based on a continuous-time observation of $X_t, t \in [0, T]$ as $T \rightarrow \infty$. The strong consistency and the joint asymptotic distribution of $\tilde{\theta}_T, \tilde{\mu}_T$ and $\tilde{\theta}_T, \tilde{\alpha}_T$ are studied. Moreover, it is obtained that the limit distribution of $\tilde{\theta}_T$ is a Cauchy-type distribution, and $\tilde{\mu}_T$ and $\tilde{\alpha}_T$ are asymptotically normal.

Fares Alazemi, Mishari Alforaih, Khalifa Es-Sebaiy
Kuwait University
fares.alazemi@ku.edu.kw, mishari.alforaih@ku.edu.kw,
khalifa.essebaiy@ku.edu.kw

PP1

PDE-Constrained Optimisation for Thin Film Flow

The objective of this poster is to present the solution of time-dependent optimal control for nonlinear partial differential equations that arise in fluid dynamics, in particular, thin-film equations. Controlling thin film flow has been novel and desirable in many applications [Kalogirou, A. and Blyth, M.G., 2021. Instabilities at a sheared interface over a liquid laden with soluble surfactant]. For example, flat-flow film stability is preferred in coating processes, while interfacial deformation could be valuable in heat and mass transfer. Our main contribution is to develop the time-dependent optimality system for controlling thin-film flow by finding the appropriate topography for the relevant system that minimises the differences between thin-fluid film flow and the preferred state. We use the approach of Optimize then Discretize, that is optimality system will be derived [van der Zee, K.G., Tinsley Oden, J., Prudhomme, S. and Hawkins-Daarud, A., 2011. Goal-oriented error estimation for Cahn-Hilliard models of binary phase transition] as coupled systems of partial differential equations and then discretised and solved. An iterative algorithm is used for the optimality system as it contains a coupled system of time-dependent PDEs, including adjoint equations which are backwards in time [Azmi, B., Kalise, D. and Kunisch, K., 2021. Optimal feedback law recovery by gradient-augmented sparse polynomial regression].

Sattam Alrashidy
UNIVERSITY OF NOTTINGHAM
sattam.alrashidy@nottingham.ac.uk

Kris van der Zee, Anna Kalogirou
University of Nottingham
kg.vanderzee@nottingham.ac.uk,
anna.kalogirou@nottingham.ac.uk

Dante Kalise
Imperial College London
d.kalise-balza@imperial.ac.uk

PP1

Dual Solution Or Double-Diffusive Mixed Convection Opposing Flow Through a Vertical Cylinder Saturated in a Darcy Porous Media Containing Gyrotactic Microorganisms

The steady mixed convection flow towards an isothermal permeable vertical cylinder nested in a fluid saturated porous medium is studied. The Darcy model is applied to observe bioconvection through porous media. The suspension of gyrotactic microorganisms is considered for various applications in bioconvection. Appropriate similar-

ity variables are opted to attain the dimensionless form of governing equations. The resulting momentum, energy, concentration, and motile microorganism density equations are then solved numerically. The resulting dual solutions are graphically visualized and physically analyzed. The results indicate that depending on the systems parameters, dual solutions exist in opposing flow beyond a critical point where both solutions are connected. Our results were also compared with existing literature.

Abdulaziz Alsenafi
Kuwait University
abdulaziz.alsenafi@ku.edu.kw

Mohammad Ferdows
Department of Applied Mathematics
University of Dhaka
ferdows@du.ac.bd

PP1

A Hybrid Amr Low-Rank Tensor Approach for Solving the Boltzmann Equation

The Boltzmann equation describes the time evolution of a particle distribution function in six-dimensional position-velocity phase space. The exponential growth in computational complexity often challenges a grid-based approach to modeling the Boltzmann equation as the dimensionality grows. To mitigate such issues, scalable low-rank tensor decomposition techniques have recently been developed with applications to high-dimensional PDEs. Despite the remarkable progress made in the community, low-rank structures in the phase-space are not evident in realistic engineering systems with complex geometries (e.g., electric propulsion systems and fusion reactors), where discontinuities, shocks, complex boundary conditions, and material-dependent physics (e.g., collisions, fusion reactions, ionization/excitation, charge-exchange processes) pose formidable challenges. In this talk, we propose a novel hybrid algorithm where quad-tree adaptive mesh refinement (AMR) is applied in real space while a low-rank approximation is applied in the velocity space. The AMR algorithm efficiently handles challenges pertaining to complex structures in real space, while the low-rank formulation targets dimensionality challenges in the velocity space. We present preliminary results on the new algorithm applied to challenging multi-dimensional gas kinetics problems.

Samuel J. Araki
Jacobs Inc.
Air Force Research Laboratory
samuel.araki.ctr@us.af.mil

William Taitano
Air Force Research Laboratory
william.taitano.1@spaceforce.mil

PP1

Using Nonlinear Domain Decomposition As Smoother in Nonlinear Multigrid

Nonlinear partial differential equations frequently arise in different fields of science. Discretization of the nonlinear problems usually leads to largenonlinear systems. Solution of such big discretized nonlinear problems needs fast, highly scalable, and parallelize solvers. Nonlinear multigrid is a well-known method for efficiently solving nonlinear boundary value problems. The full approximation

scheme (FAS) solves nonlinear problems on fine and coarse grids. To smooth the nonlinear problem a suitable nonlinear solver is needed and since a matrix-free implementation is desirable, this form of implementation of the smoother should be plausible. For this purpose, the nonlinear additive Schwarz method (NASM) seems to be an appropriate choice. NASM converges with the same rate as linear iterations applied to the linearised equation. In addition, it is inherently parallel and proper to be implemented in matrix-free form. We combine FAS and NASM to obtain hybrid NASM/FAS. The FAS solves the nonlinear problem and the NASM is the smoother of the nonlinear boundary value problem in local subdomains on each level of the multigrid method. Within the NASM, Jacobian-Free Newton Krylov method is used as a solver on each subdomain. We consider different nonlinear equations in 3D space as test problems. We investigated several parameters of the methods to have a better understanding of influence of the parameters on the efficiency of the method and its convergence rate.

Mehdi Baba Mehdi

Forschungszentrum Jülich GmbH
m.baba.mehdi@fz-juelich.de

PP1

Two New Strategies for Stock Price Prediction Based on Machine Learning Algorithms

Stock price prediction is a complicated and interesting task. Noisy trends make stock pricing sensitive and complicated while the economical motivation behind, keeps it interesting for researchers and investors. A significant complication is the more impact of later data in prediction. In computer science, incremental learning is a class of machine learning algorithms in which input data is continuously used to extend the existing model's knowledge. These algorithms help to improve the predictions and analysis in two ways: 1) there is no need to keep all the history and 2) the old and least effective data will be easily deleted. In this paper we are to outline two novel incremental ideas for stock pricing. We also test each of our suggested algorithms for predicting the price of 6 stocks from different sectors. To show the efficiency of our proposed algorithm, we compare the predicted prices with real values and also perform a back-test to verify that the annual returns based on real data and predicted price are almost the same.

Negin Bagherpour

University of Tehran
University of Tehran
negin.bagherpour@guest.ut.ac.ir

PP1

A Relaxed Low-Rank Correction for Schur Complement Preconditioners

We describe a multiplicative low-rank correction scheme for pressure Schur complement preconditioners to accelerate the iterative solution of the linearized Navier-Stokes equations. The update method is based on a low-rank approximation to the error between the identity and the preconditioned Schur complement. We introduce a relaxation parameter that can improve the updated preconditioner. The scheme is tested for a model for buoyancy-driven fluid flows described by the Boussinesq approximation which combines the Navier-Stokes equations enhanced with a Coriolis term and a temperature advection-diffusion equation. Numerical results on a cube and a shell geometry

illustrate the action of the low-rank correction on spectra of preconditioned Schur complements using well-known preconditioning techniques and demonstrate that the update technique can accelerate the convergence of iterative solvers.

Rebekka S. Beddig

Hamburg University of Technology (TUHH)
rebekka.beddig@tuhh.de

Sabine Le Borne

Hamburg University of Technology
Department of Mathematics
leborne@tuhh.de

Jörn Behrens

Dept. of Mathematics, Universität Hamburg
joern.behrens@uni-hamburg.de

PP1

Supervised Learning for Kinetic Consensus Control

The modelling of multi agent systems (MAS) can be seen as the result of a suitable combination of within population interactions and external influences. We address how to successfully condition high dimensional MAS towards a designed cooperative goal (e.g. consensus) by means of dynamic optimization. The problem reads as the minimization of a cost functional subject to individual-based dynamics; thus, its solution easily becomes unfeasible as the number of agents grows. A natural way of circumventing this difficulty is by passing from a microscopic viewpoint to a macroscopic one, that is from an agent-to-agent description of the trajectories to the evolution of the system represented as a density in space and time. Although mean field optimal control problems are designed to be independent of the number of agents, they are feasible to solve only for moderate intrinsic dimensions of the agents' space. For this reason, we propose a procedure for approaching the solution from suboptimality by means of a Boltzmann scheme. We consider the quasi-invariant limit of binary interactions as approximation of the mean field PDE governing the dynamics of the agents distribution. This considerably tackles down the numerical complexity of the original problem, which is now reduced to a collection of many 2-agents sub-systems. The need for an efficient solver of the binary OCP motivates the use of Neural Networks.

Giacomo Albi

Università degli Studi di Verona
giacomo.albi@univr.it

Sara Bicego, Dante Kalise

Imperial College London
s.bicego21@imperial.ac.uk, d.kalise-balza@imperial.ac.uk

PP1

A State Redistribution Algorithm for Moving Geometries

State redistribution (SRD) is a modern technique for stabilizing cut cells that result from finite-volume embedded boundary methods. SRD has been successfully applied to a variety of compressible and incompressible flow problems with static geometries. In this work, we present a novel extension of the state redistribution algorithm to moving geometries. We demonstrate the effectiveness of the extended

algorithm by solving the variable density, incompressible Navier-Stokes equations using a finite-volume projection method where moving geometries are represented with a cut cell / embedded boundary approach. The resulting algorithm is shown to be conservative and validated using canonical two- and three-dimensional example problems.

Matthew Blomquist
University of California Merced
mblomquist@ucmerced.edu

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

Jordan Musser
National Energy Technology Laboratory (NETL)
jordan.musser@netl.doe.gov

Deepak Rangarajan
National Energy Technology Laboratory
deepak.rangarajan@netl.doe.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

PP1

A Machine Learning Minimal Residual Method for Solving Quantities of Interest of Parametric PDEs

We propose to combine neural networks (NNs) and finite element schemes over coarse meshes to approximate quantities of interest (QoI) of parametric PDEs. Within the context of residual minimization, we add a weight to the inner product of the test space. The inner-product weight is computed as the output of an artificial neural network and it is trained with a supervised learning procedure. We use a loss function to train the neural network that compares the QoI of the coarse-mesh discrete solution with the QoI of a precomputed high-precision solution. Then, the training process minimize the loss function to make the discrete solution a better approximation (on the QoI) for any PDE parameter.

Ignacio Brevis
School of Mathematical Sciences
University of Nottingham
ignacio.brevis@nottingham.ac.uk

Ignacio Muga
Instituto de Matemáticas
Pontificia Universidad Católica de Valparaíso
ignacio.muga@pucv.cl

David Pardo
Basque Center for Applied Mathematics
University of the Basque Country (UPV/EHU)
dzubiaur@gmail.com

Oscar Rodriguez
Basque Center for Applied Mathematics
orodriguez@bcamath.org

Kristoffer van Der Zee
School of Mathematical Sciences
University of Nottingham

kg.vanderzee@nottingham.ac.uk

PP1

Co-Design of Discretizations and Matrix-Free Solvers for Large-Scale Stokes Problems with Extreme Viscosity Variations

The simulation of real-world Earth mantle convection problems requires extreme spatial resolution. This motivates massively parallel, matrix-free methods for solving the corresponding linear systems. The optimal efficiency of such methods is only attained through careful co-design of the discretization and the linear solver. In particular, we are concerned with the choice of discretizations and solvers for Stokes problems that exhibit extreme viscosity variations, leading to severely ill-conditioned linear systems. It is unclear which discretization can cope best with this, as different discretizations produce discrete problems of various sizes, nonzero patterns, and spectral properties. We consider a new enriched Galerkin scheme that promises to be a competitive choice compared to standard discretizations. On the solver side, we employ a combination of geometric multigrid and Krylov methods, e.g., the Golub-Kahan solver. Eventually, we analyze the discretization- and solver-composite with respect to its numerical and computational efficiency on hybrid tetrahedral grids.

Fabian Böhm
Friedrich-Alexander-Universität Erlangen-Nürnberg
Fabian.boehm@fau.de

Nils Kohl
University of Erlangen-Nuremberg (FAU)
nils.kohl@fau.de

Ulrich Ruede
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)
ulrich.ruede@fau.de

PP1

Efficient Upwind Finite-Difference Schemes For Wave Equations On Overset Grids

We describe an algorithm to easily and efficiently incorporate upwinding into finite-difference schemes for solving wave equations in second-order form and apply this scheme to solve problems on complex geometry using overset grids. Upwinding can be added to an existing discretization, such as a centered and dissipation-free scheme, as a modular corrector stage. This new upwind predictor-corrector scheme significantly improves the run-time performance compared to our original formulation, with typical speedups of factors of ten or more. As with the original upwind formulation, theory and numerical results show that the new algorithm remains robust and stable even for the difficult cases of overset grids with thin boundary fitted grids, where non-dissipative schemes are generally unstable. Numerical results simulating Maxwells equations in second-order form to second- and fourth-order accuracy are used to assess the run-time performance of the new scheme.

Allison Carson
Rensselaer Polytechnic Institute
carsoa@rpi.edu

Jordan Angel
NASA Ames Research Center
jordan.b.angel@nasa.gov

Jeffrey W. Banks, William Henshaw
Rensselaer Polytechnic Institute
banksj3@rpi.edu, henshw@rpi.edu

PP1

preCICE: A Sustainable and User-Friendly Coupling Ecosystem for Partitioned Simulations

What if we could build complex and efficient multi-physics simulations by easily plugging together the tools we already have at hand? Prototyping such complex and efficient simulations has been possible since longer with the free/open-source coupling library preCICE, which provides sophisticated numerical coupling methods and scalability on ten thousands of compute cores [a href="https://doi.org/10.1016/j.compfluid.2016.04.003" data-bbox="100 270 480 285">Bunger et al., preCICE: A fully parallel library for multi-physics surface coupling, *Comp&Fluids*, 2016; a]. Today, it is significantly easier to design partitioned simulations by selecting from a list of ready-to-use integrations with widely-used simulation codes, following a unified and actively maintained online documentation, and connecting with an expanding community of users and contributors, counting more than 100 research groups worldwide. This growing ecosystem of subprojects creates challenges in structuring and automating the development, documentation, testing, and continuous integration from unit to system level. This poster will present the challenges and lessons learned in growing preCICE from an as-is coupling library to a sustainable, batteries-included ecosystem [a href="https://doi.org/10.12688/openreseurope.14445.1" data-bbox="100 460 480 475">Chourdakis et al., preCICE v2: A sustainable and user-friendly coupling library [version 1; peer review: 2 approved] *Open Res Europe* 2022, 2:51; a].

