

## SIAM Conference on Computational Science

and  
Engineering



February 27-March 3, 2017  
Hilton Atlanta, Atlanta, Georgia, USA

Sponsored by the SIAM Activity Group on  
Computational Science and Engineering (CSE)

The SIAM Activity Group on Computational Science and Engineering (CSE) fosters collaboration and interaction among applied mathematicians, computer scientists, domain scientists and engineers in those areas of research related to the theory, development, and use of computational technologies for the solution of important problems in science and engineering. The activity group promotes computational science and engineering as an academic discipline and promotes simulation as a mode of scientific discovery on the same level as theory and experiment.

The activity group organizes the biennial SIAM Conference on Computational Science and Engineering and maintains a wiki, a membership directory, and an electronic mailing list. The SIAG recently established the SIAG/CSE Early Career Prize.

A new status report, Research and Education in Computational Science and Engineering, is available as a preprint from the SIAG CSE Wiki or directly from <https://arxiv.org/abs/1610.02608>. The article summarizes the status of CSE as an emerging discipline and presents the trends and challenges for the field.



### SIAM 2017 Events Mobile App

Scan the QR code with any QR reader and download the TripBuilder EventMobile™ app to your iPhone, iPad, iTouch or Android mobile device. You can also visit [www.tripbuildermedia.com/apps/siam2017](http://www.tripbuildermedia.com/apps/siam2017)



Society for Industrial and Applied Mathematics

3600 Market Street, 6th Floor

Philadelphia, PA 19104-2688 USA

Telephone: +1-215-382-9800 Fax: +1-215-386-7999

Conference E-mail: [meetings@siam.org](mailto:meetings@siam.org)

Conference Web: [www.siam.org/meetings/](http://www.siam.org/meetings/)

Membership and Customer Service:

(800) 447-7426 (USA & Canada) or +1-215-382-9800 (worldwide)

[www.siam.org/meetings/cse17](http://www.siam.org/meetings/cse17)

**IP1****Algorithmic Adaptations to Extreme Scale**

Algorithmic adaptations to use next-generation computers closer to their potential are underway. Instead of squeezing out flops the traditional goal of algorithmic optimality, which once served as a reasonable proxy for all associated costs algorithms must now squeeze synchronizations, memory, and data transfers, while extra flops on locally cached data represent only small costs in time and energy. After decades of programming model stability with bulk synchronous processing, new programming models and new algorithmic capabilities (to make forays into, e.g., data assimilation, inverse problems, and uncertainty quantification) must be co-designed with the hardware. We briefly recap the architectural constraints, then concentrate on two kernels that each occupy a large portion of all scientific computing cycles: large dense symmetric/Hermitian systems (covariances, Hamiltonians, Hessians, Schur complements) and large sparse Poisson/Helmholtz systems (solids, fluids, electromagnetism, radiation diffusion, gravitation). We examine progress in porting solvers for these kernels to the hybrid distributed-shared programming environment, including the GPU and the MIC architectures that make up the cores of the top scientific systems on the floor and on the books. How will the hierarchical solvers that lead in scalability (e.g., fast multipole, hierarchically low rank matrices, multigrid) map onto the more rigidly programmed and less reliably performant structures within a node?

David E. Keyes  
KAUST  
david.keyes@kaust.edu.sa

**IP2****Challenges for Climate and Weather Prediction in the Era of Exascale Computer Architectures: Oscillatory Stiffness, Time-Parallelism, and the Role of Long-Time Dynamics**

For weather or climate models to achieve exascale performance on next-generation heterogeneous computer architectures they will be required to exploit on the order of hundred-million-way parallelism. This degree of parallelism far exceeds anything possible in today's models even though they are highly optimized. In this talk I will discuss one of the mathematical issues that leads to the limitations in space- and time-parallelism for climate and weather prediction models oscillatory stiffness in the PDE of the form:

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{1}{\epsilon} L(\mathbf{u}) + N(\mathbf{u}, \mathbf{u}) = D(\mathbf{u}), \quad \mathbf{u}(0) = \mathbf{u}_0,$$

where the linear operator  $L$  has pure imaginary eigenvalues, the quadratic nonlinear term is  $N(\mathbf{u}, \mathbf{u})$   $D$  represents dissipation. The  $\epsilon$  is a small non-dimensional parameter. The operator  $\epsilon^{-1}L$  results in time oscillations on an order  $\mathcal{O}(\epsilon)$  time scale, and generally requires small time steps for explicit methods, and implicit methods if accuracy is required. I will discuss the case when  $\epsilon$  is finite and two algorithms: 1) a fast-converging HMM-parareal-type method and 2) a time-parallel matrix exponential.

Beth Wingate  
University of Exeter  
B.Wingate@exeter.ac.uk

**IP3****Ingredients for Computationally Efficient Solution****of Large-Scale Image Reconstruction Problems**

Image reconstruction problems provide great opportunities to bring together many scientific computing techniques to advance the state of the art in inverse problems and in disciplinary areas. First, we need to tailor the inverse problem to the application to produce the regularized solution while considering how a priori knowledge is enforced. We might choose to enforce hard constraints, like non-negativity, sparsity and/or high-contrast. Alternatively, we may choose to employ a learned dictionary or a parameterized image model that imposes those constraints directly on image space and simultaneously restrict the search space. Second, those modeling choices, which are interesting problems in and of themselves, necessitate the use of sophisticated optimization algorithms. Third, since each optimization step requires multiple forward model simulations, techniques from the multigrid, model reduction and randomization communities need to be explored to attain the maximum computational efficiency. In this talk, we provide an overview of some of these scientific computing techniques that have been successfully used in image reconstruction, and provide some combinations of techniques that have led to particularly fruitful outcomes in the context of a few applications.

Misha E. Kilmer  
Mathematics Department  
Tufts University  
misha.kilmer@tufts.edu

**IP4****Making Sense of our Universe with Supercomputers**

In computational cosmology and astrophysics we encounter some of the most complex multi-scale and multi-physics problems. In the past decades, algorithmic advances have enabled ever more realistic numerical models of a very wide range of astrophysical objects. These range from stars to galaxies, from planets to the large scale structure of the Universe, from molecular clouds to star clusters, from supernovae explosions to super-massive black holes in centers of galaxies. We routinely create three dimensional models of how our Universe may have originated, how its structure formed, how the very first stars and galaxies came about, how pulsars work, and how black holes merge and generate gravitational waves to just name a few such applications. We will highlight some examples of three particular algorithmic breakthroughs and the particular advances and insights they have enabled so far. These describe adaptive mesh refinement simulations capturing 15 orders of magnitude in length scale, adaptive ray tracing for high accuracy radiation hydro-dynamical simulations, as well as a new noise-free approach to solve the collision-less Boltzmann equation of interest in cosmology as well as in plasma physics. We will also present the scientific visualizations created from these simulations. These have been shown on various television programs, international planetarium shows and numerous print media.

Tom Abel  
Stanford University  
Hi@TomAbel.com

**IP5****Multiscale Modelling: Powerful Tool or Too Many**

## Promises

Multiscale modelling aims to provide systematic linking of different time and length scales in order to enable simulations at different levels of spatial and/or temporal resolutions. No single unique method exists or is even foreseeable and hence the choice of the most appropriate method or mapping depends on the properties of interest. The roots of multiscale simulations go back to the 1960s and 1970s, but the last decade has been brought them in the mainstream of method development and as a viable approach to model complex systems. It has been applied to viral capsids, fluid flow, crystal growth, proteins, colloids, and polymers to mention a few examples. Current techniques range from pragmatic, such as using solubilities for force field mapping to algorithmic, using equilibrium particle data for solving an inverse problem (using methods such as Inverse Boltzmann, force matching and Inverse Monte Carlo), particle-continuum coupling and using Langevin and Fokker-Planck equations and mapping procedure. Procedures such as GENERIC also extend the multiscale approach to irreversible processes. In this talk, I will discuss multiscale methods from the perspective of soft materials based on our own work, provide perspectives for future development and problems involving multiscale.

Mikko Karttunen

Eindhoven University, The Netherlands  
mkarttu@gmail.com

## IP6

### Adaptive and Multiscale Methods in Subsurface Flow Modeling

In reservoir simulation, the ratio of the largest scale to the smallest scale is typically very large. The smaller scales can not easily be upscaled or neglected as they can have (significant) impact on reservoir flow. At the same time, decision making under uncertainty requires fast and accurate flow simulations for often large ensembles of model realizations. No wonder then that the reservoir simulation community has developed many adaptive as well as multiscale methods to reduce runtime. Especially the last decade has seen a number of exciting new approaches, which I will discuss here, as well as the outstanding challenges that remain.

Margot Gerritsen

Dept of Energy Resources Engineering  
Stanford University  
margot.gerritsen@stanford.edu

## IP7

### Productive and Sustainable: More Effective Computational Science and Engineering

Computational Science and Engineering (CSE) is effective to the extent it contributes to overall scientific and engineering pursuits. Its contributions are most tangible when delivering concrete scientific and engineering results via modeling, simulation and analysis. At the same time, delivery of CSE results is impacted by how we develop and support the ecosystem that produced these results, including, in particular, software and people. While delivering results is the ultimate goal of our CSE efforts, the shortest path to results is often not the most productive and sustainable. In this presentation we discuss elements that impact the effectiveness of CSE efforts, beyond just the direct activities to produce results. We discuss how processes, tools and a holistic view of efforts can lead to more

effective CSE. We also discuss the importance of human factors in CSE activities, highlighting ways we can provide natural incentives toward more effective CSE.

Michael Heroux

Sandia National Laboratories  
maherou@sandia.gov

## IP8

### Computational Science and Engineering Achievements in the Designing of Aircraft

This presentation will give an overview of what Computational Science and Engineering has brought in design capacities these last years in the aeronautics industry. The unceasing increase in computing resources and remarkable improvements of solving methods have enabled industry to perform computations that were not conceivable several before. An emphasis will be scoped to optimization methods as actual engineering tools utilized for industrial applications, in particular for aerodynamic shape design. Numerical formulation and implementation issues will be recalled and illustrations of applications will be discussed. The study of efficient multidisciplinary approaches will be also addressed. New field of applications of Computational Science and Engineering have emerged these last years. Stochastic methods are in the process to take more and more an important place in the toolset of the design engineer and beyond. Some examples of application will be given. The presentation will end with the challenges related to Computational Science and Engineering for aeronautical industry.

Bruno Stoufflet

Dassault Aviation, France  
Bruno.Stoufflet@dassault-aviation.com

## IP9

### Stochastic Gradient Methods for Machine Learning

The stochastic gradient method has emerged as the most powerful technique for training the large-scale statistical models that form the foundation of modern machine learning systems. This talk provides an accessible introduction to the mathematical properties of stochastic gradient methods, and the intuition behind them. To set the stage, we present two case studies, one involving sparse linear models for text classification and one involving deep neural networks for image recognition. We then discuss advanced algorithms that control noise and make use of second order information. We conclude the talk with a discussion of the geometry of deep neural networks.

Jorge Nocedal

Department of Electrical and Computer Engineering  
Northwestern University  
j-nocedal@northwestern.edu

## CP1

### Theories and Algorithms of Integrated Singular Value Decomposition (iSVD)

The singular value decomposition (SVD) is an important tool in many applications. However, the computational cost of traditional algorithms for solving SVD grows rapidly as the data size increasing. Randomized SVD provides a method to randomly sketch a matrix and find its approximate low-rank SVD with lesser resources. Some

techniques can be used to improve the accuracy of a single sketching. Instead of focusing on one single sketching, iSVD is an algorithm to improve the accuracy by integrating multiple simple random sketching. The main idea of the proposed algorithm is solving a restricted optimization problem for a suitable objective function. The Kolmogorov-Nagumo-type average is used to solve this optimization problem. Some numerical results, including the application on 1000 Genomes data, will also be presented.

Ting-Li Chen  
Institute of Statistical Science  
Academia Sinica  
tlchen@stat.sinica.edu.tw

Dawei D. Chang  
Institute of Applied Mathematical Sciences  
National Taiwan University  
davidzan830@gmail.com

Su-Yun Huang  
Institute of Statistical Science  
Academia Sinica  
syhuang@stat.sinica.edu.tw

Hung Chen  
Institute of Applied Mathematical Sciences  
National Taiwan University  
hchen@math.ntu.edu.tw

Chienyao Lin  
Institute of Statistical Science  
Academia Sinica  
youyuoims94@gmail.com

Weichung Wang  
National Taiwan University  
Institute of Applied Mathematical Sciences  
wwang@ntu.edu.tw

## CP1

### An Efficient Iterative Penalization Method Based on Recycled Krylov Subspaces and Its Application to Impulsively Started Flows

The Vortex Particle-Mesh (VPM) method is well suited for solving advection dominated incompressible flows. However, the efficient and accurate handling of solid boundaries in this method is still an active domain of research. The Brinkman penalization method embeds the object in the fluid domain and enforces the velocity inside the obstacle to be  $\mathbf{u} = \mathbf{u}_b$ , where  $\mathbf{u}_b$  is the desired velocity. This additional constraint is added to the vorticity form of the incompressible Navier-Stokes equations through a Lagrange relaxation. The boundary enforcement conditions the capture of vorticity production at the wall and is thus paramount to the accuracy of the global method. Hejlesen et al. proposed to first solve the unpenalized Navier-Stokes equations and then, to impose the constraint using a Jacobi-like iterative process. In this work, we formulate the penalization problem inside a VPM method as a linear system to solve at every time step. Recovering the velocity from the vorticity (i.e. solving a Poisson problem) makes the matrix-vector product highly expensive. We use a recycling iterative solver, rBiCGStab, to solve it in order to reduce the number of iterations. This method is validated against a benchmark flow past a cylinder ( $\mathit{Re} = 550$ ) and then, we assess the computational gain with a flow past a cylinder and a plate ( $\mathit{Re} = 9500$

and  $\mathit{Re} = 1000$ , respectively).

Thomas Gillis  
Universite catholique de Louvain  
Institute of Mechanics, Materials and Civil Engineering  
thomas.gillis@uclouvain.be

## CP1

### Preconditioning Irk Methods for Time-Dependent Fluid Flow Problems

We examine block preconditioners for time-dependent incompressible Navier-Stokes problems. In some time-dependent problems, explicit time stepping methods can require much smaller time steps for stability than are needed for reasonable accuracy. This leads to taking many more time steps than would otherwise be needed. With implicit time stepping methods, we can take larger steps, but at the price of needing to solve large linear systems at each time step. We consider implicit Runge-Kutta (IRK) methods. Suppose our PDE has been linearized and discretized with  $N$  degrees of freedom. Using an  $s$ -stage IRK method leads to an  $sN \times sN$  linear system that must be solve at each time step. These linear systems are block  $s \times s$  systems, where each block is  $N \times N$ . We investigate preconditioners for such systems, taking advantage of the structure of the subblocks.

Victoria Howle  
Texas Tech  
victoria.howle@ttu.edu

## CP1

### Polynomial Preconditioned Arnoldi for Eigenvalues

Polynomial preconditioning has been explored for Krylov methods but has not become standard. In this talk, we look at the Arnoldi method for eigenvalues and give a simple choice for the polynomial preconditioner. It is shown that this approach can significantly improve the efficiency for difficult problems.

Jennifer A. Loe  
Baylor University  
jennifer\_loe@baylor.edu

Ronald Morgan  
Department of Mathematics  
Baylor University  
ronald\_morgan@baylor.edu

## CP1

### Linear Equations and Eigenvalues Using Krylov Methods on Multiple Grid Levels

We wish to solve large systems of linear equations and large eigenvalue problems. We try to combine the efficiency of using coarse grids with the power of Krylov subspaces. For linear equations, this involves a two-grid approach that deflates eigenvalues on the fine grid using approximate eigenvectors computed on the coarse grid. For eigenvalue problems, eigenvectors from coarse grids can be improved on finer grids using a variant of restarted Arnoldi. These methods are more robust than standard multigrid and sometimes are much more efficient than standard Krylov methods.

Ronald Morgan

Department of Mathematics  
Baylor University  
ronald\_morgan@baylor.edu

### CP1

#### A Factored ADI Method for Sylvester Equations with High Rank Right Hand Sides

The factored alternating direction implicit (ADI) method is a technique used to solve Sylvester equations of the form  $AX - XB = MN^T$ , where  $M$  and  $N$  are tall-and-skinny matrices. In this talk, we develop a variation of the ADI method that improves performance when  $MN^T$  is of medium to high rank, specifically for the application of solving elliptic partial differential equations. In particular, we employ it to solve elliptic partial differential equations on the disk expressed in a low rank format.

Heather D. Wilber  
Boise State University  
hdw27@cornell.edu

Alex Townsend  
Cornell University  
townsend@cornell.edu

### CP1

#### Parallel Implementations of Integrated Singular Value Decomposition (iSVD)

Integrated Singular Value Decomposition (iSVD) is an algorithm for computing low-rank approximate singular value decomposition of large size matrices. The iSVD integrates different low-rank SVDs obtained by multiple random subspace sketches and achieve higher accuracy and better stability. While iSVD takes higher computational costs due to multiple random sketches and the integration process, these operations can be parallelized to save computational time. We parallelize the algorithm for multicore/manycore hybrid CPU-GPU clusters. We propose algorithms and data structures to increase the scalability and reduce communication. With parallelization, iSVD can solve matrices with huge size, and achieve near-linear scalability with respect to the matrix size and the number of machines. We implement the algorithms in C++, with several techniques used so that many tuning decisions can be determined at compile time to reduce run-time overhead. Some huge size examples will be presented to show the performance of the implementation.

Mu Yang  
Institute of Applied Mathematical Sciences, National Taiwan  
muyang@ntu.edu.tw

Su-Yun Huang, Ting-Li Chen  
Institute of Statistical Science  
Academia Sinica  
syhuang@stat.sinica.edu.tw, tlchen@stat.sinica.edu.tw

Weichung Wang  
National Taiwan University  
Institute of Applied Mathematical Sciences  
wwang@ntu.edu.tw

### CP1

#### A Fast Direct Solver for Elliptic PDEs on Locally-

### Perturbed Domains

Many problems in science and engineering can be formulated as integral equations with elliptic kernels. In particular, in optimal control and design problems, the domain geometry evolves and results in a sequence of discretized linear systems to be constructed and inverted. While the systems can be constructed and inverted independently, the computational cost is relatively high. In the case where the change in the domain geometry for each new problem is only local, i.e. the geometry remains the same except within a small subdomain, we are able to reduce the cost of inverting the new system by reusing the pre-computed fast direct solvers of the original system. The resulting solver only requires inexpensive matrix-vector multiplications, thus dramatically reducing the cost of inverting the new linear system. Numerical results will illustrate the performance of the solver.

Yabin Zhang  
Dept. of Computational and Applied Mathematics  
Rice University, Houston, TX  
YZ89@RICE.EDU

Adrianna Gillman  
Rice University  
Department of Computational and applied mathematics  
adrianna.gillman@rice.edu

### CP2

#### Hierarchical Model Reduction for Incompressible Flows in Pipes

The Hierarchical Model Reduction (HiMod) is a novel technique for the efficient solution of fluid problems in pipes that joins computational efficiency with numerical accuracy. According to this method, the transverse dynamics is represented in terms of a generalized Fourier modal expansion, whose axial dependence is discretized via a Finite Element Method. In such a way the original problem results as a system of coupled 1D problems. The power of this technique lies in its hierarchical nature, as the accuracy can be tuned by a proper selection of the number of transverse modes. Aiming at real medical applications, we apply HiMod to scalar (Advection-Diffusion Reaction models) and vector problems (incompressible Navier-Stokes equations) on 3D cylindrical domains. The application of the technique to such geometries is non-trivial, especially in the respect of the identification of a basis function set. We numerically assess different options. Patient-specific blood-vessels geometries are handled by appropriate geometrical maps. Numerical tests point out the capability of HiMod to detect the transverse dynamics of the physical phenomenon, as opposed to other approaches that rely on averaging the transverse dynamics in a purely 1D setting. Our method tries to find a practical trade-off between the accuracy of 3D modeling and the efficiency of 1D modeling.

Sofia Guzzetti, Alessandro Veneziani  
Department of Mathematics and Computer Science  
Emory University  
sofia.guzzetti@emory.edu, avenez2@emory.edu

Simona Perotto  
MOX - Modeling and Scientific Computing  
Dipartimento di Matematica

simona.perotto@polimi.it

## CP2

### Inertial Confinement Fusion Simulations Using a Front Tracking API

The Stony Brook front tracking code FronTier has been extracted into an Application-Programming Interface (API) for easy implementation into external physics codes. Front tracking is a well validated algorithm which shows improved accuracy relative to experiment for fluid mixing applications. In this talk we detail the first use of this API through implementation in the University of Chicago code FLASH. We detail the process required for implementation and discuss the benefits of the coupling of a front tracking algorithm for fluid mixing problems such as Rayleigh-Taylor and Richtmyer-Meshkov instabilities. Our main application is the use of the front tracking API for the simulation of Inertial Confinement Fusion (ICF) capsules. We present 2D simulations in a spherical geometry and discuss the impact of front tracking on ICF simulations. We show that for coarser grids, the front tracking simulations are closer to a converged result, a key requirement for the heavy computational requirements associated with ICF simulations.

Jeremy A. Melvin  
University of Texas - Austin  
jmelvin@ices.utexas.edu

Verinder Rana  
Stony Brook University  
Stony Brook NY  
vrana@ams.sunysb.edu

Ryan Kaufman  
SUNY at Stony Brook  
rkaufman@ams.sunysb.edu

James Glimm  
Stony Brook University  
glimm@ams.sunysb.edu

## CP2

### Conservative DEC Discretization of Incompressible Navier-Stokes Equations on Arbitrary Surface Simplicial Meshes With Applications

A conservative discretization of incompressible Navier-Stokes equations over surfaces is developed using discrete exterior calculus (DEC). The mimetic character of many of the DEC operators provides exact conservation of both mass and vorticity, in addition to superior kinetic energy conservation. The employment of various discrete Hodge star definitions based on both circumcentric and barycentric dual meshes is also demonstrated. Some of the investigated definitions allow the discretization to admit arbitrary surface simplicial meshes instead of being limited only to Delaunay meshes, as in previous DEC-based discretizations. The discretization scheme is presented in detail along with several numerical test cases demonstrating its numerical convergence and conservation properties. The developed scheme is also applied to explore the curvature effects on flow past a circular cylinder.

Mamdouh S. Mohamed  
Physical Sciences & Engineering Division, KAUST,  
Jeddah, KSA  
mamdouh.mohamed@kaust.edu.sa

Anil Hirani  
Department of Mathematics  
University of Illinois at Urbana-Champaign, IL, USA  
hirani@illinois.edu

Ravi Samtaney  
KAUST  
ravi.samtaney@kaust.edu.sa

## CP2

### Stability of Oscillatory Rotating Boundary Layers

Some of the most popular applications of fluid mechanics come in aerodynamics, and methods of laminar flow control on swept wings have become increasingly important over the past few decades; especially with the global emphasis on emissions reduction. Another reason for the recent development of this field is the availability of high-performance computers, meaning calculations that would have been impossible only a decade ago can now be performed quickly on a workstation. With instability mechanisms in common with a swept wing, the rotating disk has long been considered as a valid approximation to this flow and is also far more amenable to experiments. The experimental setup for a rotating disk study requires a much smaller space and much less expensive equipment than the wind tunnel required for a swept wing experiment. For this reason, there are a wealth of experimental and theoretical studies of the rotating disk boundary layer and this talk will extend these established results. A recent study by Thomas et. al. [Proc. R. Soc. A (2011) 467, 2643-2662] discusses adding an oscillatory Stokes layer to a channel flow and shows some stabilising results. We present a similar modification to the rotating disk configuration by way of periodic modulation and provide results from both direct numerical simulations and local eigenvalue analyses showing a stabilising effect.

Scott N. Morgan  
Cardiff University  
MorganSN@cardiff.ac.uk

## CP2

### Marker Re-Distancing and Sharp Reconstruction for High-Order Multi-Material Front Dynamics

A new method for high-order front evolution on arbitrary meshes is introduced. The method is a hybrid of a Lagrangian marker tracking with a Discontinuous Galerkin projection based level set re-distancing. This Marker-Re-Distancing (MRD) method is designed to work accurately and robustly on unstructured, generally highly distorted meshes, necessitated by applications within ALE-based hydro-codes. Since no PDE is solved for re-distancing, the method does not have stability time step restrictions, which is particularly useful in combination with AMR, used here to efficiently resolve fine interface features. A high-order (implemented up to the 6th-order) level set approach is utilized for a new sharp treatment of mix elements, which reconstructs multiple-per-element solution fields (one for each material present in the mix element). Reconstruction incorporates interfacial jump conditions, which are enforced in the least-squares sense at the interfacial marker positions provided by MRD. Since no explicit differentiation across the interface is involved in the assembly of residuals for mass, momentum and energy equations, the method is capable of capturing discontinuous solutions at multi-material interfaces with high order, and without Gibbs oscillations. The method performance is

demonstrated on a number of numerical tests, including well-known benchmarks, and phase-change fluid dynamics problems relevant to the selective laser melting applications.

Robert Nourgaliev, Patrick Greene, Sam Schofield  
LLNL  
nourgaliev1@llnl.gov,                      greene30@llnl.gov,  
schofield5@llnl.gov

## CP2

### Direct Computations of Marangoni-Induced Flows Using a Volume of Fluid Method

The volume of fluid (VOF) interface tracking methods have been used for simulating a wide range of interfacial flows. An improved accuracy of the surface tension force computation has enabled the VOF method to become widely used for simulating surface tension driven flows. We present a new method for including variable surface tension in a VOF based Navier-Stokes solver. The tangential gradient of the surface tension is implemented using an extension of the classical continuum surface force model that has been previously used for constant surface tension simulations. Our method can be applied for computing the surface gradients of surface tension that is temperature or concentration dependent.

Ivana Seric, Shahriar Afkhami  
New Jersey Institute of Technology  
is28@njit.edu, shahriar@njit.edu

Lou Kondic  
Department of Mathematical Sciences, NJIT  
University Heights, Newark, NJ 07102  
kondic@njit.edu

## CP2

### Physics-Based Preconditioning for a High-Order Rdg-Based Compressible Flow Solver with Phase Change

The numerical simulation of flows associated with metal additive manufacturing processes such as selective laser melting and other laser-induced phase change applications present new challenges. Specifically, these flows require a fully compressible formulation with phase change, since rapid density variations occur due to laser-induced melting and solidification of metal powder.

We investigate the preconditioning for a recently developed all-speed compressible Navier-Stokes solver that addresses such challenges. The equations are discretized with a reconstructed Discontinuous Galerkin method and integrated in time with fully implicit discretization schemes. The resulting set of non-linear and linear equations are solved with a robust Newton-Krylov (NK) framework.

To enable convergence of the highly ill-conditioned linearized systems, we employ a physics-based operator split preconditioner (PBP), which utilizes a robust Schur complement matrix for the velocity-pressure and velocity-temperature block-systems. We investigate different options of splitting the physics (field) blocks as well as different iterative solvers to approximate the action of the inverse of the preconditioned system. We demonstrate that our PBP-NK framework is scalable and converges for high CFL/Fourier numbers on classic problems in fluid dynamics as well as for laser-induced phase change problems.

Brian Weston

University of California, Davis  
Lawrence Livermore National Laboratory  
brianweston@gmail.com

Robert Nourgaliev  
LLNL  
nourgaliev1@llnl.gov

Jean-Pierre Delplanque  
University of California, Davis  
delplanque@ucdavis.edu

## CP3

### Multigrid Preconditioned Lattice Boltzmann Method Based on Central Moments for Efficient Computation of Fluid Flows

Lattice Boltzmann (LB) Method is one of the more recent promising developments in computational fluid dynamics (CFD) for based on a local kinetic approach flow simulations. Like other classical explicit time-marching methods, the LBM can suffer from slow convergence rate to steady state, especially at low Mach numbers. This is due to the relatively large disparity between the acoustic wave and fluid convection speeds, i.e. high eigenvalue stiffness, which can be alleviated by preconditioning the LB method. Furthermore, we combine such a preconditioned LB scheme based on the efficient stream-and-collide procedure with the multigrid method based on the nonlinear full approximation storage scheme to provide convergence acceleration at an optimal cost. The collision step is formulated using central moments to provide enhanced numerical stability. In addition, we develop consistent inter-grid transfer operators based on using extended moment equilibria in the collision step involving a tunable parameter that keeps the flow properties such as viscosities, and hence flow physics, invariant across different grid levels. Finally, we validate our new multigrid preconditioned LB method for various benchmark problems and then demonstrate the significant improvements achieved in convergence acceleration, by factors of one or more orders of magnitude, for various sets of the preconditioning parameters and Reynolds numbers and Mach numbers.

Farzaneh Hajabdollahi, Kannan Premnath  
University of Colorado Denver  
farzaneh.hajabdollahi@ucdenver.edu,                      kan-  
nan.premnath@ucdenver.edu

## CP3

### Multigrid Preconditioning for Space-Time Distributed Optimal Control of Parabolic Equations

This work is concerned with designing optimal order multigrid preconditioners for space-time distributed control of parabolic equations. The focus is on the reduced problem resulted from eliminating state and adjoint variables from the KKT system. Earlier numerical experiments have shown that our ability to design optimal order preconditioners depends strongly on the discretization of the parabolic equation, with several natural discretizations leading to suboptimal preconditioners. Using a continuous-in-space-discontinuous-in-time Galerkin discretization we obtain the desired optimality.

Mona Hajghassem, Andrei Draganescu  
Department of Mathematics and Statistics  
University of Maryland, Baltimore County

mona4@umbc.edu, draga@umbc.edu

### CP3

#### Spectral Matrix Analysis of the Semi-Discrete Compressible Navier-Stokes Equations Using Large-Scale Eigensolvers

Implicit integration methods for the compressible Navier-Stokes equations rely on iterative methods for sparse linear systems. Code performance depends upon the physics, models, boundary conditions, and numerical methods. At large scale and high order of accuracy, performance is largely determined by the linear solver which in turn depends upon the base iterative method, preconditioner selection, and eigenstructure of the linear problems constructed through spatial discretization. To greater understand the impact of algorithm choices on linear solver and preconditioner performance and design, we use large-scale eigensolvers to obtain a partial spectrum (collection of eigenvalues) of the discretization matrix. We perform this analysis on a high-order entropy-stable spectral collocation method for the compressible Navier-Stokes equations on laminar and turbulent problems.

Michael A. Hansen

University of Utah  
Sandia National Laboratories  
mike.hansen.utah@gmail.com

Travis Fisher  
Sandia National Labs  
tcfishe@sandia.gov

### CP3

#### Block Triangular Preconditioners for Linearization Schemes of the Rayleigh-Bénard Convection Problem

In this work, we compare two block triangular preconditioners for different linearizations of the Rayleigh-Bénard convection problem discretized with finite element methods. The two preconditioners differ in the nested or non-nested use of a certain approximation of the Schur complement associated to the Navier-Stokes block. First, bounds on the generalized eigenvalues are obtained for the preconditioned systems linearized with both Picard and Newton methods. Then, the performance of the proposed preconditioners is studied in terms of computational time. This investigation reveals some inconsistencies in the literature that are hereby discussed. We observe that the non-nested preconditioner works best both for the Picard and for the Newton cases. Therefore, we further investigate its performance by extending its application to a mixed Picard-Newton scheme. Numerical results of two- and three-dimensional cases show that the convergence is robust with respect to the mesh size. We also give a characterization of the performance of the various preconditioned linearization schemes in terms of the Rayleigh number.

Guoyi Ke  
Texas Tech University  
guoyi.ke@ttu.edu

Eugenio Aulisa  
Department of Mathematics and Statistics.  
Texas Tech University  
eugenio.aulisa@ttu.edu

Giorgio Bornia

Dept. of Mathematics and Statistics Texas Tech  
University  
giorgio.bornia@ttu.edu

Victoria Howle  
Texas Tech  
victoria.howle@ttu.edu

### CP3

#### Comparison of Techniques for Hermitian Interior Eigenvalue Problems: Refined, Harmonic and Polynomial Filters

Polynomial and rational polynomial filtered methods have gained popularity for large interior eigenproblems. FEAST-like solvers exhibit good performance when factorizing the operators is affordable. For larger problems, recent research in Chebyshev polynomials and inexact FEAST have proposed heuristics to tune performance sensitive parameters such as the selection of the contour points, the polynomial degree, or the tolerance for the approximate solution of the linear systems. However they are still hard to apply in black-box solvers. We present a modification of the refined extraction technique of Zhongxiao Jia, that is at least as robust but far less expensive, especially when using block methods. The new extraction in combination with Jacobi-Davidson has shown good results in the context of solving difficult interior eigenvalue problems arising in the computation of singular triplets. An important advantage of Jacobi-Davidson is that can be seen as a rational polynomial filtered method, which offers a natural way to adjust dynamically the shifts and the tolerance to solve linear systems. In this talk we provide experimental results with Hermitian problems comparing: Jacobi-Davidson with the new refined extraction as included in the last version of PRIMME, 2.0; polynomial filtered Lanczos in the recent released EVSL; and inexact solution of linear solvers with the FEAST software.

Eloy Romero Alcalde  
Computer Science Department  
College of Williams & Mary  
eloy@cs.wm.edu

Andreas Stathopoulos  
College of William & Mary  
Department of Computer Science  
andreas@cs.wm.edu

### CP3

#### Fast Algorithms for Jacobi Matrices from Modification by Rational Functions

Given a Jacobi matrix for a sequence of orthogonal polynomials with respect to some measure  $d\lambda(x)$ , the goal of this project is to generate the Jacobi matrix for modification of that measure  $d\tilde{\lambda}(x)$  obtained from  $d\lambda(x)$  by multiplying by a rational function. Through partial fraction decomposition, this amounts to modification by dividing by linear or irreducible quadratic factors. The proposed method reverses the algorithm for modifying by multiplying factors, due to Golub and Kautsky, combined with a root finding iteration such as the secant method. The entire process requires only  $O(n)$  floating point operations compared to  $O(n^3)$  floating point operations for the inverse Cholesky algorithm due to Elhay and Kautsky. One application is to obtain Jacobi matrices for generalized Jacobi polynomials.