Gerasimos Chourdakis

Technical University of Munich
gerasimos.chourdakis@tum.de

Benjamin Uekermann
University of Stuttgart
benjamin.uekermann@ipvs.uni-stuttgart.de

PP1

Material Reconstruction in Epma Using the PN Model and Adjoint Gradient Based Minimization

Electron Probe Microanalysis (EPMA) determines the chemical composition of material samples based on intensity measurements of x-rays induced by electron beams. The reconstruction of the chemical composition poses an inverse problem. The spatial resolution of EPMA is currently restricted by intensity models that assume a material with either homogeneous or depth-layered structure. We propose a more sophisticated intensity model based on the spherical harmonic (P_N) approximation of the Boltzmann equation for electron transport that allows heterogeneous structure. Current reconstruction methods invert intensity measurements obtained with a single beam setup (energy and position). For higher resolution, the reconstruction method must consider intensities obtained from multiple beam setups. We propose a reconstruction method based on gradient-based minimization of the discrepancy between multiple modeled and measured intensities. We present a numerical staggered grid method for the P_N -approximation and differentiate our intensity model using a combination of the continuous adjoint method and adjoint algorithmic differentiation. The combination achieves an efficient gradient computation of the discrepancy function and thus an

efficient reconstruction with exchangeability in the material parametrization. Using reconstruction examples we motivate further research on e.g. tailored discrepancy functions, material parametrizations or uncertainty quantification.

Tamme Claus

ACoM, Research Lab for Applied and Computational Mathematics
RWTH Aachen University
claus@acom.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

Silvia Richter
GFE, RWTH Aachen University
richter@gfe.rwth-aachen.de

PP1

Travelling Waves in a Volume-Filling Model of Cell Invasion into Extracellular Matrix

Many reaction-diffusion models produce travelling wave solutions that can be interpreted as waves of invasion in biological scenarios such as wound healing or tumour growth. These partial differential equation models have since been adapted to describe the interactions between cells and extracellular matrix (ECM), using a variety of different underlying assumptions. In this work, we derive a system of population-level reaction-diffusion equations with cross-dependent diffusion terms by coarse-graining an individual-based description of cell motility and proliferation alongside ECM degradation, taking into account the impact of both cell and ECM volume-filling effects on cell motility and proliferation. We analyse the resulting travelling wave solutions both numerically and analytically across various parameter regimes using phase plane analysis, asymptotics and perturbation techniques. Subsequently, we perform a systematic comparison between the population-level behaviours observed in this model and those predicted by simpler models in the literature, which do not take into account volume-filling effects in the same way. Our study justifies the use of some of these simpler, more analytically tractable models to reproduce the qualitative properties of the solutions in the limiting cases, as well as revealing some interesting properties caused by the introduction of cell and ECM volume-filling effects, where standard model simplifications might not be appropriate.

Rebecca Crossley

University of Oxford
rebecca.crossley@wolfson.ox.ac.uk

PP1

Sequential Learning for Fiber Coating Dynamics via Pod and Neural ODEs

Thin film liquids flowing down a vertical fiber exhibit a family of complex and interesting nonlinear dynamics driven by Rayleigh-Plateau instability, including moving droplets and irregular wavy patterns. These films play an important role in many engineering applications, such as heat and mass exchangers, desalination, and vapor/particle capture systems. This type of flow has been modeled by degenerate fourth-order parabolic PDEs for the free surface. We are interested in using Neural ODEs (NODEs) to learn the long-term behavior of the fiber coating dynamics. It has been shown that Heavy Ball NODEs (HBNODEs) have

advantages over NODEs for sequential learning in terms of the number of function evaluations needed and ability to learn long-term dependencies [Baker, et al. Learning POD of Complex Dynamics Using Heavy-ball Neural ODEs]. In this poster, we use the energy and entropy structure of the PDE to design a POD-based HBNODE framework that better captures the morphological changes in the dynamics.

Benjamin C. Daniel, Hangjie Ji
North Carolina State University
bcdanie2@ncsu.edu, hangjie_ji@ncsu.edu

PP1

Extreme Learning Machines for Variance-Based Global Sensitivity Analysis

Variance-based global sensitivity analysis (GSA) can provide a wealth of information when applied to complex models. A well-known Achilles' heel of this approach is its computational cost which often renders it unfeasible in practice. An appealing alternative is to analyze instead the sensitivity of a surrogate model with the goal of lowering computational costs while maintaining sufficient accuracy. Should a surrogate be "simple" enough to be amenable to the analytical calculations of its Sobol' indices, the cost of GSA is essentially reduced to the construction of the surrogate. We propose a new class of sparse weight Extreme Learning Machines (SW-ELMs) which, when considered as surrogates in the context of GSA, admit analytical formulas for their Sobol' indices and, unlike the standard ELMs, yield accurate approximations of these indices. The effectiveness of this approach is illustrated through both traditional benchmarks in the field and on a chemical reaction network.

John Darges
North Carolina State University
Department of Mathematics
jedarges@ncsu.edu

PP1

Exponential Methods for the Anisotropic Diffusion Problem

The diffusion of Galactic cosmic rays (CRs) depend on the properties of the Galactic magnetic field. The ordered structure of the large-scale magnetic field results in preferential diffusion, of CRs, parallel to the ordered field lines, i.e. the diffusion is anisotropic. We propose to solve the diffusion equation, characterised by a diffusion tensor, using methods that directly compute the matrix exponential. Not only do these methods converge rather quickly, one can also obtain highly accurate solutions whilst taking substantially larger step sizes than what is possible for implicit (and explicit) schemes. We propose to use the methods of Leja polynomial interpolation and Krylov-based KIOPS to compute the matrix exponential. We also investigate the performance of exponential integrators, in this context, by extracting a constant coefficient from the diffusion tensor, where the stiff constant coefficient term is solved exactly and the remainder is treated using high-order explicit methods. We show in a number of test problems that these proposed methods result in significant amount of computational savings (up to 2 orders of magnitude) over the traditionally used Crank-Nicolson integrator. These methods are expected to further boost the computational performance with the inclusion of other terms in the CR transport equation such as advection, reacceleration, en-

ergy losses, etc. (i.e. when the equation is extremely stiff).

Pranab J. Deka, Lukas Einkemmer, Ralf Kissmann
University of Innsbruck
pranab.deka@uibk.ac.at, lukas.einkemmer@uibk.ac.at,
ralf.kissmann@uibk.ac.at

PP1

Eulerian Algorithms for Computing the Forward Finite Time Lyapunov Exponent Without Finite Difference Upon the Flow Map

Effective Eulerian algorithms are introduced for the computation of the forward finite time Lyapunov exponent (FTLE) of smooth flow fields. The advantages of the proposed algorithms mainly manifest in two aspects. First, previous Eulerian approaches for computing the FTLE field are improved so that the Jacobian of the flow map can be obtained by directly solving a corresponding system of partial differential equations, rather than by implementing certain finite difference upon the flow map, which can significantly improve the accuracy of the numerical solution especially near the FTLE ridges. Second, in the proposed algorithms, all computations are done on the fly, that is, all required partial differential equations are solved forward in time, which is practically more natural. The new algorithms still maintain the optimal computational complexity as well as the second order accuracy. Numerical examples demonstrate the effectiveness of the proposed algorithms.

Shaoyong Deng
University of North Carolina at Charlotte
shaodeng@uncc.edu

Guoqiao You
Nanjing Audit University
Nanjing, Jiangsu 211815, P.R. China
magqyou@nau.edu.cn

Changfeng Xue
Yancheng Institute of Technology
Yancheng, Jiangsu 224051, P.R. China
cfxue@ycit.edu.cn

PP1

Models Correction Based on Sparse Identification and Data Assimilation

Many models assumed to be able to predict the response of structural systems fail to efficiently accomplish that purpose for two main reasons. First, some structures in operation undergo localized damage that degrades their mechanical performances. To reflect this local loss of performance, the stiffness matrix associated with the structure should be locally corrected. Second, the nominal model is sometimes too coarse grained to reflect all structural details, and consequently, the predictions are expected to deviate from the measurements. In that case, the entire domain needs to be repaired; therefore, the entire structure-stiffness matrix should be corrected. In the present work, we propose a methodology to correct locally or enrich globally the models from collected data. The proposed techniques consist in the first case of an L1-minimization procedure that, with the support of data, aims at the same time period to detect the damaged zone in the structure and to predict the correct solution. For the global enrichment, instead, the methodology consists of an L2-minimization procedure with the support of measurements. The benefits and potential of such techniques are illustrated on four

different problems, showing the large generality and adaptability of the methodology. Di Lorenzo, D.; Champaney, V.; Geroso, C.; Cueto, E.; Chinesta, F. Data Completion, Model Correction and Enrichment Based on Sparse Identification and Data Assimilation. *Appl. Sci.* 2022, 12, 7458.

Victor Champaney, Daniele Di Lorenzo
Arts et Métiers Institute of Technology
victor.champaney@ensam.eu,
daniele.di_lorenzo@ensam.eu

Geroso Claudia
Instituto Tecnológico de Santo Domingo
claudia.geroso@intec.edu.do

Elias Cueto
I3A, University of Zaragoza, Spain
ecueto@unizar.es

Francisco Chinesta
ENSAM Institute of Technology
francisco.chinesta@ensam.eu

PP1

Scikit-Shape: Python Toolbox for Shape and Image Analysis

Many tasks in image processing, e.g. segmentation, surface reconstruction, are naturally expressed as energy minimization problems, in which the free variables are shapes, curves in 2d or surfaces in 3d. We typically express such problems as energies with data (or target) mismatch and geometric regularization components, to be minimized algorithmically to attain the optimal shape. To solve such problems, we have implemented a suite comprising various building blocks of such problems and algorithms to perform the minimization, including geometric regularization, statistical shape priors, adaptive geometric discretization, and fast Newton-type minimization schemes. Moreover, we have developed crucial shape analysis algorithms for statistical analysis and evaluation of the shapes computed, based on elastic shape distance framework. Our main applications are image and data analysis problems, but the infrastructure is quite general, and can be used for problems in other fields as well. All our algorithms are implemented in Python, leveraging on the NumPy/SciPy ecosystem, making them as easy to use as Matlab, also compatible with existing Python tools. Our algorithms is freely available as an open source package for the research community at: <http://scikit-shape.org>

Gunay Dogan
National Institute of Standards and Technology
gunay.dogan@nist.gov

PP1

Optimization of Back-Contact Metal-Semiconductor-Metal Perovskite Solar Cells

Currently, hybrid organic-inorganic halide perovskites are the most promising low-cost materials for solar cell fabrication with decent photovoltaic properties compared to the widely preferred crystalline silicon. Conventionally, perovskite solar cells (PSCs) have multilayer structure, in which the perovskite layer is sandwiched between two charge-transporting layers. This structure demands a high level of precision and accuracy in fabrication to achieve a reasonable level of performance. Herein we present a

theoretical analysis of back-contact metal-perovskite-metal PSC, which intends to simplify the fabrication process. The electrical and optical properties of the two-dimensional model of devices are investigated numerically in accordance with the experimental results reported by other groups. The geometrical parameters of devices such as the thickness of the perovskite film, the distance between metallic contacts, and the width of metallic contacts are varied in search of the device structure that yields the best photovoltaic performance. An unexpected drop in the device current is obtained when the width of metallic contacts is too small. The analysis of electric field distribution in the devices revealed the underlying reasons for this phenomenon. Based on the computer simulation experiments, it is found that a theoretical power conversion efficiency of 6.5% could be achieved for observed PSCs structure with the methylammonium lead iodide perovskite layer.

Ilyas T. Dossayev, Assylan Akhanuly
Nazarbayev University
ilyas.dossayev@nu.edu.kz, assylan.akhanuly@nu.edu.kz

Erik Shalenov
Satbayev University
erik.shalenov@nu.edu.kz

Askhat Jumabekov
Nazarbayev University
askhat.jumabekov@nu.edu.kz

PP1

Mathematical Modelling of the Transmission Dynamics of Monkey Pox Using the Differential Transformation Method

This work is designed to formulate and analyze a mathematical model for dynamics of Monkey Pox disease with optimal control measures under some assumption made. Necessary conditions for the existence and stability of the disease steady states are derived. The Monkey Pox disease reproduction number is determined. Based on construction of suitable Lyapunov functional, disease-free equilibrium point of the formulated model is shown to be globally asymptotically stable when the crime reproduction number is below unity, while a unique disease present equilibrium is proved to be globally asymptotically stable whenever the disease reproduction number exceeds unity. The method of differential transformation is employed to compute an approximation to the solution of the non-linear systems of differential equations for the transmission dynamics of the disease model. The differential transformation method is a semi-analytic numerical method or technique, which depends on Taylor series Sensitivity analysis is carried out to determine the relative importance of model parameters in the disease spread. Furthermore, optimal control theory is employed to assess the impact of two time-dependent optimal control strategies, including public enlightenment and vaccination.

Helen Edogbanya
Federal University Lokoja, Nigeria
helen.edogbanya@fulokoja.edu.ng

Zamurat A. Adegboye
NIGERIAN DEFENCE ACADEMY
,KADUNA,NIGERIA

zamurat.adeboye@imsp-uac.org

PP1

A New Free-Surface Incompressible Navier-Stokes Spectral Element Model for Water Waves

We present a new high-order accurate free-surface incompressible Navier-stokes model discretized using the spectral element method for the simulation of water waves suitable for offshore engineering applications, e.g. in renewable sectors such as offshore wind and wave energy. For such applications, the development of efficient models for wave propagation and wave-structure interactions is of key interest. We present details of the derivation of the model in two space dimensions and first results from numerical experiments to highlight the capabilities of the solver.

Allan P. Engsig-Karup, Anders D. Melander, Jesper J. Lisegaard, Max Bitsch
Technical University of Denmark
apek@dtu.dk, adame@dtu.dk, s184019@student.dtu.dk, maxbitsch@gmail.com

PP1

A Fourier Series-Based Approach for Estimating Time-Varying Parameters in Differential Equations

Many real-world systems modeled using differential equations involve unknown parameters. While many estimation approaches focus on constant parameters, some unobservable system parameters may vary with time with unknown evolution models. In this work, we propose a novel approximation method inspired by the Fourier series to estimate time-varying parameters in deterministic dynamical systems modeled using differential equations. We demonstrate the capability of this approach in estimating periodic parameters, in cases where the period is both known and unknown, with several computed examples.

Molly Folino, Anna Fitzpatrick, Andrea Arnold
Worcester Polytechnic Institute
mjfolino@wpi.edu, akfitzpatrick@wpi.edu, anarnold@wpi.edu

PP1

A Parallel Approach to Approximate Bayesian Inference

Despite the increasing interest in Bayesian computing, large-scale inference tasks continue to pose a computational challenge that often requires a trade-off between accuracy and computation time. We present a highly scalable approach for performing spatial-temporal Bayesian modeling based on the methodology of integrated nested Laplace approximations (INLA), combining solution strategies from the field of high-performance computing with state-of-the-art statistical learning techniques. We leverage different parallelization strategies to fully utilize the power of today's distributed compute architectures and introduce highly optimized sparse linear algebra routines to handle the computational kernel operations, enabling fast and reliable model predictions.

Lisa Gaedke-Merzhäuser
Università della Svizzera italiana
lisa.gaedke.merzhaeuser@usi.ch

Haavard Rue

King Abdullah University of Science and Technology
haavard.rue@kaust.edu.sa

Olaf Schenk
Università della Svizzera italiana
olaf.schenk@usi.ch

PP1

p4est – State of the Software

The p4est software library implements fast algorithms for large-scale distributed adaptive mesh refinement. It employs a distributed forest-of-octrees data structure and has been demonstrated to scale to 10^6 MPI processes. The management of the software has always been exceedingly conservative, prioritizing correctness and simplicity over feature extension and the evolution of C standards. In recent times, however, it has been increasingly relevant to allow for more flexible, cross-platform build options, last but not least to address demands by the scientific user base. Thus, we have added an unofficial but functional CMake build system, support for non-Linux OSes, various CI runners, and generally hardened the software for a variety of compute environments. On the algorithmic side, we have added several performance improvements of core algorithms, such as forest creation, ghost layer construction and repartitioning. Furthermore, we have added to the set of non-process-local search and geometric matching algorithms, for example to locate (ocean gauge) points and integrate over (satellite view) rays modeling actual physical measurements. In addition to mesh-based numerical simulation, new algorithmic features expand applications to ray casting and radiosity. This poster gives a general overview and some behind-the-scenes bits and pieces on the current state of the software.