Amber C. Sumner



The University of Southern Mississippi  
 amber.sumner@usm.edu

James V. Lambers  
 University of Southern Mississippi  
 Department of Mathematics  
 James.Lambers@usm.edu

### CP3

#### Estimating Matrix Bilinear and Quadratic Forms Using Krylov Subspace Methods with Recycling

Matrix bilinear forms  $C^T A^{-1} B$ , for a nonsymmetric matrix  $A$ , and quadratic forms  $B^T A^{-1} B$  or its trace, for symmetric positive definite  $A$ , frequently appear in applications. We show how to evaluate or estimate these bilinear and quadratic forms accurately and cheaply using Krylov subspace methods with recycling. We demonstrate the effectiveness of our approach on three applications: functional error estimation and mesh adaptation in CFD, tomography, and topology optimization. This is joint work with Chris Roy and Will Tyson (VT) on CFD applications, Misha Kilmer (Tufts) on Tomography, and Xiaojia Zhang and Glaucio Paulino (Georgia Tech) on Topology Optimization.

Katarzyna Swirydowicz, Eric De Sturler  
 Virginia Tech  
 kswirydo@vt.edu, sturler@vt.edu

William Tyson, Christopher Roy  
 VT  
 wtyson45@vt.edu, cjroy@vt.edu

Misha E. Kilmer  
 Mathematics Department  
 Tufts University  
 misha.kilmer@tufts.edu

Xiaojia Zhang, Glaucio Paulino  
 Georgia Tech  
 xzhang645@gatech.edu, paulino@gatech.edu

### CP3

#### Why Are So Many Matrices in Computational Science of Low Rank?

In computational mathematics, matrices and functions that appear in practice are so often of surprisingly low rank and this structure is often expertly exploited. Since random (“average”) matrices are almost surely of full rank, mathematics needs to explain the remarkable abundance of low-rank structures in computational mathematics. In this talk, we will give a new characterization of low-rank matrices, which we use to explore why (1) Droplets on a pond, (2) Non-equally sampling of functions, (3) Elliptic partial differential equations and (4) The world flags, all lead to low-rank objects.

Alex Townsend  
 Cornell University  
 townsend@cornell.edu

Gil Strang  
 MIT

gilstrang@gmail.com

### CP4

#### Computing the Ankle-Brachial Index with Computational Fluid Dynamics

Peripheral artery disease (PAD), in which narrowing and blockage of peripheral arteries reduces blood flow to the extremities of the body, is associated with a six-fold increase in mortality risk from cardiovascular disease. PAD is diagnosed by computing the ankle-brachial index (ABI), a metric relating blood pressure in the ankles and upper arms. With parallel computing, we use 3D computational fluid dynamics to simulate flow in a complete, patient-derived arterial system and compute the ABI. The simulations employ a massively parallel CFD application, HARVEY, designed for large-scale hemodynamic simulations and based on the lattice Boltzmann method. Simulations were conducted on Vulcan, a Blue Gene/Q supercomputer at Lawrence Livermore National Laboratory with 393,216 cores. We consider the dependence of ABI on simulation resolution and find adequate numerical convergence at  $50\mu\text{m}$ . The influence of body posture on ABI is investigated by incorporating gravitational forces corresponding to supine and standing body positions. Additionally, we consider the influence of an aortic coarctation, which imposes the same hemodynamic compromise on the peripheral arteries as PAD, and observe the expected decrease in ABI.

John Gounley  
 Old Dominion University  
 john.gounley@duke.edu

Erik W. Draeger  
 Lawrence Livermore Nat. lab.  
 draeger1@llnl.gov

Jane Leopold  
 Brigham and Women’s Hospital  
 leopold@partners.org

Amanda Randles  
 Duke University  
 amanda.randles@duke.edu

### CP4

#### Dynamic Mesh Adaptation for Front Evolution Using Discontinuous Galerkin Based Weighted Condition Number Relaxation

A new mesh smoothing method designed to cluster cells near a dynamically evolving interface is presented. The method is based on weighted condition number mesh relaxation with the weight function being computed from a level set representation of the interface. The weight function is expressed as a Taylor series based discontinuous Galerkin (DG) projection, which makes the computation of the derivatives of the weight function needed during the condition number optimization process a trivial matter. For cases when a level set is not available, a fast method for generating a low-order level set from discrete cell-centered fields, such as a volume fraction or index function, is provided. Results show that the low-order level set works equally well for the weight function as the actual level set. The method retains the excellent smoothing capabilities of condition number relaxation, while providing a method for clustering mesh cells near regions of inter-

est. Dynamic cases for moving interfaces will show that the new method is capable of maintaining a desired resolution near the interface with an acceptable number of relaxation iterations per time step, which demonstrates the method's potential to be used as a mesh relaxer for arbitrary Lagrangian Eulerian (ALE) methods. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Information management release number LLNL-ABS-702781.

Patrick Greene, Samuel Schofield, Robert Nourgaliev  
LLNL  
greene30@llnl.gov, schofield5@llnl.gov, nourgaliev1@llnl.gov

#### CP4

##### A DPG Method for Viscoelastic Fluids

We propose discontinuous Petrov-Galerkin (DPG) finite element method for the steady-state Oldroyd-B equations. The method is attractive because of its built-in stability—that is, no stabilization terms need to be added to the system—as well as its built-in error estimator which can be used for adaptive mesh generation. Notable, as well, is that the method will always result in a symmetric, positive definite stiffness matrix. We exploit each of these properties in our analysis and perform verification upon a confined cylinder benchmark problem. The code was written in C++ and available online with the Camellia finite element software [1]. [1] Roberts, N. V. (2014). Camellia: A software framework for discontinuous PetrovGalerkin methods. *Comput. Math. Appl.*, 68(11):15811604.

Brendan Keith  
The University of Texas at Austin  
The Institute for Computational Engineering and Sciences  
brendan@ices.utexas.edu

Philipp Knechtges  
Chair for Computational Analysis of Technical Systems  
RWTH Aachen University  
knechtges@cats.rwth-aachen.de

Nathan Roberts  
Argonne National Laboratory  
nvroberts@alcf.anl.gov

Stefanie Elgeti  
RWTH Aachen University  
elgeti@cats.rwth-aachen.de

Marek Behr  
RWTH Aachen University  
Chair for Computational Analysis of Technical Systems  
behr@cats.rwth-aachen.de

Leszek Demkowicz  
Institute for Computational Engineering and Sciences (ICES)  
The University of Texas  
leszek@ices.utexas.edu

#### CP4

##### Large Time Step HLL and HLLC Schemes

We present the Large Time Step (LTS) extension of the Harten-Lax-van Leer (HLL) and Harten-Lax-van Leer Con-

tact (HLLC) schemes. The LTS methods are a class of explicit methods whose time step is not restricted by the classical Courant-Friedrichs-Lewy (CFL) condition. The basic idea is to increase the domain of dependence by modifying the numerical flux function, and allowing linear interaction of waves from different Riemann problems. Such schemes were developed by Randall LeVeque [R.J. LeVeque. Large time step shock-capturing techniques for scalar conservation laws, *SIAM J. Numer. Anal.*, 19:1091-1109, 1982] in the nineteen eighties. The original scheme and the successive versions have been developed and applied mostly within a framework of the Godunov scheme and Roe's approximate Riemann solver. We show that it is possible to construct the LTS extension of the HLL and HLLC schemes. We apply the LTS HLL and HLLC schemes to a number of test cases for inviscid gas dynamics, such as shock tube, double rarefaction and Woodward-Colella blast wave problem. It is shown that the schemes yield results comparable to those of standard and LTS Roe scheme. In addition, we show that both LTS HLL and HLLC schemes preserve positivity for a double rarefaction test case where standard Roe and LTS Roe schemes fail and that both schemes yield entropy satisfying resolution of the rarefaction waves.

Marin Prebeg, Bernhard Müller  
Norwegian University of Science and Technology  
marin.prebeg@ntnu.no, bernhard.muller(at)ntnu.no

#### CP5

##### Multigrid Kss Methods for Time-Dependent PDEs

Krylov Subspace Spectral (KSS) methods are traditionally used to solve time-dependent, variable-coefficient PDEs. Lambers, Cibotarica, and Palchak improved the efficiency of KSS methods by optimizing the computation of high-frequency components. This talk will demonstrate how one can make KSS methods even more efficient by using a multigrid-like approach for low-frequency components. The essential ingredients of multigrid, such as restriction, residual correction, and prolongation, are adapted to the time-dependent case. Numerical experiments demonstrate the effectiveness of this approach.

Haley Dozier  
The University of Southern Mississippi  
Haley.Dozier@usm.edu

James V. Lambers  
University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

#### CP5

##### Automatic Construction of Scalable Time-Stepping Methods for Stiff Pdes

Krylov Subspace Spectral (KSS) Methods have been demonstrated to be highly scalable time-stepping methods for stiff nonlinear PDEs. However, ensuring this scalability requires analytic computation of frequency-dependent quadrature nodes from the coefficients of the spatial differential operator. This talk describes how this process can be automated for various classes of differential operators to facilitate public-domain software implementation.

Vivian A. Montiforte  
The University of Southern Mississippi  
vivian.mclain@usm.edu

James V. Lambers  
University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

#### CP5

### Multigrid Preconditioning of Linear Systems Arising in the Semismooth Newton Solution of Distributed Optimal Control of Elliptic Equations with State Constraints

The purpose of this research is to design efficient multigrid preconditioners for distributed optimal control of elliptic equations. In this talk we focus on preconditioning of linear systems arising in the semismooth Newton solution of distributed control for elliptic equations with state-constraints. This research is building upon our earlier work on the associated control-constrained problems. Analytical and numerical results are presented.

Jyoti Saraswat  
Department of Mathematics & Physics  
Thomas More College  
saraswj@thomasmore.edu

Andrei Draganescu  
Department of Mathematics and Statistics  
University of Maryland, Baltimore County  
draga@umbc.edu

#### CP5

### Inexact Algebraic Factorization Methods for the Steady Incompressible Navier-Stokes Equations

The steady Navier-Stokes problem is of fundamental importance in science and engineering. Although there has been much research on the topic through the decades, the problem remains a challenging one due to its nonlinearity and the large size of its associated linear systems. Because of its importance and difficulty, new techniques for reducing the problems computational burden remain in demand. For the time-dependent problem, there exists a class of popular methods for solving the Navier-Stokes problem efficiently known as Inexact Algebraic Factorization Methods (see e.g. [A. Quarteroni, F. Saleri, and A. Veneziani, Factorization methods for the numerical approximation of Navier-Stokes equations, *Comp. Meth. in App. Mech.*, 188 (2000), pp. 505-526]). These methods work by approximating the saddle-point system with an inexact block LU factorization. These methods exhibit good accuracy and stability properties while significantly reducing the costs associated with solving, assembling, and storing the linear system. In this work, we extend these methods to the steady problem by showing that the stiffness matrix can be used as a suitable approximation to the nonlinear term under certain conditions. Numerical results in 2D and 3D are then discussed and presented.

Alex Viguerie  
Emory University  
aviguer@emory.edu

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

#### CP5

### Investigations of Several Mhd Solvers Based on

### Discontinuous Galerkin Finite Element Method

Compressible flow with magnetic phenomena occurs in a wide variety of scientific and engineering applications. The efficient parallel numerical simulation for these problems is very important. In this paper, several 1D and 2D Hydro-Dynamics (HD) and Magneto-Hydro-Dynamic (MHD) solvers have been developed based on Discontinuous Galerkin method on Finite Element method. Their performance has been investigated and compared. These include various Riemann solvers for capturing the shocks and contact waves. Furthermore new method has been used to capture the shocks, and it has been compared to above traditional solvers. These solvers are in parallel version, benchmark and simulations will be shown, and some conclusions will be drawn from them.

Xiaohu Zhufu  
Institute of Software, Chinese Academy of Science  
xiaohu@iscac.ac.cn

Yanfei Jiang  
Center for Astrophysics, Harvard University  
yanfei.jiang@cfa.harvard.edu

Zhaoming Gan, Defu Bu, Maochun Wu  
Shanghai Astronomical Observatory, Chinese Academy of Science  
zmgan@shao.ac.cn, dfbu@shao.ac.cn,  
maochun@ustc.edu.cn

Jin Xu  
Institute of Software, ISCAS  
Chinese Academy of Science, China  
xu\_jin@iscas.ac.cn

#### CP6

### DOF-Reducing Small-Lebesgue Polygonal Spectral Basis Functions with Application to Discontinuous FEM

A closed form relation is proposed to approximate Fekete points (a Nondeterministic Polynomial (NP) problem) on a general convex/concave polyhedral. The approximate points are used to generate basis functions using the SVD of the Vandermonde matrix. Arbitrary order nodal, orthogonal and orthonormal polygonal basis functions are derived. It is shown that these basis are the best choice to enforce minimum DOF while maintaining a small Lebesgue constant when very high p-refinement is done. The proposed basis are rigorously proven to achieve arbitrary order of accuracy by satisfying Weierstrass approximation theorem in  $\mathbb{R}^d$ . The practicality of these basis are evaluated in Discontinuous Galerkin (DG) and Discontinuous Least-Squares (DLS) formulations. The accuracy, stability and DOF reduction is demonstrated for the linearized acoustics and two-dimensional compressible Euler equations on some benchmark problems including a cylinder, airfoil, vortex convection and compressible vortex shedding from a triangle.

Arash Ghasemi  
SimCenter: National Center for Computational Engineering  
ghasemi.arash@gmail.com

#### CP6

### A Hybrid Adaptive Compressible/Low-Mach-

## Number Method

Flows in which the primary features of interest do not rely on high-frequency acoustic effects, but in which long-wavelength acoustics play a nontrivial role, present a computational challenge. Integrating the entire domain with low Mach number methods would remove all acoustic wave propagation, while integrating the entire domain with the fully compressible equations would be prohibitively expensive due to the CFL time step constraint. For example, thermoacoustic instabilities might require fine resolution of the fluid/chemistry interaction but not require fine resolution of acoustic effects, yet one does not want to neglect the long-wavelength wave propagation and its interaction with the larger domain. The proposed lecture will present a new hybrid algorithm that has been developed to address these type of phenomena. In this new approach, the fully compressible Navier-Stokes equations are solved on the entire domain, while their low Mach number counterparts are solved on a subregion of the domain with higher spatial resolution. The coarser acoustic grid communicates inhomogeneous divergence constraints to the finer low Mach number grid, so that the low Mach number method allows the long-wavelength acoustics. We will demonstrate the effectiveness of the new method on practical cases such as the aeroacoustics generated by the vortex formation in an unstable low-Mach mixing layer.

### Emmanuel Motheau

Lawrence Berkeley National Lab  
Center for Computational Sciences and Engineering  
emotheau@lbl.gov

Ann S. Almgren  
Lawrence Berkeley National Laboratory  
asalmgren@lbl.gov

John B. Bell  
CCSE  
Lawrence Berkeley Laboratory  
jbbell@lbl.gov

## CP6

### **An Arbitrary High Order Imex Scheme For Extended Magnetohydrodynamics Equations Using Entropy Conservative Flux**

XMHD is an extended plasma fluid model which assumes quasineutrality and differs from the existing reduced models by the formulation of Generalized Ohm's Law. Previous models for example Ideal MHD or Hall MHD are not sufficient to describe all the dynamics encountered in general plasma due to neglected terms in their derivation. In particular to capture low density current, inclusion of electron inertia terms is necessary. The Ohm's Law for XMHD has been modified from resistive MHD Ohm's Law through the introduction of electron pressure, electron inertia and Hall term allowing ion and electron demagnetization. In this work we proposed a finite difference scheme that uses an entropy conservative flux with an appropriate numerical diffusion operator for the simulation of the fluid part and rest of the system is treated with a Local Lax Friedrich flux. For the divergence constraints of Maxwell equation to be explicitly satisfied, correction potential for  $\mathbf{B}$  and  $\mathbf{E}$  has been enforced. The IMEX idea where source term is treated implicitly and flux is treated explicitly increases the efficiency of the scheme by reducing numerical cost since using the special structure of the source term we will show that we only need to solve a system of 9 linear equations at each grid point explicitly. The scheme performs well when

two fluid effects are important.

### Chhanda Sen

Indian Institute of Technology Delhi  
chhanda27sen@gmail.com

Harish Kumar  
Indian Institute of Technology Delhi, India  
hkumar[at]maths.iitd.ac.in

## CP6

### **Energy Conservation Moment Method to Solve the Multi-Dimensional Vlasov-Maxwell-Fokker-Planck Equations**

We present a numerical method to solve the Vlasov-Maxwell-Fokker-Planck (VMFP) system using the regularized moment method proposed in [Z. Cai, R. Li, Numerical regularized moment method of arbitrary order for Boltzmann-BGK equation, SIAM J. Sci. Comput (2010), Z. Cai, Y. Fan, R. Li, Globally hyperbolic regularization of Grad's moment system in one dimensional space, Comm. Math. Sci (2013)]. In [Y. Wang, S. Zhang, CICIP (accepted)], the globally moment system to the 1D Vlasov-Poisson-Fokker-Planck (VPFP) is deduced and the numerical scheme to keep the balance law of the total momentum is provided. However, it cannot keep the conservation of total energy. In this paper, the moment model of the electromagnetic field term is derived and then, we extend the global moment system to the multi-dimensional VMFP and VPFP systems, where the electromagnetic field term and the Fokker-Planck collision term are reduced into the linear combination of the moment coefficients. The Strang-splitting method is adopted to solve the whole moment system, which is splitted into the conservation part and the MFP part. Most importantly, a special semi-implicit numerical scheme which could keep the conservation of total energy is proposed to solve the Maxwell's equations at the MFP part. The time evolution of the solutions to both 2D VMFP and VPFP systems are studied to demonstrate the stability and accuracy of the regularized moment method when applied to the VMFP system.

### Yanli Wang

Institute of Applied Physics and Computational Mathematics  
wang\_yanli@iapcm.ac.cn

## CP7

### **Adjoint-Enabled Optimization and UQ for Radiation Shield Design**

Radiation shields make commodity microelectronics practical for use in satellite and other space systems. Shield designers wish to take advantage of new materials and manufacturing processes to meet strict weight limits while protecting electronics from naturally occurring proton and electron radiation environments. Our work couples Sandia National Laboratories Dakota software (<http://dakota.sandia.gov>) with its SCEPTRE radiation transport code to automate the design exploration and reliability analysis process, enabling analysts to evaluate prospective shield materials and geometries. This talk highlights efficiency gains from pairing gradient-based optimization and uncertainty quantification algorithms in Dakota with newly implemented adjoint sensitivities in SCEPTRE.

Brian M. Adams

Sandia National Laboratories  
 Optimization/Uncertainty Quantification  
 briadam@sandia.gov

### CP7

#### Ensemble Kalman Filtering for Inverse Optimal Control

Solving the inverse optimal control problem for discrete-time nonlinear systems requires the construction of a stabilizing feedback control law based on a control Lyapunov function (CLF). However, there are few systematic approaches available for defining appropriate CLFs. We propose an approach that employs Bayesian filtering methodology to parameterize a quadratic CLF. In particular, we use the ensemble Kalman filter to estimate parameters used in defining the CLF within the control loop of the inverse optimal control. Results are demonstrated on a real-world application to mathematical biology.

Andrea Arnold  
 North Carolina State University  
 anarnold@ncsu.edu

Hien Tran  
 Center for Research in Scientific Computation  
 North Carolina State University  
 tran@ncsu.edu

### CP7

#### Interpolatory Model Reduction of Parameterized Bilinear Dynamical Systems

Interpolatory projection methods for model reduction of nonparametric linear dynamical systems have been already successfully extended to nonparametric bilinear dynamical systems. However, this is not the case for parametric bilinear systems. In this work, we aim to close this gap by providing a natural extension of interpolatory projections to model reduction of parametric bilinear dynamical systems. We introduce the conditions that projection subspaces need to satisfy in order to obtain parametric tangential interpolation of each subsystem transfer functions. These conditions also guarantee that the parameter gradient of each subsystem transfer function is matched tangentially by the parameter gradient of the corresponding reduced order model transfer function. Similarly, we obtain conditions for interpolating the parameter Hessian of the transfer function by including extra vectors in the projection subspaces. As in the linear case, for two-sided projections, the basis construction does not require computing neither the gradient nor the Hessian to be matched.

Andrea Carracedo Rodriguez  
 Virginia Polytechnic Institute and State University  
 (Virginia Tech)  
 crandrea@vt.edu

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

### CP7

#### Domain Decomposition Algorithms for Uncertainty Quantification: High-Dimensional Stochastic Systems

### tic Systems

Domain decomposition (DD) algorithms developed by Subber for uncertainty quantification of large-scale stochastic PDEs are extended using multi-level preconditioned conjugate gradient methods (PCGM) to enhance its capability to tackle high-dimensional stochastic systems. These DD algorithms will be further tuned using Anderson acceleration method to expedite the convergence. Parallel sparse matrix-vector operations are used to cut floating-point operations and memory requirements. Both numerical and parallel scalabilities of these algorithms are presented for a diffusion equation having spatially varying diffusion coefficient modeled by a non-Gaussian stochastic process.

Abhijit Sarkar  
 Associate Professor  
 Carleton University, Ottawa, Canada  
 abhijit\_sarkar@carleton.ca

Ajit Desai  
 Carleton University, Canada  
 ajit.desai@carleton.ca

Mohammad Khalil  
 Sandia National Laboratories  
 mkhalil@sandia.gov

Chris Pettit  
 United States Naval Academy, USA  
 pettitcl@usna.edu

Dominique Poirel  
 Royal Military College of Canada, Canada  
 poirel-d@rmc.ca

### CP7

#### An Asymptotic-Preserving Stochastic Galerkin Method for the Semiconductor Boltzmann Equation with Random Inputs and Diffusive Scalings

In this talk, I will introduce the generalized polynomial chaos approach based stochastic Galerkin (gPC-SG) method for the linear semi-conductor Boltzmann equation with random inputs and diffusive scalings. The random inputs are due to uncertainties in the collision kernel or initial data. We study the regularity of the solution in the random space, and prove the spectral accuracy of the gPC-SG method. We then use the asymptotic-preserving framework for the deterministic counterpart developed in [Jin] to come up with the stochastic asymptotic-preserving gPC-SG method for the problem under study, which is efficient in the diffusive regime. Numerical experiments will be presented to validate the accuracy and asymptotic properties of the method. This is a joint work with Prof Shi Jin. [Jin]: Discretization of the multi scale semiconductor Boltzmann equation by diffusive relaxation schemes, S. Jin and L.Pareschi, J. Comput. Phys., 161: 312-330, 2000.

Liu Liu  
 University of Wisconsin Madison  
 lliu84@wisc.edu

### CP7

#### Bayesian Model Reduction for Nonlinear Dynamics Using Automatic Relevance Determination

The concept of automatic relevance determination (ARD) is invoked in this study to select models of nonlinear dy-

namical systems in the form of stochastic ordinary differential equations (ODEs). The Bayesian method can provide misleading results when the prior probability distributions are assigned arbitrarily to a subset (some) of parameters for which no information is available. In such cases, ARD provides an automatic model selection scheme to identify the optimal model nested under an overly complex model. Given noisy measurement data, a relatively complicated model is envisioned to represent the dynamical system. Then a model selection problem is posed to find the best model nested under the envisioned model. This problem is transferred from the model parameter space to a hyper-parameter space by imposing a parametrized prior distribution called the ARD prior. The parameters of the prior distribution are known as hyper-parameters; in practice, they explicitly capture the relevance of model parameters. This approach regularizes the adaptation of the posterior distribution to the data so as to avoid overfitting.

Abhijit Sarkar  
Associate Professor  
carleton University, Ottawa, Canada  
abhijit\_sarkar@carleton.ca

Rimple Sandhu  
Carleton University, Canada  
rimple\_sandhu@carleton.ca

Chris Pettit  
United States Naval Academy, USA  
pettitcl@usna.edu

Mohammad Khalil  
Sandia National Laboratories  
mkhalil@sandia.gov

Dominique Poirel  
Royal Military College of Canada, Canada  
poirel-d@rmc.ca

## CP7

### A Computational Bayesian Framework to Parallelize an Adaptive Markov Chain Monte Carlo

This paper introduces a method to optimize the convergence of a new adaptive Markov Chain Monte Carlo approach (AMCMC) needed to formulate high dimensional parametric Bayesian formulations. The proposed approach relies on the sampling of parallel chains to ensure the capturing of all modes of the posterior distribution by developing a two-step synchronous sampling mechanism. The full integration of the posterior distribution applying the Bayesian paradigm using MCMC and Metropolis-Hasting (MH) algorithms is known to be computationally inefficient. AMCMC presents a deterministic tuning of the proposal distribution in a two-step process, to optimize the acceptance ratio and expedite the MCMC convergence via several independent runs. A methodology to randomly generate parallel combinations of MCMC is proposed, which aim is to search for the optimized chains acceptance ratio. A comparison between the proposed and conventional MCMC-MH algorithms is discussed, when these are applied to a forward model simulating shale gas well-production, presented with real borehole production of daily observations, required to complete a probabilistic model calibration. The proposed approach is validated by its use on different well production data, showing significant computational efficiency, but most importantly, the promise for the method to be fully Bayesian parallelized

(work in progress).

Yasser Soltanpour  
Texas A&M University  
yasser\_soltanpour@tamu.edu

Zenon Medina-Cetina  
Stochastic Geomechanics Laboratory  
Texas A&M University, TX, USA  
zenon@tamu.edu

## CP7

### Stochastic Dirichlet Boundary Optimal Control of Steady Navier-Stokes Equations

When a physical system under control includes a stochastic component, the construction, modeling, and analysis of the controls become much more difficult. As a specific case, we consider the optimal control of a system governed by the Navier-Stokes equations with a stochastic Dirichlet boundary condition. Control conditions applied only on the boundary are associated with reduced regularity, as compared to internal controls. To ensure existence of solution and efficiency of numerical simulations, the stochastic boundary conditions are required to belong almost surely to  $H^1(\partial\mathcal{D})$ , similar to the  $H$ -valued infinite dimensional Wiener process. To simulate the system, state solutions will be approximated using the stochastic collocation finite element approach, and sparse grid techniques are applied to the boundary random field. The one shot optimality systems are derived from the Lagrange functional. Error estimates are computed for the optimality almost surely using samples, and for the state equation using interpolated boundary conditions. Error estimates for the adjoint equations are derived from a duality argument, and the control equation comes via a non-conforming finite element variational crime. A numerical simulation can then be made, using a combination of Monte Carlo and sparse grid methods, which demonstrates the efficiency of the algorithm.

Wenju Zhao  
Department of Scientific Computing  
Florida State University  
wz13@my.fsu.edu

Max Gunzburger  
Florida State University  
Department of Scientific Computing  
mgunzburger@fsu.edu

## CP8

### Compact Implicit Integration Factor Method for Solving High Order Differential Equations

Due to the high order spatial derivatives and stiff reactions, severe temporal stability constraints on the time step are generally required when developing numerical methods for solving high order partial differential equations. Implicit integration method (IIF) method along with its compact form (cIIF), which treats spatial derivatives exactly and reaction terms implicitly, provides excellent stability properties with good efficiency by decoupling the treatment of reaction and spatial derivatives. One major challenge for IIF is storage and calculation of the potential dense exponential matrices of the sparse discretization matrices resulted from the linear differential operators. The compact representation for IIF (cIIF) was introduced to save the computational cost and storage for this purpose. Another chal-

lenge is finding the matrix of high order space discretization, especially near the boundaries. In this paper, we extend IIF method to high order discretization for spatial derivatives through an example of reaction diffusion equation with fourth order accuracy, while the computational cost and storage are similar to the general second order cIIF method. The method can also be efficiently applied to deal with other types of partial differential equations with both homogeneous and non-homogeneous boundary conditions. Direct numerical simulations demonstrate the efficiency and accuracy of the approach.

Sameed Ahmed, Xinfeng Liu  
University of South Carolina  
ahmed3@math.sc.edu, xliu@math.sc.edu

### CP8

#### Concept of Spectral Differentiation for Solving Corner Flow

The concept of differentiation with the first kind of spectral method is used in this presentation to solve the model of movement of flow around an upright angle. In this model, lots of fluid features and interactions and separations take a place and hence a high accurate numerical techniques is needed. Because of the singularities in some nodes, we used a conformal mapping which is used by many researchers. A numerical simulation of the fluid movements will be showed along side with some tables to validate our results with others.

Badr Alkahtani  
King Saud University  
alhaghog@gmail.com

### CP8

#### Optimal Recovery of Integral Operators and Its Applications

In this talk we present the solution to a problem of recovering a rather arbitrary integral operator based on incomplete information with error. We apply the main result to obtain optimal methods of recovery and compute the optimal error for the solutions to certain integral equations as well as boundary and initial value problems for various PDE's. In particular, to illustrate the method, we present results for boundary value problems for wave, heat, and Poisson's equations. Nevertheless, the developed method is more general and can be applied to other similar problems.

Yuliya Babenko  
Kennesaw State University  
ybabenko@kennesaw.edu

Vladyslav Babenko  
Dnepropetrovsk National University, Ukraine  
Donetsk Institute of Applied Mathematics and Mechanics, UKR  
babenko.vladislav@gmail.com

Natalia Parfinovych, Dmytro Skorokhodov  
Dnepropetrovsk National University  
xvector2016@gmail.com, dmitriy.skorokhodov@gmail.com

### CP8

#### Compatible-Strain Mixed Finite Element Methods

#### for Nonlinear Elasticity

We introduce a new class of mixed finite element methods for large deformations of solids called compatible-strain mixed finite element methods (CSFEMs). We use a Hu-Washizu-type mixed formulation for CSFEMs, and the independent unknowns are the displacement, the displacement gradient, and the first Piola-Kirchhoff stress tensor. To obtain the solution and the test spaces of CSFEMs, we discretize the nonlinear elasticity complex using the finite element exterior calculus. The nonlinear elasticity complex is a Hilbert complex that describes large deformations of solids and contains information about topological properties of bodies. Accordingly, the trial spaces of the displacement gradient automatically satisfy the classical Hadamard jump condition, which is a necessary condition for the compatibility of non-smooth displacement gradients. By solving selected problems for nonlinear elastic bodies, we demonstrate that CSFEMs have a good performance for bending problems, in the nearly incompressible regime, and for bodies with complex geometries. CSFEMs are capable of accurately approximating stresses and perform well in problems that standard enhanced strain methods suffer from the hourglass instability. Moreover, CSFEMs provide a novel and convenient framework for modeling inhomogeneities.

Mostafa Faghih Shojaei  
School of Civil and Environmental Engineering  
Georgia Institute of Technology  
mfsh@gatech.edu

Arash Yavari  
Georgia Institute of Technology  
arash.yavari@ce.gatech.edu

Arzhang Angoshtari  
Department of Civil and Environmental Engineering  
The George Washington University  
aangoshtari@email.gwu.edu

### CP8

#### An Optimal Complexity Spectral Method for Elliptic PDEs on Rectangular Domains

We present an  $\mathcal{O}(n^2 \log n \log(1/\epsilon))$  complexity spectrally accurate elliptic PDE solver for PDEs of splitting rank 2 on rectangular domains, where  $n$  is the discretization size in each spatial variable and  $0 < \epsilon < 1$  is the working precision. The method is based on a modification of the ultraspherical spectral method and the construction of certain structured Sylvester matrix equations. We present a proof of the algorithmic complexity of the solver and show that it is amenable to general boundary conditions and right-hand sides.

Dan Fortunato  
Harvard University  
dfortunato@g.harvard.edu

Alex Townsend  
Cornell University  
townsend@cornell.edu

### CP8

#### Scalable Simulation of Systems of PDEs for Modeling Tumor Growth

There are existing sophisticated mathematical models de-

signed to model tumor growth with partial differential equations (PDEs); however, finding solutions is computationally expensive so researchers have to lower resolution, simplify models, or run fewer simulations with representative sets of parameters, etc. Thus, results are less accurate, less timely, or not tailored to individual patients. We hypothesize that, a component-wise time-stepping method can be used so that greater accuracy, efficiency and scalability can be achieved.

Elyse Garon

The University of Southern Mississippi  
elyse.garon@eagles.usm.edu

James V. Lambers

University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

### CP8

#### **A High Order Accurate Direct Solution Technique for High Frequency Problems with Body Loads**

Highly oscillatory differential equations arise in many areas of science and engineering including medical imaging, material science, fluid dynamics, electromagnetics, and geophysics. When solutions are highly oscillatory, classic discretization techniques suffer from dispersion and poor conditioning, which result in accuracy issues and slow convergence for iterative solvers. Direct solvers are robust in the high frequency regime, avoiding these issues, and are excellent for problems with many right hand sides such as occur in most applications of interest. This talk presents the recently developed Hierarchical Poincaré-Steklov (HPS) method. HPS is a high order discretization technique that provides robust solutions even in the high frequency regime. The method comes with a direct (as opposed to iterative) solver which processes solves with nearly linear cost (with respect to the number of discretization points), making it ideally suited for applications when many solves are required. Numerical results will report on the performance of the method.

Peter Geldermans

Rice University  
pjpg4@rice.edu

Adrianna Gillman

Rice University  
Department of Computational and applied mathematics  
adrianna.gillman@rice.edu

### CP8

#### **Mesh Imprinting Reduction by Multi-Dimensional Slope Limiters in Hydrodynamic Simulations**

Slope limiters are a powerful tool to suppress numerical oscillations of second order finite volume schemes to solve inviscid compressible flows. On Cartesian meshes, the numerical schemes typically employ directional splitting. Split schemes are easy to implement in 2D/3D and they are computationally effective. These schemes use only one-dimensional (1D) information along the current split direction to reconstruct fluid quantities in computational cells. This reconstruction is strongly dependent on the orientation of a discontinuity with respect to a computational mesh, resulting in increased mesh imprinting in the numerical solution. We investigate the effect of the mesh imprinting on radially symmetric problems includ-

ing simulations of converging cylindrical/spherical shells on Adaptive Mesh Refinement (AMR) meshes using the Los Alamos National Laboratory (LANL) multi-material multi-physics code xRage [Gittings et al. 2008]. The imprinting is evaluated as the artificial angular dependency (anisotropy) of the solution. We suggest a correction of the 1D slopes to reduce the imprinting. The correction is making use of the local direction of the gradient of a reconstructed quantity. This method is compared to a truly multi-dimensional reconstruction using the Barth-Jespersen slope limiter [Barth, Jespersen. 1989] or Multi-dimensional Limiting Process (MLP) [Park, Kim. 2012] with respect to the symmetry preservation versus a computational cost.

Jan Velechovsky, Marianne M. Francois

Los Alamos National Laboratory  
jan@lanl.gov, mmfran@lanl.gov

Thomas Masser

Continuum Dynamics Group, CCS-2  
CCS Division, Los Alamos National Laboratory  
tmasser@lanl.gov

### CP9

#### **An Incremental Compressive Sampling Approach for Sparser Recovery of Polynomial Chaos Expansions**

Compressive sampling techniques have successfully allowed reconstruction of sparse Polynomial Chaos (PC) expansions from a small set of samples, by effectively solving an underdetermined linear system. In this presentation, we introduce an efficient sparse recovery approach that promotes the sparsity by breaking the dimensionality of the PC expansion. The proposed algorithm incrementally explores sub-dimensional expansions for a sparser recovery, and shows success when removal of uninfluential parameters allows for a measurement matrix with a lower coherence, and/or a higher order expansion to be recovered. Using several numerical examples, we demonstrate the capability of the proposed approach in both producing more accurate PC coefficient estimates, and also reducing the dimensionality of the stochastic representation.

Hadi Meidani, Negin Alemazkoor

University of Illinois at Urbana-Champaign  
meidani@illinois.edu, alemazk2@illinois.edu

### CP9

#### **Stochastic Model Reduction of Allen-Cahn Phase Field Model with High Dimensional Random Forcing**

Using the framework of *Dynamical Orthogonality* (DO) condition, proposed by Sapsis & Lermusiaux, a coupled set of evolution equations are derived for Allen-Cahn phase field transport equation. The reduced order field equations (DO equations) consist of (i) the system of Partial Differential equations (PDEs) that describes the evolution of the mean field and the orthonormal spatial basis (modes) defining the stochastic subspace where uncertainty 'lives' and (ii) the set of Stochastic Differential equations (SDEs) that describes the evolution of stochastic basis (coefficients) defining how the stochasticity will be evolved within the reduced order stochastic subspace. Numerical simulations are performed for time independent high dimensional stochastic forcing based on Karhunen-Loeve (KL) decomposition. The aim of the current work is to investigate dimension re-



duction in high dimensional random space by analyzing the number of DO modes required to resolve the transient stochastic Allen-Cahn equation. More specifically, our results demonstrates that the DO decomposition does not suffer from the *curse of dimensionality* in comparison to Polynomial Chaos(PC) methods.

Mayank Bajpayi  
Brown University  
mayank\_bajpayi@brown.edu

Hessam Babaee  
MIT  
babaee@mit.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

### CP9

#### An Information Theoretic Approach to Use High-Fidelity Codes to Calibrate Low-Fidelity Codes

In this presentation, we discuss an information theoretic approach to employ high-fidelity codes to calibrate low-fidelity codes used for design optimization or control implementation. The objective is to employ a limited number of high-fidelity code evaluations as data for Bayesian calibration of the low-fidelity code. We employ the mutual information between parameters and designs to determine input values to the high-fidelity code, which maximize the available information. For computationally expensive codes, surrogate models are used to approximate the mutual information. The framework is illustrated for examples arising in nuclear power plant design.

Ralph C. Smith  
North Carolina State Univ  
Dept of Mathematics, CRSC  
rsmith@ncsu.edu

Kayla Coleman  
North Carolina State University  
kdcolem2@ncsu.edu

Brian Williams  
Los Alamos National Laboratory  
brianw@lanl.gov

### CP9

#### Mesh Discretization Error and Uncertainty Quantification: A Variational Multiscale Approach

When trying to quantify the uncertainty of a given simulation, one has to identify the sources of such uncertainty. These sources can be classified into input uncertainty, numerical approximation uncertainty, or modelling uncertainty. In this work we focus on the numerical approximation uncertainty, but it is well known that the accuracy of the numerical approximation methods depends on the input parameters. For example, considering convection-diffusion-reaction problems, uncertainty may be associated to the direction of advection or the magnitude of the diffusion coefficient. In this work we firstly assess the behaviour of some existing Variational Multiscale (VMS) error estimators for the Uncertainty Quantification (UQ) of Mesh Discretization Error (MDE) under uncertain input param-

eters. The VMS error estimators are typically described in a Finite Element (FE) framework and, since they involve second order spatial derivatives, they are hardly extensible to Finite Volume (FV) methods. In order to overcome this difficulty, we also propose new VMS error estimators that can be easily used in a FE or FV context. The VMS error estimators are analysed for canonical convection-diffusion-reaction problems and then extended to turbulent incompressible flow problem.

Oriol Colomes  
International Center for Numerical Methods in Engineering  
oriol.colomes@duke.edu

Guglielmo Scovazzi  
Duke University  
guglielmo.scovazzi@duke.edu

### CP9

#### Model Reduction from Partial Observations

Our contribution takes place within the context of model reduction for parametric partial differential equations (PPDE). When the solution of the PPDE has to be evaluated for many different values the parameters, the computational effort may become prohibitive. To circumvent this issue, model reduction intends to simplify the resolution of the PPDE by (typically) constraining the solutions to belong to some low-dimensional subspace. Many techniques have been proposed in the literature to identify such subspaces: Taylor or Hermite expansions, proper orthogonal decomposition (POD), balanced truncation, reduced basis techniques, etc. All the methods mentioned above presuppose some refined knowledge of the solution manifold  $\mathcal{M}$ . Unfortunately, in practice a refined knowledge of  $\mathcal{M}$  may not always be available. The question addressed in our work is therefore as follows: can we still build a “good” approximation subspace if  $\mathcal{M}$  is imperfectly known but some partial measurements of the latter are available? We propose practical procedures to deal with this problem and provide theoretical results to analyze their behavior.

Cédric Herzet  
Inria, Center of Rennes  
cedric.herzet@inria.fr

Angélique Drémeau  
ENSTA Bretagne  
angelique.dremeau@ensta-bretagne.fr

Patrick Héas  
Inria, Center of Rennes  
patrick.heas@inria.fr

### CP9

#### Transported PDF Approaches for Propagating Parameter Uncertainty: A Reacting Flow Case Study

In propagating parameter uncertainty forward through a system of PDEs, reduced representations are often sought to deal with the computational complexity and expense of the forward PDE system. The holy grail is an approach that is computationally feasible, can handle large dimensionality (of input parameters and/or solution variables) and, most importantly, can stay faithful to the PDE system for arbitrarily complex computational domains and boundary conditions. Transported probability density function (T-PDF) approaches, developed primarily in the context

of turbulent reacting and non-reacting flows, have some of these very desirable properties. In the T-PDF approach one solves for the transport equations of joint pdfs of the solution variables as implied by the PDE system. In this talk I will give a brief overview of the T-PDF approaches and how they are readily amenable to the forward propagation problem. I will describe the advantages and limitations of T-PDF approaches, and discuss numerical methods for solving them. In particular I will focus on a hybrid Eulerian-Lagrangian methodology that is advantageous for dealing with large dimensionality. I will conclude with preliminary results from a simple convective-reactive problem with a prescribed distribution of the uncertain reaction rate parameter.

Hemanth Kolla

Sandia National Laboratories  
hnkolla@sandia.gov

Xinyu Zhao

University of Connecticut  
Department of Mechanical Engineering  
xinyu.zhao@uconn.edu

Habib N. Najm

Sandia National Laboratories  
Livermore, CA, USA  
hnnajm@sandia.gov

## CP9

### A Stochastic Galerkin Method for the Boltzmann Equation with High Dimensional Random Inputs Using Sparse Grids

We propose a stochastic Galerkin method using sparse grids for the Boltzmann equation with high dimensional random inputs. The method uses locally supported piecewise polynomials as an orthonormal basis of the random space. By a sparse grid technique, only a moderate number of basis functions are required to achieve good accuracy in high dimensional random spaces. We discover a sparse structure of a set of basis-related coefficients, which allows us to accelerate the computation of the collision operator. Regularity of the solution of the Boltzmann equation in the random space and an accuracy result of the stochastic Galerkin method are proved in multidimensional case. The efficiency of the method is illustrated by numerical examples with uncertainties from the initial data, boundary data and collision kernel.

Ruiwen Shu

University of Wisconsin-Madison  
rshu2@math.wisc.edu

Shi Jin

Shanghai Jiao Tong University, China and the  
University of Wisconsin-Madison  
sjin@wisc.edu

Jingwei Hu

Purdue University  
jingwei.hu@purdue.edu

## CP9

### Impact of Parametric Uncertainty on Estimation of the Energy Deposition into an Irradiated Brain

## Tumor

We analyze the effect of parametric uncertainty on the total energy deposited in a brain tumor through X-ray irradiation. Both the location of the region over which a dose-enhancing, iodinated contrast agent spreads out after injection into the tumor, and the agent's concentration, are allowed to be uncertain. We model this problem via a probabilistic approach in which the coordinates of the center of the contrast agent region, as well as the effective atomic number in this area (which depends on the agent's concentration), are represented as mutually independent, uniformly distributed random variables. Employing the stochastic collocation (SC) method, we estimate statistical moments of the deposited energy as a function of the mean and/or variance of the random inputs. We find that in most cases, the coefficient of variation of the uncertain parameters is amplified by the nonlinearity of the problem, yielding a larger coefficient of variation for the energy deposition. As the stochastic dimension increases, the magnitude of the predictive uncertainty in the energy deposition, as measured by its standard deviation, approaches that of the prediction (mean energy deposition) itself. This demonstrates that accurate prediction of the energy deposition requires a proper treatment of even small parametric uncertainty. Our analysis also reveals that SC outperforms standard Monte Carlo, with the largest difference in efficiency occurring for the case of a single uncertain parameter.

Soren Taverniers, Daniel M. Tartakovsky

University of California, San Diego  
staverni@ucsd.edu, dmt@ucsd.edu

## CP10

### Numerical Methods of Rational Form for Solving Partial Differential Equations

The purpose of this study is to investigate selected numerical methods that demonstrate good performance in solving PDEs. We adapt alternative method that involve rational polynomials. Padé time stepping (PTS) method, which is highly stable for the purposes of the present application and is associated with lower computational costs, is applied. Furthermore, PTS is modified for our study which focused on diffusion equations. Due to the rationality form of PTS, some numerical error occurs and then some kind of control is imposed. Finally, numerical runs are conducted to obtain the optimal local error control threshold.

Said Algarni

Assistant Professor - Department of Mathematics  
King Fahd University of Petroleum and Minerals  
garnis@kfupm.edu.sa

## CP10

### Multi-Dimensional Sublinear Sparse Fourier Algorithm

We discuss how to develop a deterministic sublinear Fourier algorithm on sparse data in high-dimensional space. We develop an algorithm for both noisy and noiseless data introducing new concepts called "partial unwrapping method" and "tilting method". These two methods overcome difficulties encountered when we extend one-dimensional algorithm presented in "Adaptive Sub-linear Time Fourier Algorithm" by D. Lawlor, Y. Wang and A. Christlieb (2012). Furthermore we use error-

correction to make the algorithm work for noisy data in high-dimensional space.

Bosu Choi  
Michigan State University  
choibosu@math.msu.edu

### CP10

#### Indifference Pricing of a GLWB Option in Variable Annuities

We investigate the valuation problem of variable annuities with Guaranteed Lifelong/Lifetime Withdrawal Benefit (GLWB) options, which give the policyholder the right to withdraw a specified amount as long as he/she lives, regardless of the performance of the investment. We apply the principle of equivalent utility to find the indifference price for a variable annuity with a GLWB contract with an equity indexed death benefit. Using an exponential utility function, Hamilton-Jacobi-Bellman (HJB) type partial differential equations (PDEs) are derived for the pricing functions. We first assume the mortality is deterministic, and the pricing PDE is solved numerically using a finite difference method. The effects of various parameters are investigated, including the age at inception of the policyholder, the withdrawal rate, the risk-free rate, and the volatility of the underlying asset. We also consider a roll-up option and analyze the effect of delaying the start of the withdrawals. Another pricing PDE is derived with a stochastic mortality, when the force of mortality is modeled with a stochastic differential equation. A finite difference method is used again to solve the pricing PDE numerically, and the sensitivities of the GLWB contracts with respect to the withdrawal rate and the risk-free rate are explored.

Jungmin Choi  
East Carolina University  
choiju@ecu.edu

### CP10

#### Polynomial Particular Solutions for Solving Elliptic Partial Differential Equations

In the past, polynomial particular solutions have been obtained for certain types of partial differential operators without convection terms. In this paper, a closed-form particular solution for more general partial differential operators with constant coefficients has been derived for polynomial basis functions. The newly derived particular solution is further coupled with the method of particular solutions (MPS) for numerically solving a large class of elliptic partial differential equations. In contrast to the use of Chebyshev polynomial basis functions, the proposed approach is more flexible in selecting the collocation points inside the domain. The polynomial basis functions are well-known for yielding ill-conditioned systems when their order becomes large. The multiple scale technique is applied to circumvent the difficulty of ill-conditioning problem. Some numerical examples are presented to demonstrate the effectiveness of the proposed algorithm.

Thir R. Dangal  
The University of Southern Mississippi  
Hattiesburg, MS 39406  
thir.dangal@usm.edu

C.S. Chen  
The University of Southern Mississippi

cs.chen@usm.edu

Ji Lin  
Hohai University - Jiangning Campus Nanjing, CHINA  
linji861103@126.com

### CP10

#### An Efficient Numerical Scheme for Space-Fractional Fitzhugh-Nagumo Model

Fractional models proved to be effective in describing the heterogeneity and complex connectivity of electrophysiological dynamics in biological tissue. A strongly stable second order exponential time differencing scheme for solving the space-fractional Fitzhugh-Nagumo model with non-smooth data is presented. Stability, efficiency and reliability of the proposed scheme is discussed. A comparison with existing methods in the literature shows the usefulness and robustness of the scheme. Furthermore, the scheme is shown to be easily applicable to large systems of reaction-diffusion equations.

Olaniyi S. Iyiola  
University of Wisconsin-Milwaukee, WI, USA  
osiyiola@uwm.edu

Bruce Wade  
Department of Mathematical Sciences, UW-Milwaukee  
Milwaukee, Wisconsin 53201-0413  
wade@uwm.edu

Khaled Furati  
Department of Mathematics & Statistics  
King Fahd University of Petroleum & Minerals  
kmfurati@kfupm.edu.sa

Abdul Khaliq  
Middle Tennessee State University  
abdul.khaliq@mtsu.edu

### CP10

#### A Comparison of Adaptive Mesh Refinement Techniques for Poisson's Equation

The technique of Adaptive Mesh Refinement provides efficient solution of a PDE problem by selecting appropriate regions of the physical domain to be refined. This results in reduced degrees of freedom, thus the efficiency. In here the local mesh refinement is based on designed a-posteriori error estimators. We compare two different estimators for the Poisson problem discretized with finite element method. One of them is taken from literature, and the other one is proposed by our research group. The performance of these two methods is studied in terms of  $L_2$  and  $H_1$  error norms. In the numerical validations part, the proposed error estimator performs better in terms of reduced number of degree of freedoms for given tolerance.

Shih-Yu Lee  
Texas Tech University  
shihyu.lee@ttu.edu

Eugenio Aulisa  
Department of Mathematics and Statistics.  
Texas Tech University  
eugenio.aulisa@ttu.edu

Guoyi Ke

Texas Tech University  
guoyi.ke@ttu.edu

### CP10

#### The Computational Analysis of Problems on Domains with Small Holes

A new computational method is proposed to address meshing and convergence challenges that can arise when the problem domain has small holes in it. The method combines analytic singularities of the solution with finite element approximation of its smooth components. Theoretical and numerical results are provided, to establish the efficacy of the method, both in the energy norm and in extracting a representative quantity of interest. The method converges both with respect to the size of the holes and the mesh discretization parameter, and provides a more accurate alternative to using the asymptotic limit.

Ana Maria Soane  
U.S. Naval Academy  
Mathematics Department  
soane@usna.edu

Ivo Babuska  
ICES, The University of Texas at Austin  
babuska@ices.utexas.edu

Manil Suri  
University of Maryland Baltimore County  
suri@umbc.edu

### CP10

#### Sub-Linear Time Discrete Sparse Fourier Transform Algorithm

During the past 10 years, a lot of sparse Fourier Transform algorithms (random and deterministic) have been built. However, they all required that the algorithm have the access to the signal anywhere they want (or in another word, the signal is continuous). This requirement is hard to reach since in real life people usually only have access to limited or discrete samples. In our algorithm we designed an interpolation method which can be used along with some deterministic SFT algorithms to capture the significant frequencies in sub-linear time. The advantage of our algorithm is that we only need  $N$  equal-spaced samples from the signal ( $N$  is the bandwidth) and we can find all the significant frequencies in sub-linear time ( $O(k^a(\log N)^b)$ ). We proved (both theoretically and numerically) that our interpolation method can be used along with some existing SFT algorithms. In the poster, we will give both the theoretical and numerical results of our algorithm.

Ruochuan Zhang  
Michigan State University  
zhangr12@msu.edu

### CP11

#### Bayesian Calibration for Parameters of Jwl Equation of State in Cylinder Test

In detonation physics, the cylinder test is the benchmark experiment to evaluate output energy of the explosive and calibrate parameters of equation of state for explosive products. The traditional calibration method is usually only based on radius data in the radial direction, which is smooth and regular functional data. The Gaussian pro-

cess (GP) was used to build surrogate models for radius data, and the Bayesian method was used for calibration. As the improvement of test technology, both the radius and velocity can be observed in cylinder test. The above traditional calibration method may not make the simulation results of the velocity match the data well. In this talk, we want to use both the radius and velocity data to improve the prediction accuracy of the calibrated model. The GP model is hard to applied to the velocity data because it is very irregular functional data. So we propose that the GP and wavelet methods are used to build surrogate models for the radius and velocity data respectively, and then calibrate parameters of Jones-Wilkins-Lee (JWL) state of equation. Our method decrease the prediction uncertainty of calibrated models.

Hua Chen, Haibing Zhou, Guozhao Liu  
Institute of Applied Physics and Computational Mathematics  
chen\_hua@iapcm.ac.cn, zhou\_haibing@iapcm.ac.cn, liu\_guozhao@iapcm.ac.cn

Zhanfeng Sun  
Laboratory for Shock Wave and Detonation Physics Research,  
Institute of Fluid Physics, CAEP  
sunzf7695@163.com

Shudao Zhang  
Institute of Applied Physics and Computational Mathematics  
zhang\_shudao@iapcm.ac.cn

### CP11

#### An Efficient Stochastic Quasi-Newton Mcmc Method for Stochastic Inversion of Linear Elasticity Problem Using Kernel Pca, Automatic Differentiation and Adjoint Method

We address a linear elasticity inversion in the framework of Bayesian inference. The nonlinear mapping between the observables and parameters leads to non-Gaussian posteriors even with the Gaussian priors. The computational cost of the forward model makes it intractable using traditional methods such as Markov chain Monte Carlo (MCMC) and the Polynomial Chaos Expansion (PCE). We propose a novel stochastic inversion framework, where we first derive continuous adjoint partial differential equations (PDEs) that facilitate efficient computation of the gradient of an objective functional with respect to parameters. Next, we use a conceptual model representing prior knowledge to generate a series of realizations of the complex structural model. Following that, we construct a low-dimensional feature space where the non-linearly correlated parameters are represented in terms of independent Gaussian random variables using a kernel principle component analysis (KPCA) and an iso-probabilistic transformation. Given the gradient of the objective functional with respect to parameters, we use automatic differentiation to derive an adjoint model of KPCA-based iso-probabilistic transformation and obtain the gradient of the objective functional with respect to the low-dimensional feature variables. Finally, we devise an efficient quasi-Newton MCMC scheme to sample the posteriors of the random variables in the feature space then obtain the posteriors of the high-dimensional parameter space.

Xiao Chen  
Lawrence Livermore National Laboratory  
Center for Applied Scientific Computing

chen73@llnl.gov

Wenju Zhao  
Department of Scientific Computing  
Florida State University  
wz13@my.fsu.edu

Charanraj Thimmisetty  
Lawrence Livermore National Lab  
thimmisetty1@llnl.gov

Joshua A. White, Charles Tong  
Lawrence Livermore National Laboratory  
white230@llnl.gov, tong10@llnl.gov

### CP11

#### **Bayesian, Multi-Fidelity, Optimization under Uncertainty**

This paper is concerned with the optimization of high-dimensional, complex models in the presence of uncertainty. As many other uncertainty quantification tasks, this is hindered by the significant cost associated with each forward model evaluation and the large number of random and design variables. We reformulate the problem as one of probabilistic inference which enables a uniform treatment of both sets of variables as well as the assessment of the objective's sensitivity. In order to alleviate the computational burden we employ alternative models that are of lower-fidelity but much cheaper to evaluate. The use of such surrogates introduces an additional source of (epistemic) uncertainty. That is, regardless of the amount of training data, the lower-fidelity model is incapable of predicting exactly the output of interest of the reference, high-fidelity model. This additional uncertainty can be readily introduced in the probabilistic framework advocated and its effects on the results can be quantified. We couple the aforementioned formulation with stochastic Variational inference tools that are capable of providing accurate approximations to the solution. Furthermore they can identify regions in the input space where additional, high-fidelity runs are needed in order to most efficiently refine the estimates obtained.

Maximilian Koschade  
Technical University of Munich  
maximilian.koschade@tum.de

Phaedon S. Koutsourelakis  
Technical University Munich  
p.s.koutsourelakis@tum.de

### CP11

#### **Bayesian Coarse-Graining in Atomistic Simulations: Adaptive Identification of the Dimensionality and Salient Features**

We present a Bayesian formulation to coarse-graining (CG) of atomistic systems using generative probabilistic models. A crucial component of all CG schemes is the definition of the coarse variables. The key component of the formulation advocated is the definition of a probabilistic coarse-to-fine mapping. This enables the implicit definition of CG variables as latent generators of the fine-scale configurations. Appropriate prior specifications on the coarse-to-fine mapping provide insights into the physical meaning of the coarse variables. We demonstrate how such an approach can lead to adaptive determination of the dimension of the

coarse variables. A second critical question pertains to the parametrization of the CG model and the type of interactions between the CG variables. We propose the sequential addition of features by employing an information theoretic metric based on the Kullback-Leibler divergence between the exact and approximate description. In addition, the use of sparsity-enforcing priors leads to a computationally efficient coarse model with as few features as necessary. The methodology proposed is demonstrated in the context of coarse-graining of peptides. The approach advocated enables the significant increase in the time and space scales that can be explored by peptide simulations.

Markus Schoeberl  
Technical University of Munich  
m.schoeberl@tum.de

Nicholas Zabarar  
University of Notre Dame  
nzabarar@gmail.com

Phaedon S. Koutsourelakis  
Technical University Munich  
p.s.koutsourelakis@tum.de

### CP11

#### **Method Coupling Harmonic Decomposition and Polynomial Chaos for Seismic Wave Propagation in Uncertain Medium**

A surrogate is proposed to study seismic wave propagation in uncertain medium. The surrogate is based on a double decomposition of the signal: a damped harmonic decomposition coupled with a polynomial chaos (PC) representation of the four coefficients of each harmonic term (amplitude, decay constant, pulsation, and phase). An efficient PC representation of the coefficients are obtained through non-intrusive spectral projections. It requires the resolution of a nonlinear least squares problem for each integration point of the sparse grid. The implementation of the surrogate is illustrated on applications to layered soils with uncertainties in the geological data (geometry, wave velocities, damping factor). Computational tests show that the stochastic signal can be efficiently represented with a low-order PC representation leading to the use of a low-level sparse grid integration. For each test case, a global sensitivity analysis is performed in time and frequency domains to investigate the relative impact of the random parameters.

Pierre Sochala, Florent De Martin  
BRGM  
p.sochala@brgm.fr, f.demartin@brgm.fr

### CP11

#### **An Importance Sampling Approach to Risk Estimation**

Estimation of risk functions of computationally expensive quantities of interest is an important task in uncertainty quantification and risk-averse optimization under uncertainty for PDEs with random inputs. Some of the commonly used risk functions, such as conditional value-at-risk and mean-upper-semideviation, depend only on samples of the quantity of interest that lie above given quantiles. A natural approach for efficient estimation of such risk functions is importance sampling. However, application of importance sampling is a non-trivial task. The fundamental issue in importance sampling is the choice of biasing distri-

bution. In this talk we propose an approach for estimating finite coherent risk functions based on their conjugate duality properties. Utilizing the concept of risk identifiers we construct the biasing distribution that naturally incorporates the information about the quantity of interest into the sampling process. The proposed approach allows us to considerably reduce the number of samples required to estimate risk functions of the quantities of interest, making this approach preferable for the problems which involve costly simulations, such as PDE-constrained optimization problems under uncertainty.

Timur Takhtaganov  
Rice University  
Houston, TX  
tat3@rice.edu

Drew P. Kouri  
Optimization and Uncertainty Quantification  
Sandia National Laboratories  
dpkouri@sandia.gov

Denis Ridzal  
Sandia National Laboratories  
dridzal@sandia.gov

Matthias Heinkenschloss  
Department of Computational and Applied Mathematics  
Rice University  
heinken@rice.edu

## CP11

### **A Weighted Kernel PCA-Based Method for Goal-Oriented Optimization and Uncertainty Quantification with Applications to Elastic Parameter Inversion**

Viewing instantiations of a channelized medium as realizations of a random field, we demonstrate that methods based on machine learning can be successfully applied to inverse problems in the context of linearly elasticity through the use of the adjoint method coupled with kernel PCA-based optimization. In addition, we show how the convergence of the linear variant of the method can be accelerated by weighting realizations of the random field by the value of the objective function and its gradient with respect to control variables in a low-dimensional feature space. As linear machine learning methods are inherently limited in their ability to represent features of non-Gaussian stochastic random fields, we extend this goal-oriented, weighted method using kernel PCA to capture nonlinear spatial relationships and multipoint statistics so as to properly characterize complex features such as the tortuosity of channelized media. We demonstrate how this Weighted Kernel PCA (WKPCA) based optimization both captures nonlinear correlations and exhibits accelerated convergence to a target instantiation of the random field. We demonstrate how a response surface can be constructed to furnish weights using Bayesian methods, demonstrate an application of the method to learning the spatial distribution of material parameter values in the context of linear elasticity, and discuss further extensions of the method to stochastic parameter inversion and uncertainty quantification.

Charles Talbot  
Lawrence Livermore National Laboratory  
University of North Carolina at Chapel Hill  
ctalbot@live.unc.edu

Chen Xiao, Charles Tong  
Lawrence Livermore National Laboratory  
chen73@llnl.gov, tong10@llnl.gov

Charanraj Thimmisetty  
Lawrence Livermore National Lab  
thimmisetty1@llnl.gov

Wenju Zhao  
Department of Scientific Computing  
Florida State University  
wz13@my.fsu.edu

## CP11

### **High-Dimensional Intrinsic Interpolation Using Gaussian Process Regression**

We introduce a regression procedure for intrinsic variables constrained onto a manifold embedded in an ambient space. The procedure is meant to sharpen high-dimensional interpolation by introducing constraints delineated from within the data being interpolated. Our method augments manifold learning procedures with a Gaussian process regression (GPR). The proposed method first identifies, using diffusion maps, a low dimensional manifold embedded in an ambient high dimensional space associated with the data. It relies on the diffusion distance associated with this construction to define proximity between points on the manifold. This distance is then used to compute the correlation structure of a Gaussian process that describes the dependence of quantities of interest in the high dimensional ambient space. The proposed method is applicable to arbitrarily high-dimensional datasets and is demonstrated in the context of characterization of the 3D seismic properties from the vertical seismic profile (VSP) logs. Lawrence Livermore National Laboratory: LLNL-ABS-702635-DRAFT This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Charanraj Thimmisetty  
Lawrence Livermore National Lab  
thimmisetty1@llnl.gov

Roger Ghanem  
University of Southern California  
Aerospace and Mechanical Engineering and Civil Engineering  
ghanem@usc.edu

Christina Morency  
Lawrence Livermore National Laboratory  
morency1@llnl.gov

Xiao Chen  
Lawrence Livermore National Laboratory  
Center for Applied Scientific Computing  
chen73@llnl.gov

Joshua A. White, Charles Tong  
Lawrence Livermore National Laboratory  
white230@llnl.gov, tong10@llnl.gov

## CP12

### **A Realizability Preserving Discontinuous Galerkin Method for Radiative Transport**

The  $M_1$  radiative transfer model plays an important role

in describing radiative transport in physics and engineering applications. The  $M_1$  model is a non-linear hyperbolic system that describes the distribution of radiative particles (e.g. photons) as they interact with matter. In this talk I will introduce a high order discontinuous Galerkin method for solving the  $M_1$  model. I will introduce the concept of realizability, discuss numerical challenges posed by the  $M_1$  model and show the modifications to the standard DG method that are necessary in order to guarantee stability and accuracy of the method. I will conclude with numerical experiments to demonstrate the effectiveness of the DG method.

Prince Chidyagwai  
Loyola University Maryland  
Department of Mathematics and Statistics  
pchidyagwai@loyola.edu

Florian Schneider  
Kaiserslautern University of Technology  
schneider@mathematik.uni-kl.de

Benjamin Seibold  
Temple University  
seibold@temple.edu

Martin Frank  
RWTH Aachen University  
Center for Computational Engineering Science  
frank@mathcces.rwth-aachen.de

### CP12 Adaptive Methods for Multidimensional Cubature

Computation of  $\int_{\mathbb{R}^d} f(\mathbf{x}) \varrho(\mathbf{x}) d\mathbf{x}$  arises in finance, high energy physics and other applications. Monte Carlo, quasi-Monte Carlo, and Bayesian cubature are among the methods used to evaluate such integrals. This talk describes recent theoretically justified adaptive versions of these algorithms, which choose the sample sizes to automatically satisfy user-defined error.

Fred J. Hickernell  
Illinois Institute of Technology  
Department of Applied Mathematics  
hickernell@iit.edu

### CP12 Dual Least-Squares Finite Element Methods for Hyperbolic Problems

Least-squares finite element methods have demonstrated good potential for solving problems of elliptic and parabolic type. We further study this type of methods for hyperbolic balance and conservation laws. Hyperbolic problems result in different challenges and possess solutions of very low regularity as they can have jump discontinuities. This research concentrates on negative norm and  $LL^*$  least-squares methods. The negative norm methods provide natural formulations in view of the definition of a weak solution. However, they only possess coercivity in a norm strictly weaker than the  $L^2$ -norm. Therefore,  $LL^*$  methods are used to obtain stronger coercivity in the  $L^2$ -norm. This is different from previous uses of  $LL^*$  methods. Their use, e.g., for elliptic problems is to improve the  $L^2$ -norm error (e.g. as part of a hybrid method) of the minimizer of a certain least-squares functional for problems that can already be approximated in a stronger norm. The purpose here is to obtain  $L^2$  coercivity for methods that are coercive only

in a strictly weaker norm.

Delyan Z. Kalchev, Thomas Manteuffel  
University of Colorado  
delyan.kalchev@colorado.edu, tmanteuf@colorado.edu

### CP12 Higher Order Numerical Schemes for Non-Conservative Hyperbolic Equations on 2D Unstructured Grids

Standard finite volume schemes consider the solution of hyperbolic equations in conservative form. Similar to for example the shallow water equations, many hyperbolic moment models that are used for the simulation of rarefied gases cannot be given in conservative form. This requires special non-conservative numerical schemes to accurately resolve shocks and reduce runtime. We describe high-order numerical methods for the solution of general non-conservative hyperbolic equations that have been implemented on unstructured quadrilateral grids. Especially the reconstruction and the slope limiting procedure are discussed. The schemes are compared to their first order versions with respect to accuracy and runtime for two-dimensional test cases of hyperbolic moment equations.

Julian Koellermeier, Manuel Torrilhon  
RWTH Aachen University  
koellermeier@mathcces.rwth-aachen.de,  
mt@mathcces.rwth-aachen.de

### CP12 Sobolev Discontinuous Galerkin Methods

The Discontinuous Galerkin (DG) method is a popular polynomial-based method for solving wave propagation problems known for its spectral accuracy and geometric flexibility. Despite its nice properties, DG suffers from a very restrictive time step size, at least inversely proportional to the order ( $p$ ) of the DG method. For a fixed order explicit time-stepping algorithms, the time step size ( $\Delta t$ ) is determined by the following Courant-Friedrichs-Lewy (CFL) condition:  $a(\Delta t/\Delta x) \leq C/p^2$ , where  $a$  is the maximum wave speed,  $\Delta x$  is the mesh size, and  $C$  is a method-specific constant. In this talk, we present Sobolev DG, a novel DG method that allows for order-independent CFL condition for sufficiently accurate time evolution. Sobolev DG tests with low degree polynomials while maintaining the accuracy of high order polynomials by using different variational equations. As a consequence, Sobolev DG can take at least  $p$  times larger time steps compared to traditional DG methods. We will describe Sobolev DG method and present computational results illustrating its excellent time stepping properties.