Gautam Ghosh
Universität Bonn
ghosh@ins.uni-bonn.de

Tim Griesbach
Universität Bonn, Germany
Institute for Numerical Simulation
tim.griesbach@uni-bonn.de

Hannes Brandt
University of Bonn
brandt@ins.uni-bonn.de

Michael Hirsch
Boston University
mhirsch@bu.edu

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

PP1

Discovering Shared Causal Drivers from Highly-Corrupted Time Series

Many experimental measurements arise from subsystems that implicitly share a common driving signal. In systems biology, examples include targets of transcription factors, regulation of circadian rhythms, and descending control in animal nervous systems. Previous theoretical work on attractor reconstruction suggests that partial information present in each subsystem can reveal subregimes within

the unseen driving signal, introducing the possibility of fully recovering the driver given sufficient data. Here, we show that these results motivate a new type of manifold learning algorithm based on persistent homology. We show empirical results demonstrating the ability to reliably detect a common driving signal from complex time series, even in the presence of many noisy, incomplete measurements. We show applications to real-world datasets including species abundances in microbial ecosystems and simultaneous recordings of neurons.

William Gilpin

The University of Texas at Austin
wgilpin@utexas.edu

PP1

An Integral Equation Method for a Broad Class of Elliptic PDEs on Surfaces

Elliptic PDEs on a surface embedded in three dimensions occur frequently in computer graphics, including in shape analysis and surface interpolation problems. They also occur in many areas of physics, including molecular dynamics and fluid dynamics. They are particularly useful in magnetostatics and plasma problems, where they are used to construct the Hodge decompositions of tangential vector fields. In this poster, we present the first published method for converting a *broad class of elliptic PDEs* on a *general smooth surface* Γ into second kind Fredholm integral equations on Γ . Doing so ensures the equations are well conditioned and makes it possible to construct high-order numerical solvers. The structure of the integral equation also enables the use of hierarchical compression schemes, which can be used to solve the resulting linear system in a comparable speed to non-integral equation methods. To derive the integral equation, we extend the known method for the Laplace-Beltrami problem [Kropinski and Nigam, 2014] on a sphere to a broader class of equations on general smooth surfaces. The new method uses the observation that the Greens function of a corresponding PDE in \mathbb{R}^2 is a parametrix (an approximate Greens function) for the PDE on Γ to derive an integral equation form of the PDE.

Tristan Goodwill, Mike O'Neil

New York University
tristan.goodwill1@gmail.com, oneil@cims.nyu.edu

PP1

Geometric Shape Optimization for Reflexive Optics

Geometrical optics describes the propagation of light in optical media. In many physical situations, light can undergo successive reflections on objects like mirrors. In telescopes, designing of mirror shapes aims at providing an enlarged and good-quality image of an observed light source like a star or Earth. Traditionally, optical design makes use of parametric shape optimization to compute the design of mirrors in optical instruments. However attractive mathematical techniques like shape differentiation in the spirit of Hadamard's boundary variation partially alleviates the use of shape parametrization and potentially provides better designs for optical criteria. Numerous mathematical models coexist for the description of light rays. Here we choose a kinetic Liouville PDE supplemented with reflexive boundary conditions which offer a unified framework to deal with reflections and geometric shape optimization. The aim of this poster is to present a new work on the boundary variation method to the aforementioned PDE model. We will highlight the mathematical difficulties

which arise with this equation in our shape optimization application such as regularity issues, non local boundary conditions, lack of variational formulation and numerical computations in the phase space. Although covering many possible applications, our research is motivated by the design of optical telescope and some numerical illustrations will be presented in this context.

Luca Gorini

Centre de Mathématiques Appliquées, CMAP,
École Polytechnique de Paris, 91118 Palaiseau, France
luca.gorini@polytechnique.edu

Samuel Amstutz

Centre de Mathématiques Appliquées, CMAP,
École Polytechnique de Paris, 91128 Palaiseau, France
samuel.amstutz@polytechnique.fr

PP1

Line-Search Methods for Unconstrained Optimization with Inexactness Arising from Reduced Order Models

I discuss line-search algorithms for the solution of smooth unconstrained optimization problems that allow the use of approximate objective function and gradient information. These algorithms are motivated by the need to rigorously incorporate reduced order models (ROMs) into the solution of large-scale optimization problems governed by partial differential equations. Problems of this nature are common in many science and engineering applications. The developed algorithms are error aware and use on the fly updates to error tolerances in order to reduce the accuracy requirements on the ROM approximations. The considered algorithms require no explicit information about the underlying true model and operate entirely using error bounds on the objective and its gradient. These algorithms are implementable and provide convergence guarantees subject to some reasonable assumptions on the underlying optimization problem. Numerical results show that the proposed line-search methods, combined with ROMs, converge to local minima of the original optimization problem at a fraction of the computational cost required by traditional Newton CG algorithms.

Dane Grundvig

Rice University
dsg12@rice.edu

Matthias Heinkenschloss

Department of Computational and Applied Mathematics
Rice University
heinken@rice.edu

PP1

Accelerating Parallel-in-Time Methods

To reduce the computation time of numerically solving partial differential equations, time-parallel methods have received increasing interest in recent years. Most parallel-in-time methods are iterative, and their parallel efficiency depends on fast convergence of the iterative scheme. In this poster we will discuss approaches, e.g., using machine learning, to improve the performance of parallel-in-time methods.

Sebastian Götschel

Hamburg University of Technology

sebastian.goetschel@tuhh.de

PP1

Representation Theory Based Algorithm to Compute Boltzmann's Bilinear Collision Operator in the Irreducible Spectral Burnett Ansatz Efficiently

Numerically solving the Boltzmann equation is computationally expensive in part due to the number of variables the distribution function depends upon. Another contributor to the complexity of the Boltzmann Equation is the the quadratic collision operator describing changes in the distribution function due to colliding particle pairs. Solving it as efficiently as possible has been a topic of recent research, e.g. [Cai & Torrilhon, "On the Holway-Weiss debate: Convergence of the Grad-moment-expansion in kinetic gas theory", 2019], (Wang & Cai, "Approximation of the Boltzmann collision operator based on hermite spectral method", 2019), (Cai & Fan & Wang, "Burnett spectral method for the spatially homogeneous Boltzmann equation", 2020)]. In this paper we exploit results from representation theory to find a very efficient algorithm both in terms of memory and computational time for the evaluation of the quadratic collision operator. With this novel approach we are also able to provide a meaningful interpretation of its structure.

Andrea Hanke, Manuel Torrilhon
RWTH Aachen
hanke@aices.rwth-aachen.de, mt@acom.rwth-aachen.de

PP1

Pore-Scale Simulation of Hydrogen Transport in Porous Media

Underground Hydrogen Storage (UHS) in porous formations is a promising technology for large-scale (TWh) energy storage. To ensure the efficiency of the storage operation, multiscale modeling and simulation strategies are essential, in which the micro-scale physics are studied to derive input parameters for the continuum-scale dynamic models. The present study develops a dynamic pore-network modeling (D-PNM) approach to simulate the immiscible two-phase flow of hydrogen and water through representative digital network models of different porous structures. The model input parameters are based on the experimentally obtained static and dynamic wettability analyses as presented in the literature. As for the rock, digital networks are constructed based on 3D X-ray images of porous samples. The topology of the pore space geometry is translated into a representative pore-network model. To preserve the simulation stability, the developed D-PNM solves the transient multi-phase Stokes equations fully implicitly, for pressure and phase volume concentration. Through several test cases, we analyze the transport characteristics of hydrogen/water interface, especially the fingering and spreading physics. These results shed new lights on how a representative continuum-scale model should be created to study the process at field scale.

Leila Hashemi
Technical University of Delft
Department of Applied Mathematics
l.hashemi@tudelft.nl

Rainer Helmig
IWS, University of Stuttgart, Germany
Dept. of Hydromechanics
rainer.helmig@iws.uni-stuttgart.de

Cornelis Vuik
Technical University of Delft
Delft Institute of Applied Mathematics
c.vuik@tudelft.nl

Hadi Hajibeygi
Delft University of Technology
The Netherlands
h.hajibeygi@tudelft.nl

PP1

Iterative Regularization Schema for Volume Reconstruction in Single Particle Analysis

In this contribution, we study discrete inverse problems $Ax = b$ arising in cryo-electron microscopy single-particle analysis (SPA). SPA aims to reconstruct the three-dimensional high-resolution structure of biological macromolecules from a set of its two-dimensional projections. Structural information is vital to understanding biological life processes and can be used, for example, in drug discovery. Inverse problems arising in SPA are challenging especially due to the large dimensionality of the measured data, suffering from the presence of extreme amounts of noise and missing information. Computationally efficient methods and strong regularization, therefore, need to be applied in order to obtain reliable approximation of a molecular structure. Here, we focus on an inverse problem of volume reconstruction from a given set of measured particle projections with estimated viewing angles. We start with a brief overview of the problem formulation and present sample realistic data. Further, we focus on a discretization schema that yields a highly structured large-scale approximation problem. An inner-outer iterative solver suitable for the solution of the studied problem will be presented and its highly-parallel GPU implementation will be discussed. Finally, we present results obtained with realistic data and address future work.

Eva Havelková
Charles University, Prague, Czech Republic.
eva.havelkova@karlin.mff.cuni.cz

PP1

Romnet: Learning PDE Dynamics from Data with Reduced Order Model Neural Networks

Data-driven modeling of dynamical systems is an active area of research. However, current techniques very often require extensive prior knowledge of the governing equations, or are limited to linear or first-order equations. In this work, we propose a neural net-based approach for learning the dynamics of systems described by Partial Differential Equations (PDEs), without requiring any prior knowledge of the system. Specifically, we propose a novel deep learning framework, called Reduced Order Model Network (ROMNet), that consists of three modules responsible for (i) learning a lower dimensional representation of the data, (ii) learning the dynamics and advancing the solution in the reduced latent space, and (iii) mapping the advanced solution from the latent space to the original space. We demonstrate the effectiveness of ROMNet for learning PDE dynamics on complex simulated and real-world data, showing that our model accurately learns unknown linear and nonlinear PDEs (in 2D and 3D). We compare our approach to conventional numerical schemes and find that ROMNet advances the dynamics considerably faster and more efficiently in addition to having comparable accuracy. Our results showcase the implications of deep learning models

(such as ROMNet) in learning complex PDEs and the potential to significantly enhance current numerical methods for large systems as well as improving the analysis of systems with limited prior knowledge.

A. Ali Heydari
University of California, Merced
aheydari@ucmerced.edu

PP1

Data-Driven Diffusion Coefficient Estimation in Marine Lakes

Marine lakes are located near the ocean and their geological location allows exchanges between the lake and the sea to occur. This grants marine lakes the potential to develop near-isolated marine ecosystems. We wish to quantify the mixing within the lakes by estimating the diffusion coefficient, D , of quantities such as temperature and salinity from direct measurements. We set up our system as a PDE-constrained optimization problem, typically a Helmholtz-type equation, and minimize a cost function consisting of a misfit term between data and state solution and a regularization term for D . We deploy Finite Element Method (FEM) and Newton-type algorithm for discretization and optimization, respectively. We obtain good convergence towards the optimal diffusion coefficient. We derive a method to systematically select the regularization parameter. Through empirical study, we also show the importance of accurately modeling the exchanges with the ocean. We demonstrate the capability of our method on synthetic data as well as real measurement data.

Alex Ho
University of California Merced
aho38@ucmerced.edu

Francois Blanchette, Michael Dawson
University of California, Merced
fblanchette@ucmerced.edu, mdawson@ucmerced.edu

PP1

Measurement of Gas Bubble Size in Fluid Flow

We describe a method to identify, in photographs of gas bubbles suspended in liquid flow, the location and diameter of the bubbles. Challenges include discriminating in-focus from out-of-focus bubbles; densely packed, overlapping bubbles; and deformed (non-spherical) bubbles. Recent progress in algorithm development and comparison to human analysis is presented.

Thomas Hoft, Thomas Shepard
University of St. Thomas
hoft@stthomas.edu, thomas.shepard@stthomas.edu

PP1

Reflectionless Discrete Analytic Perfectly Matched Layers for Higher-Order Finite Difference Schemes

We propose discrete-analytic Perfectly Matched Layers (PMLs) for high-order finite difference (FD) discretizations of the scalar wave equation which, unlike PDE-based PMLs, produce no numerical reflections (up to machine precision) at the PML interface. We prove that our numerical solutions decay exponentially within the PML domain as its width increases. Our approach generalizes the ideas put forth by Chern in [1] for the standard second-order FD

method, to arbitrary high-order schemes at the cost of introducing additional localized PML variables that account for the larger stencils used. The merits of the method are demonstrated on a variety of numerical examples including waveguide problems where high-order schemes are often needed to mitigate undesired numerical dispersion errors.

Vicente Hojas
Pontificia Universidad Católica de Chile
vahojas@uc.cl

Carlos Pérez
University of Twente
c.a.perezarancibia@utwente.nl

Manuel Sánchez
Pontificia Universidad Católica de Chile
manuel.sanchez@ing.puc.cl

PP1

A Massively Parallel Approach to Forecasting Electricity Prices

With the ongoing energy crisis in Europe, accurate forecasting of electricity price levels and volatility is essential to planning grid operations and protecting consumers from extreme prices. We present how massively parallel stochastic optimal power flow models can be deployed on modern many-core architectures to efficiently forecast power grid configurations in real time. Processing of stochastic weather and economic scenarios is optimized on many-core CPUs to achieve maximal throughput and minimize latency from the receipt of weather data to the output and interpretation of model results.

Timothy Holt
Università della Svizzera italiana (USI)
timothy.holt@usi.ch

Juraj Kardos
Institute of Computing
Università della Svizzera italiana
juraj.kardos@usi.ch

Olaf Schenk
Università della Svizzera italiana
olaf.schenk@usi.ch

PP1

Parallel Zolotarev-SVD for the Analysis of Rocket Combustion Data

Singular value decomposition (SVD), also known as principal component analysis (PCA), is one of the oldest and most fundamental, but also very powerful mathematical tools in data science. This poster is concerned with the application of modern high-performance SVD algorithms to problems arising in, e.g., the context of rocket combustion at the German Aerospace Center (DLR). We report on our implementation within the PyTorch- and mpi4py-based HPC-data analytics software HEAT (Helmholtz Analytics Toolkit) developed at DLR, JSC, and KIT (Gtz et al., 2020 IEEE International Conference on Big Data, pp. 276-287). The core of our SVD implementation is formed by the highly parallelizable Zolotarev-SVD algorithm proposed by Y. Nakatsukasa and R.W. Freund (2016, SIAM Review, Vol. 58, No. 3, pp. 461-493). We present numerical experiments on the respective scaling behavior.

Moreover, depending on the concrete problem type, other well-known techniques, such as randomized or incremental algorithms, are shown to allow for further reduction of the computational costs.