Adeline Kornelus, Daniel Appelo  
University of New Mexico  
kornelus@unm.edu, appelo@math.unm.edu

Thomas M. Hagstrom  
Southern Methodist University  
Department of Mathematics  
thagstrom@smu.edu

### CP12 Life and Its Many Layers: Eigenfunctions of 1-D Differential Operators with Piecewise Constant Co-

**efficients**

Today's computers are able to perform simulations at higher resolutions. However, as the spatial resolution increases, using current simulation methods, such as time-stepping methods for solving PDEs, the time step must decrease which drastically increases computation time. Mathematically speaking, the time-stepping methods that are available today are not scalable when applied to stiff problems. This research project will be exploring solutions of time-dependent PDEs in media consisting of more than two materials. The real world applications of this project will be to simulate physical phenomena such as waves propagating through air, water, and solids. Due to the discontinuity at the interface between different materials, it is far more challenging to find the natural frequencies that simplify simulation because they change between materials. The Uncertainty Principle provides a geometrically intuitive approach to obtain these natural frequencies and will help in finding an initial approximation of them. Life is very layered and the algorithms we will develop can more accurately and precisely account for it.

Sarah D. Long

The University of Southern Mississippi  
sarah.d.long@usm.edu

James V. Lambers

University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

**CP12****Monotone Local Projection Stabilization for Non-linear Hyperbolic Systems**

We propose a stabilized finite element scheme for hyperbolic systems on unstructured meshes with continuous Galerkin (CG) discretization in space. Starting with any hyperbolic system, we introduce a CG semi-discrete scheme with low order artificial dissipation that satisfies the local extremum diminishing (LED) condition for systems. To be high resolution, we propose a gradient based element limiting strategy that limits any anti-diffusion added to the low order scheme, while preserving the LED condition and also being linearity preserving. To tackle any possibility of phase errors or terracing, we introduce linearity preserving high order dissipation in the antidiffusive correction. The limiting procedure being at the semi-discrete level allows for an LED semi-discrete scheme which can be discretized in time using any time integrator, only requiring a CFL type of condition for explicit schemes. We demonstrate the robustness of this numerical scheme through several examples including: linear transport, nonlinear Burgers equation, KPP problem, and several Euler equations hydrodynamics simulations. For time discretization, we use Crank-Nicholson, backward Euler, SSP Runge-Kutta, and SDIRK schemes. The next stage is to consider multifluid plasma physics applications, especially Z-pinch implosions. **collaboration with John Shadid and Dmitri Kuzmin**

Sibusiso Mabuza

University of Houston  
smabuza@sandia.gov

**CP12****A Godunov-Type Finite-Volume Solver for Non-hydrostatic Euler Equations with a Time-Splitting****Approach**

A two-dimensional nonhydrostatic (NH) model based on the compressible Euler system has been developed in the Cartesian  $(x, z)$  domain. The spatial discretization is based on a Godunov-type finite-volume (FV) method employing dimensionally split fifth-order reconstructions. The model uses the explicit strong stability-preserving Runge-Kutta scheme and a split-explicit method. The time-split approach is generally based on the split-explicit method, where the acoustic modes in the Euler system are solved using small time steps, and the advective modes are treated with larger time step. However, for the Godunov-type FV method this traditional approach is not trivial for the Euler system of equations. In the present study, a new strategy is proposed by which the Euler system is split into three modes and a multi-rate time integration is performed. The computational efficiency of the split-scheme is compared with the explicit scheme using the FV model with various NH benchmark test cases.

Farshid Nazari

NCAR Earth System Laboratory  
fnazari@ucar.edu

Ram Nair

NCAR  
Institute for Mathematics Applied to Geosciences  
rnair@ucar.edu

**CP13****Efficient Multilevel Methods for Optimal Control of Elliptic Equations with Stochastic Coefficients**

A common strategy for solving optimal control of stochastic PDEs relies on stochastic collocation, which reduces the problem to multiple solves of optimal control problems constrained by deterministic PDEs. In this work we investigate an alternative approach where we use a stochastic Galerkin formulation and discretization of the PDE prior to solving the optimal control problem. Ultimately this requires solving a potentially very large linear system, which we then solve using specially designed multilevel algorithms.

Sumaya Alzuhairy

UMBC  
alsum1@umbc.edu

Andrei Draganescu

Department of Mathematics and Statistics  
University of Maryland, Baltimore County  
draga@umbc.edu

Bedrich Sousedik

University of Maryland, Baltimore County  
sousedik@umbc.edu

**CP13****Reduced Order Models for Symmetric Eigenvalue Problems**

Reduced order models are becoming an indispensable tool in large scale numerical simulation. The ever-increasing demand for accuracy entails overwhelming demands on computational resources. In model reduction much smaller models are generated that admit fast and accurate approximations to the original problem. Here we are concerned with eigenvalue problems  $Ax = \lambda x$  where  $A = A(s)$  is a



symmetric matrix depending on a parameter  $s$ . The matrix  $A$  may be obtained from the approximation of an elliptic operator by finite elements,  $s$  may represent physical parameters. We are particularly interested in shape optimization where  $s$  represents in some way the shape of the computational domain. Large eigenvalue problems have to be solved for many values of  $s$ . We investigate reduced order models that admit to cheaply compute accurate approximations of the most important eigenvalues of  $A(s)$ . Reduced order models allow a fast scan of (parts of) the parameter space in order to cheaply localize optima of some objective function. A standard optimization procedure or a zoom into the region of the optimum of may follow. In our investigation the objective is to determine a shape  $s$  such that  $A(s)$  has some prescribed (smallest) eigenvalues.

Peter Arbenz  
ETH Zurich  
Computer Science Department  
arbenz@inf.ethz.ch

Zlatko Drmac  
University of Zagreb  
Department of Mathematics  
drmac@math.hr

### CP13

#### Parallel-In-Time Gradient-Type Method For Optimal Control Problems

We present a new parallel-in-time gradient-type method for time-dependent optimal control problems. When the classic gradient method is applied to time-dependent optimal control problems, in each optimization iteration, before control update, the forward solution of state equations and the backward solution of adjoint equations are required, which consumes most of the computation time. This new parallel-in-time gradient-type method introduces parallelism by splitting the time domain into  $N$  time subdomains and executes the forward and backward computation in each time subdomain in parallel using state and adjoint variables at time subdomain boundaries from last optimization iteration as the initial value so that each gradient-type iteration takes roughly  $1/N$  of the gradient iteration time. For Linear-Quadratic optimal control problems, the resulting iteration can be interpreted as a  $(2N-1)$ -part splitting iteration scheme. Convergence proof is given. Numerical examples on Linear-Quadratic problems and general nonlinear problems, such as water flooding optimization in oil reservoir recovery, demonstrate significant speedup. Strong scaling is achieved in some cases for a moderate number of time subdomains.

Xiaodi Deng  
Rice University  
xiaodi.deng@rice.edu.

Matthias Heinkenschloss  
Department of Computational and Applied Mathematics  
Rice University  
heinken@rice.edu

### CP13

#### Mode Reduction Methods for Data Assimilation: Subspace Projection Using Koopman Operators

In order to reduce the computational burden of assimilating data into large-scale systems, spectral decomposition

methods are used to to define a subspace that reduces the dimension of the problem. In this talk we use a recent decomposition technique based on the Koopman operator and present how it applied to data assimilation methods. We use the eigenmodes defined by the Koopman operator that represent the non-linear behavior of a dynamical system for the assimilation of data into a shallow water model, and compare its performance with other subspace projections methods.

Humberto C. Godinez  
Los Alamos National Laboratory  
Applied Mathematics and Plasma Physics  
hgodinez@lanl.gov

Nick Hengartner  
Los Alamos National Laboratory  
nickh@lanl.gov

### CP13

#### Uncertainty Quantification in High Frequency Electromagnetics Using the Low Rank Tucker Decomposition

The analysis of high frequency components, e.g. waveguides, is based on the numerical solution of Maxwell's equations. Finite difference and finite element schemes are typically employed for time- and frequency-domain analysis. Despite their accuracy, these approaches assume the precise knowledge of the geometry and are therefore incapable of quantifying uncertainties due to manufacturing. These uncertainties should not be neglected, as hyperbolic equations can be very sensitive to geometric deviations. The stochastic collocation (SC) method (Xiu/Hesthaven, 2005) is one of the most reliable uncertainty quantification (UQ) methods available. In the multi-parametric case, a tensor-product SC scheme is employed. The resulting tensors are ideal for efficient algebraic computations, but their storage and computational complexities grow exponentially w.r.t. the number of parameters. The current state-of-the-art in reducing such complexities is based on sparse grids (SG). Recently, tensor decomposition methods (Kolda/Bader, 2009) arised as an alternative to SG, especially in the presence of many parameters. In this work, the Tucker decomposition is employed for a waveguide model with geometrical uncertainties. While a complexity reduction better than SG-SC is not expected, the Tucker format is the basis of hierarchical decompositions, which are suitable for high dimensional problems. Such approaches shall be investigated in an extension of this work.

Dimitrios Loukrezis  
Institut Theorie Elektromagnetischer Felder, TU Darmstadt  
Graduate School for Computational Engineering, TU Darmstadt  
loukrezis@temf.tu-darmstadt.de

Ulrich Roemer  
Institut fuer Theorie Elektromagnetischer Felder  
Technische Universitaet Darmstadt  
roemer@temf.tu-darmstadt.de

Herbert De Gerssem  
Institut Theorie Elektromagnetischer Felder, TU Darmstadt  
Graduate School Computational Engineering, TU Darmstadt

degersem@temf.tu-darmstadt.de

### CP13

#### A Dictionary Learning Strategy for Bayesian Inference

Inference of model parameters or spatially-distributed properties is an important and difficult step in many engineering processes. Information often comes from only a few spatially-localized measurements, nonlinearly related to the quantity of interest, resulting in a very challenging ill-posed problem. A typical example is the inference of subsurface resources from scarce surface measurements. Fortunately, these high-dimensional quantities are often sparse in properly chosen bases so that dedicated recovery algorithms such as compressed sensing can be employed. We introduce a novel approach which does not rely on an a priori choice of basis for approximating the parameter field. From a prior set of realizations, one seeks a basis (dictionary) such that each realization is likely to admit a sparse representation. A critical additional aspect is that this basis has to be *observable by the sensors*, so that the basis modes can indeed be informed by the data. In the particular context of Bayesian inference with a linear model, this problem formulates via constructive sparse Bayesian learning and leads to lower variance estimates. Our approach will be illustrated with several examples, such as the recovery of the inhomogeneous diffusivity field from a few sensors. Thanks to the enforced observability property, it will be shown to significantly outperform current methods, such as those based on K-SVD.

Lionel Mathelin  
LIMSI - CNRS, France  
mathelin@limsi.fr

### CP13

#### Simultaneous Estimation of Material Parameters and Neumann Boundary Conditions in a Linear Elastic Model by PDE-Constrained Optimization

We describe a formulation of and solution strategy for a simultaneous inversion problem in linear elasticity. Given point measurements of displacement, we concurrently estimate spatially-varying shear modulus and traction boundary condition fields by posing the inverse problem as a PDE-constrained optimization problem. Regularization is used to relax modeling assumptions (i.e. assumed parametric forms for the solution fields) that are often made when solving inverse problems. We present results obtained by application of our technique to both synthetic and experimental data.

Daniel T. Seidl, Bart G. Van Bloemen Waanders,  
Timothy Wildey  
Sandia National Laboratories  
dtseidl@sandia.gov, bartv@sandia.gov,  
tmwilde@sandia.gov

### CP13

#### Multilevel Monte Carlo Methods for Eigenvalue Problems

We consider multilevel Monte Carlo methods for accelerating the computation of eigenvalue problems arising from PDEs in the presence of uncertainty. Eigenvalue problems are important in a number of applications where randomness in the underlying system has a significant impact on

the response, such as structural vibration and buckling. Monte Carlo methods are simple for quantifying the impact of uncertainties on eigenvalue problems, but are considered to be too expensive in many practical cases. We propose employing multilevel Monte Carlo methods to accelerate the computation of the statistics of eigenvalue problems. This involves sampling of 'coarse grid' representations of the problem of interest. We will argue that the efficiency gains of the multilevel method make the approach competitive with some less flexible methods that have been developed as alternatives to Monte Carlo methods for eigenvalue problems, and specifically for structural vibration.

Juliette Unwin  
University of Cambridge  
hjt2@cam.ac.uk

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

Nathan Sime  
University of Cambridge  
njcs4@cam.ac.uk

### CP14

#### Fluctuating Hydrodynamic Methods for Manifolds: Applications to Microstructures Within Curved Fluid Interfaces

We develop stochastic computational methods using the exterior calculus of differential geometry for hydrodynamic interactions of microstructures embedded within curved fluid interfaces. For manifolds of spherical topology we develop spectral methods for the hydrodynamic equations based on hyperinterpolation and use of Lebedev quadrature. We show how geometric and topological effects can significantly augment hydrodynamic flow within curved interfaces. We then present results showing how these effects contribute to the individual and collective motions of particle inclusions and their drift-diffusion. Overall, we present general fluctuating hydrodynamic computational methods for manifolds that can be used for investigations of many-body systems involving diffusion and hydrodynamic coupling within thin fluid interfaces having curved geometries.

Paul J. Atzberger  
University of California-Santa Barbara  
atzberg@gmail.com

### CP14

#### Space-Time Least-Squares Petrov-Galerkin Nonlinear Model Reduction

Reduced-order models (ROMs) of nonlinear dynamical systems are essential for enabling computational high-fidelity models (HFMs) to be used in many-query and real-time applications such as uncertainty quantification and design optimization. Majority of existing ROMs explore only spatial domain, not temporal domain. Although there are some space-time ROMs that explore both space and time domains, they depend on space-time HFM formulation, which is not readily available in known commercial and academic simulation codes. We introduce a POD projection-based space-time ROM that does not rely on space-time HFMs. Instead it exploits left as well as right singular vectors to form a spatio-temporal state basis. We derive error bounds that show the method exhibits slower

time-growth of the error as compared to typical spatial-projection-based ROMs. Numerical results show that the proposed technique enables additional computational savings to be realizable with similar accuracy as compared to typical spatial-projection-based ROMs.

Youngsoo Choi, Kevin T. Carlberg  
Sandia National Laboratories  
ychoi@sandia.gov, ktcarl@sandia.gov

#### CP14

##### **A Robust Discontinuous Galerkin Scheme for Ten-Moment Gaussian Closure Equations**

Euler equation for compressible flow treats pressure as a scalar. However, for several applications this description of pressure is not suitable. Several extended model based on the higher moments of Boltzmann equations are considered to overcome this issue. One such model is Ten moment Gaussian closure equations, which treats pressure as symmetric tensor and allow anisotropic effects. In this work, we develop a higher-order, positivity preserving discontinuous Galerkin scheme for Ten-moment Gaussian closure equations. The key challenge is to preserve positivity of density and pressure tensor. This is achieved by constructing a positivity limiter to ensure positivity of density and pressure tensor. In addition to preserving positivity, the scheme also ensures the accuracy of the approximation. The theoretical results are then verified using several numerical experiments. Furthermore, these results can be extended to achieve positivity of pressure tensor in other extended models like 13-moment equations.

Asha K. Meena  
Indian Institute of Technology Delhi  
ashameena01@gmail.com

Praveen Chandrashekar  
TIFR Centre for Applicable Mathematics  
Bangalore  
praveen@tifrbng.res.in

Harish Kumar  
Indian Institute of Technology Delhi  
hkumar@maths.iitd.ac.in

#### CP14

##### **Accelerating Multiscale Coupling Methods in Heterogeneous Adjacent Multidomains with Shear Flow**

We examine the numerical treatment of dissipative particle dynamics (DPD) and momentum-conserving molecular dynamics, which are significantly more expensive than continuum methods (e.g., Navier–Stokes) in multiscale coupling. The largest integration stepsize of traditional DPD methods is limited due to the discretization error as well as the presence of nonequilibrium perturbations (e.g., shear flow), either of which can be thought of excess energy pumped into the system preventing it from maintaining the desired temperature. We demonstrate that it is possible to improve both the accuracy and the stability of DPD by employing a pairwise adaptive Langevin thermostat that precisely matches the dynamical characteristics of DPD simulations (e.g., autocorrelation functions) while automatically correcting thermodynamic averages using a negative feedback loop. We also derive the associated stationary distribution when the system is subject to a general class of nongradient external perturbations, which provides theoretical founda-

tions for nonequilibrium modelling. Our findings are verified by numerical experiments.

Xiaocheng Shang  
University of Edinburgh  
x.shang@brown.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

#### CP14

##### **Solution of PDEs for Photobleaching Kinetics Using Krylov Subspace Spectral Methods**

We show how to solve the first order photobleaching kinetics partial differential equations with prebleach steady state initial conditions using a time-domain method known as a Krylov Subspace Spectral method (KSS method). KSS are explicit methods for solving time-dependent variable-coefficient partial differential equations (PDEs). KSS methods are advantageous compared to other methods because of its high resolution and its superior scalability. We will apply Gaussian Quadrature rules in the spectral domain developed by Golub and Meurant to solve PDEs. We present a simple rough analytical solution, as well as a computational solution that is first-order accurate. We then use this solution to examine short and long time behaviors.

Somayyeh Sheikholeslami  
Ph.D Student  
s.sheikholeslami@usm.edu

James V. Lambers  
University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

#### CP14

##### **Algebraic Analysis and Numerical Illustration on Artificial Odd-Even Grid Oscillation and Its Presence in Domain Decomposition Computation**

Although artificial odd-even grid oscillation associated with central difference has been a classic topic in research, still there is lack of a direct and complete explanation on its onset and clear understanding of its behaviors. This presentation makes a systematic analysis in conjunction with numerical illustration on such oscillation in solutions of a model problem, presenting a plain, straightforward explanation as well as criteria on its presence and behavior. Two types of odd-even grid oscillation are identified; one comes from dual-mode patterns in numerical solutions, and the other results from over-imposing of boundary conditions. The first type of the oscillation decays with grid spacing, while the second one tends to remain regardless of its size. As a consequence of their presence in single-domain solutions, the two kinds of fluctuation also occur in computation by domain decomposition, but they are affected by algorithms of the decomposition. Further analysis demonstrates that the fluctuation inherited in the model problem also leads to zig-zag forms of solutions for more complicated nonlinear flow problems whether they are solved in a single domain or two subdomains using central difference.

Hansong Tang

Dept. of Civil Eng., City College, City Univ. of New York  
htang@ccny.cuny.edu

Wenbin Dong  
Dept. of Civil Eng., City College of New York, CUNY  
wdong000@citymail.cuny.edu

#### CP14

##### Curvature-Augmented Numerical Methods for Interfaces and Surfaces

Whether considering surface tension at a fluids interface, bending resistance of a surface, electric potential on a surface, or diffusion across a surface, having robust numerical methods to handle interfaces and surfaces is key to simulating a wide variety of phenomena. In the case of smooth interfaces/surfaces, curvature usually plays an important role in the governing dynamics, either directly or indirectly. Here, the curvature is also used to augment numerical methods for tracking interfaces/surfaces (level set method) and for solving PDEs on a surface (closest point method). For the level set method, using a velocity extension is vital in maintaining the signed-distance property. The traditional velocity extension is augmented with curvature information, leading to substantially improved mass conservation. For the closest point method, the discretization of the embedded PDE is augmented with a variable-coefficient Laplacian stencil, which depends on the curvature. The new approach has better accuracy and sparsity than the traditional method. To tie these two augmented methods together, some results from vesicle simulations are presented.

Chris Vogl  
University of Washington  
cvogl@uw.edu

#### CP14

##### Patient-Specific 3D Reconstruction of Biore-sorbable and Metallic Stent: A Critical Step for CFD Analysis of Hemodynamics in Stented Coronary Artery

This work provides a pipeline to perform patient-specific quantitative analysis of coronary arteries treated with Biore-sorbable Vascular Stents (BVS). The struts of BVS, which are thicker than that of regular metallic devices, may trigger local alterations in the blood stream. Using Optical Coherence Tomographies (OCT) and Angiographies, we are able to reconstruct the 3D geometry of deployed BVS in vivo. We illustrate the automatic strut detection from OCT, the identification of the different stent components, the 3D registration of the stent curvature using OCT wire trajectory derived from Angiogram, and the final stent volume reconstruction. The semi-automatic stent reconstruction is validated against clinical cases and a virtual phantom. The reconstruction method can also be modified to adapt the metallic stent reconstruction, and the current progress on metallic stent reconstruction is also presented. Our methodology allows reliable and timely reconstruction of patient-specific BVS using multimodality image analysis. This work is the first step toward a fully automated reconstruction and simulation procedure. We have successfully used the reconstructed to conduct computational fluid dynamics simulations for 16 patients, some of which will be presented. The reconstructed geometries enable an extensive quantitative analysis of the impact of BVS struts

on hemodynamics via computational fluid dynamics.

Boyi Yang  
Emory University  
Department of Mathematics & Computer Science  
byang8@emory.edu

#### CP15

##### Multiscale Methods for Flow and Transport in Porous Media

Large scale simulations are performed to evaluate various scenarios such as CO<sub>2</sub> sequestration, contaminant plume migration as well as hydrocarbon recovery in subsurface porous media. The reservoir properties relevant to flow and transport modeling is often available at fine scale either from well logs or inferred from geological models. However, incorporating this fine scale data is computationally expensive and prohibitive in evaluating a scenario. Additionally, modeling complex physical processes result in restrictive time-step sizes due to convergence issues associated with non-linear iterations. In this work, we present multiscale approaches, both in the spatial and temporal sense, to alleviate these issues in favor of computational efficiency. We use a two-scale homogenization approach for upscaling single phase, slightly compressible, flow and transport in porous media from a fine to a coarse spatial scale. Furthermore, an adaptive, time discrete, domain decomposition scheme is proposed which allows for different time-step sizes in different domains. A lowest-order mixed finite element scheme is used for spatial discretization with rectangular parallelepiped, RT0 elements. We verify our approach by comparing the numerical results against a known analytical or benchmark fine scale solution for accuracy and computational efficiency. Numerical experiments are also presented for extensions to field scale studies for evaluating scenarios of practical interest.

Yerlan Amanbek  
The Institute for Computational Engineering and Sciences  
yerlan@ices.utexas.edu

Gurpreet Singh  
The University of Texas at Austin  
Center for Subsurface Modeling  
gurpreet@ices.utexas.edu

Mary F. Wheeler  
Center for Subsurface Modeling, ICES  
University of Texas at Austin  
mfw@ices.utexas.edu

#### CP15

##### Abstract Framework for Separable FWI Problems

We present an abstract framework for separable optimization problems and illustrate its use in solving Full Waveform Inversion (FWI) problems in seismology. This framework is part of the Rice Vector Library (RVL), a C++ package that provides mathematics-emulating abstract classes for expression of linear algebra and nonlinear optimization algorithms. The qualifier "separable" describes least-squares problems in which the simulation operator is a linear function of a subset of model parameters. For example, parameter perturbations and sources (right-hand sides in wave equations) are linear parameters in linearized ("Born") inversion and source-medium parameter estimation respectively. We define the abstract interface LinOpValOp (LOVOp) representing operators whose prod-

uct domain is split into nonlinear and linear subsets. This LOVOp class enables abstract expressions of algorithms for separable problems, such as the Variable Projection Method, while allowing for efficient implementation. As a subclass of the Operator (vector valued nonlinear function) class, LOVOp can be used in any setting in which an Operator can be used, thus opening up many other algorithmic possibilities. In this presentation, we will briefly describe the software structure of RVL for separable problems, and illustrate its use in the solution of some typical inverse problems from seismology.

Mario Bencomo

Department of Computational and Applied Math  
Rice University  
mario.j.bencomo@rice.edu

William Symes  
Rice University  
symes@caam.rice.edu

### CP15

#### Optimal Reconstruction of Constitutive Relations for Porous Media Flows

Comprehensive full-physics models for flow in porous media typically involve convection-diffusion partial differential equations whose parameters are unknown and have to be reconstructed from experimental data. Quite often these unknown parameters are coefficients represented by space-dependent, sometimes correlated, functions, e.g. porosity, permeability, transmissibility, etc. However, special complexity is seen when the reconstructed properties are considered as state-dependent parameters, e.g. the relative permeability coefficients  $k_{rp}$ . Modern petroleum reservoir simulators still use simplified approximations of  $k_{rp}$  as single variable functions of  $p$ -phase saturation  $s_p$  given in the form of tables or simple analytical expressions. This form is hardly reliable in modern engineering applications used, e.g., for enhanced oil recovery, carbon storage, modeling thermal and capillary pressure relations. Thus, the main focus of our research is on developing a novel mathematical concept for building new models where  $k_{rp}$  are approximated by multi-variable functions of fluid parameters, namely phase saturations  $s_p$  and temperature  $T$ . Reconstruction of such complicated dependencies requires advanced mathematical and optimization tools to enhance the efficiency of existing engineering procedures with a new computational framework generalized for use in various earth science applications.

Vladislav Bukshtynov

Florida Institute of Technology  
Department of Mathematical Sciences  
vbukshtynov@fit.edu

### CP15

#### A Finite Element Flow and Transport Model in Porous Media for Enhanced Oil Recovery Tests at Laboratory Scale

In this work, a general flow and transport model in porous media to simulate, analyze and interpret enhanced oil recovery tests at core scale under laboratory conditions is presented. The flow model is based on the oil phase pressure and total velocity formulation, in which the capillary pressure, relative permeabilities, the effects of gravity and the dynamic porosity and permeability modification are allowed. Whereas, the transport model includes physical-

chemical phenomena such as advection, diffusion, dispersion, adsorption-desorption, and reactions. For the numerical solution is applied a finite element method and its computational implementation was carried out in Python using FeniCS project. From the methodological point of view, each stage of model development (conceptual, mathematical, numerical and computational) is described. Finally, the resulting coupled flow and transport model is numerically validated in a case study.

Martin A. Diaz-Viera

INSTITUTO MEXICANO DEL PETROLEO  
mdiazv@imp.mx

Eduardo Linares-Perez, Mario Noyola-Rodriguez  
Posgrado en Ciencias de la Tierra  
Universidad Nacional Autonoma de Mexico  
eduardo\_linares@comunidad.unam.mx, estadua-  
gro@hotmail.com

### CP15

#### Domain-Specific Abstractions for Full-Waveform Inversion

Full-waveform inversion is a PDE-constrained optimisation problem involving massive amounts of data (petabytes) and large numbers of unknowns ( $O(10^9)$ ). This well known compute-intensive and data-intensive is extremely challenging for several reasons. First, there is the complexity of having to handle extremely large data volumes with metadata related to experimental details in the field, and the discretization of the unknown earth parameters and approximate physics. Second, reduced or adjoint-state methods call for computationally intensive PDE solves for each source experiment (of which there are thousands) for each iteration of a gradient-based optimization scheme. The talks will give an overview how carefully chosen layers of abstraction can help manage both the complexity and scale of inversion while still achieving the high degree of computational performance required to make full-waveform a practical tool. Specifically, the presentations will focus on domain specific stencil language for time-stepping methods to solve various types of wave equations and on abstracts for large-scale parallel optimization frameworks.

Felix J. Herrmann

Seismic Laboratory for Imaging and Modeling  
The University of British Columbia  
fherrmann@eos.ubc.ca

### CP15

#### Adaptive BDDC for Flow in Heterogeneous Porous Media

We present a method based on Balancing Domain Decomposition by Constraints (BDDC) for a numerical solution of a single-phase flow in heterogenous porous media. The method solves for both flux and pressure variables. The fluxes are resolved in three steps: the coarse solve is followed by subdomain solves and last we look for a divergence-free flux correction and pressures using conjugate gradients with the BDDC preconditioner. Our main contribution is the application of an adaptive algorithm for selection of flux constraints. Performance of the method is illustrated on benchmark problems including SPE10.

Bedrich Sousedik

University of Maryland, Baltimore County

sousedik@umbc.edu

## CP16

### Fluidity Based Approach to Modeling Ice Sheets

This talk describes a nonlinear full-Stokes method for modeling ice sheets that avoids the difficulty of infinite viscosity introduced by Glen's Flow law. The model is framed in a First-order System Least-squares (FOSLS) type formulation and solved using a Nested Iteration (NI), Newton-FOSLS-AMG approach in which the majority of the work is done on coarse grids. A couple of test problems are presented to show the effectiveness of this NI-Newton-FOSLS-AMG approach.

Jeffery M. Allen  
University of Colorado at Boulder  
jeffery.allen@colorado.edu

## CP16

### Accurate and Stable Time Stepping in Ice Sheet Modeling

We introduce adaptive time step control for simulation of evolution of ice sheets. The discretization error in the approximations is estimated using "Milne's device" by comparing the result from two different methods in a predictor-corrector pair. The time step is changed in an efficient way that the velocity field equation is only solved once per time step. The stability of the numerical solution is maintained and the accuracy is controlled by keeping the local error below a given threshold using PI-control. Our method takes a shorter time step than an implicit method but with less work in each time step and the solver is simpler. The method is analyzed theoretically with respect to stability and applied to the simulation of a 2D ice slab and a 3D circular ice sheet. The stability bounds in the experiments are explained by and agree well with the theoretical results.

Gong Cheng  
Department of Information Technology, Uppsala University  
cheng.gong@it.uu.se

Per Lötstedt, Lina von Sydow  
Uppsala University  
per.lotstedt@it.uu.se, Lina.von.Sydow@it.uu.se

## CP16

### Numerical Shape Optimization to Decrease Failure Probability of Ceramic Structures

Ceramic is a material frequently used in industry because of its favorable properties. Due to its popularity shape optimization is needed in these industrial applications. Common approaches in shape optimization for ceramic structures concerning tensile loading aim to minimize the stress acting on the component, as it is the main indicator for failure. In contrast to this we follow a more natural approach by minimizing the component's probability of failure. For this purpose and in full consideration of fracture mechanics matters, the objective functional describing this probability of failure for a given ceramic component under one tensile load was recently established. With this objective functional the problem is stated as a PDE constrained optimization problem. To solve the minimization problem, we choose a gradient based method combined with a

first discretize then optimize approach. For discretization finite elements are used. By converting the objective functional with the help of the Lagrange function we are able to calculate the shape gradient. The implementation was verified by comparison of it with a finite difference method applied to a minimal 2d example. This demonstrated that the calculated shape gradient actually points towards the direction of the optimal shape in terms of the failure probability. In this talk we present the objective functional as well as our implementation and some numerical examples.

Camilla Hahn, Matthias Bolten  
Universität Kassel  
chahn@mathematik.uni-kassel.de,  
matthias.bolten@mathematik.uni-kassel.de

Hanno Gottschalk  
University of Wuppertal  
FBC - AG Stochastics  
hanno.gottschalk@uni-wuppertal.de

## CP16

### Hurricane Uncertainty Propagation for Real-Time Storm Surge Forecasting

In recent decades, computational hurricane storm surge models have become increasingly accurate due to improvements in numerical methods, improved model parameters, and implications of advancements in high performance computing, e.g. the ability to more finely resolve spatial domains. Unfortunately, when storm surges are forecasted in real-time, many uncertainties are introduced due to the uncertainties in the hurricane (i.e. wind) forecast. To aid in emergency response, an ensemble of hurricane scenarios is used to determine a range of possibilities of storm surge inundation. However, the likelihood of each scenario remains unclear. Here, we use the uncertainties associated with various storm parameters to estimate these likelihoods. We also investigate the evolution of the likelihoods as the uncertainty in the hurricane forecasts is reduced.

Talea Mayo  
Princeton University  
talea.mayo@ucf.edu

## CP16

### An Implicit Approach to Phase Field Modeling of Solidification for Additively Manufactured Materials

We develop a fully-coupled, fully-implicit approach to phase field modeling of solidification for additively manufactured materials. Predictive simulation of solidification in pure metals and alloys remains a significant challenge in the field of materials science, as micro-structure formation during the solidification of a material plays an important role in the properties of the solid material. Our approach consists of a finite element spatial discretization of the fully-coupled nonlinear system of partial differential equations at the microscale, which is treated implicitly in time with a preconditioned Jacobian-free Newton-Krylov (JFNK) method. The approach allows timesteps larger than those restricted by the traditional explicit CFL limit on structured and unstructured 2D and 3D meshes, is algorithmically scalable and efficient due to an effective preconditioning strategy.

Chris Newman, Marianne M. Francois  
Los Alamos National Laboratory

cnewman@lanl.gov, mmfran@lanl.gov

#### CP16

##### A Fast Direct Solver for Boundary Integral Equations in Evolving 2D Domains

Recently, fast direct solver has emerged as a promising technique to solve the dense linear systems resulting from discretizing the boundary integral equations. One advantage it has over traditional techniques is its efficiency in handling multiple right hand sides. However, many engineering applications require solving a large number of problems with slightly different geometries, instead of different right hand sides. This talk is going to introduce a new method that recycles the structural information and hence reduces the cost of rebuilding the direct solver for new geometries. A preliminary analysis based on kernel free FMM will be presented to justify the method. Finally, We will compare the method to several alternative options.

Tianyu Qiu  
Rice University  
tq4@rice.edu

Adrianna Gillman  
Rice University  
Department of Computational and applied mathematics  
adrianna.gillman@rice.edu

#### CP17

##### Comparison of Different Algorithms for Biomolecule Simulations

Biomolecule simulation emerges with the fast developments of computational biology in 1990s. It now investigates the structure and mutual interaction between biomolecules through mathematical models and numerical modeling. The most popular model for biomolecule simulation is described by the nonlinear Poisson-Boltzmann (PB) equation, and it has many challenges in numerical simulation, such as 3D complex protein geometry, strong nonlinearity, and singular charge source, etc. In this talk, the methods of CG and DG have been used to solve the PB equation. They have both been realized on nodal based FEM. The methods are explained first, and then detailed comparisons including stability, time step, speed, etc have been conducted. Some conclusions on the future work will be given.