Fabian Hoppe
Deutsches Zentrum für Luft- und Raumfahrt e.V. - DLR
Institut für Softwaretechnologie
fabian.hoppe@dlr.de

Philipp Knechtges
German Aerospace Center (DLR)
philipp.knechtges@dlr.de

Alexander Rüttgers
Deutsches Zentrum für Luft- und Raumfahrt (DLR)
Institut für Softwaretechnologie
alexander.ruettgers@dlr.de

PP1

An Adaptive Sparse Proper Generalized Decomposition for Real-Time Structural Health Monitoring

Real-time structural health monitoring represents an added-value service for airlines. It consists in implementing an automated procedure to determine the health status of a structure at regular intervals. Recently, data-driven modelling approaches have been the focus of increasing interest due to their capacity to handle the large amount of data acquired through sensors networks. However, one of the challenges is the lack of sufficiently predictive physical models for damage identification. For this purpose, we consider the use of so-called reduced order models relying on an offline/online strategy. First, a data compression is performed during the offline (or learning) phase. Synthetic damaged structural responses generated through high-fidelity simulations and design-of-experiment samplings are collected to construct a low-dimensional approximation subspace based on the Sparse Proper Generalized Decomposition (sPGD). However, the classical sPGD fails to capture the influence of damage localization on the solution. To alleviate the just referred difficulties, our work proposes an adaptive sPGD. First, a change of variable is carried out to place all the damage areas on the same reference region, where an adapted interpolation can be done. Finally, during the online use, an optimization algorithm is employed to estimate model parameters which allow to defining the health state of the structure.

Maurine Jacot
PIMM Lab, Arts et Metiers Institute of Technologie
ENSAM
maurine.jacot@ensam.eu

Champaney Victor, Chinesta Francisco
PIMM Lab, Arts et Metiers Institute of Technologie
ENSAM
victor.champaney@ensam.eu,
francisco.chinesta@ensam.eu

Cortial Julien
Safran Tech, Department of Digital Sciences and
Technologies
julien.cortial@safrangroup.com

PP1

A Mimetic Finite Difference Based Quasi-Static Magnetohydrodynamic Solver for Force-Free Plas-

mas in Tokamak Disruptions

In this work, we propose a regularized quasi-static MHD model for force-free plasmas in tokamak disruptions and a corresponding mimetic finite difference (MFD) algorithm, which is targeted at applications such as simulating the cold vertical displacement event (VDE) of a major disruption in an ITER-like tokamak reactor. In the case of whole device modeling, we further consider the two sub-domains of the plasma region and wall region and their coupling through an interface condition. We develop a parallel, fully implicit, and scalable MFD solver based on PETSc and its DMStag structure for the discretization of the five-field quasi-static perpendicular plasma dynamics model on a 3D structured mesh. The MFD spatial discretization is coupled with a fully implicit second-order DIRK scheme. The MFD algorithm exactly preserves the divergence-free condition of the magnetic field under a generalized Ohms law. The preconditioner employed is a four-level field-split preconditioner, which is created by combining separate preconditioners for individual fields, that calls multigrid or direct solvers for sub-blocks or exact factorization on the separate fields. The numerical results confirm the divergence-free constraint is strongly satisfied and show the good performance of the fieldsplit preconditioner and overall algorithm. The simulation of ITER VDE cases over the actual plasma current diffusion time is also demonstrated.

Zakariae Jorti
Los Alamos National Laboratory
zjorti@lanl.gov

Qi Tang
Los Alamos National Lab
qtang@lanl.gov

Konstantin Lipnikov, Xianzhu Tang
Los Alamos National Laboratory
lipnikov@lanl.gov, xtang@lanl.gov

PP1

Certifying Stability in Runge-Kutta Methods Via Semidefinite Programming

In the development of new time-integration schemes, one often seeks to impose desirable stability properties (e.g., A-stability, G-stability, B-stability etc.), or at the very least, verify stability of newly constructed schemes. It has been known since the 1970's and 80's that various numerical stability conditions for Runge-Kutta or multistep methods can be formulated as feasibility problems involving semi-definite matrices. In practice, utilizing these algebraic conditions to obtain rigorous stability certificates may fail for various reasons. In this poster we present new theory for several semi-definite algebraic feasibility conditions. These theoretical observations have the practical implication of enabling the rigorous certification of several types of numerical stability through the solution of semi-definite programming (SDP) problems. We highlight the approach by using SDPs to provide rigorous certificates of numerical stability for a range of linear time stepping methods (e.g., Runge-Kutta and multistep methods), including rigorous bounds on α for the $A(\alpha)$ stability of several schemes published in the literature.

Austin T. Juhl, David Shirokoff
New Jersey Institute of Technology

aj659@njit.edu, david.shirokoff@njit.edu

PP1

Acceleration of Convergence of Fixed-Point Iterations

We propose a convergence accelerating method for fixed-point iterations that can be applied to $q + 1$ times differentiable iteration functions. In this method, we consider given q linearly independent fixed-point iteration functions to solve a nonlinear equation, each of which generates a sequence of iterates that may be either convergent or divergent. We select a linear combination function of the q iteration functions, whose first $q - 1$ derivatives are zeros at the true root. The resulting combined fixed-point iteration has order of convergence higher than or equal to q . We apply this method to obtain a quartically convergent algorithm for the square root of a positive number.

Chansoo Kim

Korea Institute of Sci. and Tech. & Univ. of Sci. and Tech.

AI/R Lab.

eau@kist.re.kr

Young Sang Choi

Healthcare AI Team, National Cancer Center

ychoi@ncc.re.kr

Byoungseon Choi

Department of Economics, Seoul National University

bschoi12@snu.ac.kr

PP1

Towards Explainable Neural Network Models of Mosquito Abundance

Vector-borne disease outbreaks are closely tied to vector abundance, which makes knowledge of population dynamics useful in preventing future outbreaks. Here we present Aedes-AI, a collection of neural network models of Aedes aegypti mosquito abundance [Kinney, A., Current, S., and Lega, J., Aedes-AI: Neural network models of mosquito abundance, PLoS Comp. Bio. (2021)]. The models are trained on synthetic data generated from a mechanistic model, in contrast to other models of mosquito abundance that rely on noisy, real world trap data for training. We assess the impact of model architecture and input data oversampling on the ability of a model to replicate the dynamics of the mechanistic model, and we show the neural networks can learn the spatiotemporal features of mosquito populations. The Aedes-AI model framework presents an opportunity to study and advance interpretability methods because the models are trained on data generated from known dynamics. Here, we focus on extracting feature importance and representing the learned latent features of the hidden layers. We conclude with a discussion on how the Aedes-AI models are appealing from a public health perspective. The models are computationally efficient compared to the mechanistic model and are designed to be user friendly. Further, the present approach to developing AI models of abundance for Aedes aegypti mosquitoes should be replicable for other potent disease vectors.

Adrienne C. Kinney

University of Arizona Applied Mathematics

akinney1@math.arizona.edu

PP1

Design and Development of Intro to HPC Program Targeting Underrepresented Undergraduates

In high-performance computing (HPC), traditionally underrepresented groups (including people of color, women, and first-generation scholars) engage at significantly lower rates than their representation in the general population. Often for these groups, exposure to introductory HPC material at the undergraduate level is needed to bridge gaps in foundational computational skills in order to increase the size and diversity of the workforce. In order to directly address this skill gap, we have created a two-pronged approach. First, we are developing a program for upper-level undergraduate students to gain these computational skills, preparing them for more advanced internships, research opportunities, and next steps in the field of high-performance computing and scalable artificial intelligence. Building off a foundational curriculum developed by Sustainable Horizons Institute and using the expertise of subject matter experts from across the national laboratory system, we designed an introductory course tailored to use national laboratory HPC resources to learn fundamental HPC concepts. Second, using the insights obtained during a listening session with professors at minority serving institutions we are working to create a more sustainable, long-term absorption of HPC and scalable AI content to enhance their existing curriculum.

Sreeranjani Ramprakash

Argonne National Laboratory, U.S.

ramprakash@anl.gov

Paige Kinsley

Argonne National Laboratory

pkinsley@anl.gov

Linda Akli, Mary Ann E. Leung

Sustainable Horizons Institute

akli@shinstitute.org, mleung@shinstitute.org

Lois Curfman-McInnes

Argonne National Laboratory

curfman@anl.gov

PP1

Added Technical Aspects of P4est: Alternative Quadrant Representation and MPI-3 Shared Memory

Parallel adaptive mesh refinement (AMR) is a key technique when simulations are required to capture time-dependent and/or multiscale features. A forest of octrees is a data structure to represent the recursive adaptive refinement of an initial, conforming coarse mesh of hexahedra. This poster presents several recent enhancements to the p4est software for forest-of-octrees AMR. The first introduces new ways of encoding quadrants as atomic objects, which vary both the in-memory binary format and the associated algorithms. We present a 128-bit AVX version and an optimized long integer format, respectively. The second enhancement exploits MPI-3 shared memory windows to eliminate redundancy of quadrant and metadata storage within each shared memory node. In conclusion, we demonstrate how different approaches to shared memory use affect performance, along with the comparison of

runtimes for various quadrant implementations and representative simulation pipelines.

Mikhail Kirilin, Carsten Burstedde
Universität Bonn
kirilin@ins.uni-bonn.de, burstedde@ins.uni-bonn.de

PP1

Exascale Ready Adaptive Mesh Refinement for Hybrid Meshes

t8code is a versatile open source library for parallel adaptive mesh refinement on hybrid meshes. [1] It is exascale-ready and capable of efficiently managing meshes with up to a trillion elements distributed on a million of cores as already shown in a peer-reviewed research paper. [2] On the top-level, t8code uses forests of trees to represent unstructured meshes with complex geometries. Space-filling curves index individual elements within a forest, which requires only minimal amounts of memory allowing for efficient and scalable algorithms of mesh management. In contrast to existing solutions, t8code has the capability to manage an arbitrary number of tetrahedra, hexahedra, prisms and pyramids within the same mesh. With this poster we want to present the first official release (v1.0) of our software and give a quick overview over its main features. Besides presenting the core algorithms of t8code, we give application scenarios on how our library integrates into major simulation frameworks for weather forecasting, climate modeling and engineering; and how they benefit from our approach to do AMR. [1] <https://github.com/DLR-AMR/t8code> [2] <https://epubs.siam.org/doi/abs/10.1137/20M1383033>

David Knapp
Universität Bonn
david.knapp@dlr.de

Johannes Holke
German Aerospace Center (DLR)
johannes.holke@dlr.de

Johannes Markert
German Aerospace Center
johannes.markert@dlr.de

PP1

Finite-Element Based Computational Methodologies for Non-Collinear Magnetism and Spin-Orbit Coupling in Real-Space Density Functional Theory

Spin-Orbit Coupling (SOC) is a relativistic effect that plays a critical role in many experimental phenomena, including magnetic anisotropy, phosphorescence, and spin-orbit torque. SOC is also vital in other active fields, such as spintronics, low-dimensional materials, and topological insulators. Extensions of pseudopotential Density Functional Theory (DFT), a widely used first principles material simulation tool, to account for non-collinear magnetism and SOC are known to predict various material properties successfully. The traditional implementations of DFT commonly employ either a plane wave (PW) basis or an atom-centered orbital (AO) basis set. PW basis sets restrict the simulation domains to periodic boundary conditions while the AO basis is not systematically convergent. Further, both these basis sets suffer from poor scalability on massively parallel computing architectures. Recently, computational methodologies for DFT calculations based on a finite-element basis, incorporated in open-source code DFT-FE, have demonstrated reduced computational pre-

factor delaying the onset of cubic scaling till system sizes of 30,000 electrons while demonstrating accuracy comparable to existing PW implementations. In this work, we will discuss the real-space formulation and an efficient, scalable finite-element-based implementation methodology of non-collinear magnetism and SOC in the framework of DFT-FE on both multinode CPU and GPU architectures.

Nikhil Kodali
Indian Institute of Science, Bangalore, India
nikhilkodali@iisc.ac.in

Phani S. Motamarri
Indian Institute of Science Bangalore
phanim@iisc.ac.in

PP1

Time Adaptive Quasi-Newton Waveform Iteration

We consider methods for coupled problems, in particular partitioned solvers for fluid structure interaction where different sub solvers are used for the fluid and solid domain. More specifically we want to create detailed simulation of a jackdaw's feather to investigate how feathers generate lift and ultimately answer the question to how birds evolved flight. Another point of interest are partitioned solvers for thermal transfer problems with two subdomains that share a boundary. Interface Quasi-Newton is one of the standard solvers used for the coupling of the two subdomains in fluid structure interaction. It has recently been combined with waveform iterations to extend interface Quasi-Newton to work with multirate time stepping, where the two sub solvers use different timesteps. We further extend the Quasi-Newton waveform iteration to the time adaptive case, where both of the sub solvers use an adaptive time stepping scheme. We also provide some insight to how one should select tolerances and termination criteria within time adaptivity, the waveform iteration and possible inner solvers for the subproblems.

Niklas Kotarsky
Centre for Mathematical Sciences at Lund University
niklas.kotarsky@math.lu.se

Philipp Birken
Lund University
philipp.birken@math.lu.se

PP1

Pattern Formation in a Nonlocal Lotka-Volterra System

We analyse the stability of the spatially-uniform, coexisting steady state of a diffusive nonlocal Lotka-Volterra (NLLV) system for two species in one spatial dimension, and include terms that model both nonlocal intraspecific and local interspecific competition. We find that the coexisting state of the system can lose stability once nonlocality is introduced. We use asymptotic and numerical analysis to find the neutral curve and also derive amplitude equations using weakly nonlinear analysis, which show that the bifurcation behaviour close to the neutral curve is consistent with numerical results. Finally, we construct the asymptotic solution in the limit of weak diffusivity, and find that the leading order periodic solution consists of disjoint regions where either one or the other species is absent.

Jia Yin Lee
University of Nottingham, U.K.

Jia.Lee@nottingham.ac.uk

John Billingham, Stephen Coombes
University of Nottingham
John.Billingham@Nottingham.ac.uk,
Stephen.Coombes@nottingham.ac.uk

PP1

Learning Sparse Approximate Inverse Preconditioners with Graph-Conditioned Variational Auto-Encoders

Preconditioning is an essential ingredient for iterative algorithms, such as the conjugate gradient method, to achieve fast and robust convergence when solving large and sparse linear systems arising from mesh-based discretizations of partial differential equations (PDEs). However, generating reliable and efficient preconditioners can be both challenging and expensive. We propose a deep learning based generative technique which models the probability distribution of the exact inverse of sparse system matrices generated from finite element discretizations of linear elliptic PDEs. Our model, based upon a graph-conditioned variational auto-encoder architecture, is capable of generating sparse approximate inverse (SPAI) preconditioners which yield competitive condition numbers compared with those obtained through brute force SPAI search methods. We are aiming to generalise this approach to allow training on coarser discretizations than those considered in our test set.

Mou Li

University of Leeds
School of Computing
scmli@leeds.ac.uk

He Wang, Peter K. Jimack
School of Computing
University of Leeds
h.e.wang@leeds.ac.uk, p.k.jimack@leeds.ac.uk

PP1

A Symmetrized Parametric Finite Element Method for Anisotropic Surface Diffusion

We deal with a long-standing problem about how to design an energy-stable numerical scheme for solving the motion of a closed curve/surface under *anisotropic surface diffusion* with a general anisotropic surface energy $\gamma(\mathbf{n})$, where \mathbf{n} is the outward unit normal vector. By introducing a novel surface energy matrix $\mathbf{Z}_k(\mathbf{n})$ which depends on the Cahn-Hoffman $\boldsymbol{\xi}$ -vector and a stabilizing function $k(\mathbf{n})$, we first reformulate the equation into a conservative form, and derive a new symmetrized variational formulation for anisotropic surface diffusion with weakly or strongly anisotropic surface energies. Then, a semi-discretization in space for the variational formulation is proposed, and its area conservation and energy dissipation properties are proved. The semi-discretization is further discretized in time by an implicit structural-preserving scheme (SP-PFEM) which can rigorously preserve the enclosed area in the fully-discrete level. Furthermore, we prove that the SP-PFEM is unconditionally energy-stable for almost any anisotropic surface energy $\gamma(\mathbf{n})$ under a simple and mild condition on $\gamma(\mathbf{n})$. For several commonly-used anisotropic surface energies, we construct $\mathbf{Z}_k(\mathbf{n})$ explicitly.