Weishan Deng, Xiaohu Zhufu  
Institute of Software, Chinese Academy of Science  
417408020@qq.com, xiaohu@iscac.ac.cn

Jin Xu  
Institute of Software, ISCACS  
Chinese Academy of Science, China  
xu.jin@iscas.ac.cn

#### CP17

##### Influence of Scaffold Permeability on Tissue Growth in a Perfusion Bioreactor

As the field of tissue engineering evolves, there is growing interest in determine how features of the porous scaffold used in many perfusion bioreactors can be tuned to achieve desired results. In this talk we present a mathematical model to analyze how scaffold permeability may be engineered to control outcomes. A general model is proposed in which tissue proliferates on an underlying scaffold

of spatially-varying permeability. Several scenarios of possible experimental relevance are considered.

Daniel Fong  
United States Merchant Marine Academy  
Department of Mathematics and Science  
fongd@usmma.edu

Linda Cummings  
New Jersey Institute of Technology  
linda.cummings@njit.edu

Jeff Pohlmeier  
New Jersey Institute of Technology  
Department of Mathematical Sciences  
jp99@njit.edu

#### CP17

##### Mathematical Modeling and Numerical Simulation of Drug Delivery in the Coronary Arterial Wall Using Drug Eluting Stent

In recent years, mathematical modeling of cardiovascular drug delivery systems has become an effective tool to gain deeper insights into the cardiovascular diseases like atherosclerosis. In the case of coronary biodegradable stent which is a tiny expandable metallic mesh tube covered by biodegradable polymer, it leads to deeper understanding of drug release mechanisms from polymeric stent into the arterial wall. A coupled non-Fickian model of a cardiovascular drug delivery system using a biodegradable drug-eluting stent is proposed in this talk. Energy estimate using variational formulations is used to study the qualitative behavior of the model. The influence of arterial stiffness in the sorption of drug eluted from the stent is analyzed. The numerical results are obtained using finite element method in the weak form by COMSOL Multiphysics. The results in this presentation will open new perspectives to adapt the drug delivery profile to the needs of the patient.

Jahed Naghipoor  
Buaahuas University of Weimar  
Jahed.naghipoor@uni-weimar.de

Jose Augusto Ferreira  
CMUC, Department of Mathematics, University of Coimbra  
Coimbra, Portugal  
ferreira@mat.uc.pt

Lino Goncalves  
Department of Cardiology, Hospital of the University of Coimbra, Coimbra, Portugal  
lgoncalv@ci.uc.pt

Timon Rabczuk  
Institute of Structural Mechanics, Bauhaus University of Weimar, Weimar, Germany  
timon.rabczuk@uni-weimar.de

Paula de Oliveria  
CMUC, Department of Mathematics, University of Coimbra  
Coimbra, Portugal  
poliveir@mat.uc.pt

#### CP17

##### Educational Modules on HPC Bioinformatics Al-

**gorithms**

In 2015, Fulbright Specialists Angela Shiflet (computer science and mathematics) and George Shiflet (biology) participated in a three-week collaborative project at University “Magna Græcia” of Catanzaro in Italy, in the Dept. of Medical and Surgical Sciences, hosted by Mario Cannataro. The three, along with Pietro Guzzi and NSF-funded Blue Waters Interns Daniel Couch and Dmitriy Kaplan, started a project to develop educational modules on high performance computing (HPC) bioinformatics algorithms. The undergraduate interns have written sequential and HPC programs and performed timings to accompany modules, such as “Aligning Sequences—Sequentially and Concurrently,” available at <http://www.wofford.edu/ecs/>. Covering the necessary biological background, the named module develops the sequential Needleman-Wunsch Algorithm (NWA) for aligning DNA sequences. Then, after motivating the need for HPC, the module discusses HPC pipeline versions of NWA along with timings. To aid students, the module contains fifteen Quick Review Questions, many with multiple parts; nine exercises; and five projects. Completed sequential and parallel C with MPI programs are available upon request for instructors. The talk will also detail experiences using the material in a bioinformatics course at University “Magna Græcia” of Catanzaro. Besides being appropriate for such a course, the module can provide a meaningful application for a high performance computing or a data structures class.

Angela B. Shiflet

McCalla Professor of Math. & CS, Dir. of Computational Sci.  
Wofford College  
shifletab@wofford.edu

George W. Shiflet, Daniel Couch  
Wofford College  
shifletgw@wofford.edu, couchds@email.wofford.edu

Pietro Guzzi, Mario Cannataro  
University Magna Græcia of Catanzaro, Catanzaro, Italy  
hguzzi@unicz.it, cannataro@unicz.it

**CP17****Numerical Simulation of a Nonlocal Variational Problem**

A ternary inhibitory system is a three-component system characterized by two properties: growth and inhibition, which motivated by the triblock copolymer theory is studied as a nonlocal geometric variational problem. Theoretical analysis provides the existence of a stable stationary point of the free energy functional. Numerically a phase field model is proposed here to track the sharp interface. The method combines a semi-implicit scheme in time and a Poisson solver with Neumann Boundary Condition. To reduce the computational complexity, a reduced FFT-based fast algorithm is also discussed here.

Chong Wang

George Washington University  
chongwang@gwu.edu

Xiaofeng Ren  
The George Washington University  
Department of Mathematics  
ren@gwu.edu

Yanxiang Zhao  
Department of Mathematics  
George Washington University  
yxzhao@email.gwu.edu

**CP17****Models for Principled Characterization of Dynamic, Spatially Embedded, Multiscale Networks**

Advances in neuroimaging techniques have made it possible to reconstruct whole-brain networks composed of structural (physical fiber tracts) and functional (statistical relationships) connections among brain regions. Analysis of these static networks has revealed a host of non-random attributes, including highly connected hubs, modular architecture, and rich clubs. In this talk I will discuss two recent advances in brain network modeling. First, I will introduce the multi-layer network model for characterizing time-varying functional brain networks. I will cover recent work showing that the brains flexibility the extent to which its multi-layer modular organization is stable across time can be used to predict an individuals learning rate, is associated with executive function and state of arousal, and is also altered in psychiatric diseases such as schizophrenia. Second, I will discuss the role that the brains intrinsic geometry plays in shaping its network architecture. I will cover two recent studies one in which we show that simple wiring rules can explain a wide variety of the brains topological features and another in which we modify classical community detection tools to uncover space-independent community structure in brain networks.

Richard betzel

Dept. Bioengineering  
University of Pennsylvania  
rbetzel@seas.upenn.edu

**CP18****Mathematical Modeling and Analysis of a Type-3 DFIG/DFAG Wind Turbine Dynamics: Numerical, Analytical, and Simulation Results**

In this research, a type-3 DFIG/DFAG wind turbine generator is considered for modeling and analysis. The main blocks of the model in transfer function domain is taken from the literature and are described and translated into a system of differential equations with algebraic constraint. Time domain analysis for the system is provided by computing the steady states as functions of the wind speed and the grid parameters, eigenvalues sensitivity to the wind speed, and stability in parameter space. Analytical proofs of boundedness, existence, and uniqueness under control limits are provided to give more assurance about numerical exploration using this model. Simulations for time dependent wind speed and terminal voltage along with system responses are presented as well.

Sameh Eisa, William Stone, Kevin Wedeward

New Mexico Tech  
sameh.eisa@student.nmt.edu, william.stone@nmt.edu,  
kevin.wedeward@nmt.edu

**CP18****Optimal Parametric Model Reduction in H2-L2 Norm**

Reduced order models play a significant role in simulation, design and optimization as they are able to reduce the



computational complexity drastically while retaining accuracy. Interpolatory model reduction is one of the widely used methods where transfer function of the reduced model interpolates that of the original at carefully selected frequency domain interpolation points. Indeed, for linear non-parametric dynamical system, Iterative Rational Krylov Algorithm chooses these points optimally in the  $\mathcal{H}_2$  norm. Even though interpolatory methods have been extended to parametric systems, except for special cases, there exists no optimal selection strategy for frequency and parameter interpolation points jointly for a combined error measure. The usual approach is to employ a greedy search in the parameter domain followed by a search in the frequency domain. In this talk, we will introduce a new framework where the frequency and parameter interpolation points are selected jointly to minimize a global tensor  $\mathcal{H}_2 - L_2$  norm,  $\mathcal{H}_2$  in the frequency,  $L_2$  in the parameter domain. We will provide the first-order conditions in this  $\mathcal{H}_2 - L_2$  norm for an optimal parametric reduced model. Various numerical examples will be presented to illustrate the proposed method.

Alexander R. Grimm  
Department of Mathematics  
Virginia Tech  
alex588@vt.edu

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

Christopher A. Beattie  
Virginia Polytechnic Institute and State University  
beattie@vt.edu

#### CP18

##### **Fluctuating Hydrodynamic Methods for Manifolds: Particle Dynamics Within Curved Fluid Interfaced and Membranes**

We present computational methods for simulating particle drift-diffusion and hydrodynamic coupling within fluid interfaces. Motivated by the kinetics of proteins within lipid bilayer membranes, we develop theory and methods based on the exterior calculus of differential geometry to take into account the hydrodynamics within two curved leaflets of bilayer membranes, intermonolayer slip between the leaflets, and coupling with the surrounding bulk solvent fluid. In 1975, Saffman and Delbruck introduced a hydrodynamic theory for infinite flat sheets that is still a widely used theory for the diffusivity of membrane proteins. For finite curved membrane sheets, we investigate how geometric and topological effects can significantly augment the protein diffusivity and hydrodynamic coupling. We present results showing how these effects contribute to the individual and collective motions of particle inclusions in vesicles having different shapes with spherical topology. We also present general fluctuating hydrodynamic methods for manifolds that can be used for general investigations of many-body systems involving diffusion and hydrodynamic transport within thin curved fluid interfaces.

Ben J. Gross, Paul Atzberger  
UC Santa Barbara  
bgross09@yahoo.com, atzberg@ucsb.edu

#### CP18

##### **Influence of Storm Characteristics on Hurricane**

#### Surge

The potential wind damage impact of impending hurricanes has traditionally been measured by the Saffir-Simpson scale, a simple model to provide weather forecasters and emergency planners with a Category 1-5 rating. However, reliance on this scale as an indicator of storm surge, the primary destructive force during a hurricane, leads to misconceptions by the public and scientists alike of the impending danger. In this talk, it is shown that the size of a hurricane windfield, the intensity, and a newly proposed potential kinetic energy are the much more influential/indicative of resulting storm surge. We present results of a suite of synthetic storms impacting the Galveston Bay, TX area that demonstrates this phenomena. The simulations are performed on the Texas Advanced Computing Centers supercomputers using the highly accurate and robust Advanced Circulation (ADCIRC) finite element numerical model.

Jennifer Proft  
University of Texas at Austin  
jennifer@ices.utexas.edu

Clint Dawson  
Institute for Computational Engineering and Sciences  
University of Texas at Austin  
clint@ices.utexas.edu

#### CP18

##### **Approximating Optimal Local Subspaces for Localized Model Order Reduction By Random Sampling**

Recently optimal local reduced spaces for the Generalized Finite Element Method in [Babuska, Lipton 11] and static condensation or substructuring methods in [Smetana, Patera 16] have been proposed. Those spaces can be approximated by, say, the Finite Element Method and allow for an exponential convergence to the true solution. However, computing an approximation requires first the computation of many local solutions of the PDE and subsequently the solution of an eigenvalue problem, which can become very costly. Extending ideas in [Buhr, Engwer, Ohlberger, Rave 15] and [Eftang, Patera 13] we propose to build the local reduced spaces from local solutions of the PDE with random boundary conditions. Here, the number of required local solutions of the PDE equals approximately the dimension of the desired optimal local reduced spaces. We prove that our proposed reduced spaces have nearly the same convergence behavior as the optimal local reduced spaces in [Babuska, Lipton 11] and [Smetana, Patera 16]. Numerical experiments confirm the excellent convergence rate and show that the suggested spaces can be constructed in a much shorter computational time, requiring significantly less memory.

Kathrin Smetana  
Department of Mechanical Engineering  
Massachusetts Institute of Technology  
kathrin.smetana@www.de

Kathrin Smetana, Andreas Buhr  
Institute for Computational and Applied Mathematics  
University of Muenster  
kathrin.smetana@www.de, andreas@andreasbuhr.de

#### CP18

##### **Numerical Methods for Stochastic Differential**

### Equations with Highly Nonlinear Coefficients

We consider numerical methods for nonlinear stochastic differential equations with highly nonlinear coefficients. The coefficients can grow at most polynomially at infinity but can be locally Lipschitz continuous or Holder continuous. Both explicit and implicit schemes are explored. Theoretical and numerical study of convergence is presented.

Zhongqiang Zhang  
Worcester Polytechnic Institute  
zzhang7@wpi.edu

### CP19

#### Dynamics of a Two-Vector, Two-Pathogen, Two-Host Model

In this talk, the speaker will present recent theoretical results from the dynamics of a two-vector, two-pathogen, two-host model. A system of ordinary differential equations is used to model the dynamics of two vector-borne pathogens (*Rickettsia parkeri* and *Rickettsia amblyommii*) that are increasingly found within tick populations of Virginia spread by two species of ticks (*Amblyomma maculatum* and *Dermacentor variabilis*), within a dual host system. Three methods of transmission are included in the model: vector-borne, transovarial, and co-feeding. Results of numerical simulations are presented and determine a range of parameter values which lead to coexistence of the two pathogens and values which lead to the extinction of one pathogen and persistence of the other.

Caleb L. Adams  
Radford University  
cadams5@radford.edu

### CP19

#### Development of a Dynamical System for the Simulation of Immune Cells and Antigens

A mathematical model has been developed to simulate a dynamical system including non-specific immune cells, HIV, and tuberculosis. The dynamical system for immune cells, HIV, and tuberculosis has been defined by three ordinary differential equations. The model formulation included growth, decay, and interaction of non-specific immune cells, HIV, and tuberculosis. The equilibrium solutions have been developed for the investigation of the system as time goes to infinity. A Jacobian matrix has been evaluated to examine the stability of the system based on the equilibrium solutions of the dependent variables. In addition to the stability, phase plane diagrams have been developed to perform the bifurcation analysis of the dynamical system. For different initial conditions and different values of the model parameters, the bifurcation analyses have been used to establish the stable and unstable systems. The growth and decay of the immune cells and antigens have been observed by examining the phase plane diagrams. Finally, the unsteady solutions have been generated to examine the time-varying effects of immune cells on the antigens.

Rumana Ahmed  
Graduate Student/Research Assistant  
rumed1992@gmail.com

Mahbubur Rahman  
University of North Florida

mrahman@unf.edu

### CP19

#### Robust Regulation of Hepatic Pericentral Amination by Glutamate Dehydrogenase Kinetics

Impaired glutamate dehydrogenase (GDH) sensitivity to its inhibitors causes excessive insulin secretion by pancreatic beta-cells and defective ammonia metabolism in the liver. These are commonly associated with the hyperinsulinism/hyperammonemia syndrome (HI/HA), causing recurrent hypoglycaemia in early infancy. Hepatic localization of GDH amination and deamination activities linked with the urea cycle is known to be involved in ammonia metabolism and detoxification. Although deamination activities of hepatic GDH in the periportal zones of liver lobules and its connection to urea cycle have exhaustively been investigated, physiological roles of GDH amination activity observed at pericentral zones has often been overlooked. Using kinetic modelling approaches, here we report a new role of hepatic GDH amination kinetics for maintaining ammonia homeostasis under excess intrahepatic input of ammonium. We have shown that alpha-ketoglutarate substrate inhibition kinetics of GDH robustly control the ratio between glutamate and ammonium under a wide range of intracellular substrate variation. Dysregulation of this activity under pericentral nitrogen insufficiency contributes to breaking down of ammonia homeostasis and thereby can significantly affect HI/HA syndrome.

Soumen Bera  
P.hD Research Scholar  
soumenmath4189@gmail.com

### CP19

#### Scanning Electron Microscope Simulation with Adaptive Finite Elements

There are important industrial measurement applications for scanning electron microscope (SEM) images of samples of mixed insulators and conductors, for example in nanomaterials or in integrated circuits, where oxides insulate transistors and their current-carrying interconnections from one another. The SEM's electron beam charges insulators. Quantitative inferences of feature positions and dimensions require accounting for the effect of the resulting electric fields. Our model contains two coupled processes: a Monte Carlo method to deal with electron scattering, secondary electron production, and detection, and a finite element analysis to determine the electric field that affects the electron paths and resulting image. The finite element analysis requires a highly nonuniform mesh to resolve small-scale structures in the materials and local variations in the electric field due to the charge distribution. Although one can manually design a reasonable mesh from knowledge of the structures in the materials, the region being scanned, and the decay rate of the electric potential away from the charged region, adaptive mesh refinement can produce a more optimal mesh, and also allows for a dynamically changing mesh as the charge distribution evolves. In this talk we will present an adaptive mesh refinement strategy and a posteriori error estimator tailored for the SEM simulation, and demonstrate its performance on sample problems.

William F. Mitchell  
NIST, Gaithersburg, MD  
william.mitchell@nist.gov

John Villarrubia  
NIST  
john.villarrubia@nist.gov

### CP20

#### An Experimental Investigation into the Practical Performance of Lattice Reduction Algorithms on Ideal Lattices

Lattice-based cryptography is considered to be a strong candidate for post-quantum cryptography. In practice, lattices with a special algebraic structure, called ideal lattices, are regularly implemented because their structure allows for more efficient storage and computation. It is commonly assumed for security purposes that problems in ideal lattices are not easier than in the general case. The most well-known hard lattice problem is the shortest vector problem (SVP), for which lattice reduction algorithms are the standard approach. In 2008, Gama and Nguyen published an extensive experimental study regarding the practical performance of several lattice reduction algorithms on general lattices. We present preliminary results from an analogous experimental study performed to assess the practical performance of lattice reduction algorithms on ideal lattices. By comparing our results to the Gama and Nguyen study, we may determine if lattice reduction algorithms perform better on ideal lattices than general lattices, allowing us to reasonably conjecture whether or not ideal lattices offer the same security. Our results may also be used to predict the performance of these algorithms on ideal lattices, and identify their limitations.

Scott C. Batson  
North Carolina State University  
scott.batson@navy.mil

Jamie R. Lyle, Bryan Williams  
SPAWAR Systems Center Atlantic  
jamie.lyle@navy.mil, bryan.l.williams1@navy.mil

Kayla Capitan  
Palmetto Scholars Academy  
kaylacapitan125@gmail.com

Georgianna Campbell  
University of Georgia  
georgeltcampbell@gmail.com

Tuwaner Lamar  
Morehouse College  
tuwaner.lamar@morehouse.edu

Hemant Pendharkar  
Worcester State University  
hemant.pendharkar@worcester.edu

### CP20

#### Domain Management for Future Scale Parallelism

One of the main challenges in numerical simulation of real world phenomena on modern high performance clusters is the decomposition and distribution of the overall work load – especially for growing problem sizes and larger systems with increasing heterogeneity. In this presentation, we introduce our CFD Framework which has been successfully deployed on up to 140,000 cores on two of Germany's top-tier supercomputing systems. For domain decomposition we rely on non-overlapping block-structured Cartesian

Grids gained through hierarchical space-tree refinement, perfectly suited for a custom tailored parallel multigrid-like solver. While domain management is centralised – using a dedicated management server, aware of the complete topological view of the data structure –, load balancing is realised using space-filling curves. This approach is, however, limited by the available memory of these management servers. As remedy, we propose a novel approach currently under development, where every sub-domain has only a local view of its spatial and hierarchical neighbours. Load balancing is then realised by a diffusion process that will be solved using our multigrid-like approach across the complete hierarchy.

Christoph M. Ertl  
Chair of Computation in Engineering, Prof. Rank  
Technische Universität München  
christoph.ertl@tum.de

Christoph M. Ertl  
Technische Universität München  
Chair for Computation in Engineering  
christoph.ertl@tum.de

Jérôme Frisch  
Institute of Energy Efficiency and Sustainable Building  
E3D  
RWTH Aachen University  
frisch@e3d.rwth-aachen.de

Ralf-Peter Mundani  
TUM, Faculty of Civil, Geo and Environmental  
Engineering  
Chair for Computation in Engineering  
mundani@tum.de

Ernst Rank  
Technische Universität München  
Chair for Computation in Engineering  
rank@bv.tum.de

### CP20

#### Mapping of Language to a Logic Model Checker (LMC)

In this paper we present an automated logic model checker (LMC) named Meth8 to accept input of expressions with literals, modifiers, and connectives. The modifiers are the negation operator of Not and the modal operators of the Necessity box and Possibility lozenge. Logical input is assigned from the components of natural language such as parts of speech (POS), sentences, paragraphs, and documents. The approach is an end-to-end evaluation of requirements through the correctness of components. The goal is to assure there are no mistakes in the requirements, the implementation units, or in the configuration items as deliverables. Verification and validation (VV) of the models is checked by logical proof from Meth8. The approach is bivalent, hence avoiding modal mu calculus and temporal logic.

Colin James III  
ESMC [Ersatz Systems Machine Cognition, LLC]  
info@cec-services.com

### CP20

#### Mathematical Modelling and Analysis for Making Career Decisions in Football Using Data Mining

## and Statistical Techniques

With the advent of more and more talents in the field of football it is imperative for football club managers to negotiate transfer deals effectively to ensure that the clubs finances are not harmed. This paper aims to propose a model for professional football players and club managers with basic guidelines for making better career decisions that would result in the collective benefit of the entire team. The research is based on a case study of twenty forward players from the football clubs participating in various European Football Leagues. For every player, the total time played, numbers of goals, assists, shots on target, shots off target are taken into account and subsequently calibrated points are calculated accordingly using simulation and optimization techniques. Results show that the combination of two strikers of the same level will not benefit the team compared to a good striker paired with an average one. Also, players tend to lose confidence when they get less playing time which affects their performance, and consequently, the performance of the team. This reversible process slowly weakens a player and leads to a change of clubs which can have potential repercussions on the players career. The paper also introduces a concept of spikes which are quantitative indicators that can be used to help players as well as club-managers in making important career decisions.

Parardha Kumar, Aakar Dwivedi  
Indian Institute of Technology (Indian School of Mines)  
Dhanbad  
parardha@ap.ism.ac.in, aakardwivedi@am.ism.ac.in

## CP20

### Performance Portability and Analysis of a Finite Element Multi-Physics Code

In recent years, large, high performance computing clusters have become increasingly more heterogeneous as part of an ongoing push towards exascale computing. Often times, a code or algorithm has to be rewritten or optimized for a particular computing architecture which leads to an enormous amount of time being spent in porting codes to next generation hardware. Kokkos is a library-based programming model which seeks to eliminate this issue by providing an easy to use framework for performance portability across multi-core, many-core or GPU architectures. In this talk, I will discuss our efforts in developing a performance portable implementation of the finite element code called Albany by utilizing the Kokkos library to obtain high performance parallelism with minimal effort. I will provide strong and weak scalability studies on MPI+OpenMP and MPI+GPU frameworks for climate and other applications implemented within the Albany multi-physics code. In particular, I will show how the results highlight the performance capabilities of the Aeras atmospheric model dynamical core and the FELIX land-ice solver.

Jerry Watkins  
Stanford University  
Sandia National Labs  
watkins2@stanford.edu

Irina Tezaur  
Sandia National Laboratories  
ikalash.sandia.gov

## CP20

### An Interoperable Workflow Platform for Multi-

## Disciplinary Simulations

We present an interoperable workflow platform the open Distributive Interoperable Executive Library (openDIEL) - to facilitate the collaboration, exploration, and execution of multi-disciplinary modeling projects suited for a diversified research community on emergent large-scale parallel computing platforms. It does so by providing a managing executive, a workflow configuration input file, and two sets of native communication protocols. OpenDIEL allows users to plug in their individual science codes (modules), prescribe the interactions between those modules, and schedule communications between them. OpenDIEL encapsulates user codes in a single MPI executable and executes them as arranged in the configuration file.

Kwai L. Wong  
Joint Institute for Computational Science  
University of Tennessee/ORNL  
kwong@utk.edu

Tanner Curren  
Maryville College  
tanner.curren@mymaryvillecollege.edu

## CP20

### Optimizing Structured-Grid Halo-Exchange Communications for Accelerator-Resident Computing on Multicore Platforms with Multiple MICs

Accelerator-resident computing is widely used in numerical simulations on accelerator-based platforms, with possible benefits of reduced host-accelerator data transfer. However, since data fields are now resident on accelerators, communications such as halo-exchange become rather expensive. In this talk, we present our work on optimizing structured-grid halo-exchange communications for typical applications on multi-CPU multi-MIC platforms. Using a combined approach of a data-centric communication model and accelerator-oriented message scheduling, we improve the performance of halo-exchange communication of FDTD and molecular dynamics simulation to nearly an order of magnitude.

Zhang Yang, Yang Yang, Aiqing Zhang  
Institute of Applied Physics and Computational Mathematics  
yang\_zhang@iapcm.ac.cn, yang\_yang@iapcm.ac.cn,  
zhang\_aiqing@iapcm.ac.cn

## CP20

### Systematic Modeling of Structured-Grid Numerical Applications for Automatic Parallelization

The ever-increasing complexity of the parallel programming stack imposes significantly challenge on high performance numerical application development. To address this challenge in structured-grid numerical simulations, we explore a model-oriented approach. The main idea of our approach is to systematically model the common patterns of numerical computation on various types of structured grids, and form a series of domain specific parallel computing models. These models can be efficiently implemented on modern high performance computing platforms and support domain experts implementing parallel structured-grid applications easily and quickly.

Aiqing Zhang  
Institute of Applied Physics and Computational

Mathematics  
zhang\_aiqing@iapcm.ac.cn

Zeyao Mo  
CAEP Software Center  
for High Performance Numerical Simulations  
zeyao\_mo@iapcm.ac.cn

Xiaolin Cao, Qinkai Liu, Xiaowen Xu  
Institute of Applied Physics and Computational  
Mathematics  
xiaolincao@iapcm.ac.cn, liuqk@iapcm.ac.cn,  
xwxu@iapcm.ac.cn

### CP21

#### Analysis of DMA Instrumentation for Viscoelastic Materials Using the DPG Methodology

Current inverse models used by DMA instrumentation will be compared to a calibration procedure that uses the force measurements from the instrument and simulations of classical viscoelasticity in the time-harmonic regime. This is done in the context of characterization of epoxy-like materials used in medium-voltage form-wound coils lying in large electromachinery. Numerical simulations are performed with an in-house software that uses the DPG methodology. The numerical method is proved and thoroughly verified to be stable at the expected convergence rates for different polynomial orders. Moreover, the methodology carries a natural and automatic residual-based h-adaptivity scheme used to determine the fine solution features.

#### Federico Fuentes

Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
federico@ices.utexas.edu

Leszek Demkowicz  
Institute for Computational Engineering and Sciences  
(ICES)  
The University of Texas  
leszek@ices.utexas.edu

Aleta Wilder  
The University of Texas at Austin  
a.wilder@mail.utexas.edu

### CP21

#### Improving the Computational Efficiency of Multi-scale Simulations with Filtering

In multiscale simulations that concurrently couple continuum approximations with first-principles based atomistic representations, the atomistic component of the model typically dominates the computational cost of the overall method. In this talk we demonstrate the effectiveness of a spectral filter for improving the accuracy of noisy continuum data obtained from atomistic simulations. The reduction of errors due to noise enables running less expensive atomistic simulations to achieve a desired accuracy, thereby lowering the primary cost of the model and leading to faster simulations. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-ABS-702605

David J. Gardner

Lawrence Livermore National Laboratory  
gardner48@llnl.gov

Daniel R. Reynolds  
Southern Methodist University  
Mathematics  
reynolds@smu.edu

### CP21

#### Risk Analysis Using Monte Carlo Simulation in Bowtie Models

Risk assessment and analysis, originally a standard process in drilling and mining industries, has recently gained popularity in a variety of domains: health industry, transportation, handling hazardous materials, environment etc. Quantitative risk assessment of a critical event (accident) is essentially the relationship (often modeled as the product) between the probability of the event and the severity of its consequences. While the mathematical model may look straightforward, in practice, risk analysis is a very complicated process that includes making decisions based on uncertain events. An exact risk analysis is, in general, not possible due to the large number of the parameters in the model. We use Monte Carlo simulation to perform risk analysis in a bowtie model. For different probability distributions of causes and different severities of consequences in the model we perform all computations and produce different outcomes of risk assessment values. The Monte Carlo simulation in a bowtie model can assist risk analysts in their decision-making process over a range of possibilities. It shows the sensitivity of the outcome to input changes, the extreme situations, and all intermediate control values along the ways from the critical events causes to its consequences.

#### Ionut E. Iacob

Department of Mathematical Sciences  
Georgia Southern University  
ieiacob@georgiasouthern.edu

### CP21

#### Arbitrary-Shaped Walls with Controllable Surface Roughness in Dissipative Particle Dynamics Simulations

We present a boundary method for dissipative particle dynamics (DPD) simulations involving arbitrary-shaped wall boundaries. By introducing a parameter of boundary volume fraction (BVF) associated to each fluid particle, the solid boundary is detected on-the-fly by the fluid particles according to local particle configuration. Therefore, it becomes no necessary to predefine the boundary geometry for DPD simulations, which makes it possible to read-in the geometry of fluid systems directly from experimental images or AutoCAD files. Quantitative validity of this boundary method is performed with Poiseuille flow, Couette flow and Wannier flow being compared to the analytical solutions. It is verified that the proposed boundary method yields controllable surface roughness on walls and gives correct no-slip boundary and negligible fluctuations of density and temperature in the vicinity of solid boundary. Moreover, we use an example to explicitly show the strategy of generating a DPD system from an experimentally scanned image and simulating the system with the proposed boundary method. Since this boundary method uses local information within a neighbor list and satisfies desired boundary conditions for the velocity on-the-fly, it is not only valuable for static arbitrary-shaped boundaries, but also the mov-

ing boundaries and deformable boundaries in the problems involving fluid-solid interactions.

Zhen Li

Division of Applied Mathematics  
Brown University  
Zhen\_Li@brown.edu

Xin Bian, Yu-Hang Tang

Brown University  
xin.bian@tum.de, yuhang\_tang@brown.edu

George E. Karniadakis

Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

## CP21

### A Parallel Approach to Viscoelastic Polymer Gel Modeling

Stimuli-responsive hydrogels have attracted a great deal of attention due to their tunable mechanical properties. These transient polymer networks respond to environmental stimuli such as pH, UV, or temperature, making them ideally suited for a variety of applications. Recent advances include the creation of multiresponsive interpenetrated self-assembled polymer networks (IPSAN) formed from the mixing of two different stimuli responsive hydrogels. However, mathematical models, algorithms and numerical simulations for the investigation of IPSANs remain largely unexplored. In this contribution, we present an elastic dumbbell chain model that leverages the parallel processing power of graphics processing units (GPU) to create a unique micro-macro scale driven design. This platform incorporates non-linear viscoelastic behavior as well as the stochastic processes that describe the breaking and reforming of entanglement points. The added fidelity of this approach allows a recreation of not only the coupling between microstructure and flow within a single hydrogel, but also captures the emergent properties of their mixture. In this talk we discuss the numerical and computational challenges arising from this study and demonstrate the model's ability to efficiently recreate measured data from the mixing of a UV-responsive network formed by copolymers based on poly-ethylene oxide (tPEO) and a pH-responsive network formed by copolymers based on poly-acrylic acid (tPAA).

Erik Palmer

Department of Mathematics  
University of South Carolina  
etpalmer@math.sc.edu

## CP21

### A Partition of Unity Method for Generalized Eigenvalue Problems with Applications in Electronic Structure Calculations

In this presentation we are concerned with the numerical treatment of the generalized eigenvalue problem

$$Kx = \lambda Mx$$

by Partition of Unity Methods (PUM). We focus on the construction of appropriate enrichment functions to attain a high approximation accuracy with very small numbers of degrees of freedom. Moreover, we are interested in the construction of an orthogonal basis of the employed PUM

space to attain  $M = \mathbb{I}$  to reduce the generalized eigenvalue problem to a classical eigenvalue problem. The fundamental ideas employed in these constructions are based on the techniques presented in [?, ?] which essentially allow for the construction of a stable (block-)orthogonal basis for arbitrary enrichments with a computational effort that scales linearly with the number of employed partition of unity functions. The overall goal of this work is to provide a parallel PUM simulation environment with linear scaling for generalized eigenvalue problems esp. in the context of electronic structure calculations. We discuss the current state of our software framework PUMA and present first numerical results obtained with PUMA for some reference problems in this context.

Marc A. Schweitzer

Institut für Numerische Simulation, Universität Bonn and Fraunhofer SCAI  
schweitzer@ins.uni-bonn.de

Constanze Klaar, Albert Ziegenhagel

Fraunhofer SCAI  
constanze.klaar@scai.fraunhofer.de, albert.ziegenhagel@scai.fraunhofer.de

## CP21

### An Adaptive Coupled Level Set and Moment-of-Fluid Method for Simulating the Solidification Process in Multimaterial Systems

An adaptive hybrid level set moment-of-fluid method is developed to study the material solidification of static and dynamic multiphase systems. The main focus is on the solidification of water droplet, which may undergo normal or supercooled freezing. We model the different regimes of freezing such as supercooling, nucleation, recalescence, isothermal freezing and solid cooling accordingly to capture physical dynamics during impact and solidification of water droplets to solid surfaces. The numerical simulations are validated by comparison to analytical results and experimental observations. The present simulations demonstrate the ability of the method to capture sharp solidification front, and handle contact line dynamics, and the simultaneous impact, merging and freezing of a drop. Parameter studies have been conducted, which show the influence of the Stefan number on the regularity of the shape of frozen droplets. Also, it is shown that impacting droplets with different sizes create ice shapes which are uniform near the impact point and become dissimilar away from it. In addition, surface wettability determines whether droplets freeze upon impact or bounce away.