Yifei Li

Department of Mathematics, National University of Singapore
e0444158@u.nus.edu

Weizhu Bao
National University of Singapore
Department of Mathematics
matbaowz@nus.edu.sg

Wei Jiang
Wuhan University
jiangwei1007@whu.edu.cn

PP1

Data-Driven Solver Selection for Sparse Linear Matrices at Scale

Sparse linear systems sit at the core of many computational problems, and their solution strongly correlates to overall execution time. However, with the constant increase of the number of linear solver and preconditioner implementations available across a plethora of numerical libraries, choosing the most efficient combination for a given problem (in terms of time-to-solution) is a challenging task. Indeed, even selecting a numerically stable combination may seem to be an unsurmountable endeavor, especially for a novice user. This is particularly evident when observing performance at scale, where the HPC system architecture and even the number of cores used can make a drastic difference. In this work, we compare previous machine learning approaches to the solver-preconditioner selection problem and develop a performance model for a selection of Krylov solver implementations and preconditioners in the PETSc framework over the SuiteSparse matrix collection. We then use the developed model to address the selection of an optimal solver for a given input matrix, showing that the current approach performs better than the black-box approach on a broad range of systems.

Hayden Liu Weng, Felix Dietrich
Technical University of Munich
h.liu@tum.de, felix.dietrich@tum.de

Hans Joachim Bungartz
Technical University Munich
bungartz@tum.de

PP1

Gauss' Law Preserving Methods for the Multi-Fluid Plasma Model

Multi-fluid plasma models are useful for representing a broad range of physics. These models use a fluid representation for electrons as well as other charged and neutral species. The models also include the full set of Maxwell's equations. A numerical method for these models is presented which satisfies Gauss' law. A compatible exact sequence spatial discretization is used. A nodal discontinuous Galerkin (H-Grad) basis is used to discretize the fluids, and a Raviart-Thomas (H-Div) basis and Nédélec (H-Curl) basis are used to discretize Maxwell's equations. Temporally this method is discretized with an implicit-explicit (IMEX) Runge Kutta scheme. This time integration method allows implicit evolution of fast time scales in a stable manner, while also efficiently explicitly solving slower time scales. This discretization is shown to implicitly satisfy Gauss' law. However, if the fluid representation is modified in a non-charge density preserving manner, e.g.

slope limiting, then a correction to the electric field is required to preserve Gauss' law. Two such corrections are presented and compared in this study. The first is a global Poisson solve, similar to the Boris(1970) correction used in particle in cell (PIC) algorithms. The second is local elementwise Poisson solve, which takes advantage of the DG discretization. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Caleb D. Logemann
Sandia National Labs
clogema@sandia.gov

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratories
eccyr@sandia.gov

Duncan A. McGregor
Sandia National Laboratories
damcgre@sandia.gov

PP1

Full Waveform Inversion Using Fourier Neural Operators and An Adversarial Regularization Network

We propose a workflow to solve Full Waveform Inversion problems with regularization to detect defects in concrete structures. An actuator creates an ultrasound wave that travels through the structure and changes its speed when hitting defects inside it. Then, a sensor array on the outside of the structure observes the waves propagating back. From these measurements, the internal defects in the structure can be inferred. Traditionally, optimization techniques in combination with classical numerical methods are used to solve this inverse scattering problem. Wave propagation in the domain of interest is simulated in high resolution several times for different proposed defects. Significant savings in computational resources can be realized with a suitable surrogate for the traditional numerical solver. Here, we train Fourier Neural Operators to map actuator data to the wave inside the structure at certain time points. Inspired by work in computer vision for computer tomography in medicine, we also train a regularizer term for the outer inverse problem, using an Adversarial Regularization Network. Compared to the standard inverse problem solution with a numerical solver and classical regularization, our approach can resolve defects with higher on-line efficiency and much higher fidelity.

Rahul Manavalan, Yiming Zhang
Chair for Scientific Computing in Computer Science
Technical University of Munich
rahul.manavalan@tum.de, ge59daw@mytum.de

Felix Dietrich
Technical University of Munich
felix.dietrich@tum.de

PP1

Reconstructing Unobserved Cardiac Excitations from Experimental Recordings using Data Assimilation

In typical cardiac experiments, optical mapping is used to measure spatial and temporal dynamics of excitation across the surface(s) of the heart. The reconstruction of excitation patterns through the unobserved depth of the tissue is

essential to realizing the potential of computational models in cardiac medicine. We have interpolated experimental recordings of cardiac excitation on the epicardium and endocardium surfaces of a slice of cardiac tissue, which are then treated as observations in an local ensemble transform Kalman Filter (LETKF) scheme with observations having a prescribed uncertainty. We demonstrate that by including explicit information about the stimulation protocol we can improve the accuracy of the ensemble reconstruction and improve the reliability of the assimilation over time. Additionally, we find that adapting uncertainty estimates based on the phenomenological features of the excitation waves can lead to improvements in the reconstruction.

Christopher Marcotte, Elizabeth Cherry
School of Computational Science and Engineering
Georgia Institute of Technology
christopher.marcotte@durham.ac.uk,
elizabeth.cherry@gatech.edu

Flavio H. Fenton
Georgia Institute of Technology
flavio.fenton@physics.gatech.edu

Matthew J. Hoffman
Rochester Institute of Technology
mjhsma@rit.edu

PP1

Material Changes As Fictitious Heat Sources in Inverse Heat Transfer Problems

Often, in inverse heat transfer problems, we seek an estimate of some localized changes in material properties, in zones of the domain that are not directly observable. If we assume to do experiments with infrared thermographic data as measurements on the boundary, a nonlinear inverse problem arises. A commonly used method is to define a quadratic cost function involving the prediction error of temperatures at the boundary, obtained by a conveniently parameterized PDE model, and to use some regularized nonlinear optimization method. We show that the gradient of the prediction error, often used in the literature, is problematic, while fictitious source term estimation behaves much better. Moreover, by exploiting the well-known maximum principle, we derive an algorithm to detect and quantify hidden material changes. To show the effectiveness of this approach, we present some results in two case studies: the hidden corrosion estimation problem, arising in nondestructive testing of metal structures, and the estimation of leavening in bread making processes.

Fabio Marcuzzi
University of Padova
marcuzzi@math.unipd.it

PP1

SIAGE/Supercomputing Initiatives: Raising Awareness of HPC Opportunities and Impact Growing the HPC Community

The SIAM Activity Group on Supercomputing (SIAG/SC, <https://siag-sc.org>) provides a forum for computational mathematicians, computer scientists, computer architects, and computational scientists to exchange ideas on mathematical algorithms and computer architecture needed for high-performance computer systems. SIAG/SC promotes the exchange of ideas by focusing on the interplay of an-

alytical methods, numerical analysis, and efficient computation. This poster provides an overview of new SIAG/SC initiatives that aim to raise awareness of opportunities and impact in high-performance computing (HPC) and grow the HPC community. A focus is the new Supercomputing Spotlights webinar series, featuring short presentations that highlight the impact and successes of HPC throughout our world. Presentations, emphasizing achievements and opportunities in HPC, are intended for the broad international community, especially students and newcomers to the field. We welcome your ideas and contributions. Join us! SIAG/SC Officers 2022-23: Lois Curfman McInnes (chair), Hatem Ltaief (vice chair), Michael Bader (program director), Rio Yokota (secretary)

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@anl.gov

Hatem Ltaief
KAUST, Saudi Arabia
hatem.ltaief@kaust.edu.sa

Michael Bader
Technical University of Munich
bader@in.tum.de

Rio Yokota
Tokyo Institute of Technology
rioyokota@gsic.titech.ac.jp

PP1

Matlab Implementation of the Finite Element Method in Electro-Viscoelasticity

In this work, we describe a numerical method for 2D non-smooth contact problems with Tresca friction. First, we study the antiplane frictional contact models for electro-viscoelastic materials. The material is assumed to be electro-viscoelastic and the friction is modeled with Tresca's law and the foundation is assumed to be electrically conductive. First we establish the existence of a unique weak solution for the model. Moreover, the Proof is based on arguments of evolutionary inequalities. Comparison is made with results obtained using a finite element program, MATLAB. Selected numerical examples of application of the algorithm are presented here.

Dalah Mohamed
University FMC-Constantine, Faculty of Exact Sciences
Department of Mathematics, Constantine/Algeria
mdalah.ufmc@gmail.com

Megrous Amar
Ecole Supérieure de Comptabilité
et de Finances à Constantine/Algeria
megrous.amar@gmail.com

derbazi ammar
University BBA
BORDJ-BOUARRRIDJ/Algeria
aderbazi@yahoo.fr

fadlia besma
University FMC-Constantine, Faculty of Exact Sciences
Department of Mathematics, Constantine/Algeria

besmafadlia@gmail.com

PP1

Challenges and Chances of Task Parallelism on Gpus

The algorithms behind many scientific applications do not provide sufficient data-parallelism to exploit the massive hardware parallelism of current CPUs and GPUs. The task-parallel programming paradigm has been introduced to enable the full use of the degree of concurrency that these algorithms provide. While task-parallelism benefits applications on CPUs, the concept is deemed unsuitable for GPUs since they are not designed for synchronization-heavy software. The latter gets especially challenging for fine-grained task-parallelism. We present an execution model for fine-grained task parallelism on GPUs. The poster provides a quantitative analysis of current architecture trends in CPUs and GPUs and outlines their influence on programming models. Further, it describes programming model features that facilitate task parallelism on GPUs. Based thereon, it outlines a mapping between CPU and GPU features by means of their concurrent processing capabilities. The poster describes a CUDA-based execution model with persistent threading that enables queue-based task scheduling on GPUs. In addition, it outlines diverse performance optimizations for the proposed task queues. Further, an overview of the task-parallel programming library Eventify and its porting to GPUs via the proposed execution model is presented. Concluding, we evaluate this approach by comparison to a data-parallel OpenACC implementation with a fast multipole method for molecular dynamics as use case.

Laura Morgenstern
Durham University
Department of Computer Science
laura.morgenstern@durham.ac.uk

Ivo Kabadshow
Research Centre Juelich
i.kabadshow@fz-juelich.de

Matthias Werner
Chemnitz University of Technology
matthias.werner@informatik.tu-chemnitz.de

PP1

Kernel-Based Approximation of Koopman Generator for Coarse-Grained Stochastic Dynamical Systems

Much attention has been paid to dimensionality reduction and model discovery for complex systems. An extension of Extended Dynamic Mode Decomposition (EDMD) has been introduced in [Klus et al., *Physica D* (2020)] to approximate the Koopman generator for system identification. However, the selection of basis functions upon which the generator is approximated is not an easy task. By taking advantage of kernel methods introduced in [Klus et al., *Entropy* (2020)], we develop a kernel-based data-driven method to approximate the Koopman generator of specifically nondeterministic dynamical systems via Galerkin projection in reproducing kernel Hilbert spaces. The method allows us to identify a stochastic differential equation governing the coarse-grained model of a high-dimensional system. Dominant dynamics and metastabilities of the system in the reduced-order space can be obtained by the eigen-decomposition of Galerkin approximation of the projected

generator. We numerically analyze the method using toy models governed by overdamped Langevin dynamics.

Vahid Nateghi

Max Planck Institute for Dynamics of Complex Technical Systems
nateghi@mpi-magdeburg.mpg.de

Feliks Nüske

Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
nueske@mpi-magdeburg.mpg.de

PP1

Open-Source Tool for Gas Composition Tracking in Pipeline Networks

Utilizing hydrogen in energy sources constitutes a key strategy towards the transition to renewable energies. Energy surplus from solar and wind sources can produce green hydrogen which is proportionally induced in existing natural gas pipelines for transport. This however requires reliable monitoring and control of gas mixtures at the entry nodes of the gas network and a detailed study of its transport. A multi-scale simulation of gas networks, that can predict the dynamics of gas mixture composition can effectively aid operational and predictive maintenance. With this aim, a scalable model is utilized to compute the flow of gas mixtures along a pipeline network. The network flow model is framed into a system of equations and finite-element discretization is utilized to solve them. The flow solution is coupled with species transport to track different components of the gas along the network. An open-source software is developed to study the non-linear aspects of flow and composition for practical scenarios of pipeline networks involving discontinuous supply and demand. The tool serves to gain insights in system response for different conditions to enable dynamic control of hydrogen composition within pipeline networks.

Ashwin S. Nayak

MPI Magdeburg
anayak@mpi-magdeburg.mpg.de

Sara Grundel

Max Planck Institute for Dynamics of Complex Technical Syst
grundel@mpi-magdeburg.mpg.de

PP1

Towards the Smarter Tuning of Molecular Dynamics Simulations

The large computational cost of pairwise force calculations within Molecular Dynamics requires the use of specialist algorithms, such as Linked Cells or Verlet Lists, as well as efficient ways of parallelising such algorithms. There is, however, no 'silver bullet' best algorithm for all simulations, and the best algorithm can change over the course of a simulation. AutoPas is a node-level particle simulation library that aims to dynamically select the most optimal algorithm, vectorisation strategy, and shared memory parallelism for a given metric, such as time for force calculation [F. Gratl et al, N ways to simulate short-range particle systems: Automated algorithm selection with the node-level library AutoPas, 2022]. In multi-node HPC systems, each node has their own AutoPas container, making it's own tuning decisions. Practically, this autotuning requires trialling algorithms during the course of the simulation, how-

ever trialling slow algorithms can provide significant overhead, and so smart tuning strategies must be developed that can select optimal, or close to optimal, performance, with minimal overhead. In this poster, we will discuss how statistical techniques such as Bayesian Optimisation, Gaussian Process Models, and Reinforcement Learning can be adapted into smart tuning strategies within AutoPas. To support our claims, we present results using our smart tuning strategies applied to the field of Molecular Dynamics.

Samuel J. Newcome, Fabio A. Gratl

Technical University of Munich
samuel.newcome@tum.de, f.gratl@tum.de

Philipp Neumann

Helmut Schmidt University, Germany
philipp.neumann@hsu-hh.de

Hans-Joachim Bungartz

Technical University of Munich
bungartz@cit.tum.de

PP1

Modelling and Simulating Integrated Energy Networks

To slow down the global warming process due to the emission of greenhouse gases, one can replace fossil fuel usage with renewable energy, hydrogen gas and cogeneration units within the current energy networks that provide electricity, gas and heat. The replacement of fossil fuel is currently ongoing and is also known as energy transition. The implementation of these components introduces interaction between different energy networks. To implement these components efficiently, a mathematical model is required that can incorporate these interactions and compute the optimal load flow of the integrated energy network. Numerical and modelling challenges arise even for medium sized integrated energy networks with more than 10^3 unknowns and equations. The main goal is to scale this to large sized integrated energy networks, such as the Dutch energy network. Solvability and scalability are two main challenges that one faces when modelling and simulating integrated energy networks, specifically for coupling between electricity, gas and heat networks. In this poster presentation, these challenges will be explained. Moreover, one will gain insight on how to tackle these challenges.