Mehdi Vahab

University of California Davis  
Department of Applied Science  
mvahab@math.fsu.edu

Mark Sussman, M. Yousuff Hussaini

Department of Mathematics  
Florida State University  
sussman@math.fsu.edu, yhussain@fsu.edu

## CP22

### Simd-Acceleration of a Sparse Analytical Chemical Kinetic Jacobian Code Via Python

Reacting flow simulations are intensely computationally demanding due to the interaction of physical processes over large ranges of spatial and temporal scales. Typ-

ical codes employ operator splitting schemes, resulting in tens of thousands to millions of separate chemical kinetic systems of non-linear ordinary differential equations (ODEs); usually evaluated in parallel via traditional multi-threaded solutions, e.g. OpenMP/MPI. The implicit integration techniques commonly used to solve these ODEs rely on repeated evaluation and factorization of a Jacobian matrix. Speedups can be achieved by analytical evaluation of the Jacobians compared to finite difference methods and sparse linear algebra techniques. Single Instruction, Multiple Data (SIMD) processing is a parallel computing paradigm found on central processing units as well as graphics processing units and other computing co-processors. SIMD provides accelerations well suited for the embarrassingly parallel nature of chemical kinetic ODEs, and can be combined with traditional parallelization methods. Historically SIMD-acceleration has been highly platform dependent and difficult to optimize for various architectures, however recently developed tools simplify this task. This talk will demonstrate the use of the python packages SymPy and Loo.py to develop SIMD-accelerated OpenCL code for evaluating chemical kinetic Jacobians for multiple platforms and languages.

Nicholas Curtis

University of Connecticut  
Department of Mechanical Engineering  
nicholas.curtis@uconn.edu

## CP22

### Hybrid Computing In Large-Scale Multiphysics Simulation: Tabulated Properties and Particle-Cell Interpolations

This work is directed toward addressing two computationally expensive segments in multiphysics reacting flow simulations. The first segment is the gas phase simulation including solving the gas-dynamics equations coupled with a model to describe the reaction. In this work, these are described by pre-computed tabulated data, and it is shown that executing the tabulated functions across different dimensionality and interpolant order on the GPU instead of the CPU platform decreases the computational cost significantly. The second bottleneck is the interactions of the gas and the particle phase that are governed by Eulerian and Lagrangian formulations, respectively. These interactions, consisting of particle-cell interpolations, can take up to 20% of the simulation time. Regarding the number of particles and gas mesh size, performing particle-cell interpolation calculations on the GPU instead of the CPU can achieve up to 50x speedup.

Babak Goshayeshi

University of Utah  
b.goshayeshi@gmail.com

Tony Saad

Institute for Clean and Secure Energy  
Department of Chemical Engineering, University of Utah  
tony.saad@utah.edu

James C. Sutherland

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

## CP22

### A Graphical Programming Approach for Parallel

## Particle Based Simulation Development

Particle based simulations have become an important tool in area of scientific computing. A major challenge lays in the rapid design and implementation of parallel particle based applications. In this paper, we introduce a graphical programming approach to address this challenge. A new semi predefined data structure is proposed to facilitate the definition of particle variable. Some user-configurable components are designed to organize parallel data communication and computation. The program architecture is auto-generated together with upper definition and configuration. These codes are combined with sequential programs written in Fortran to construct the full application. The approach is implemented to ease parallel programming for domain experts. Real applications demonstrate that the approach for developing particle based numerical applications is both practical and efficient.

Li Liao, First Name Jing

Institute of Applied Physics and Computational Mathematics  
liliao@iapcm.ac.cn, jing.cuiping@iapcm.ac.cn

## CP22

### Adding Plasticity to the Oof Finite-Element Solver for Materials Science

The NIST-developed Object-Oriented Finite Element code (OOF) is a long-standing project 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. Recently, the development team has added crystal plasticity to a development version of this tool, allowing it to model systems which acquire a permanent deformation in response to loads. Being history-dependent, this expanded the OOF project beyond its initial scope of what were essentially divergence equations. The usual approach to this problem taken by the computational mechanics community posed several challenges to the OOF scheme, which aspires to a high level of generality, ease of use, and good computational performance. The team's approaches to these challenges will be described.

Andrew Reid

Center for Theoretical and Computational Materials Science  
National Institute of Standards and Technology  
andrew.reid@nist.gov

Shahriyar Keshavarz

Theiss Research  
shahriyar.keshavarzhadad@nist.gov

Stephen Langer

Information Technology Laboratory  
National Institute of Standards and Technology  
stephen.langer@nist.gov

## CP22

### Saving Time and Energy with Oversubscription in Ab Initio Calculations for Large Chemical Systems

Although significant progress has been made in performing ab initio calculations for very large chemical systems in the last 20 years, the time to complete the calculations

still remains prohibitively long, often on the order of days or months even on today's most powerful computing platforms. In this talk, we investigate the use of oversubscription of processing units to deal with this problem, as well as to reduce the energy consumption of the calculations. We chose to test oversubscription on two MP2 algorithms within the widely used NWChem computational chemistry package. In particular, our results show that oversubscription reduces total time to solution for semi-direct MP2 energy calculations by 25–45% and the total CPU and DRAM energy consumed by 10–15% on the Intel platform. Linear regression analysis indicates a strong correlation between time to solution and total energy consumed.

Masha Sosonkina  
Old Dominion University  
Ames Laboratory/DOE  
masha@scl.ameslab.gov

Ellie Fought  
Iowa State University  
foughtel@iastate.edu

Vaibhav Sundiyal  
Old Dominion University  
vsundiya@odu.edu

Theresa Windus  
Iowa State University  
twindus@iastate.edu

## CP22

### A Quadratic High-Order Method for Mesh Generation Inspired by LBWARP

The ability to construct high-order meshes that conform to the boundary of curved geometries is a limiting factor in the adoption of high-order computational methods for solving partial differential equations. In order to leverage these high-order methods, we must pair them with a high-order mesh. In this work, we propose a method for generating curved meshes of second degree. The approach consists of the following three steps. First, the initial mesh is modified by inserting nodes at the midpoint of each line segment. Second, the midpoints that fall along the boundary are displaced to the prescribed boundary. Third, the final positions of the internal nodes are solved for based on the boundary deformation. We will present some numerical examples that demonstrate the viability of our method, as well as discuss the quality of the elements generated by our method.

Mike Stees, Suzanne M. Shontz  
University of Kansas  
mstees@ku.edu, shontz@ku.edu

## CP22

### A Fast Treecode Algorithm for Stokes Flow in 3D

A large number of problems in fluid dynamics are modeled as many-particle interactions in Stokes flows, for example, simulations of falling jets of particles in viscous fluids, microfluidic crystals, and vesicle flows. The formulation is often based on fundamental solutions. The Stokeslet and the Stresslet are the kernels in the single and double layer potentials, respectively. Many situations (e.g., through superposition or discretization of boundary integrals) involve sums of Stokeslets and Stresslets, which is an example of an  $N$ -body problem and the direct sum requires  $O(N^2)$

operations. This can make the numerical calculation prohibitively expensive. A variable-order Barnes-Hut Cartesian treecode algorithm is developed for speeding up the computation. The particles are restructured recursively into a tree, and the particle-particle interactions are replaced with particle-cluster interactions computed by either a far-field expansion or a direct summation. Numerical results exhibit the promising performance of the algorithm.

Lei Wang  
University of Wisconsin-Milwaukee  
Department of Mathematical Science  
wang256@uwm.edu

Robert Krasny  
university of michigan, ann arbor  
krasny@uwm.edu

Svetlana Tlupova  
Farmingdale State College  
stlupova@gmail.com

## CP22

### Porting a Real-World Molecular Dynamics Application to Xeon Phi Coprocessors

We ported MOASP, a real-world short-ranged MD code mainly targeting metallic materials, to Xeon Phi. The main porting effort involves parallelization of 3 key parts of the application: force calculation, neighbor list builder and particle sort. We also vectorized the parallelized code and did some scalar code optimizations to further improve the performance. When compared to the performance of the original highly parallel CPU version on a 12-core CPU, the force calculation achieved a 2.08X speedup, the neighbor list builder part achieved a 1.29X speedup, and the particle sort part constitutes less than 1% of the total time. For total computation time, the Xeon phi version achieved a 1.8X speedup, about the same performance of 21 CPU cores.

Yang Yang, Aiqing Zhang  
Institute of Applied Physics and Computational Mathematics  
yang\_yang@iapcm.ac.cn, zhang\_aiqing@iapcm.ac.cn

Letian Kang  
Hunan University  
kangleitian2014@163.com

Han Wang  
CAEP Software Center for Numerical Simulation  
wang\_han@iapcm.ac.cn

Wenjie Hong  
Hunan University  
hwj2015@hnu.edu.cn

## CP23

### Education for Simulation Sciences and HPC at JSC

Fostering a sound education for students and young researchers at bachelor, master and PhD level in simulation sciences and high-performance computing (HPC) is an essential task of the Jülich Supercomputing Centre (JSC). Education in applied mathematics and computer science plays a crucial part in these activities. This talk will give an overview of the joint degree programmes with nearby uni-



versities and informs on guest student programmes and the Joint Laboratory for Extreme Scale Computing (JLESC) tailored for master and PhD students with interest in HPC.

Johannes Grotendorst  
Forschungszentrum Juelich  
j.grotendorst@fz-juelich.de

### CP23

#### New Algorithms for Solving Multimaterial Diffusion Problems on Meshes Non-Aligned with Material Interfaces

The use of indirect Arbitrary-Lagrangian-Eulerian (ALE) methods for modeling of 2D multimaterial flows, where the mesh is not moving with the fluid, leads to appearance of cells containing a number of different materials, which might have very different physical properties. If diffusion or heat-conduction problems are also to be solved on such meshes, where material interfaces do not align with the mesh, interface reconstruction methods are usually used first. This results in so-called multimaterial cells, which are partitioned into single-material polygons, and the reconstructed material interfaces are generally discontinuous on interfaces between cells. We present a new family of methods called static condensation that uses special local coarsening algorithms to solve multimaterial diffusion problems on described meshes. Results of numerical simulations for some test problems are shown to compare the accuracy of the new methods to standard homogenization type algorithms that are typically used in existing codes.

Evgeny Kikinon, Mikhail Shashkov  
Los Alamos National Laboratory  
kikinon@lanl.gov, shashkov@lanl.gov

Yuri Kuznetsov  
University of Houston  
Department of Mathematics  
kuz@math.uh.edu

### CP23

#### Adaptive Sampling of Carrier Population Data for Multiscale Semiconductor Modeling

Recent advances have shown that adaptive sampling (AS) techniques, combined with surrogate modeling, can reduce computational cost of multiscale simulations by a factor of 1000 or more, rendering calculations computable that would otherwise require years on modern supercomputers. Among the outstanding challenges for AS is that many material models are based on multiple outputs as well as inputs. This is especially challenging in simulations of time-dependent phenomena, in which many state variables are passed forward in time. We explore this problem in the context of modeling optically-active semiconductors. Our base simulation is a time-dependent Monte Carlo (MC) simulation that computes population vectors of electrons holes and polarization, driven by optical illumination, that propagate to successive time steps. AS is needed to extend these calculations to nano-structured materials, with the ultimate goal of enabling multiscale simulation of photonic devices that employ and/or emulate high-fidelity models at the fine scale. The surrogate models we are developing are influenced by time-dependent surrogate models discussed in [Conti and OHagan, J. Stat. Plan. Inference 140:640 (2010)]. Because the population vectors have hundreds of states, we explore different approaches to surrogate mod-

eling, including accounting for the noise inherent in MC simulation outputs. We conclude with estimates of the computational time needed for a simple test problem.

Brent Kraczek  
US Army Research Laboratory  
brent.e.kraczek.civ@mail.mil

### CP23

#### Mathematical Modelling of Love-Type Wave in a Double Layered Piezoelectric Structure

A mathematical model has been developed to study the propagation characteristics of Love-type wave in a piezoelectric layer bonded perfectly to an isotropic heterogeneous layer overlying an unbounded isotropic homogeneous substrate. The closed form expressions of the dispersion relation for both the cases of electrically open and short conditions have been established by taking into account Debye asymptotic expansion. The emphatic impacts of various affecting parameters viz. wave number, piezoelectric constant of the uppermost piezoelectric layer, heterogeneity parameter of the intermediate layer and width ratio of the layers on the phase velocity of Love-type wave has been remarkably traced out. In order to analyze these effects, numerical computation and graphical illustrations have been accomplished. It is worth mentioning that piezoelectric constant of the uppermost piezoelectric layer, heterogeneity parameter of the intermediate layer and width ratio of the layers reduce the phase velocity of Love-type wave substantially for both the cases of electrical conditions. The computational results manifest that the aforementioned affecting parameters enable Love-type wave to propagate along the surface of the layers, leading to the confinement of wave for a longer duration. In view of these facts, the consequences of the study can be widely utilized in the design of surface acoustic wave devices and Love wave sensors to enhance their efficiency and achieve high performance.

Santan Kumar, Abhishek Singh  
Indian Institute of Technology  
(Indian School of Mines), Dhanbad  
ksantan.ism@gmail.com, abhi.5700@gmail.com

### CP23

#### Modeling Chemical Reactions in Classical Molecular Dynamics Using Data Analytics

We present a method of using data analytics and modeling to accelerate the molecular dynamics simulation of bond exchange reactions. Bond exchange reactions in covalent adaptable network (CAN) polymers are similar to a single displacement chemical reaction,  $AB + C = AC + B$ . In order to accurately model these reactions in large scale classical molecular dynamics (MD), ABC are temporarily bonded together. After a structure equilibration, the molecule, B or C, that will cause higher local energy has its bond to A broken. The structure is equilibrated again, and the process repeats. In this study, local molecular topology information was computed as a feature vector for each time ABC were bonded together. After structure equilibration, a binary target vector was gathered for the sample based on whether B or C bonded. We will present different regression techniques used to develop a model to determine final bonding condition based on an initial local topology sample. We are able to eliminate the temporary ABC bonding and equilibration step in the traditional simulation, replace it with our regression model, and reduce

the simulation time significantly.

Alexander Lohse

Georgia Institute of Technology  
alohse@gatech.edu

Edmond Chow

School of Computational Science and Engineering  
Georgia Institute of Technology  
echow@cc.gatech.edu

Karl Jacob

Georgia Institute of Technology  
karl.jacob@mse.gatech.edu

### CP23

#### Investigation of Stiffness Detection Metrics for Chemical Kinetics ODEs

Many simulations of combustion and reactive flows exhibit numerical stiffness in the equations governing chemical kinetics. Explicit solvers for these equations offer low computational expense, but typically cannot handle stiff systems of differential equations. In contrast, implicit methods demand greater expense but offer unconditional stability—as a result, most combustion simulations rely on these methods by default. However, if chemical kinetics systems exhibit low-to-moderate stiffness, explicit or stabilized explicit methods can instead be used to reduce the computational expense (while remaining stable and accurate). This study therefore aims to determine a stiffness quantification metric capable of efficiently and robustly determining the appropriate category of integrator required. Literature stiffness quantification methods will be surveyed and investigated as metrics to measure the stiffness of chemical kinetics states, including methods based on eigendecomposition or the spectral radius of the Jacobian matrix, error estimations, conditioning parameters, and computational cost estimations. These methods will be applied during the solution of hydrogen and methane autoignition with different initial conditions, and evaluated in terms of effectiveness and computational efficiency.

Kyle E. Niemeyer

Oregon State University  
Kyle.Niemeyer@oregonstate.edu

### CP23

#### Computation of Derived Variables for the Eddy Current Maxwell's Equations

In second order form, the eddy current Maxwell's equations involves the calculation of the electric field,  $\mathbf{E}$ , on edges as the primary variable. However, analysts are often interested in current density,  $\mathbf{J}$ , defined on cell centers or on element faces. The former is of interest for use in visualization packages. while the latter is for measuring current flow. We will detail a mathematically justifiable means of computing these variables and demonstrate their accuracy on a number of test problems.

Christopher Siefert

Sandia National Laboratories  
csiefer@sandia.gov

Logan Meredith

University of Rochester  
logan.meredith@rochester.edu

Duncan A. McGregor

Sandia National Laboratories  
Computational Multiphysics  
damcgre@sandia.gov

### CP23

#### Piecewise Linear Polynomials in Position Dependent Random Maps

Let  $\tau = \{\tau_1, \tau_2, \dots, \tau_r; p_1, p_2, \dots, p_r\}$  be a random map, where  $\tau_1, \tau_2, \dots, \tau_r$  are nonsingular Markov maps defined from  $[0, 1]$  to itself and  $p_1, p_2, \dots, p_r$  are probabilities such that at each iteration a map  $\tau_i$  is selected by the probability  $p_i$ ,  $1 \leq i \leq r$ . When such a probability depends upon position of iterations of a point, the map is called a position dependent random map. We use piecewise linear polynomials in maximum entropy principle to approximate invariant densities of the Frobenius-Perron operator related to position dependent random maps. The  $L^1$  errors between the exact and approximated invariant densities are also depicted.

Tulsi Upadhyay

The University of Southern Mississippi  
tulsi.upadhyay@usm.edu

### MS2

#### A Data Scalable Hessian/KKT Preconditioner for Large Scale Inverse Problems

Hessian or KKT solves are often the primary computational bottleneck for efficient solution of large scale inverse problems, deterministic or Bayesian. Current state-of-the-art Hessian and KKT preconditioners tend to scale poorly with increasingly informative observational data. In this talk we present a new KKT preconditioner that performs well regardless of the information content in the data, as long as the regularization is chosen appropriately. That is, as long as the regularization does not over-penalize highly informed parameter modes and does not under-penalize uninformed modes.

Nick Alger

The University of Texas at Austin  
Center for Computational Geosciences and Optimization  
nalger225@gmail.com

Umberto Villa

University of Texas at Austin  
uvilla@ices.utexas.edu

Tan Bui, Omar Ghattas

The University of Texas at Austin  
tanbui@ices.utexas.edu, omar@ices.utexas.edu

### MS2

#### Title Not Available

Abstract not available

George Biros

The Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
biros@ices.utexas.edu

### MS2

#### Large Scale Fusion of Energy Resolved Compton

## Scatter and Attenuation-Based X-Ray Data for Materials Characterization

Motivated by the need to determine the distribution of materials in luggage checked at airports, here we consider the problem of mapping mass density and the photoelectric absorption coefficient from severely limited view but energy resolved Compton scatter and absorption data. While the limited view absorption imaging problem is well studied, less attention has been given to the use of scattered X-rays for problems such as this. For the system of interest here where relatively few views are present, the Compton scattered photons significantly augments the diversity of ray paths associated with the data set. After providing an overview of the physics associated with Compton scatter, we shall describe approaches to recovering the two quantities of interest, both of which are needed for materials characterization. Technically, we focus on (a) novel regularization methods in which the density (which is, in a sense easier to estimate) is used to guide the recovery of the photoelectric absorption and (b) the computational challenge associated with solving the variational problem defining these estimates. Indeed, despite the limited view nature of the problem, the data set is still large owing to both the density of detectors and their energy resolving characteristics. Examples will be provided both based on simulation as well as real data collected from a testbed under development.

Eric L. Miller  
Tufts University  
elmiller@ece.tufts.edu

Hamideh Rezaee  
PhD Student, Department of Electrical and Computer Engineering, Tufts University  
hamideh.rezaee@tufts.edu

Brian Tracey  
Tufts University  
btracey@eecs.tufts.edu

Abdulla Desmal  
Tufts University  
Dept. of ECE  
desmal@eecs.tufts.edu

## MS2

### Sketched Ridge Regression: Optimization and Statistical Perspectives

Previous work has applied matrix sketching to speed up the least squares regression (LSR) on the  $n \times d$  data. Theoretical analysis of the sketched LSR is has been well established and refined. How the results extend to the sketched ridge regression is yet unclear. In our recent work, we study two types of the sketched ridge regression—the classical sketch and the Hessian sketch—from two perspectives the optimization perspective and statistical perspective and draw many useful conclusion. The optimization analysis shows that the sketched solutions can be nearly as good as the optimal; in contrast, the statistical analysis clearly indicates that the two sketched solutions significantly increases bias or variance. Our conclusion is that the practical usefulness of the sketched ridge regression may be very limited. We also propose a simple method which we call the model averaging to improve the quality of the sketched solution, both theoretically and empirically. We argue that model

average has several very useful applications in practice.

Shusen Wang  
University of California, Berkeley  
shusen@berkeley.edu

Alex Gittens, Michael Mahoney  
UC Berkeley  
gittens@icsi.berkeley.edu, mmahoney@stat.berkeley.edu

## MS3

### Practical Heteroskedastic Gaussian Process Modeling

We present a unified view of likelihood based-inference Gaussian process regression for simulation experiments exhibiting input-dependent noise. Replication plays an important role in that context, however previous methods leveraging replicates have either ignored the computational savings that comes from such design, or have short-cut full likelihood-based inference to remain tractable. Starting with homoskedastic processes, we show how multiple applications of a well-known Woodbury identity facilitates inference for all parameters under the likelihood (without approximation) bypassing the typical full-data sized calculations. We then borrow a latent-variable idea from the machine learning literature to address heteroskedasticity, but adapt that method to work within the same thrifty inferential framework, thereby simultaneously leveraging the computational and statistical efficiency of designs with replication. The result is a method that can be characterized as single objective function, complete with closed form derivatives, for rapid library-based optimization. We provide open source software and illustrations on two synthetic data sets, and two real-world simulations experiments from manufacturing and the management of epidemics.

Mickaël Binois  
Mines Saint-Etienne and Renault  
mbinois@uchicago.edu

Robert Gramacy  
Virginia Tech  
rbg@vt.edu

Mike Ludkovski  
University of California at Santa Barbara  
ludkovski@pstat.ucsb.edu

## MS3

### A Novel Voronoi Piecewise Approach to Solve Stochastic PDEs

In this talk, we present a novel approach to Solve Stochastic partial differential equations. Our approach is based on an implicit Voronoi piecewise domain decomposition that avoids the curse of dimensionality. We discretize the UQ space using implicit Voronoi cells around a set of random uniformly-distributed points (seeds). Next, we solve the given stochastic PDE independently in the neighborhood of each seed. A low order expansion is constructed in a box around each seed and the local solutions are then concatenated to form a credible global piecewise solution.

Mohamed S. Ebeida  
Sandia National Laboratories  
msebeid@sandia.gov

Ahmad A. Rushdi

University of California, Davis  
aarushdi@ucdavis.edu

Eric Phipps  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

Marta D'Elia  
Sandia National Laboratories  
mdelia@sandia.gov

### MS3

#### Local Approximate Gaussian Processes for Large Computer Experiments

We provide a new approach to approximate emulation of large computer experiments. By focusing expressly on desirable properties of the predictive equations, we derive a family of local sequential design schemes that dynamically define the support of a Gaussian process predictor based on a local subset of the data. We further derive expressions for fast sequential updating of all needed quantities as the local designs are built-up iteratively. Then we show how independent application of our local design strategy across the elements of a vast predictive grid facilitates a trivially parallel implementation. The end result is a global predictor able to take advantage of modern multicore architectures, GPUs, and cluster computing, while at the same time allowing for a non stationary modeling feature as a bonus. We demonstrate our method on examples utilizing designs sized in the tens of thousands to over a million data points. Comparisons are made to the method of compactly supported covariances, and we present applications to computer model calibration of a radiative shock and the calculation of satellite drag.

Robert Gramacy  
Virginia Tech  
rbg@vt.edu

### MS3

#### An Adaptive Method for Solving Stochastic Equations by Reduced Order Models Over Voronoi Cells

An adaptive collocation-based surrogate model is developed for solving stochastic equations with random parameters. The method is based on a surrogate model of the solution constructed through a Voronoi tessellation of the samples of the random parameters with centers chosen to be statistically representative of these samples. We investigate the use of various interpolants over Voronoi cells in order to formulate the surrogate. For example, an interpolant based on first order Taylor expansions in the probability space is considered which matches solutions and their gradients at the centers of the Voronoi cells. Unlike Monte Carlo solutions, relatively small numbers of deterministic calculations are needed to implement the surrogate models. These models can be used to generate large sets of solution samples with a minimum computational effort. The method is applied to solve stochastic partial differential equations in materials science and random vibration problems of interest in wind engineering. Efficiency and accuracy of this model are demonstrated through numerical examples.

Wayne Isaac T. Uy  
Center for Applied Mathematics  
Cornell University

wtu4@cornell.edu

Haoran Zhao  
School of Civil and Environmental Engineering  
Cornell University  
hz289@cornell.edu

Mircea Grigoriu  
Cornell University  
mdg12@cornell.edu

### MS4

#### Spacey, Super Spacey, and Regenerative Spacey Random Walks

Recent work on eigenvalues of hypermatrices and tensors has generated an algebraic analogue of the stationary distribution vector for a Markov chain. We show that this tensor eigenvector corresponds to the stationary distribution of a new stochastic process called a spacey random walk. Our insight provides a solid probabilistic foundation for these tensor eigenvectors, their interpretation, and their application to data problems with higher-order structure. We then discuss two variants of the process: the super spacey random walk for handling data sparsity and the regenerative spacey random walk for additional uniqueness guarantees.

Austin Benson  
Stanford University  
arbenso@stanford.edu

### MS4

#### Symmetric Orthogonal Approximation to Symmetric Tensors and Its Applications to Image Reconstruction

In this talk, we consider an approximation of symmetric tensors by symmetric orthogonal decomposition. The advantage of using orthogonal decomposition is to represent tensors without containing redundant information in the approximation. Here we would give an iterative algorithm to determine a symmetric orthogonal approximation. The convergence of the proposed algorithm also be discussed. We would show its applications in image reconstruction.

Junjun Pan  
Hong Kong Baptist University  
junjpan@hkbu.edu.hk

### MS4

#### On the Uniqueness of the $Z_1$ -Eigenvector of Transition Probability Tensors

Transition probability tensors of order 3 in dimension 3 and of order 4 in dimension 2 are studied. In both cases, we prove that an irreducible symmetric transition probability tensor has a unique positive  $Z_1$ -eigenvector.

Kelly Pearson  
Murray State University  
kpearson@murraystate.edu

### MS4

#### Transition Tensor of Spacey Random Walk for The

### Stationary Distribution of Polya Urn Process

We study the connections between the recently proposed spacey random walk and the generalized Polya Urns. The standard Polya Urn can be modeled as a special case of the second-order spacey random walk with certain transition tensor. On the other hand, the second-order spacey random walk can be used to generalize the Polya Urn process. We present two types of generalizations based on the transition tensors: one type with all vectors being stationary distributions; and the other type is the Polya Urns have a unique stationary distribution  $e/n$ . We give necessary and sufficient conditions on the transition tensor for the above two types of Polya Urns. Lastly we show that higher-order spacey random walk can be used to further generalize the Polya Urn process with multiple draws each run.

Tao Wu

Purdue University  
wu577@purdue.edu

### MS5

#### ADVENTURE: Scalable Solvers for Industrial Applications

Exa-scale supercomputers will appear around 2020–2022. To obtain high intra-node performance, efficient utilization of processor cache memory should be considered. The traditional memory access-intensive approach, which prefers less computing and more storage on main memory, might not be effective for supercomputers in near future. The Domain Decomposition Method (DDM) is one of the effective parallel finite element schemes. We have been developing an FE-based parallel structural analysis code, ADVENTURE.Solid, based on DDM, with the Balancing Domain Decomposition (BDD) pre-conditioners. The re-design of the subdomain local FE solver part, which is a performance sensitive kernel in the DDM code, is required. Here in this work, an “on-cache” iterative solver based on the DDM framework is developed. The subdomain local FE solver of the DDM code is implemented using CG solvers with various types of pre-conditioners, such as diagonal scaling, SSOR and ICT. These iterative solvers are parallelized using OpenMP, so that each subdomain can be solved by multiple cores. By adjusting the subdomain size so that the footprint fits within the last-level cache of a processor, this DDM code can be considered as a kind of an “on-cache” iterative solver. Performance benchmark results are shown on various kinds of HPC platform, such as Haswell, Knights Corner, RIKEN K Computer and Fujitsu PRIMEHPC FX100.

Hiroshi Kawai

Tokyo Univ. of Science-Suwa  
kawai@rs.tus.ac.jp

Masao Ogino  
Nagoya University  
masao.ogino@cc.nagoya-u.ac.jp

Ryuji Shioya  
Toyo University  
shioya@toyo.jp

Tomonori Yamada, Shinobu Yoshimura  
The University of Tokyo

yamada@race.u-tokyo.ac.jp, yoshi@sys.t.u-tokyo.ac.jp

### MS5

#### Parallel Tensor Arithmetic Applied to Multigrid Methods

Solutions of high-dimensional problems can be approximated as tensors in the data-sparse Hierarchical Tucker Format, if the dependency on the parameters fulfills some low rank property. We use parallelized algorithms which perform arithmetic like the addition or the dot product of two tensors, as well as the application of an operator to a tensor, directly in the Hierarchical Tucker Format. The parallel algorithms result in a time to solution which grows like  $\log(d)$ , where  $d$  is the tensor dimension (i.e. number of parameters + remaining variables of each fixed-parameter-solution). The tensor is supposed to be distributed over  $2d - 1$  compute nodes. We present results on using parallel tensor arithmetic for iterative methods like multigrid in order to solve parameter dependent problems directly in the data-sparse format.

Christian Loebbert

RWTH Aachen  
loebbert@igpm.rwth-aachen.de

### MS5

#### Accelerating Groundwater Flow Simulation

Abstract not available

Zeyao Mo

CAEP Software Center  
for High Performance Numerical Simulations  
zeyao\_mo@iapcm.ac.cn

### MS5

#### Identifying Structures in the Human Skin Using High Performance Shape Optimization Algorithms

In many processes, which are modeled by PDEs, there is a small number of materials with distinct properties involved. The aim is to determine the contours of their spatial distribution based on measurements. In this talk we present an algorithm which combines techniques from high performance computing and shape optimization. We utilize this method in order to determine the shape of permeability parameter distributions in diffusive processes in the human skin with respect to data measurements. The main focus is on the scalability and the applicability on supercomputers. We therefore discuss a combination of multigrid and quasi Newton techniques for shape identification problems. It is also demonstrated how the quality of the underlying finite element meshes can be maintained through the entire optimization process using suitable inner products in the space of feasible shapes. In a short outlook we discuss how to utilize our algorithm for more general problem settings with a wide range of applications. As an example we investigate the optimization of the stiffness of cellular materials by deforming the shape of inclusions.

Martin Siebenborn

Department of Mathematics  
University of Trier  
siebenborn@uni-trier.de

Volker H. Schulz  
University of Trier

Department of Mathematics  
Volker.Schulz@uni-trier.de

### MS6

#### Fast Reaction-Diffusion of Defects in Cadmium Telluride

In this talk, we discuss the asymptotic behavior of a reaction-diffusion system with extremely fast reactions. We use a perturbation method to derive a limiting equation for a simple reaction network and present numerical results for a more complicated model with a scheme derived from a bimolecular kinetic reaction model. Finally, we briefly discuss asymptotic preserving schemes for Cadmium Telluride Solar Cells. Cadmium Telluride is the leading silicon alternative for consumer solar cells. The behavior of the device can be modeled as a reaction-diffusion equation for charged carriers with a self-consistent coupled electric field. In the course of device production additional elements (such as Chlorine and Copper) are added to the device. These molecules enter the device as diffusing interstitial defects, but can react nearly instantaneously with the CdTe lattice. To better understand the fabrication of such devices, we must first understand this physical process.

Daniel Brinkman  
Department of Mathematics  
San Jose State University  
daniel.brinkman@sjsu.edu

Christian Ringhofer  
Arizona State University  
ringhofer@asu.edu

### MS6

#### Electron Transport in Nanostructures: Physical Models and Numerical Methods

Research on electron devices scaling at the 5 nm gate-length has triggered a growing interest on two-dimensional (2D) materials, even beyond carbon-based electronics. Among them, transition metal dichalcogenides (TMDs), silicene, germanene, phosphorene, stanene, and topological insulators. Here we will focus on the use of pseudopotentials to treat the electronic structure of systems at the nanometer scale (2D crystals, nanowires, nanoribbons), their stability and properties, and to study electronic transport, both semiclassically (mobility, Monte Carlo) and from a quantum perspective (open systems). The former task requires fast matrix-free eigenvalue solvers and I will briefly show how our use of the Residual Minimization Method-Direct Iteration in Iterative Subspace (RMM-DIIS) and fast-Fourier transforms has been used to calculate the tunneling current between large graphene nanoribbons. Pseudopotential-based quantum transport, on the other hand, requires not only the implementation of sophisticated physical transport models, but also the solution of very large linear systems. We will present our formulation of the problem and results of the simulation of field-effect transistors with graphene nanoribbons and Si nanowires as active channels. Finally, we will discuss the mathematical formulation of a related problem, namely, the calculation of the dielectric response (or, simply put, of the position-dependent dielectric constant) of the nanostructures of interest.