Buu-Van Nguyen

Technical University of Delft
buuvan.nguyen@gmail.com

PP1

Multilevel Basis Reduction

In the last two decades, model order reduction has been established as an important tool for the solution of high-dimensional parametrized partial differential equations, particularly in the many-query contexts of optimization, control, uncertainty quantification, and inverse problems. However, even with the development and success of new methods that exploit machine learning tools, the problem of offline sampling — in which data from observations or simulations are used to train the surrogate model — remains. Most methods still rely on a random sampling of the parameter space, which typically necessitates large amounts of training data especially in the case of high-dimensional parameter spaces. In this work we exploit the dependence on the variance of the solution and explore the

use of variance reduction and multi-level techniques to reduce the number or expensive training samples required to train a reduced basis model. We test our proposed method on a high-dimensional PDE.

Evie Nielen

Eindhoven University of Technology
e.t.e.nielen@tue.nl

PP1

Point Data in the Differentiable Code Generation System, Firedrake

Firedrake is a code generation system which enables straightforward, scalable and differentiable model development using finite element methods. We present new capabilities which allow the evaluation, manipulation, and assimilation of equation systems which contain *point data* i.e. sets of values which are defined at particular points in space. Point data turn up regularly when finding the value of a field at a particular point or when assimilating point measurement data into a model. We will discuss the necessary data structures, operators and code generation pathways, as well as exciting new Domain Specific Language (DSL) language features. Our work is compatible with the Automatic Differentiation (AD) system `dolfin-adjoint/pyadjoint` which can straightforwardly generate code for a models discrete-adjoint and perform second derivative calculations. We are therefore able to solve PDE constrained optimisation problems involving point data and demonstrate the advantages of this over alternative approaches with a straightforward data assimilation example.

Reuben W. Nixon-Hill

Department of Mathematics
Imperial College London
reuben.nixon-hill10@imperial.ac.uk

Daniel Shapero
Polar Science Center
Applied Physics Lab
shapero@uw.edu

Colin J. Cotter
Imperial College London
Department of Mathematics
colin.cotter@imperial.ac.uk

Stephan Kramer
Department of Earth Science and Engineering
Imperial College London
s.kramer@imperial.ac.uk

David Ham
Department of Mathematics
Imperial College London
david.ham@imperial.ac.uk

PP1

Negative Binomial Optimization for Novel Structural Variant Detection

Structural variants (SVs) – such as insertions, deletions, and duplications of an individuals genome – represent an important class of genetic mutations which have been associated with both genetic diseases (e.g. cancer) and promotion of genetic diversity. Common approaches to detect

SVs in an unknown genome require sequencing fragments of the genome, comparing them to a high-quality reference genome, and predicting SVs based on identified discordant fragments. However, inferring SVs from sequencing data has proven to be a challenging mathematical and computational problem because true SVs are rare and prone to low-coverage noise. We developed a computational method which seeks to improve existing SV detection methods in three ways: First, we generalize previous work by implementing an optimization approach consisting of a negative binomial log-likelihood objective function. Second, we use a block-coordinate descent approach to simultaneously predict if an SV is homozygous (SV is on two chromosomes) or heterozygous (SV is on one chromosome) given genomic data of related individuals. Third, we model a biologically realistic scenario where variants in the child are either inherited and therefore must be present in the parent or novel. We present results on simulated data, which demonstrate improvements in predicting SVs and uncovering true SVs from false positives.

Jocelyn Ornelas-Munoz, Erica M. Rutter

University of California, Merced
jornelasmunoz@ucmerced.edu, erutter2@ucmerced.edu

Mario Banaelos

California State University, Fresno
mbanaelos22@csufresno.edu

Roummel F. Marcia

University of California, Merced
rmarcia@ucmerced.edu

PP1

Statistical Modeling of Covid-19 Outbreak

In epidemiology, statistical modelling helps to understand the mechanisms that influence the spread of infectious diseases, and it suggests prevention and control strategies. It can be used to comprehend how a virus spreads across a large region or country hence, this work is devoted to large scale analysis and application of different statistical tools such as machine learning and deep learning to model the COVID-19 pandemics so as to predict and forecast the evolution of the disease at the population level. COVID-19 cases prediction models are currently divided into three categories: theoretical models, single artificial intelligence models, and decomposition integration models. The majority of statistical methods are used in the case study of COVID-19 to make prediction of the pandemic evolution at different phases. We present in this work some useful quantitative statistical tools for analysing epidemiological data. The current study demonstrated the ability of machine learning models and deep learning to accurately identify major components of viral disease such as COVID-19 and several determinants by retro-predicting the spread of the disease.

Kayode Oshinubi

Université Grenoble Alpes
Université Grenoble Alpes
kayode.oshinubi@univ-grenoble-alpes.fr

PP1

Enhanced Parametric Level Set Methods for Tracking Evolving Objects

We consider the time series reconstruction of piecewise constant objects using a new parametric level-set method

called PaLEnTIR (Parametric level-sets enhanced to improve reconstruction). PaLEnTIR is a significantly enhanced PaLS model relative to the current state-of-the-art which requires only a single level-set function to recover a scene with piecewise constant objects possessing multiple unknown contrasts. We first demonstrate the ability of PaLEnTIR to recover static images in 2D and 3D. We then show how to incorporate time into our PaLEnTIR image model, and demonstrate the ability of the resulting time-varying PaLEnTIR to recover and track the change of shape and position of objects over time.

Ege Ozsar
Tufts University
ege.ozsar@tufts.edu

Misha E. Kilmer
Mathematics Department
Tufts University
misha.kilmer@tufts.edu

Eric Miller
Department of Electrical and Computer Engineering
Tufts University
eric.miller@tufts.edu

PP1

An Efficient Hardware-Aware Matrix-Free Implementation for the Finite-Element Discretized Operator Action on Multi-Component Vectors

The finite-element (FE) discretization of a partial differential equation usually involves construction of a FE discretized operator and computing its action on trial FE discretized fields for the solution of a linear system of equations or eigenvalue problems and is traditionally computed using global sparse-vector multiplication modules. However, recent hardware-aware algorithms for evaluating such matrix-vector multiplications suggest that on-the-fly matrix-vector products without building and storing the cell-level dense matrices reduce both arithmetic complexity and memory footprint and are referred to as matrix-free approaches. These approaches exploit the tensor-structured nature of the FE polynomial basis for evaluating the underlying integrals and the current state-of-the-art matrix-free implementations deal with the action of FE discretized matrix on a single vector. These are neither optimal nor readily applicable for matrix-multivector products involving large number of vectors. We discuss a computationally efficient and scalable matrix-free implementation procedure to compute the FE discretized matrix-multivector products on both multi-node CPU and GPU architectures. The performance of the proposed implementation procedure is thoroughly assessed on representative FE discretized problems corresponding to the Poisson, Helmholtz equations with multiple forcing vectors and the solution of the Kohn-Sham DFT eigenvalue problem arising in quantum modeling of materials.

Gourab Panigrahi
Indian Institute of Science
pgourab@iisc.ac.in

Nikhil Kodali
Indian Institute of Science, Bangalore, India
nikhilkodali@iisc.ac.in

Phani S. Motamarri
Indian Institute of Science Bangalore

phanim@iisc.ac.in

PP1

Finite Element Method for P-Laplace and Infinity Laplace Equations on Surfaces

We seek to construct convergent numerical methods for the p-Laplace and infinity Laplace equations posed on surfaces. On the plane, finite element methods have been shown to converge for non-smooth viscosity solutions where finite difference schemes fail to succeed. We implement a surface finite element method. In the limit as $p \rightarrow \infty$, the problem becomes ill-conditioned and so requires careful treatment. We present results for the Dirichlet problem on smooth surfaces lying in \mathbb{R}^2 and \mathbb{R}^3 and outline image processing applications.

Hannah Potgieter, Steve Ruuth, Razvan Fetecau
Simon Fraser University
hpa59@sfu.ca, steven_ruuth@sfu.ca, razvan_fetecau@sfu.ca

PP1

Evaluating Differences Between #BlackLivesMatter and #AllLivesMatter: Discourse and Interpretations

Viral hashtags often emerge on Twitter to discuss social and political topics. In response, others regularly surface to express divergent stances. Among the most widespread and influential of this phenomenon is the case of #BlackLivesMatter and #AllLivesMatter. Although neither hashtag is formally associated with any particular identity, both the content of the tweets and the hashtag attached can serve as nuanced communicative tags, marking tweets with contextual information possessing varying implicatures for different audiences. To address variation between the two hashtags, we sought to (1) qualitatively record variations in discourses corresponding to each respective hashtag, and (2) correlate differences with resulting interpretations. We constructed our dataset by scraping #BlackLivesMatter and #AllLivesMatter tweets from 2020, then performed analysis with word embeddings to understand qualitative differences present between the two hashtags. We additionally created, employed, and collected survey data to determine interpretations of the two hashtags as signals. We then identified correlations between evaluations of stimuli (tweets), characteristics of tweets, and evaluator demographics. We find that the two hashtags differ greatly in terms of diversity of conversation, yet possess notable consistency in topics. Additionally, the political orientation of participants proved to be the most significant predictor of evaluations.

Maia Powell
University of California Merced
mpowell2@ucmerced.edu

Arnold D. Kim
Applied Mathematics Department
University of California, Merced
adkim@ucmerced.edu

Paul Smaldino
University of California, Merced
Cognitive and Information Sciences Department

psmaldino@ucmerced.edu

PP1

Large Scale Randomized Iterative Least Squares

Solving the least squares problem is fundamental to many typical predictive techniques of today, such as 4D-Var and GLMs. Unfortunately, our increasing desire for more accurate predictions requires that we use finer grids and more data, which increases the computational difficulty of solving these problems. This increase arises from the high cost of moving memory at a large scale. Using Krylov methods or Incremental QR ameliorates the issues arising from these memory costs when the matrix has either a high row or column dimension. Unfortunately, these methods fail to be acceptable solutions when the system has both a high row and column dimension. In this case, Iterative Random Sketching (IRS) appears to be a good solution because of its ability to leverage random sketching to approximate large systems with low-dimensional ones. However, for IRS to work efficiently, its progress must be cheaply tracked and stopped. This presentation will outline a technique for estimating and stopping the progress of such Iterative Random Sketching methods. It will establish theoretically and experimentally that our method is able to track progress with high accuracy and stop progress with a low failure rate. Using this progress tracking method, we will finish the presentation by showing how IRS is able to facilitate solving both 4D-Var and GLM problems of about 0.76 TB in size using only 100 MB of memory in scenarios where state-of-the-art methods fail.

Nathaniel Pritchard

University of Wisconsin-Madison
npritchard@wisc.edu

Vivak Patel

University of Wisconsin Madison
vrpatel6@wisc.edu

PP1

Sparse Machine-Learning Proxy Models for Geothermal Reservoir Simulation

Uncertainty quantification of highly-parameterised subsurface reservoir models over multiple geological realisations can be computationally prohibitive. The use of a low-fidelity subsurface proxy model has been shown to maximise use of information inherently contained in high-dimensional data, while significantly reducing computational demands necessary for accurate statistical evaluation. However, a principal concern in proxy model applications is loss of physical and geological realism within the reduced-order model, subsequently compromising our understanding of reservoir dynamics and resulting in loss of model prediction power. This work explores the framework of Machine Learning (ML)-based subsurface proxy models for prediction of reservoir production performance. The Sparse Identification of Nonlinear Dynamics (SINDy) algorithm [Brunton et al., 2016] is used for discovery of parsimonious physical models from spatiotemporal data. Previous work on SINDy illustrates a SINDy autoencoder method [Champion et al., 2019] for discovery of nonlinear reduced coordinates along with associated nonlinear governing equations for dynamics in a joint optimization. We investigate the possibility of extending the SINDy autoencoder method to establish a parsimonious ML-based subsurface proxy model. For these applications, synthetic training data is generated using the multi-physics Delft Ad-

vanced Research Terra Simulator (DARTS) framework.

Farah Rabie

Heriot-Watt University
fr2007@hw.ac.uk

Gabriel Serrao Seabra
Technische Universiteit Delft
Petrobras Brazil
g.serraoseabra@tudelft.nl

Vasily Demyanov
Heriot-Watt University
v.demyanov@hw.ac.uk

Denis Voskov
Delft University of Technology
Stanford University
d.v.voskov@tudelft.nl

PP1

Efficient and Scalable Finite-Element Based Computational Methodologies for Large-Scale Ab-Initio Modelling of Energy Storage Materials

The predictive capability offered by ab-initio quantum mechanical modelling of materials opens a gateway to understand the underpinning mechanisms that govern the efficiency and safety of energy storage materials. This demands large-scale density functional theory simulations beyond the current high-throughput calculations routinely done by employing plane-wave or atomic-orbital based approaches which are restrictive in terms of the nature of boundary conditions one can employ or the accuracy achieved. To this end, the recently proposed finite-element-based methods for DFT (DFT-FE) provide an HPC-centric framework that overcomes these limitations. We first discuss an efficient, scalable computational methodology that leverages this DFT-FE framework for extracting chemical bonding information of large-scale material systems using a projected Hamiltonian population analysis approach. This has implications in screening for alloying elements to enhance energy storage properties. Subsequently, we discuss the real-space formulation and efficient finite-element-based implementation strategy to incorporate projector augmented wave (PAW) formalism in the DFT-FE framework on hybrid CPU-GPU architectures. This FE-based PAW framework allows us to reduce the degrees of freedom required to achieve the required chemical accuracy, thereby enabling much large length scales and longer time scales crucial for addressing complex problems in the design of energy storage devices.

Kartick Ramakrishnan

Department of Computational & Data Science
Indian Institute of Science, Bangalore
kartickr@iisc.ac.in

Sambit Das
Department of Mechanical Engineering
University of Michigan Ann Arbor
dsambit@umich.edu

Phani S. Motamarri
Indian Institute of Science Bangalore

phanim@iisc.ac.in

PP1

pyMOR - Model Order Reduction with Python

pyMOR (<https://pymor.org>) is a free and open source model reduction software library for the Python programming language. Originally created with the application of Reduced Basis methods to large-scale problems in mind [Ohlberger, Rave et al., 2014], it has been designed from the ground up for seamless integration with external PDE solvers by expressing all algorithms in terms of operations on VectorArray, Operator and Model interface classes [Milk, Rave et al., 2016]. Since its inception in 2012, pyMOR has grown significantly beyond its original scope and now offers a wide selection of both Reduced Basis and system-theoretic algorithms. Developers from both fields [Balicki, Mlinaric et al., 2019], [Mlinaric, Rave et al., 2021] drive the continuing evolution of the library as an open source project. Recent additions include data-driven algorithms such as Dynamic Mode Decomposition or neural-network based approaches, structure-preserving methods as well as randomized numerical linear algebra algorithms. With this poster we will give an overview on pyMOR's design and features. We will also discuss our current and future development goals.

Linus Balicki
Virginia Tech
Blacksburg, USA
balicki@vt.edu

René Fritze
University of Münster
rene.fritze@wwu.de

Petar Mlinaric
Virginia Tech
Blacksburg, USA
mlinaric@vt.edu

Stephan Rave
University of Münster
stephan.rave@wwu.de

Jens Saak
Max Planck Institute for Dynamics of Complex Technical
Syste
saak@mpi-magdeburg.mpg.de

Felix Schindler
University of Muenster
Applied Mathematics Muenster
felix.schindler@wwu.de

PP1

Narrowband Transmit Beam Pattern in Medical Ultrasound: a Stochastic Approach to Delays Optimization

Medical Ultrasound is the most widespread real-time non-invasive diagnostic tool: it exploits the ability of human tissue to reflect ultrasound signals to reconstruct images of the organs or quantify physical and physiological parameters. The performance of those different modalities is greatly affected by the transmit Beam pattern (TBP), whose shape and intensity depends on several parameters such as transmit pulse shape, central frequency, focal depth

and so on. TBP optimization is an important task in some advanced applications, which work with narrow band signals. In this work we present a model for narrowband beam patterns in time frequency domain and propose an optimization method to overcome typical limitations of standard Beam Patterns, like non-uniform beam width over depth, presence of significant side lobes and quick energy drop out after the focal depth. We propose to no longer fix a single focal depth determining a whole-time delay curve but to consider each transmit delay as a free variable. The optimization problem is formulated as a Least Square problem by considering a prescribed TBP. Since the obtained problem is non-convex we tackle its solution using Particle Swarm Optimization. The obtained results have been compared with a set of standard TBP, showing an overall improvement of desired features, thus demonstrating the effectiveness of the proposed approach. Moreover, to allow a quantitative evaluation, we introduce a novel set of metrics.

Chiara Razzetta

University of Genoa
razzetta@dima.unige.it

Marco Crocco
Esaote S.p.A.
marco.crocco@esaote.com

Federico Benvenuto
Università di Genova
benvenuto@dima.unige.it

PP1

Training Large-Scale Neural Networks with a Newton Conjugate Gradient Method

Training deep neural networks consumes tremendous computational resources - be it for brute-force hyperparameter search (1) or optimizer search (2). Our approach is (1) to enhance this by a *second-order* optimization methods with *fewer hyperparameters* for large-scale neural networks and (2) to perform a survey of the performance optimizers for specific tasks to suggest users the best one for their problem. We introduce a novel second-order optimization method that requires the effect of the Hessian on a vector only and avoids the huge cost of explicitly setting up the Hessian for large-scale networks. The Newton equation is solved by a conjugate gradient method (with regularization and cheap preconditioning). We compare the proposed second-order method with two state-of-the-art optimizers on six representative neural network problems, including regression, bayesian neural networks and very deep networks from computer vision, variational autoencoders or transformers with $\geq 20M$ weights. We efficiently parallelized the optimizers with Horovod and applied it to a 8 GPU NVIDIA A100 (DGX-1) machine with $\geq 60\%$ parallel efficiency.

Severin Reiz, Tobias Neckel
Technical University of Munich
s.reiz@tum.de, neckel@in.tum.de

Hans Joachim Bungartz
Technical University Munich

bungartz@tum.de

PP1

Low-Rank Methods for Iga with Multiple Patches

Isogeometric analysis is an important tool in the discretization of PDEs on complex domains. In this poster, we illustrate that low-rank techniques based on tensor decompositions can be applied when the domain consists of multiple patches. For this, we utilize the tensor train format in combination with the Block-AMEN iterative solver. We further show that this methodology can also be used when a PDE-constrained optimization problem has to be solved.

Tom-Christian Riemer

Technische Universität Chemnitz
tom.riemer@mathematik.tu-chemnitz.de

PP1

Effects of Prey Capture on the Swimming and Feeding Performance of Choanoflagellates

Choanoflagellates, eukaryotic predators of bacteria in aquatic ecosystems, are used as a model system to study the evolution of animals from protozoan ancestors. The choanoflagellate, *Salpingoeca rosetta*, has a complex life cycle that includes unicellular and multicellular stages, provides a model system to study the consequences of different cell morphologies, being free-swimming vs. sessile, or being a single cell vs. a multicellular colony. A unicellular *S. rosetta* has an ovoid cell body and a single flagellum surrounded by a collar of microvilli. The cell swims by waving its flagellum, creating a water current that brings bacteria to the collar of prey-capturing microvilli. One measure of the performance of a suspension-feeding organism is the volume of fluid that it can move into its collar during a beat cycle. The inward flux of fluid acts as a proxy for the rate of bacterial capture. While this is a good measure of uptake of dissolved nutrients, it is only an approximate measure of prey capture. Here we use a regularized Stokeslet framework to model the hydrodynamics of a unicellular choanoflagellate, the captured bacterial prey of non-zero volume, and their effect on swimming performance and clearance rate. We compare the model predictions with high-speed microvideography. Moreover, we will discuss current model assumptions, and future model improvements that, together with coordinated lab experiments, will help us probe this intriguing biophysical system.

Emma C. Ross, Hoa Nguyen

Trinity University
eross1@trinity.edu, hnguyen5@trinity.edu

M.A.R. Koehl
University of California Berkeley
Integrative Biology
cnidaria@berkeley.edu

Lisa J. Fauci
Tulane University
Department of Mathematics
fauci@tulane.edu

Ricardo Cortez
Tulane University
Mathematics Department

rcortez@tulane.edu

PP1

Adaptive Discontinuous Galerkin Methods for 2D Unsteady Convection-Diffusion Problems on Moving Mesh

We propose and analyse a semi-Lagrangian method for the convection-diffusion equation. An a posteriori error estimation for interior penalty discontinuous Galerkin semi-discretization of the nonstationary problem is derived to link together the principles of semi-Lagrangian method and adaptive mesh refinement. These error criteria give an explicit expression of their dependence on each parameter of the simulation. They are finally tested through tailored test cases.

Ezra Rozier

Universität Hamburg (UHH)
ezra.rozier@uni-hamburg.de

Jörn Behrens

Dept. of Mathematics, Universität Hamburg
joern.behrens@uni-hamburg.de

PP1

The Pseudo-Spectrum of the Parareal Parallel-in-Time Iteration

Parareal is one of the most widely studied parallel-in-time algorithms. For linear problems, it can be written and analyzed as a stationary linear iteration. Its iteration matrix is nil-potent, owing to the well-known fact that Parareal always converges after a finite number of iterations. Therefore, the spectral radius is always zero and thus not useful to assess whether Parareal converges monotonically and quickly enough to provide parallel speedup. The poster will instead look at the pseudo-spectrum of the iteration matrix and explore if it provides useful information about how Parareal convergence. In particular, we will analyze if the pseudo spectral-radius can predict the non-monotonic convergence behavior that Parareal can show when applied to problems with imaginary eigenvalues or hyperbolic partial differential equations.

Daniel Ruprecht

Institute of Mathematics
Hamburg University of Technology
ruprecht@tuhh.de

PP1

Scalable I/O for Firedrake and PETSc

One important aspect of scientific and engineering simulations is scalable input/output of solutions. In this work we enhanced HDF5 input/output capabilities of Firedrake [Florian Rathgeber and David A. Ham and Lawrence Mitchell and Michael Lange and Fabio Luporini and Andrew T. T. McRae and Gheorghe-Teodor Bercea and Graham R. Markall and Paul H. J. Kelly, *Firedrake: automating the finite element method by composing abstractions*. ACM Trans. Math. Softw., Vol.43, pp. 24:124:27, 2016] and PETSc to allow for saving/loading finite element solutions in association with the mesh of the computational domain efficiently in parallel. Particularly, our new implementation allows for using different number of MPI processes for saving and for loading. Transparent interfaces for mesh extrusion and timestepping problems are also pro-

vided.

Koki Sagiyama
Imperial College London
k.sagiyama@imperial.ac.uk

Václav Hapla
ETH Zurich
vaclav.hapla@erdw.ethz.ch

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Lawrence Mitchell
Nvidia
wence@gmx.li

David Ham
Department of Mathematics
Imperial College London
david.ham@imperial.ac.uk

PP1

A Finite-Element-Solver for Coupled Domains in Rust

The Finite-Element-method (FEM) is one of the most used numerical techniques to solve partial differential equations (PDEs). Over the last decades, many FEM-solvers have been published, free software as well as proprietary. However, most of them are written in 'classical' languages like C/C++, Fortran or Julia, and some of them with a Python-interface. FEniCS, deal.ii, Gridap.jl and Comsol come to mind, to list a few. Rust, however, is a language primarily designed as a system programming language and is rarely used in scientific computing. However, in our opinion, features like the borrow-checker, built-in support for parallelization and its type system make it also an excellent choice for involved scientific software. Furthermore, solving coupled problems with multiple domains (which may or may not have different dimensions) is not straight forward to implement efficiently with most other solvers to the author's knowledge. Our project serves as a proof of concept regarding scientific software written in Rust, and additionally provides a nice way to solve coupled problems.

Alireza Selahi
WIAS Berlin
selahi@wias-berlin.de

PP1

Unbiased Stochastic Optimization for Gaussian Processes on Finite Dimensional RKHS

Gaussian Processes (GPs), while being a powerful probabilistic framework which has been applied to a wide range of learning applications, scales poorly, especially when it comes to optimizing the marginal likelihood, one of the key advantages of GPs over kernel methods. As a consequence, approximations are required for any modern application which involves Gaussian Processes on big data. Arguably in big data regimes it is important to optimize the marginal likelihood with algorithms that are based on stochastic mini-batches, and this is especially true in so-called Deep Gaussian Processes, in which the covariance function is based on a deep neural network. Current

methods for stochastic inference of GPs are based on approximations (e.g., computing biased stochastic gradients, or stochastic variational inference), and as such are not guaranteed to converge to a stationary point of the true marginal likelihood. In this work, we propose an algorithm for exact stochastic inference of GPs with kernels that induce a RKHS of moderate dimension which is a common case in deep kernels. We show experimentally that our method can be also useful for in the infinite case by using random features for kernel approximation.

Neta Shoham, Avron Haim
Tel Aviv University
shohamne@gmail.com, haimav@tauex.tau.ac.il

PP1

Jax-Based Grey-Box Modelling Framework for Building Energy Models

Grey-box modelling has many advantages over white- and black-box modelling and is therefore widely applied in the field of building energy management. However, it still requires significant expert knowledge for creating the model. In this research work, we set up a JAX-based [Schoenholz et. al., Jax] grey-box modelling framework. The usage of JAX comes with just-in-time compilation, automatic differentiation, and a rapidly growing ecosystem in the area of machine learning. This enables the usage of exact gradients and Hessians for optimization. Furthermore, it can be leveraged for hybrid modelling approaches, that is, an unknown component in a grey-box model is described by a learned model, e.g., a neural network. This idea originates from the solver in the loop paper [Kiwon UM et al., Solver-in-the-loop], with the difference that we want to learn a modelling error due to missing physics and not due to discretization. The potential of the framework is shown on a single-zone building model, where the thermal dynamics can be described with a linear state space model, however, the loads and the energy consumption can be non-linear. To minimize the modelling effort while keeping interpretability, we combine RC-Models [Yanfei Li, Grey-box modeling and application for building energy simulations] with machine learning, where the latter is used to provide corrections to the unknown physics. This sets the ground for a future highly generic system identification framework.

Felix M. Sievers
Siemens
felix.sievers@siemens.com

Dirk Hartmann
Siemens AG
Corporate Technology
hartmann.dirk@siemens.com

Hans Joachim Bungartz
Technical University Munich
bungartz@tum.de

PP1

Distributed-Memory Algorithms for Tucker Tensor Completion

Using Tucker decomposition to predict the missing entries of a tensor is considered as a more suited method as compared to CP decomposition. However, there hasn't been any work that compares the scalability and accuracy of these algorithms for distributed-memory Tucker completion. We extend our software infrastructure for tensor

completion which enables easy implementation of Tucker completion algorithms, including Riemannian Tucker completion, and analyze their performance over large synthetic and real-world datasets on distributed-memory architecture.

Navjot Singh
University of Illinois at Urbana Champaign
navjot2@illinois.edu

PP1

Preconditioning Phase Field Equations for Solar Cells

Phase field models play a big role in computational science and engineering as a tool to describe processes from materials science to imaging. We here discuss the numerical solution of a system of partial differential equations arising from modeling the morphology formation of an organic solar cell. In particular, we discuss a preconditioning strategy aimed at resolving the large scale linear systems via a block Schur-complement preconditioner.

Martin Stoll
Department of Mathematics,
Technische Universitaet Chemnitz
martin.stoll@math.tu-chemnitz.de

PP1

Nonlinear Inversion Using Data-Driven Reduced Order Models Applied to the Helmholtz Impedance Boundary Value Problem

Recently, reduced order model (ROM) techniques have received attention and have been applied to Dirichlet and Neumann inverse boundary value problems. With this poster, we give an introduction to the ROM inversion method applied to the Helmholtz impedance problem. We also discuss the limitations of this approach due to the nature of the governing boundary condition. Additionally, we propose a non-linear optimization method and a data-assimilation method to solve the inverse scattering problem based on the ROM approach.

Andreas Tataris
Utrecht University
a.tataris@uu.nl

Tristan Van Leeuwen
Centrum Wiskunde & Informatica (CWI)
t.van.leeuwen@cwi.nl

PP1

An Exploration of Optimal Tensor Contraction Strategies for Vector Inner Product in Tensor-Train Format

Tensor decompositions are proving to be a powerful tool widely used in many high-dimensional problems in quantum computing, machine learning, and scientific computing, as they prevent the inherent exponential data size, also called “the curse of dimensionality”, arising in these problems. Tensor decompositions reduce the data size to a polynomial in the problem’s dimensionality, and enable carrying out tensor operations in this compressed form. In this study, we explore effective computational strategies for the inner product of two vectors in tensor-train format, widely arising in the solution of high-dimensional systems. This

operation involves a series of tensor contractions whose order plays a significant role in the overall computational cost. Though optimizing this order is an NP-hard problem for a general tensor network, we investigate ordering algorithms specific to the network topology manifesting in the inner product of two vectors in tensor-train format. We present various contraction strategies together with their experimental evaluation on tensor-train objects of various sizes and dimensions.

Atte Torri, Brice Pointal, Oguz Kaya
Université Paris-Saclay
atte.torri@universite-paris-saclay.fr,
brice.pointal@universite-paris-saclay.fr,
oguz.kaya@universite-paris-saclay.fr

Laércio Lima Pilla
Université de Bordeaux, CNRS, Bordeaux INP, Inria,
LaBRI
laercio.lima-pilla@labri.fr

Olivier Coulaud
INRIA
olivier.coulaud@inria.fr

PP1

Fast Optimization of Viscosities for Frequency-Weighted Damping of Second-Order Systems

We consider frequency-weighted damping optimization for vibrating systems described by a second-order differential equation. The goal is to determine viscosity values such that eigenvalues are kept away from certain undesirable areas on the imaginary axis. We present two complementary techniques. First, we propose new frameworks using nonsmooth constrained optimization problems, whose solutions both damp undesirable frequency bands and maintain stability of the system. These frameworks also allow us to weight which frequency bands are the most important to damp. Second, we also propose a fast new eigensolver for the structured quadratic eigenvalue problems that appear in such vibrating systems. In order to be efficient, our new eigensolver exploits special properties of diagonal-plus-rank-one complex symmetric matrices, which we leverage by showing how each quadratic eigenvalue problem can be transformed into a short sequence of such linear eigenvalue problems. The result is an eigensolver that is substantially faster than existing techniques. By combining this new solver with our new optimization frameworks, we obtain our overall algorithm for fast computation of optimal viscosities. The efficiency and performance of our new methods are verified and illustrated on several numerical examples.