Massimo V. Fischetti  
Department of Materials Science and Engineering

University of Texas at Dallas  
max.fischetti@utdallas.edu

William Vandenberghe  
The University of Texas at Dallas  
wxv101020@utdallas.edu

Maarten Van de Put  
University of Antwerp and imec  
vandeput.maarten@gmail.com

Jingtian Fang  
Department of Physics and Astronomy, Vanderbilt University  
jxf114930@utdallas.edu

Shanmeng Chen  
The University of Texas at Dallas  
sxc142530@utdallas.edu

### MS6

#### Modeling Spin-Dependent Phenomena for New Device Applications

To enhance performance and to reduce energy consumption, electron spin is investigated as replacement of electron charge for storage and computation in future electronic devices. In contrast to charge, spin is not conserved but relaxing to its equilibrium zero value. In the case of semiconductor films, technologies are needed to boost the spin lifetime to make the technology feasible. In silicon films the spin lifetime is determined by intervalley scattering. Applying stress appropriately decreases this undesired scattering effect, which should reduce the spin relaxation [Sverdlov and Selberherr, Phys. Rep. 585, 2015]. To confirm, we solve numerically the quantum mechanical description of the behavior of spin by calculating the wave functions and the scattering matrix elements via solving a kp Hamiltonian with spin-orbit interaction. For the computation of spin relaxation, the momentum dependence of the wave functions due to spin-orbit coupling responsible for spin admixture must be preserved, which significantly increases the demands for computational resources and requires extensive code parallelization. We predict a significant spin lifetime increase with appropriate viable stress. A universal behavior of the spin lifetime on the spin injection orientation is also obtained, with even a two times spin lifetime enhancement for relevant cases. These results confirm that electron spin can be used in electronic devices in a competitively to electron charge superior manner.

Viktor Sverdlov  
Institute for Microelectronics  
TU Wien  
sverdlov@iue.tuwien.ac.at

Josef Weinbub  
Christian Doppler Laboratory for High Performance TCAD  
Institute for Microelectronics, TU Wien  
weinbub@iue.tuwien.ac.at

Siegfried Selberherr  
Institute for Microelectronics  
TU Wien

selberherr@iue.tuwien.ac.at

## MS6

### Computational and Numerical Challenges in Semiconductor Process Simulation

Semiconductor process simulation enables to predict critical manufacturing steps of semiconductor devices and circuits. In particular, it allows for so-called structure prototyping used in the critical exploratory phases of new semiconductor technologies. The ever-ongoing reduction of feature sizes and the required move towards increasingly intricate device structures requires efficient methods to enable fast, accurate, and stable modeling solutions. This is particularly challenging for necessary three-dimensional modeling approaches due to the involved computational effort and the limited acceptable simulation run-times, dictated by the fast-pacing progress in electronics. In this talk, current challenges in semiconductor process simulation, with a focus on computational and numerical aspects, will be discussed along with suggested approaches to tackle those. Among the topics are high performance visibility computation for plasma etching, three-dimensional growth-rate modeling for silicon carbide oxidation, parallel level set re-distancing for surface tracking, and efficient spatial discretizations for surface representations. The talk will also show the inter-disciplinary nature of the involved topics, enabling to introduce synergistic interfaces to other fields of research. The financial support by the Austrian Federal Ministry of Science, Research and Economy and the National Foundation for Research, Technology and Development is gratefully acknowledged.

Paul Manstetten  
Christian Doppler Laboratory for High Performance TCAD  
Institute for Microelectronics, TU Wien  
manstetten@iue.tuwien.ac.at

Vito imonka  
CD Laboratory for HPTCAD, Institute for Microelectronics  
TU Wien  
simonka@iue.tuwien.ac.at

Georgios Diamantopoulos, Lukas Gnam, Alexander Makarov  
Christian Doppler Laboratory for High Performance TCAD  
Institute for Microelectronics, TU Wien  
diamantopoulos@iue.tuwien.ac.at,  
gnam@iue.tuwien.ac.at, makarov@iue.tuwien.ac.at

Andreas Hössinger  
Silvaco Europe Ltd.  
andreas.hoessinger@silvaco.com

Josef Weinbub  
Christian Doppler Laboratory for High Performance TCAD  
Institute for Microelectronics, TU Wien  
weinbub@iue.tuwien.ac.at

## MS7

### Walk-Based Centrality and Communicability Measures: Algorithms and Applications

In this talk I will review some network centrality and com-

municability measures based on walks. These measures can be given an elegant closed form in terms of functions of the adjacency matrix. I will introduce the notion of total communicability of a network as a measure of network connectivity and robustness and show that it can be computed very quickly even for large graphs. Finally, I will discuss efficient edge modification strategies (including edge removal, addition, and rewiring) that can be used to obtain networks with desirable communicability properties. The talk is based on recent work in collaboration with Christine Klymko (LLNL) and Francesca Arrigo (Strathclyde).

Michele Benzi  
Department of Mathematics and Computer Science  
Emory University  
benzi@mathcs.emory.edu

## MS7

### Multilevel Model Reduction for Dynamic Power Grid Networks

An electric power grid system can be described by a system of differential-algebraic equations defined on a graph. A relevant task for the power grid community is to develop accurate model reduction for these systems, especially given the large size and stochasticity of the emerging grid. In this talk, we will look at multi-level approaches to constructing the reduced model. The basis of these approaches is the information contained in the weighted graph Laplacian generated by the admittance matrix of the grid. We describe this basis and examine some subtle challenges that these multi-level approaches face in order to achieve accurate simulation of dynamical power grid systems.

Barry Lee  
Department of Mathematics  
Southern Methodist University  
barryl@mail.smu.edu

## MS7

### Dynamic Processes Over Information Networks: Representation, Modeling, Learning and Inference

Large-scale and high resolution data from dynamic processes over networks are becoming increasingly available nowadays from online social platforms, such as Twitter and Facebook. Such data provide great opportunities for understanding and modeling both macroscopic (network level) and microscopic (node-level) patterns in human dynamics. Such data have also fueled the increasing efforts on developing methods to address the challenges arising from understanding, predicting, controlling and distilling knowledge from these dynamic processes over networks, and answer query such as "who will do what and when?" To tackle these challenges, I will present a framework based on point processes for representing and modeling such data, and performing learning, inference and control over dynamic processes over networks.

Le Song  
Georgia Tech  
lsong@cc.gatech.edu

## MS7

### Influence Prediction for Continuous-Time Information Propagation on Networks Using Graph-Based Fokker-Planck Equation

We consider the problem of predicting influence, defined as

the expected number of infected nodes, resulted from information propagating from any given set of source nodes on a network. We develop a novel and transformative framework that adaptively aggregates the activation states of the network according to the number of active nodes, leading to the construction of a system of differential equations that governs the time evolution of the state probabilities. This system is analogous to the Fokker-Planck equation in continuous space, and the solution readily yields the desired influence. This approach gives rise to a class of novel and scalable algorithms that work effectively for large-scale and dense networks. Numerical results on a variety of synthetic and real-world networks will be presented.

Xiaojing Ye  
Department of Mathematics & Statistics  
Georgia State University  
xye@gsu.edu

## MS8

### Exploiting Low Tensor-Rank Structures in the Solutions to High-Dimensional Parametric PDE Problems

We will discuss the intrinsic low-rank structures of numerical solutions to control and optimization problems for unsteady partial differential equations (PDEs) with possibly uncertain parameters and their exploitation using tensor-based techniques. These solutions often exhibit low-rank tensor structures after a separation of variables, which in our case will be time, spatial coordinates, and (stochastic/uncertain) parameters. This is essentially the main idea used in proper generalized decomposition (PGD) methods, but other techniques can be derived based on the same principle. We will highlight two examples: the first can be considered as generalized proper orthogonal decomposition (genPOD), the second builds upon the tensor train (aka matrix product states) format and can be used to implement the numerical solution of optimality systems for PDE-constrained optimization problems with a fraction of the memory requirement and computational cost it would need to solve these systems in the usual sparse matrix formats obtained from classical discretizations. Both approaches are illustrated for optimal control of flow problems, described by Burgers, Stokes, Stokes-Brinkman, and incompressible Navier-Stokes equations.

Peter Benner  
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany  
benner@mpi-magdeburg.mpg.de

Sergey Dolgov  
University of Bath  
Department of Mathematical Sciences  
s.dolgov@bath.ac.uk

Jan Heiland  
TU Berlin, Germany  
heiland@mpi-magdeburg.mpg.de

Akwum Onwunta  
Max Planck Institute, Magdeburg, Germany  
onwunta@mpi-magdeburg.mpg.de

Martin Stoll  
Max Planck Institute, Magdeburg

stollm@mpi-magdeburg.mpg.de

## MS8

### Structure-Preserving Nonlinear Model Reduction for Finite-Volume Models

Finite-volume discretizations are commonly used for numerically solving systems governed by conservation laws. Such techniques numerically enforce conservation laws over cells by discretizing the *integral form* of the conservation laws. However, such models can become very large-scale for when high spatial resolution is required, leading to simulation times that can exceed weeks on a supercomputer; this renders such models impractical for many-query and real-time scenarios. Reduced-order models (ROMs) have been developed to mitigate this burden. While some techniques have been successfully applied to finite-volume models, these techniques do not guarantee that conservation laws are enforced over any subset of the computational domain. As a result, instability and inaccurate responses—characterized by spurious generation/dissipation of quantities that should be conserved—are common. To address this, we propose a novel nonlinear model-reduction technique that *explicitly enforces* conservation laws over subdomains of the problem. This guarantees that, even when the dimensionality of the model is reduced, the most important structure intrinsic to the finite-volume model—the conservation laws—are enforced over subdomains. Furthermore, we equip both the objective function and constraints with hyper-reduction via gappy POD; this ensures computational efficiency in the presence of nonlinearities.

Kevin T. Carlberg, Youngsoo Choi  
Sandia National Laboratories  
ktcarlb@sandia.gov, ychoi@sandia.gov

Syuzanna Sargsyan  
University of Washington  
ssusie@uw.edu

## MS8

### Optimal Model Reduction of Systems with Quadratic Nonlinearity

Input-independent, optimal model reduction techniques have been successfully extended from linear to bilinear dynamical systems, a special case of weakly nonlinear systems, thus producing a reduced model which is uniformly good for a wide range of input functions. In this talk, we extend these developments to an important class of nonlinear systems, namely quadratic-bilinear (QB) dynamical systems. First, we define a truncated  $\mathcal{H}_2$ -norm for QB systems and derive the first-order necessary conditions for an optimal approximation, minimizing the truncated  $\mathcal{H}_2$ -norm. This leads to an iterative model reduction algorithm, which upon convergence yields a reduced-order system that *approximately* satisfies the derived optimality conditions. We illustrate the efficiency of the proposed method by means of several nonlinear partial differential equations, and show its competitiveness with the existing model reduction schemes for QB systems such as moment-matching methods and balanced truncation.

Peter Benner  
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany  
benner@mpi-magdeburg.mpg.de

Pawan Goyal



Max Planck Institute Magdeburg for  
Dynamics of Complex Technical Systems  
goyalp@mpi-magdeburg.mpg.de

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

## MS8

### Surrogate Modelling by Nonsymmetric Greedy Kernel Approximation

Many applications in CSE have as essential modular component one or several mappings from some input space to some output space which needs to be evaluated repeatedly and rapidly. If this mapping is computationally costly, e.g. requiring the solution of a PDE, this may be prohibitive for real-time situations and data-based surrogate models may be a promising solution. In this presentation, we focus on non-symmetric greedy kernel approximation schemes for construction of such surrogate models. Based on a large sample set of input-output data pairs, we incrementally construct approximating models by a greedy selection procedure. This nonsymmetric approach is more versatile than the symmetric greedy approximation, since it comprises it as a special case and, moreover, we also allow the kernel centers to be different from the data-sites. We discuss different selection criteria for the test and trial spaces, stability and approximation properties, and relate to some existing methods. Experimentally, we demonstrate the potential on artificial examples as well as on real world biomechanical applications.

Bernard Haasdonk, Gabriele Santin  
University of Stuttgart  
haasdonk@mathematik.uni-stuttgart.de,  
gabriele.santin@mathematik.uni-stuttgart.de

## MS9

### $\mathcal{H}$ -Matrix Accelerated Second Moment Analysis for Second Order PDEs

The efficient solution of operator equations with random right hand sides is considered. The solutions two-point correlation can efficiently be computed by means of a sparse grid or a low-rank approximation if the two-point correlation of the right hand side is sufficiently smooth. Unfortunately, the problem becomes much more involved in case of rough data. However, the rough data and also the inverse operators can efficiently be represented or approximated by means of  $\mathcal{H}$ -matrices. This enables us to solve the correspondent  $\mathcal{H}$ -matrix equation in almost linear time by the use of the  $\mathcal{H}$ -matrix arithmetic. Numerical experiments stemming from partial differential equations with random input data discretized by the finite element and the boundary element method are provided. Shape calculus allows the efficient treatment of PDEs on random domains by the presented method.

Juergen Doelz  
Universitaet Basel  
Departement of Mathematics and Computer Science  
juergen.doelz@unibas.ch

Helmut Harbrecht  
Universitaet Basel  
Department of Mathematics and Computer Science  
helmut.harbrecht@unibas.ch

Michael Peters  
University of Basel  
michael.peters@unibas.ch

Christoph Schwab  
ETH Zuerich  
SAM  
christoph.schwab@sam.math.ethz.ch

## MS9

### Hierarchical Matrices in Scattered Data Approximation

A scattered data interpolation problem might be stated as follows: Given data  $\{(x_i, y_i) \mid 1 \leq i \leq N\}$  with data sites  $x_i \in R^d$  and measurements  $y_i \in R$ , find a function  $s \in S_X$  such that  $s(x_i) = y_i$ ,  $i = 1, \dots, N$ . Here,  $S_X$  denotes an  $N$ -dimensional linear space of (continuous) functions that depend on the data sites. Given a basis  $\{b_1, \dots, b_N\}$  of  $S_X$ , the coefficient vector  $c \in R^N$  of the interpolant  $s(x) = \sum_{i=1}^N c_i b_i(x)$  may be computed as the solution of the linear system  $Bc = y$  with  $B = (b_{ij}) = (b_j(x_i))$ . While properties of the matrix  $B$  obviously depend on the function space  $S_X$  and its basis, several of the most commonly used approaches yield ill-conditioned, dense matrices  $B$ , resulting in a challenge to solve the linear system  $Bc = y$  which is the topic of this talk. In particular, we study the application of techniques from the  $\mathcal{H}$ -matrix framework both for the approximation of the system matrix  $B$  itself as well as for the construction of preconditioners.  $\mathcal{H}$ -matrices provide a data-sparse matrix format that permits storage and matrix arithmetic in complexity  $\mathcal{O}(N \log^\alpha N)$  for moderate  $\alpha$ . It turns out that several typical sets of basis functions from the (scattered data) literature lead to matrices  $B$  that fit into this framework, yielding a cost-effective approximation scheme to be illustrated in this talk.

Sabine Le Borne  
Hamburg University of Technology  
Department of Mathematics  
leborne@tuhh.de

## MS9

### Stable Recurrences for Orthogonal Transformation of Givens-Weight and Block Givens-Weight Representations

This talk describes a formal framework for the derivation and error analysis of algorithms that perform various orthogonal transformations of a rank structured matrix represented using the Givens-weight parameterization of Delvaux and Van Barel. The use of swapping in Givens-weight algorithms can be formalized with a set of swapping recurrences that can be specialized to give multiple algorithms for matrices with quasiseparable structure, including a fast  $QR$  factorization and an associated system solver. A general error analysis of the swapping recurrences translates to a backward stability proof of the class of orthogonal algorithms based on the recurrences. A joint Givens-weight parameterization can be used to extend the recurrences so that they can be applied to matrices with rank structured blocks.

Michael Stewart  
Georgia State University  
Department of Mathematics

mastewart@gsu.edu

troy.butler@ucdenver.edu

**MS9****Fast Bidirectional Solver for the High-Frequency Lippmann-Schwinger Equation**

In this talk we present a fast iterative solver for Lippmann-Schwinger equation for high-frequency waves scattered by a smooth medium with a compactly supported inhomogeneity. The solver is based on the sparsifying preconditioner and the method of polarized traces. The iterative solver has two levels, the outer level in which a sparsifying preconditioner for the Lippmann-Schwinger equation is constructed, and the inner level, in which the resulting sparsified system is solved fast using an iterative solver preconditioned with a bi-directional preconditioner. The complexity of the construction and application of the preconditioner is  $O(N)$  and  $O(N \log N)$  respectively, where  $N$  is the number of degrees of freedom. Numerical experiments in 2D indicate that the number of iterations in both levels depends weakly on the frequency resulting in a method with an overall  $O(N \log N)$  complexity.

Leonardo Zepeda-Nunez  
University of California, Irvine  
lzepeda@math.mit.edu

Hongkai Zhao  
University of California, Irvine  
Department of Mathematics  
zhao@math.uci.edu

**MS10****From Inverses to Predictions: End-to-End UQ With MC Methods**

We use a measure-theoretic framework to describe a full end-to-end quantification of uncertainties by solving stochastic inverse problems in order to inform predictions of a computational model. We use Monte Carlo and other sample based methods to describe the solutions, which are often complicated non-parametric probability measures. Furthermore, the ultimate goal for solving these problems can often be written as an integral over the probability spaces (e.g., as expected values of certain random variables on either input or prediction spaces). We then offer a measure-theoretic perspective for Monte Carlo integration that can be extended to most numerical integration schemes. The basic idea is that samples implicitly define a Voronoi tessellation of a space, and we can use the geometric properties of both the Voronoi tessellation and response surfaces to inform solutions to both forward and inverse quantification of uncertainty problems. We use additional Monte Carlo schemes for interrogating these geometric properties so that explicit construction of Voronoi cells and almost everywhere evaluation of response surfaces can be avoided in solving these problems. Examples are given where Monte Carlo methods are used at every stage of the formulation and solution of an end-to-end UQ problem including (1) the design of experiments leading to inverse problems, (2) the solution of these inverse problems, and (3) the propagation of the inverse solution to a prediction space.

Troy Butler  
University of Colorado Denver

**MS10****A Fast Algorithm to Compute Flow Ensembles**

Repeated computations of flow equations with varying parameters are commonly seen in many engineering and geophysical applications as an effort to deal with inherent uncertainties. These computations are generally treated as independent tasks. While parallel computing can save computational time in this setting, no savings are realized in terms of total computational cost. In this talk, we will describe a new way to perform multiple simulations efficiently, in terms of both storage and computational cost. The proposed algorithm computes all realizations at one pass by adopting an ensemble time-stepping scheme, which results in the same coefficient matrix for all realizations. This reduces the problem of solving multiple linear systems to solving one linear system with multiple right-hand sides, for which many efficient methods, e.g., block CG, block QMR, block GMRES, have been developed to significantly save the computation cost.

Nan Jiang  
Missouri University of Science and Technology  
jiangn@mst.edu

**MS10****Efficient Algorithms for MHD Simulation and Ensemble Calculations**

An efficient algorithm is proposed and studied for computing flow ensembles of incompressible magnetohydrodynamic (MHD) flows under uncertainties in initial or boundary data. The key ideas are the use of Elsasser variable formulation and adapting a clever idea for ensemble averaging of Jiang et. al to the MHD setting. At each time step, the exact same linear systems arise for each of the  $J$  system solves. Further, the Elsasser variable formulation is exploited, following a recent idea of Trenchea et. al to stably decouple each MHD system into two Oseen problems. We prove stability and convergence of the algorithm, and test it with numerical experiments.

Leo Rebholz  
Clemson University  
Department of Mathematical Sciences  
rebholz@clemson.edu

Muhammad Mohebujjaman  
Clemson University  
mmohebu@clemson.edu

**MS10****Efficient Sampling Strategies for the Consistent Bayesian Approach for Solving Stochastic Inverse Problems**

Uncertainty is ubiquitous in computational science and engineering. Often, parameters of interest cannot be measured directly and must be inferred from observable data. The mapping between these parameters and the measurable data is often referred to as the forward model and the goal is to use the forward model to gain knowledge about the parameters given the observations on the data. Statistical Bayesian inference is the most common approach for incorporating stochastic data into probabilistic descriptions of the input parameters. We have recently developed

an alternative Bayesian solution to the stochastic inverse problem. We use measure-theoretic principles to prove that this approach produces a posterior density that is consistent with the model and the data in the sense that the push-forward of the posterior through the model will match the observed density on the data. Our approach requires approximating the push-forward of the prior through the computational model, which is fundamentally a forward propagation of uncertainty. Samples from the posterior can be selected from the samples used to compute the push-forward of the prior using a basic rejection strategy. This does not require MCMC or additional model evaluations to generate the initial set of samples, but it may not yield sufficient samples from the posterior. In this presentation, we explore approaches to generate additional samples from the posterior using the consistent Bayesian approach.

Tim Wildey

Optimization and Uncertainty Quantification Dept.  
Sandia National Laboratories  
tmwilde@sandia.gov

Troy Butler

University of Colorado Denver  
troy.butler@ucdenver.edu

John D. Jakeman

Sandia National Labs  
jdjakem@sandia.gov

## MS11

### A Unified Approach to Periodization of Fast Algorithms for Laplace and Stokes Problems

We present a spectrally-accurate approach to periodic boundary conditions built on standard fast algorithms for the free-space kernels. The basic idea is to split the free-space image sum into near neighbors summed via the FMM, plus a far-field remainder which is approximated by a small number of auxiliary sources whose strengths are found by dense linear algebra. Physical boundary conditions such as pressure drops are applied directly, avoiding special choices for non-convergent lattice sums. The scheme is simple, applies to arbitrary unit cells, inherits the adaptivity of the FMM (no uniform meshes are needed), is essentially dimension- and PDE-independent, and couples to existing spectral close-evaluation quadratures. We showcase applications in large-scale high-accuracy integral equation simulations of the effective permeability of composite media, microfluidic design, and vesicle dynamics, in two and three space dimensions and various dimensions of periodicity.

Alex H. Barnett

Dartmouth College and Simons Foundation  
ahb@math.dartmouth.edu

Gary Marple, Shравan Veerapaneni

Department of Mathematics  
University of Michigan  
gmarple@umich.edu, shravan@umich.edu

Lin Zhao

Intech, Princeton, NJ  
lzhao@intechjanus.com

## MS11

### Contact-Aware Simulations of Particulate Stoke-

## sian Suspensions

We present an efficient, accurate, and robust method for the simulation of dense suspensions of deformable and rigid particles immersed in Stokesian fluid. We use a well-established boundary integral formulation for the problem as the foundation of our approach. This type of formulations, with a high-order spatial discretization and an implicit and adaptive time discretization, have been shown to be able to handle complex interactions between particles with high accuracy. Yet, for dense suspensions, very small time-steps or expensive implicit solves as well as a large number of discretization points are required to avoid non-physical contact and intersections between particles, leading to numerical instability. In our framework, the interference-free configuration is ensured by introducing contact constraints into the system. Our method maintains the accuracy of previous methods at a significantly lower cost for dense suspensions. While interference-free constraints are unnecessary in the continuous formulation, in the discrete form of the problem they make it possible to eliminate catastrophic loss of accuracy by preventing contact.

Abtin Rahimian, Libin Lu

Courant Institute of Mathematical Sciences  
New York University  
arahimian@acm.org, ll1488@nyu.edu

Denis Zorin

Computer Science, NYU  
zorin@cims.nyu.edu

## MS11

### An Integral Equation Method for 3D Surfactant-Covered Drops

We consider picoliter sized drops covered with an insoluble surfactant; at these small scales the inertial effects are negligible and the dynamics are governed by the Stokes equations. Surface active agents (surfactants) modify the surfaces tension forces and can be used to stabilize the system, i.e. preventing the droplets from coalescing. A boundary integral formulation for Stokes flow is used, where an accurate representation of each surface is based on a spherical harmonic expansion. In this context, special attention must be given to the numerical integration, to accurately handle both singular (self-self interaction) and nearly-singular (multiple drops interaction) integrals. When simulating drops, a strong deformation of the point distribution might appear and this can cause aliasing errors and possibly numerical instabilities. For this reason we propose a new reparametrization method which takes into account the whole system (drop/surfactant), updating coherently the point distribution and the surfactant concentration taking advantage of the spectral representation of both.

Chiara Sorgentone

La Sapienza University of Rome  
sorgento@kth.se

## MS11

### A New Boundary Integral Equation Formulation for Vesicle Electrohydrodynamics

The electrohydrodynamics (EHD) of vesicle suspensions is characterized by studying their pairwise interactions in applied DC electric fields in two dimensions. In this talk, a

boundary integral equation (BIE) based formulation for vesicle EHD is introduced, followed by a solution scheme based on Stokes and Laplace potential theory. In the dilute limit, the rheology of the suspension is shown to vary nonlinearly with the electric conductivity ratio of the interior and exterior fluids. The prolate-oblate-prolate transition and other transitional dynamics observed in experiments and previously confirmed via numerical simulations is further investigated here for smaller reduced areas. When two vesicles are initially un-aligned with the external electric field, three different responses are observed when the key parameters are varied: (i) chain formation they self-assemble to form a chain that is aligned along the field direction, (ii) circulatory motion they rotate about each other, (iii) oscillatory motion they form a chain but oscillate about each other.

Bowei Wu  
University of Michigan  
boweiwu@umich.edu

Shravan Veerapaneni  
Department of Mathematics  
University of Michigan  
shravan@umich.edu

## MS12

### Higher-Order Cut Cell Finite Volume Discretizations: Theory and Applications

Embedded boundary (cut cell) approaches for solving PDE's have reduced cost for grid generation and computation, but introduce additional challenges to achieve high accuracy, numerical stability, and other important discretization qualities. We present a quick survey of cut cell techniques and some mathematical and software frameworks, along with some examples of model problems and science applications we are developing. The biggest open questions involve tradeoffs, including stability (theory of "small cells") and reliability (complex geometry, software) issues, but in some cases these are outweighed by the benefits of conservation, accuracy, and performance.

Hans Johansen  
Lawrence Berkeley National Laboratory  
Computational Research Division  
hjohansen@lbl.gov

## MS12

### A Moving Boundary Cut Cell Approach for Incompressible Euler with a Free Surface

We present an adaptive solver for single-phase incompressible flows with free surface motion, which uses an adaptive, higher-order cut-cell discretization. Our approach evolves the free surface in a narrow region at the finest resolution with a gradient-augmented level set method. We reconstruct the interface from the level set using an extremal-limited Hermite interpolant, which improves accuracy and volume preservation. The fluid solver is an adaptive resolution projection scheme that uses higher-order cut cells to represent fluid boundaries. We use an approximate projection scheme that exactly matches the fluid boundary conditions but enforces the divergence-free constraint only to the order of truncation error. Because the fluid is single-phase, we have introduced a novel algorithm to extend the velocity field outside the fluid domain. We show results of validation against idealized non-linear surface waves, and discuss progress towards using the algorithm for simulat-

ing fluid-structure interactions in wave energy conversion devices.

Curtis Lee  
Duke University  
calee181@gmail.com

Dan Graves  
Lawrence Berkeley Laboratory  
dtgraves@lbl.gov

Hans Johansen  
Lawrence Berkeley National Laboratory  
Computational Research Division  
hjohansen@lbl.gov

Amneet Pal S. Bhalla  
Lawrence Berkeley National Laboratory  
apbhalla@lbl.gov

John Dalbow  
Duke University  
john.dalbow@duke.edu

## MS12

### Inverse Lax-Wendroff Procedure for Numerical Boundary Conditions of Convection-Diffusion Equations

We consider numerical boundary conditions for high order finite difference schemes for solving convection-diffusion equations on arbitrary geometry. The two main difficulties for numerical boundary conditions in such situations are: (1) the wide stencil of the high order finite difference operator requires special treatment for a few ghost points near the boundary; (2) the physical boundary may not coincide with grid points in a Cartesian mesh and may intersect with the mesh in an arbitrary fashion. For purely convection equations, the so-called inverse Lax-Wendroff procedure, in which we convert the normal derivatives into the time derivatives and tangential derivatives along the physical boundary by using the equations, have been quite successful. In this talk, we will discuss our extension of this methodology to convection-diffusion equations. It turns out that this extension is non-trivial, because totally different boundary treatments are needed for the diffusion-dominated and the convection-dominated regimes. We design a careful combination of the boundary treatments for the two regimes and obtain a stable and accurate boundary condition for general convection-diffusion equations. We provide extensive numerical tests for one- and two-dimensional problems involving both scalar equations and systems, including the compressible Navier-Stokes equations, to demonstrate the good performance of our numerical boundary conditions

Jianfang Lu  
Applied and Computational Mathematics  
Beijing Computational Science Research Center  
jflu@csrc.ac.cn

Jinwei Fang  
Department of Mathematics  
Harbin Institute of Technology  
gerrard.fang@gmail.com

Sirui Tan  
Division of Applied Mathematics  
Brown University

sirui\_tan@alumni.brown.edu

Chi-Wang Shu  
Brown University  
Div of Applied Mathematics  
shu@dam.brown.edu

Mengping Zhang  
Department of Mathematics  
University of Science and Technology of China  
mpzhang@ustc.edu.cn

### MS12

#### Application of a High-Order Adaptive Cut-Cell Method to Complex Moving Deformable Boundary Problems

Previously we presented a high-order adaptive Cartesian cut-cell method for simulation of embedded boundaries in a compressible viscous flow. The main idea is a high-order Central Essential Non-Oscillatory (CENO) scheme with a novel cell clustering algorithm that addresses the small cell problem common to all cut-cell methods. In this talk, we examine the applicability of the cut-cell method to more challenging engineering applications such as flow past moving deformable boundaries and flow through narrow gaps. Numerical analysis of conservation properties, accuracy, and stability of the scheme in such rigorous conditions will be presented. We also extend the cell clustering algorithm to when a new fluid cell emerges or an existing fluid cell turns into a solid cell due to boundary motion. In particular, this is shown to satisfy the geometric conservation laws especially during the cell vanishing and emerging events. Corrections to fluxes to enforce conservation when an embedded boundary crosses a coarse/fine mesh refinement interface are introduced, and we assess the impact of this correction on the numerical properties of the scheme.

Balaji Muralidharan  
Georgia Institute of Technology  
bmuralidharan3@gatech.edu

Suresh Menon  
Georgia Inst. of Technology  
Sch. of Aerospace Engineering  
suresh.menon@aerospace.gatech.edu

### MS13

#### Numerical Method for the Dirac Equation in the Nonrelativistic Limit Regime

Dirac equation, proposed by Paul Dirac in 1928, is a relativistic version of the Schrödinger equation for quantum mechanics. It describes the evolution of spin-1/2 massive particles, e.g. electrons. Due to its applications in graphene and 2D materials, Dirac equations has drawn considerable interests recently. We are concerned with the numerical methods for solving the Dirac equation in the non-relativistic limit regime, involving a small parameter inversely proportional to the speed of light. We begin with commonly used numerical methods in literature, including finite difference time domain and time splitting spectral, which need very small time steps to solve the Dirac equation in the non-relativistic limit regime. We then propose and analyze a multi-scale time integrator pseudospectral method for the Dirac equation, and prove its uniform convergence in the non-relativistic limit regime.

Yongyong Cai

Purdue University  
yongyong.cai@gmail.com

### MS13

#### Numerical Analysis of Finite Temperature DFT

We study finite dimensional approximations of the Mermin-Kohn-Sham equation, which is derived from the finite temperature DFT model. For general numerical discretization methods, we prove the convergence of the finite dimensional approximations and derive the optimal a priori error estimates. We also provide numerical simulations for several molecular systems that support our theory.

Huajie Chen  
Beijing Normal University  
hjchen@lsec.cc.ac.cn

### MS13

#### A Fast Deterministic Solver for the Schrodinger-Quantum Boltzmann System

We develop a fast deterministic solver for the Schrodinger-quantum Boltzmann system. This is a coupled system to describe the kinetics of Bose-Einstein condensate at finite temperature, in which the dynamics of the condensate is described by a generalized Gross-Pitaevskii equation, and the thermal cloud is represented by a generalized quantum Boltzmann equation. We will demonstrate the efficiency and accuracy of the solver using a series of examples.

Jingwei Hu  
Department of Mathematics, Purdue University  
jingwei.hu@purdue.edu

### MS13

#### Towards a Mathematical Understanding of Surface Hopping Algorithms

We develop a surface hopping algorithm based on frozen Gaussian approximation for semiclassical matrix Schrödinger equations. The algorithm is asymptotically derived from the Schrödinger equation with rigorous approximation error analysis. The resulting algorithm can be viewed as a path integral stochastic representation of the semiclassical matrix Schrödinger equations. Our results provide mathematical understanding to and shed new light on the important class of surface hopping methods in theoretical and computational chemistry. Also, I would like to report our recent progress on the improved surface hopping algorithm with various numerical tests.

Zhennan Zhou  
Duke University  
zhennan@math.duke.edu

Jianfeng Lu  
Mathematics Department  
Duke University  
jianfeng@math.duke.edu

### MS14

#### Radiation Hydrodynamics in GenASiS

GenASiS (*General Astrophysical Simulation System*) is a code under development initially and primarily for the core-collapse supernova problem, and as such neutrino radiation transport is a key part of the physics being de-

veloped within it. Classes representing phase space (position space plus momentum space) and particle distributions thereon are key pieces of infrastructure, and are being developed in a way that will be extensible from energy-dependent angular moment formalisms in the near term to full Boltzmann transport in the long term, including the potential for adaptive mesh refinement in both position space and momentum space. The spatial streaming operator is handled in conservative form in a time-explicit manner similar to fluid dynamics, while the collision operator is handled implicitly; these substeps are combined in an implicit-explicit Runge-Kutta scheme. Both finite-volume and discontinuous Galerkin versions are contemplated, as are gradual implementation of relativistic effects and increasingly sophisticated interactions with matter.

Christian Cardall

Oak Ridge National Laboratory  
cardalcy@ornl.gov

#### MS14

##### **Positivity Limiters on the Filtered $P_N$ Method for Linear Transport Equations**

We analyze and compare the properties and performance of several positivity limiters on the recently proposed filtered  $P_N$  ( $FP_N$ ) method for linear transport equations. The original  $FP_N$  method is known to suffer from the occurrence of (unphysical) negative particle concentrations, which originates from the fact that  $FP_N$  spherical harmonic approximations are not always positive at the kinetic level. These limiters enforce positivity of the  $FP_N$  approximations on a finite set of pre-selected points. With a proper PDE solver, this ensures positivity of the particle concentration at each step in the time integration. We give error estimates for the positive approximations produced by these limiters, and verify the estimates with numerical consistency tests. We simulate problems of various regularities using the  $FP_N$  method with several limiters, and report the efficiency of these limiters. The numerical results give a guideline on selecting positivity limiters for solving problems with different regularities.