Matea Ugrica
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg
ugrica@mpi-magdeburg.mpg.de

Nevena Jakovcevic Stor
Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture, University of Split, Croatia
nevena@fesb.hr

Tim Mitchell
Queens College / CUNY
tim.mitchell@qc.cuny.edu

Zoran Tomljanovic

University of Osijek
ztomljan@mathos.hr

PP1

An Efficient Numerical Method for the Maxey-Riley Equation

The Maxey-Riley Equation (MRE) models the motion of a finite-sized, spherical particle moving in a fluid. Applications using the MRE are, for example, the study of the spread of Coronavirus particles in a room, the formation of clouds and the so-called marine snow. The MRE is a second-order, implicit integro-differential equation with a singular kernel at initial time. For over 35 years, researchers used approximations and numerical schemes with high storage requirements or ignored the integral term, although its impact can be relevant. A major break-through was reached in 2019, when Prasath et al. mapped the MRE to a time-dependent Robin-type boundary condition of the 1D Heat equation, thus removing the requirement to store the full history. They provided an implicit integral form of the solution by using the so-called Fokas method that could be later solved with numerical scheme and a nonlinear solver. While Prasath et al.s method can deliver numerical solutions of very high accuracy, the need to evaluate nested integrals makes it computationally costly and it becomes impractical for computing trajectories of a large number of particles. In the poster, we present a finite differences approach that it is not only storage efficient but also much faster. We will compare our approach to both Prasath et al.s method as well as a to the numerical schemes developed by A. Daitche in 2013 for direct integration of the original Maxey-Riley equation with integral term.

Julio Urizarna Carasa
Technical University Hamburg
julio.urizarna@tuhh.de

Daniel Ruprecht
Institute of Mathematics
Hamburg University of Technology
ruprecht@tuhh.de

Leon Schegel
Technical University Hamburg
leon.schegel@tuhh.de

PP1

Mathematical Modelling for All-Solid-State Batteries

Polycrystalline solid-state electrolyte (PSSE) such as LLZO plays an essential role in augmenting electro-chemo-structuro-mechanical stability of the next-generation batteries at cell-level, as well as in diminishing flammability and leakage at product-level. However, such PSSE exhibits microscopic grain boundaries, which tempts imperfectly dendrites and consequently crack propagation. Systematically investigating misorientation and patterns of grain within LLZO gives rise to mechanical and electrochemical properties of the PSSE. Approach: By setting Lam coefficients within structural tensors spatially dependent, the infinitesimal-strain elastic deformation partial differential equation (PDE) helps to solve the misorientation grainy problem of PSSE. Starting point is the linear momentum

PDE used in stress analysis:

$$\partial_t \mathbf{u}^{(s)} + \nabla \cdot \left(\mathbb{C}^{f_{\text{allocation}}}(\lambda, \mu, \mathbf{d}_{G_i}^R, i=1, \dots, N, \mathbf{d}^E; \mathbf{x}) \nabla \mathbf{u}^{(s)} \right) + \rho \mathbf{b} = \mathbf{f}, \quad (4)$$

where $\mathbf{d}_{G_i}^R$ is the local grain direction i , N number of observable grains, \mathbf{d}^E globally uniform electric potential, and λ, μ the two Lam constants characterizing solid properties. Furthermore, coupled problem is taken into consideration. Finally, Griffith criterion can be computed directly based on strain energy, used for dendrite nucleation.

Tuan Vo
RWTH Aachen University
Applied and Computational Mathematics (ACoM)
vo@acom.rwth-aachen.de

PP1

Monolithic Algebraic Multigrid Preconditioning of the Stokes Equations

Advanced discretizations and complex meshes are becoming increasingly commonplace in a range of coupled physics and engineering applications, from weather prediction to fluid flow to graph problems. These discretizations also place increased demands on the underlying solvers, focusing on key properties of the PDE models - e.g., underlying conservation laws. While algebraic multigrid methods effectively tackle some PDEs and some discretizations on unstructured meshes, the presence of higher-order basis functions and multiple coupled unknowns makes the algebraic coarsening a challenging endeavor. This talk will highlight several new monolithic algebraic multigrid preconditioning approaches targeting high-order and non-standard discretizations of Stokes flow problems on structured and unstructured meshes. We will demonstrate how the robust monolithic AMG preconditioners can be constructed to precondition higher-order Taylor-Hood and Scott-Vogelius discretizations.

Alexey Voronin
Department of Computer Science
University of Illinois Urbana-Champaign
voronin2@illinois.edu

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

Raymond Tuminaro
Sandia National Laboratories
rstumin@sandia.gov

PP1

Matrix-Free Hyperparameter Optimization for Gaussian Processes

Gaussian processes (GPs) are a crucial tool in machine learning and their use across different areas of science and engineering has increased given their ability to quantify the uncertainty in the model. The covariance matrices of GPs arise from kernel functions, which are crucial in many learning tasks and the matrices are typically dense and large-scale. Depending on their dimension even comput-

ing all their entries is challenging and the cost of matrix-vector products scales quadratically with the dimension, if no customized methods are applied. We present a matrix-free approach that exploits the computational power of the non-equispaced fast Fourier transform (NFFT) and is of linear complexity for fixed accuracy. With this, we cannot only speed up matrix-vector multiplications with the covariance matrix but also take care of the derivatives needed for the gradient method avoiding Hadamard products of the Euclidean distance matrix and the kernel matrix. This arises when differentiating kernels as the squared-exponential kernel with respect to the length-scale parameter in the denominator of the exponential expression. Our method introduces a derivative kernel which is then well suited for multiplying with the Hadamard product. By applying our NFFT-based fast summation technique, fitting the kernel and the derivative kernel will allow for fast tuning of the hyperparameters.

Theresa Wagner, Franziska Nestler
TU Chemnitz
theresa.wagner@math.tu-chemnitz.de,
franziska.nestler@math.tu-chemnitz.de

Martin Stoll
Department of Mathematics,
Technische Universitaet Chemnitz
martin.stoll@math.tu-chemnitz.de

PP1

Moving Mesh Virtual Element Methods

The use of polygonal discretization techniques for solving partial differential equations (PDEs) has been gaining substantial traction in recent years primarily due to the improved flexibility in representing complex geometries and the computational ease with which a polygonal mesh can be refined and agglomerated. The literature is rich with works that generalise h- and p-adaptive algorithms to polygonal meshes but only recently has there been interest in developing polygonal moving mesh methods, sometimes referred to as r-adaptivity. Developing moving polygonal mesh methods raises additional challenges beyond standard finite element approaches on simplicial meshes. Such challenges include the construction of mappings between moving polygonal elements, the time-dependent nature of discrete spaces, and the robustness of the method. In this poster, we present a Velocity-based Moving Mesh Virtual Element Method for non-linear diffusion free boundary problems. Here, polygonal meshes can be exploited to represent moving boundaries and interfaces with a minimal number of degrees of freedom and to produce efficient local mesh refinement when a change in mesh connectivity is required. Empirical results are presented for the method using the two-dimensional porous medium equation as a benchmark problem. Extensions of the method are then presented that exploit the mesh generality to tackle more geometrically complex problems, such as contact problems.

Harry Wells, Matthew Hubbard
University of Nottingham
harry.wells@nottingham.ac.uk,
pmzmeh@exmail.nottingham.ac.uk

Andrea Cangiani
SISSA

acangian@sissa.it

PP1

Stable Nodal Projection Method on Octree Grids

We present a novel projection solver for the incompressible Navier-Stokes equations with arbitrary boundaries, where all variables are collocated at the nodes of non-graded octree grids. Both the viscosity and projection steps are discretized using supra-convergent finite difference approximations with sharp boundary treatments. The resulting projection operator is stable, as our analysis demonstrates. We verify the stability of our method on uniform Cartesian grids both analytically and numerically, as well as provide a framework for proving the stability on adaptive grids. On adaptive grids, we verify that the operator is stable for highly non-graded grids and arbitrary boundary conditions. We further demonstrate the accuracy and capabilities of our solver with several canonical two- and three-dimensional simulations of incompressible fluid flows. Overall, our method is second-order accurate, allows for dynamic grid adaptivity with arbitrary geometries, and reduces the overhead in code development through data collocation.

Scott R. West
Department of Applied Mathematics
University of California, Merced
swest10@ucmerced.edu

Matthew Blomquist
University of California Merced
mblomquist@ucmerced.edu

Cole Cooper, Dustin Kleckner
Department of Physics
University of California, Merced
ccooper10@ucmerced.edu,

Maxime Theillard
University of California, Merced
mtheillard@ucmerced.edu

PP1

Accelerating Randomized Algorithms for Massive-Scale Zeroth-Order Optimization

Rapid advances in computing, sensing, and other data-generating mechanisms are spawning new classes of challenging optimization problems across computational science and engineering. We address problems where an objective function is available only by querying a zeroth-order oracle. In contrast to gradient-based algorithms, algorithms for such derivative-free problems have necessarily focused on dozens or fewer decision variables. We outline practical randomized strategies for solving problems with hundreds to thousands of decision variables while still preserving theoretical guarantees in idealized cases.

Stefan Wild
Argonne National Laboratory
LANS, Mathematics and Computer Science
wild@lbl.gov

Kwassi Joseph Dzahini
Argonne National Laboratory
kdzahini@anl.gov

Xiaoqian Liu
University of Texas MD Anderson Cancer Center
xiaoqian.liu1025@gmail.com

PP1

Computing p -Harmonic Descent Directions and Their Limits for Shape Optimization

Shape optimization constrained to partial differential equations is a vivid field of research with high relevance for industrial grade applications. Recent development suggests that using a p -harmonic approach to determine descent directions is superior to classical Hilbert space methods, but features the solution of a vector-valued p -Laplace problem with a boundary force in each iteration. We present numerical results in a fluid dynamic setting based on the extension of an algorithm for scalar Dirichlet problems. The approach does not require an iteration over the order p and thus enables the efficient computation of higher-order descent directions. A general requirement on the transformations of the computational domain is to keep it of Lipschitz type. While solutions for finite p yield approximations in $W^{1,p}$, analytically only descent directions in $W^{1,\infty}$ are admissible. However, this is challenging since the limit of the p -Laplace problem features in general non-unique solutions, in particular arising from a change of sign in the force term required to fulfill geometric constraints in shape optimization. Therefore, we make progress towards an algorithm for computing admissible descent directions, that are the limit of the corresponding p -harmonic descent directions. The resulting deformations then still preserve the quality of the underlying mesh, which is crucial for applications.

Henrik Wyszka
Universität Hamburg
henrik.wyszka@uni-hamburg.de

Martin Siebenborn
University Hamburg
Mathematics Department
martin.siebenborn@uni-hamburg.de

Winnifried Wollner
University of Hamburg
Department of Mathematics
winnifried.wollner@uni-hamburg.de

PP1

Deep Learning with Gaussian Processes

By integrating Gaussian processes (GPs) with deep neural networks (NNs), three types of deep models are discussed and compared for multivariate function approximation in this work. The first model is referred to as an *NN-equivalent GP* (Lee et al., 2017), formulated based on the fact that a prior of a Bayesian NN can be equivalent to GPs under certain assumptions. An NN-induced kernel is hence defined as the GP covariance. *Deep kernel learning* (Wilson et al., 2016) is considered a second model, which embeds a deep NN architecture into the kernel of a GP and thus improves the GP's nonlinear expressive power. The third is termed a *deep GP* model (Damianou & Lawrence, 2013), in which GPs are adopted instead of NN layers to reduce computational complexity, and variational inference is utilized to approximate the posterior distribution of model outputs. Numerical experiments are conducted on three benchmark problems with increased input dimensionality, and the test accuracy and Kullback-Leibler divergence are

compared among the three models employing the same NN topology and GP kernel. It has been discovered that the deep kernel learning model provides superior performance. In addition, we summarize the differences among the examined models and provide an explanation of why the deep kernel learning model produces preferred results.

Weihao Yan, Mengwu Guo
University of Twente
w.yan@utwente.nl, m.guo@utwente.nl

PP1

A Linearly Implicit Global Energy Preserving Reduced-Order Model for Cubic Hamiltonian Systems

We present a linearly implicit multi-symplectic model reduction technique for multi-symplectic PDEs with cubic invariants. The proposed approach preserves the global energy of the reduced-order models. Thus, it is suitable for the long-time integration of reduced-order Hamiltonian systems. The proposed method is tested on the wave equation, Korteweg-de Vries equation and Camassa-Holm equation. The numerical results verify that preserving the multi-symplectic structure of PDEs in the reduced-order models yields stable reduced-order models.

Süleyman Yildiz
Max Planck Institute for Dynamics of Complex Technical System
yildiz@mpi-magdeburg.mpg.de

Pawan Goyal
Max Planck Institute for
Dynamics of Complex Technical Systems Magdeburg
goyalp@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

PP1

Droplets Evaporation on Chemically Patterned Surfaces

The evaporation of fluids on the different patterned surfaces is omnipresent in nature. A comprehensive study of the evaporation process coupling with the wetting effect through modeling will give us a complete understanding of the underlying mechanisms and help us construct a digital twin, enabling us to control the whole system. The poster is divided into two parts. Firstly, based on the idea of minimum surface, a theoretical model is established to describe the three-dimensional droplet shape with straight edges and sharp corners on a polygon-patterned substrate in quasi-equilibrium state. This kind of setup is widely used in droplet sampling for high-throughput screening of live cells and chemical reactions. We relate the volume of the shaped-droplet to its height, aiming to address the challenge of measuring the volume of evaporation droplets with usual experimental techniques. The proposed model is compared with phase-field simulation and experiments. Secondly, a Cahn-Hilliard phase field model is utilized to describe the diffusion dominated evaporation process of multi droplets. Through this model, we investigate the effect of the key parameters including the humidity, volume, droplet position/distance/numbers, liquid type/concentration etc. on the evaporation process.

Our aim is to identify an optimal condition for culturing cells and sample preparation on Droplet Microarray (DMA) through the digital twin system.

Hongmin Zhang
Karlsruhe Institute of Technology, Germany
hongmin.zhang@kit.edu

tic PDEs but has extensions to a much broader class of finite element problems.

Jennifer E. Zvonek
Cornell University
jez34@cornell.edu

PP1

Topology-Based Comparison of Population Activities in General Neural Networks

Real-world data are often encoded in high-dimensional representations. Moreover, it is often unclear which coordinates and metrics can be meaningfully justified. Topological properties are well-suited for characterizing the structure of such high-dimensional data: they are generalized to high-dimensional surfaces, and they are invariant under different coordinates and robust to the choice of metrics. The proposed approach is motivated by emerging open problems in neuroscience to analyze the high-dimensional collective activity of a population of neurons in response to specific stimuli (neural population response). A crucial gap in related works is that they have not considered how these neural population responses can be appropriately compared, which is key to understanding neural representations. We develop an approach based on persistent homology and p-Wasserstein distance to quantitatively compare between neural population responses arising from artificial and biological neural networks. This approach further allows one to perform statistical inference on a distribution of topological signatures for the respective neural population responses. We demonstrate the approach on a variety of neural population data from experiments in different regions of the brain and from numerical simulations in different artificial neural networks. This approach thus provides a scalable way to compare the topology of population activities across general neural networks.

Liu Zhang
Princeton University
Princeton University
lz1619@princeton.edu

Fei Han
National University of Singapore
mathanf@nus.edu.sg

Kelin Xia
School of Physical and Mathematical Sciences
School of Biological Sciences, NTU
xiakelin@ntu.edu.sg

PP1

Multi-Objective Adaptive Mesh Refinement Using Reinforcement Learning

The accuracy and efficiency of finite element methods depend immensely on the choice of mesh. However, choosing a mesh that appropriately balances accuracy with computational efficiency is an ongoing challenge. We propose an algorithm for learning a single, multi-purpose adaptive mesh refinement (AMR) policy that produces a family of Pareto optimal solutions. In particular, we adapt a generic multi-objective Pareto front learning algorithm to the context of reinforcement learning for AMR. We demonstrate that this approach is considerably faster than training individual policies for several different target errors. Our approach is targeted at hp refinement techniques for ellip-