Ming Tse P. Laiu

Department of Electrical and Computer Engineering  
University of Maryland - College Park  
mtlaiu@umd.edu

Cory Hauck

Oak Ridge National Laboratory  
hauckc@ornl.gov

#### MS14

##### **Multi-Dimensional Boltzmann-Neutrino-Radiation-Hydrodynamic Simulations in Core Collapse Supernovae**

Core collapse Supernovae (CCSNe) are catastrophic explosion phenomena in our universe. They are not only astrophysically interesting events but also important experimental sites for high energy physics. Unfortunately, however, astrophysicists have not achieved to comprehend the explosion mechanism yet since the physics in CCSNe is quite complex. Numerical modeling of CCSNe would be strong (and may be unique) tool to determine what physical ingredient(s) is (are) essential. Therefore, the progress of CCSNe community is strongly related with developing computational science. The major impediment to build the theory of CCSNe is related with the matter of radiation-

hydrodynamics. Neutrinos play a fundamental role for CCSNe, and they are, in general, not in equilibrium with matter. Therefore, we have to treat their dynamics as the 7-dimensional phase space problems i.e., we need to solve Boltzmann equation. It should be noted, however, that there have been, thus far, many technical and computational problems. Our group have tackled these issues and eventually overcame all of them. Indeed, we recently succeeded to carry out 2D CCSNe simulations by using newly developed code. In this conference, I will talk about methods, results of recent simulations and further progress in our project.

Hiroki Nagakura

California Institute of Technology  
hirokin@caltech.edu

Wakana Iwakami

Yukawa Institute for Theoretical Physics  
wakana@heap.phys.waseda.ac.jp

Shun Furusawa

Frankfurt Institute for Advanced Studies,  
J.W.Goethe University  
furusawa@fias.uni-frankfurt.de

Kohsuke Sumiyoshi

Numazu College of Technology  
sumi@numazu-ct.ac.jp

Shoichi Yamada

Advanced Research Institute for Science & Engineering  
Waseda University  
shoichi@heap.phys.waseda.ac.jp

Hideo Matsufuru

High Energy Accelerator Research Organization  
hideo.matsufuru@kek.jp

Akira Imakura

University of Tsukuba  
imakura@ccs.tsukuba.ac.jp

Sherwood Richers

California Institute of Technology  
srichers@tapir.caltech.edu

Christian Ott

Caltech  
cott@tapir.caltech.edu

#### MS14

##### **Two-Moment Neutrino Transport Coupled to Relativistic MHD for Stellar Core Collapse**

Collapse and explosion of massive stars depend critically on a wide variety of physics including relativistic gravity and hydrodynamics, potentially magnetic fields, the nuclear equation of state, and the transport of neutrinos and their interaction with matter. Our simulation code couples special relativistic magnetohydrodynamics and two-moment neutrino transport, which is based on the evolution equations for the energy density and fluxes of the neutrinos that are closed by a local algebraic closure. The transport equations are written in a frame comoving with the fluid and include velocity-dependent terms and, additionally, corrections due to gravity. We will present our implementation and results of multi-dimensional simula-

tions of rotating, magnetized high-mass stars for which explosions of different types are launched, depending on the importance of neutrino heating and magneto-rotational effects.

Martin Obergaullinger, Miguel Angel Aloy  
University of Valencia  
martin.obergaullinger@uv.es, miguel.a.aloy@uv.es

#### MS15

##### **Spectral Element Simulations of Turbulence on ALCF Theta**

We present performance results for the spectral-element-based Navier-Stokes code, Nek5000, on the Argonne Leadership Computing Facility Theta, which features Intel's KNL many-core nodes. We describe on- and off-node performance as well as ongoing optimization efforts.

Paul Fischer  
University of Illinois at Urbana-Champaign  
fischerp@illinois.edu

#### MS15

##### **H-to-P Efficiently: The Use of Collections with Accelerators Within Nektar++**

As the HPC community looks towards the use of many-core technology to enable the next generation of supercomputers, there is a need to understand how this new paradigm maps to existing codes and algorithms that can already effectively utilise current generation HPC facilities. In this presentation, we present the approaches currently being undertaken inside the Nektar++ spectral/hp element framework to tackle this challenging issue. Nektar++ provides solver developers with a toolbox of techniques that can be used to employ the use of high-order spectral elements for various solvers, with application areas including incompressible and compressible fluid dynamics, cardiac electrophysiology and advection-diffusion-reaction problems. Spectral elements provide increased compactness and memory efficiency over the more traditionally-used lower-order methods, meaning they have far greater potential to exploit many-core processors. The presentation will demonstrate our existing techniques for decreasing runtimes across a range of polynomial orders, by exploiting the mathematical framework of the method to admit a number of different collective operators that are key to the method. We will show how this collective nature provides a route through which we can tackle the challenge of many-core hardware in our existing codebase.

David Moxey  
Department of Aeronautics  
Imperial College London  
d.moxey@imperial.ac.uk

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

Mike Kirby  
University of Utah  
School of Computing  
kirby@cs.utah.edu

Spencer Sherwin  
Imperial College London

s.sherwin@imperial.ac.uk

#### MS15

##### **Performance Evaluation of a Very Efficient High-Order Accurate Imex Solver for Incompressible Fluid Dynamics on Intel Xeon Phi**

Abstract not available

Matteo Parsani  
KAUST  
matteo.parsani@kaust.edu.sa

#### MS15

##### **PyFR and GiMMiK on Intel KNL: Performance, Scalability, and Applications**

PyFR is a Python based framework for solving the compressible Navier-Stokes equations on mixed unstructured grids using the high-order flux reconstruction (FR) approach. Through the use of run-time code generation, enabled via the use of a bespoke domain specific language, the framework is capable of achieving portability across a range of architectures. Within an FR time step it is possible to cast many of the operations in the form of block-by-panel type matrix multiplications. In the case of PyFR these operations are offloaded to the GiMMiK library. In this talk we will show how PyFR and GiMMiK were adapted to run on the Intel KNL platform. Performance and scalability results will be presented and compared with our existing results on both the first generation Xeon Phi and traditional Xeon platforms. Further, we will also outline our experiences with how best to manage the MCDRAM from Python and how this can be integrated into the NumPy ecosystem. Special attention will also be given towards the issue of the start-up time of PyFR on KNL.

Freddie D. Witherden  
Department of Aeronautics and Astronautics  
Stanford University  
fdw@stanford.edu

Jin Seok Park  
Department of Aeronautics  
Imperial College London  
jin-seok.park@imperial.ac.uk

Alexander Heinecke  
Parallel Computing Laboratory  
Intel Corporation, Santa Clara, CA, USA  
alexander.heinecke@intel.com

Paul Kelly  
Imperial College London  
p.kelly@imperial.ac.uk

Peter E. Vincent  
Department of Aeronautics  
Imperial College London  
p.vincent@imperial.ac.uk

Antony Jameson  
Professor, Department of Aeronautics & Astronautics  
Stanford University

jameson@baboon.stanford.edu

## MS16

### Hessian Approximation Using Operator Symbols

In this talk we discuss a method to derive the Hessian for shape optimization with the help of operator symbols based on Fourier analysis. The derivation of the symbol is performed analytically for the Stokes equations and shows good agreements with the Hessian obtained with finite differences for small Reynolds flows. From the resulting symbol one can deduce that the Hessian is a pseudo-differential operator of order one. As a consequence the symbol can be used to smooth the search direction, potentially leading to a faster convergence while preserving a smooth design. To assure an efficient computation of the resulting preconditioner the pseudo-differential Hessian operator is approximated by an even order differential operator defined on the shape. To mimic linear scaling behavior for significant frequencies, we make use of a windowed Fourier transform. The preconditioner has been applied for drag minimization and shows a faster convergence compared to classical Sobolev smoothing for Reynolds numbers up to 80.

Jonas Kusch

RWTH Aachen University

jkusch@mathcces.rwth-aachen.de

Nicolas R. Gauger

TU Kaiserslautern

nicolas.gauger@scicomp.uni-kl.de

Stephan Schmidt

University Würzburg

stephan.schmidt@mathematik.uni-wuerzburg.de

## MS16

### Procedurally Generated Shape Hessians and Their Computation Using Fenics

A novel approach using shape hessian information to solve optimisation problems governed by the Stokes equations is considered. The shape hessian is used within the SQP framework to solve the nonlinear shape optimisation problem. This is either done by solving a reduced or full size KKT System to gain a decent direction for the change of the domain. Standard gradient type line search algorithms have to face the problem of slow convergence and determination of a suitable step length. Newton type optimisation approaches in shape optimisation have the ability to overcome these issues, however they introduce several challenges which will be addressed within this presentation. The shape hessian can be divided into a geometric and dynamic part, whereas the dynamic part is the dominating contribution. Computation of the dynamic term implies to solve a linearized state equation for the local shape derivative and to evaluate its normal derivative, which leads to a lack of regularity within practical implementation. The Finite Element software FEniCS was used to implement the proposed shape hessian approach. Numerical difficulties in the mesh update and mesh quality preservation together with possible measures to address these problems will be presented. Having a stable algorithm, the gained results and their superiority over gradient based methods will be shown.

Edwin Mai

Universität der Bundeswehr

München

edwin.mai@unibw.de

Stephan Schmidt

University Würzburg

stephan.schmidt@mathematik.uni-wuerzburg.de

## MS16

### Different Taping Strategies for AD-based Adjoints Analyzed for SU2

Code generation for adjoints can be automatized via tools for Algorithmic Differentiation (AD). There are two ways how intermediate values of the initial (primal) simulation can be stored for the adjoint or reverse sweep for efficient derivative computation, namely Jacobi taping and primal value taping. Jacobi taping is easy to apply and implement but yields an overall higher memory consumption. Primal value taping is more involved in the implementation and more things have to be considered when it is used to generate adjoint code. However, the advantage is that the memory consumption is lower compared to Jacobi taping. Both approaches can be supported by different memory management techniques which result in different ways to generate adjoints. We compare in total four techniques, which differ in the required memory, the computational complexity and the maintenance effort. The analysis of the different techniques is done for the SU2 suite, which is an open-source collection of software tools written in C++ for performing Partial Differential Equation (PDE) analysis and solving PDE-constrained optimization problems. The toolset is designed with computational fluid dynamics and aerodynamic shape optimization in mind, but has been extended to treat arbitrary sets of governing equations. We will show an in-depth analysis of all four adjoint approaches realized in SU2 by the AD tool CoDiPack, and finally show their potential for the effective use in aerodynamic shape optimization.

Max Sagebaum, Nicolas R. Gauger, Tim Albring

TU Kaiserslautern

max.sagebaum@scicomp.uni-kl.de,

nicolas.gauger@scicomp.uni-kl.de,

tim.albring@scicomp.uni-kl.de

## MS16

### Automatic Generation of Shape Derivatives and Hessians

A fully automatic framework for generating directional shape derivatives as well as shape Hessians is discussed. To this end, the mathematical problem is formulated in the Uniform Form Language (UFL) and then processed by a semantic analysis, which automatically applies the formal differentiation rules of shape calculus. Based on regularity assumptions of the problem provided by the user, the analysis step then, on its own, chooses or omits tangential calculus, such as the divergence theorem in tangent spaces, and generates the surface representation of the shape derivative if applicable. Otherwise, volume expressions can also be generated. Repeated application of the analysis step results in the automatic derivation of the shape Hessian of the problem. As the directional shape derivative is generated within the UFL language, the resulting UFL expression can next be seamlessly processed by the FEniCS environment, thereby automatically generating the complete optimization loop, including the primal and dual solver.

Stephan Schmidt



University Würzburg  
stephan.schmidt@mathematik.uni-wuerzburg.de

### MS17

#### Advances in Generalised Metropolis-Hastings Algorithms

A recent generalization of the Metropolis-Hastings algorithm allows for parallelizing a single chain using existing MCMC methods (Calderhead, PNAS, 2014). The construction involves proposing multiple points in parallel, then defining and sampling from a finite-state Markov chain on the proposed points such that the overall procedure has the correct target density as its stationary distribution. In this talk I'll discuss this algorithm and some of the most recent advances employing this approach.

Ben Calderhead  
Imperial College London  
b.calderhead@imperial.ac.uk

### MS17

#### Estimating Large-Scale Chaotic Dynamics

Obvious likelihood approaches are not available for chaotic systems, due to the sensitivity of the trajectories to any perturbations in the calculations. For large systems, such as used for weather predictions, Ensemble Prediction Systems (EPS) are used to quantify the uncertainty. Here we extend EPS, with essentially no additional CPU costs, to online estimation by perturbing the model parameters as well. The estimation can be performed both by a covariance update process using importance weights, or by employing evolutionary (DE) optimisation. Here we emphasize the use of DE for multiple cost function situations.

Heikki Haario  
Lappeenranta University of Technology  
Department of Mathematics and Physics  
heikki.haario@lut.fi

### MS17

#### Bayesian Inverse Problems with Infinitely-Divisible Priors

The main modelling step in a Bayesian inverse problem is the choice of the prior measure that reflects our prior information about the solution. In this talk we discuss the well-posedness of Bayesian inverse problems with non-Gaussian priors on Banach spaces. Given a forward map and noise model, our goal is to identify the classes of prior measures that result in a well-posed inverse problem. We observe that well-posedness depends on a balance between the growth of the log-likelihood function and the tail behavior of the prior measure. Then we apply our well-posedness theory to special cases such as additive noise models, linear problems and infinitely-divisible prior measures. In each case we derive simple conditions that ensure the well-posedness of the inverse problem. Finally, we study some practical aspects of Bayesian inverse problems such as their consistent approximation and present concrete examples of well-posed Bayesian inverse problems with infinitely-divisible prior measures.

Bamdad Hosseini  
Department of Mathematics  
Simon Fraser University

bhossein@sfu.ca

### MS17

#### Triple Model Reduction for Efficient Bayesian Inversion

In this work, we present an approach to solve large-scale Bayesian inverse problems in high dimensional parameter spaces. We combine a goal-oriented model reduction (using various flavors of empirical interpolation methods) approach for state, data-informed reduction for parameter, and randomized misfit approach for data reduction. The method is designed to mitigate the bottlenecks due to large scale PDEs, high dimensional parameter space, and an ever-increasing data space.

Vishwas Rao  
ICES, UT Austin  
visrao@vt.edu

Tan Bui-Thanh  
The University of Texas at Austin  
tanbui@ices.utexas.edu

### MS18

#### Practical Algorithms for Large Tensor Computations

This talk reviews the algorithms contained in libtensor, a general-purpose library for tensor computations. The library is not specific to any particular domain of applications, however its main strength is efficient handling of tensor symmetry and sparsity found in many-body quantum chemistry theory, which remains its primary application to date. The central data structure in libtensor is the block tensor, a tiled representation of the tensor with blockwise specification of symmetry and sparsity, i.e. symmetry is defined through relations between blocks, which can be zero or non-zero. This representation is particularly suitable for handling large tensors ( $10^{10}$  entries and more) by storing them out-of-core on shared-memory computers or in-core on distributed architectures. Algorithms that take advantage of symmetry and sparsity in block tensors will be discussed in this presentation. Programming of tensor expressions with libtensor is enabled through a C++ internal domain-specific language. Tensor expressions are converted into an abstract syntax tree, which allows pre-evaluation analysis, symbolic manipulation, and optimization of the computation at runtime. This approach allows the programmer to focus on the correctness of problem setup leaving the details of tensor evaluation to runtime.

Evgeny Epifanovsky  
Q-Chem, Inc.  
Pleasanton, CA 94588  
epif@q-chem.com

### MS18

#### Tensors and the Power of Imagination

Tensor contraction (TC) is an important kernel in many scientific computing applications (SCAs), such as quantum chemistry. Due to the fundamental similarity of TC to matrix multiplication (MM) and to the availability of optimized implementations such as the BLAS, tensor operations have traditionally been implemented in terms of BLAS operations, incurring both a performance and a storage overhead as tensors are converted to matrix form. In-

stead, we employ our imagination and simply pretend that the tensors are matrices during the bulk of the MM algorithm. This is possible using the BLIS framework, as implemented in C++, using duck typing and operator overloading. When actual tensor elements are required, a novel matrix-centric tensor layout—the block scatter matrix—is utilized. We show that this approach yields a TC algorithm with essentially no overhead relative to the equivalent matrix multiplication. We present performance results for this algorithm and also discuss how such an algorithm is best leveraged and extended for use in actual quantum chemistry applications.

Devin Matthews

The University of Texas at Austin  
dmatthews@utexas.edu

**MS18**

**Loop Fusion to Optimize Sequences of Tensor Contractions**

Integral transformation is a fundamental and computationally demanding calculation used in many computational chemistry suites such as NWChem. A commonly used variant transforms a four-dimensional tensor through a sequence of four tensor contractions that each contract a four-dimensional tensor with a two-dimensional transformation matrix. Differing degrees of permutation symmetry in the intermediate and final tensors in the sequence of contractions cause intermediate tensors to be much larger than the final tensor and can limit the number of electronic states that can be modeled on a parallel computer system. Loop fusion, in conjunction with tiling, can be very effective in reducing the total space requirement, as well as data movement. However, there are a large number of possible choices for loop fusion and tiling, and for data/computation distribution across a parallel system. We develop an approach to choice of fusion/tiling configuration, using lower bounds modeling of data movement complexity. We establish relationships between available aggregate physical memory in a parallel computer system and ineffective fusion configurations, enabling their pruning and consequent identification of effective choices and the development of a more effective fused implementation than the existing versions in NWChem.

P. (Saday) Sadayappan, Samyam Rajbhandari  
Ohio State University  
saday@cse.ohio-state.edu, samyamrb@gmail.com

Karol Kowalski, Sriram Krishnamoorthy  
Pacific Northwest National Laboratory  
karol.kowalski@pnnl.gov, sriram@pnnl.gov

Fabrice Rastello  
INRIA  
fabrice.rastello@inria.fr

**MS18**

**TCCG: Tensor Contraction Code Generator**

We present “GEMM-like Tensor-Tensor multiplication” (GETT), a novel approach to tensor contractions that mirrors the design of a high-performance general matrix-matrix multiplication (GEMM). The critical insight behind GETT is the identification of 3 index sets, involved in the contraction, which enable us to systematically reduce an arbitrary tensor contraction to loops around a highly tuned “macro-kernel”. This macro-kernel operates on suitably

prepared (“packed”) sub-tensors that reside in a specified level of the cache hierarchy. In contrast to previous approaches to tensor contractions, GETT exhibits desirable features such as unit-stride memory accesses, cache-awareness, as well as full vectorization, without requiring auxiliary memory. To compare our technique with other approaches, we integrated GETT alongside the so called Transpose-Transpose-GEMM-Transpose and Loops-over-GEMM approaches into an open source “Tensor Contraction Code Generator” (TCCG). The performance results suggest that GETT has the potential of becoming the method of choice: While GETT exhibits excellent performance across the board, its effectiveness for bandwidth-bound tensor contractions is especially impressive, outperforming existing approaches by up to 12.3x. More precisely, GETT achieves speedups of up to 1.42x over an equivalent-sized GEMM for bandwidth-bound contractions while attaining up to 91.3% of peak floating-point performance for compute-bound contractions.

Paul Springer

AICES  
RWTH Aachen University  
paul.springer@rwth-aachen.de

Paolo Bientinesi  
AICES, RWTH Aachen  
pauldj@aices.rwth-aachen.de

**MS19**

**Sedimentation and Deposition Simulation with LibMesh**

Turbidity currents are particle-laden, gravity-driven underflows in which the particles are largely or wholly suspended by fluid turbulence. The particles can be carried for long distances and eventually they will settle, being responsible for sediment deposits generating geological formations of considerable interest for the oil and gas industry. In this work, we present a finite element variational multi-scale formulation applied to the numerical simulation of particle-laden flows. We employ a Eulerian framework to describe the flows in which the mathematical model results from the incompressible NavierStokes equation combined with an advection-dominated transport equation. Parallel adaptive mesh refinement and coarsening (AMR) provided by libMesh is used as an optimal strategies for tackling large-scale sedimentation simulations. Numerical experiments have shown that the present formulation captures most of the relevant turbulent flow features with reasonable accuracy, when compared with highly resolved numerical simulations and experimental data.

Jose J. Camata  
Federal University of Rio de Janeiro  
camata@gmail.com

Andre Rossa  
Dept of Civil Engineering  
COPPE/Federal University of Rio de Janeiro, Brazil  
andre\_rossa@hotmail.com

Adriano Cortes  
UNIRIO  
Brazil  
adrimacortes@gmail.com

Alvaro Coutinho  
Dept. of Civil Engineering

COPPE/Federal University of Rio de Janeiro  
alvaro@nacad.ufrj.br

### MS19

#### Scalable Parallel Implementation of a Coupling Scheme Using LibMesh

In this article, we will present a framework to solve coupled equation systems, based on the libMesh library. This framework adds an abstraction layer that allows the user to couple different models at different scales and with different meshes. The coupling scheme itself is an implementation of the Arlequin method. This method couples the models, defined over different and non-conformal meshes, using a volumetric term, which is calculated from the intersections between the meshes. The coupled system is solved using the LATIN method. Each iteration of this method is divided into a coupled and a decoupled stages, and different solvers can be used in the former stage for each model. As an application of this framework, we will present an implementation of a computational homogenization method applied to polycrystalline materials, which allows us to identify the properties of the macro scale model of these materials knowing those of their heterogeneous and non-linear micro scale model. This work benefited from French state funding managed by the National Research Agency under project number ANR-14-CE07-0007 CouEst.

Regis Cottereau  
CNRS, CentraleSupélec, Université Paris-Saclay  
regis.cottereau@centralesupelec.fr

Thiago Milanetto Schlittler  
CentraleSupélec  
thiago.milanetto-schlittler@centralesupelec.fr

### MS19

#### A Monolithic Approach to Numerical Simulation of Airflow in the Lungs

We present a novel monolithic preconditioner for iterative solution of the geometric multiscale problem of a finite element Navier-Stokes model coupled to a reduced dimensional model. Using a monolithic approach as opposed to the more widely used partitioned approach allows for more robust simulations with larger timesteps. This monolithic preconditioner has been developed using a Schur complement approach that allows for the use of different Navier-Stokes preconditioners depending on the chosen application. We investigate this preconditioner in the context of airflow in the lungs using code developed using the scalable high-performance libraries libMesh and PETSc. We show that using this preconditioner the number of linear solver iterations is not mesh dependent and is largely independent of the number of coupling boundaries between the Navier-Stokes model and reduced dimensional model. We also show that by using a larger timestep we can perform simulations that are more than twice as fast as competing methods and are scalable to millions of unknowns.

James P. Mbewu  
University of Oxford  
james.mbewu@gmail.com

### MS19

#### LibMesh: Past, Present, and Future

An introduction to the design and development process

of the libMesh Finite Element Library. Beginning with the history of libMesh collaboration, this talk discusses the software development challenges posed by an ever-increasing variety of use cases, growing user base, and distributed collaboration between developers with separate goals at a variety of institutions. Interactions with the wider software ecosystem of third party dependencies and dependent applications and middleware are also explored. Many of the design decisions in libMesh and their evolution are discussed, and the talk will conclude with a few examples of upcoming library development directions and features.

Roy Stogner  
University of Texas at Austin  
roystgnr@ices.utexas.edu

### MS20

#### Data Compression for Checkpoint Optimization

Data movement is a major computational bottleneck in many HPC applications. Future HPC systems are being designed with limited bandwidth to the parallel file system, remaining at about 1 TB/s. Thus, reducing the volume of data needed for transfer to the parallel file system could mitigate the stagnate growth in file system bandwidth. One approach is through compression; however, traditional lossless compression techniques do not significantly reduce the size of HPC floating-point data. Lossy compression, on the other hand, is proving to be a viable tool for effectively reducing the size of checkpoint and visualization files transferred to the parallel file system by allowing small controllable amounts of error into the data. Yet, a critical aspect of utilizing lossy compression in practice is control of the resulting error. In this talk, selection of lossy compression error tolerances based on program on knowledge of the truncation error of the problem is discussed. The efficacy of this method on model problems is presented along with two production level HPC codes: PlasComCM and Nek5000.

Jon Calhoun  
Department of Computer Science  
University of Illinois at Urbana-Champaign  
jccalho2@illinois.edu

Franck Cappello  
ANL  
cappello@mcs.anl.gov

Luke Olson  
Department of Computer Science  
University of Illinois at Urbana-Champaign  
lukeo@illinois.edu

Marc Snir  
University of Illinois at Urbana-Champaign  
snir@illinois.edu

### MS20

#### Using Error Estimates In Resiliency

Data corruption may arise from a wide variety of sources from aging hardware to ionizing radiation, and the risk of corruption increases with the computation scale. Corruptions may create failures, when execution crashes; or they may be silent, when the corruption remains undetected. We studied solutions to failures and silent data corruptions for numerical integration solvers, which are particularly

sensitive to corruptions. Numerical integration solvers are step-by-step methods that approximate the solution of a differential equation. Corruptions are not only propagated all along the resolution, but the solution could even diverge. In numerical integration solvers, approximation error can be estimated at a low cost. We use these error estimates for two applications in fault tolerance. Concerning silent data corruptions, we demonstrated a new lightweight detector for solvers. We rigorously showed that all corruptions affecting the accuracy of a simulation are detected by our method. Concerning failures, the classic checkpointing-restart mechanism can be a bottleneck because of data movement costs. Lossy compression constitutes a promising solution, because it reduces IO bandwidth needs. However, it is unclear what level of error from lossy compression is acceptable for a give application. Using the error estimates determined above from integration solvers is one possible way to guarantee that the compression loss is no worse than the error of the solver.

Pierre-Louis Guhur

École normale supérieure de Cachan  
pierre-louis.guhur@ens-cachan.fr

Franck Cappello  
ANL  
cappello@mcs.anl.gov

Emil M. Constantinescu  
Argonne National Laboratory  
Mathematics and Computer Science Division  
emconsta@mcs.anl.gov

## MS20

### Soft Error Resilience in Big Data Kernels: An Algorithmic Approach

The shrinking processor feature and operating voltages of processor circuits are making them increasingly vulnerable to soft faults, which calls for fault resilience techniques at both the software and hardware levels under the big data context. In this talk, we propose a low cost algorithmic approach to significantly improve soft error resilience for Big Data kernels.

Lu Peng

Department of Electrical and Computer Engineering  
Louisiana State University  
lpeng@lsu.edu

Sui Chen  
Department of Electrical and Computer Engineering  
Louisiana  
csui1@lsu.edu

Travis LeCompte, Walker Legrand  
Department of Electrical and Computer Engineering  
Louisiana State University  
tlecom3@lsu.edu, wlegra2@lsu.edu

## MS20

### A Self-Correcting Connected Components Algorithm

We present an algorithmic design principle that we refer to as *self-correction*. In this approach, we consider the possible states of the algorithm and divides them into *valid* and *invalid* states. A valid state is one from which the algo-

gorithm will converge to a correct result assuming a fault-free execution. By contrast, from an invalid state the algorithm instead produces an incorrect result or diverges. A hardware fault may bring the algorithm to an invalid state. A *self-correcting algorithm* detects that it is in an invalid state, and using the previously known valid state, brings itself back to *some* valid state, and thus ensuring eventual convergence. To illustrate the principle, we present a self-correcting algorithm for the problem of computing the connected components of a graph. Our algorithm derives from a highly parallel but non-resilient algorithm, which is based on the technique of *label propagation* (LP). Our self-correcting algorithm, Ft-LP, has relatively small storage and computation overheads: in empirical tests on a variety of input graphs, we observe execution time overheads of 10-35% in Ft-LP compared to LP even at high fault rates, with the computation overhead increasing gracefully as fault rates increase.

Piyush Sao, Richard Vuduc  
Georgia Institute of Technology  
piyush3@gatech.edu, richie@cc.gatech.edu

## MS21

### Challenges in Managing Scientific Workflows in Heterogeneous Environments

Today, complex scientific computations are often structured as workflows. To support the execution of these workflows, new knowledge must be gained in how to deliver high-performance and distributed computing resources to the scientists desktop in an accessible, reliable and scalable way. This talk will describe the challenges arising when managing scientific workflows in such distributed environments. These issues include data management, resource provisioning, computation scheduling, and provenance tracking. Additional considerations include performance optimization and reliability.

Ewa Deelman

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

## MS21

### Towards Executing Dynamic and Adaptive Application Workflows on Heterogeneous Infrastructure Using Pilot-Abstractions

Scientific application workflows are comprised of heterogeneous tasks that are constructed dynamically and vary in the resources required. In addition to efficient execution there is also the need to adapt the (tasks of the) workflows, e.g. because of fluctuating resource availability. We present an abstraction-based and pilot-based software stack that enables the efficient execution of workflows on diverse resources. In choosing the right abstractions for bi-directional information traversal this also enables application workflow adaptivity.

Mark Santcroos, Shantenu Jha  
Department of Electrical and Computer Engineering  
Rutgers University  
mark.santcroos@rutgers.edu, shantenu.jha@rutgers.edu

## MS21

### The Uintah - VisIt Coupled Workflow

We present recent work with the Uintah framework that is

part of the developer and application scientist workflow, which is composed of three layers. The first, the runtime layer, computer scientists are tasked with developing and maintaining simulation controllers to parallelization schemes and require workflow tools that allow them to visualize and analyze changes in the memory to the mpi scheduling. The second, the application layer, domain scientists are tasked with developing and maintaining the physics to the meshing schemes and require workflow tools that allow them to visualize and analyze changes in the solvers to the mesh elements. Finally, the third, the user layer, requires workflow tools that allow users to visualize and analyze not only everything previously described but also parameters changes to the simulation results. To accomplish this task we have coupled Uintah's framework with the VisIt framework to allow in-situ interactive parallel visualization and analysis. For the runtime and application layers some of this coupling is performed through a series of lightweight structures giving access to both runtime and application settings not normally available except as hard to examine text files. At the same time Uintah's DataWarehouse and other components are completely exposed to the user for visualization and analysis. All together this work forms an interactive runtime simulation dashboard that has been incorporated into the Uintah workflow.

Allen Sanderson  
University of Utah  
allen@sci.utah.edu

Alan Humphrey  
Scientific Computing and Imaging Institute  
University of Utah  
ahumphrey@sci.utah.edu

John A. Schmidt  
SCI Institute  
University of Utah  
jas@sci.utah.edu

Martin Berzins  
Scientific Computing and Imaging Institute  
University of Utah  
mb@sci.utah.edu

## MS21

### Implementing Multi-Scale Simulation Workflows with COMPS

COMPS is a coarse grain, task-based programming model aiming at the implementation of dynamic computational workflows. COMPS supports the execution of workflows in clusters, private, public and hybrid clouds, including platforms such as the Chameleon cloud. The talk will present recent work on the extensions of COMPS to support MPI simulations as tasks and the features to deal with efficient multi-scale simulations, which have been designed having in mind the Human Brain Project requirements.

Raül Sirvent  
Barcelona Supercomputing Center, Spain  
raul.sirvent@bsc.es

Rosa M. Badia  
Barcelona Supercomputing Center / CSIC

rosa.m.badia@bsc.es

## MS22

### Model Error in CO2 Retrieval from Satellite Measurements

NASA's OCO-2 mission collects space based measurements of atmospheric CO<sub>2</sub>. Data are collected with high spatial and temporal resolution and the data product includes both an estimate of column averaged CO<sub>2</sub> dry air mole fraction ( $X_{CO_2}$ ) and an estimate of uncertainty. The OCO-2 instrument measures reflected sunlight and uses a physical model and Bayes Theorem to estimate  $X_{CO_2}$ . However, computational shortcuts are taken to obtain an estimate of the posterior mode ( $\hat{X}$ ) and posterior variance ( $\hat{S}$ ). Even though the forward model is not linear, users usually treat the posterior distribution as Gaussian with mean  $\hat{X}$  and variance  $\hat{S}$ . Also, uncertainty due to several uncertain parameter inputs and model discrepancy is not taken into account. A UQ group within the OCO-2 mission has developed a testbed, where a surrogate model (simplified, but physically realistic) can be used to study various aspects of the retrieval. In this talk we will discuss a few testbed experiments, such as full exploration of the posterior via MCMC, effects of model discrepancy and parameter uncertainty.

Jenny Brynjarsdottir  
Case Western Reserve University  
Dept. of Mathematics, Applied Mathematics and Statistics  
jxb628@case.edu

Jonathan Hobbs, Amy Braverman  
Jet Propulsion Laboratory  
California Institute of Technology  
jonathan.m.hobbs@jpl.nasa.gov,  
amy.j.braverman@jpl.nasa.gov

## MS22

### Bayesian Calibration of Simulators with Structured Discretization Uncertainty

When computational constraints prohibit model evaluation at all but a small ensemble of parameter settings, a dimension-reduced emulator of the system can be constructed and interrogated at arbitrary parameter regimes. The choice of the emulator model is a critical aspect of calibration. Existing approaches treat the mathematical model as an unknown, deterministic response surface. However, in many cases the underlying mathematical model, or the simulator approximating the mathematical model, are stochastic. In this paper, we propose a Bayesian statistical calibration model for stochastic simulators. The approach is motivated by two applied problems: a deterministic mathematical model of intra-cellular signaling whose implementation on computer nonetheless has discretization uncertainty, and a stochastic model of river water temperature commonly used in hydrology.

Matthew T. Pratola, Oksana A. Chkrebtii  
The Ohio State University  
mpratola@stat.osu.edu, oksana@stat.osu.edu

## MS22

### Model Error and Statistical Calibration of Physical

## Models

The predictive fidelity of computations of physical systems is often limited due to structural errors in model specification. Increasingly, there is interest in, and ongoing work on, the modeling of model error, and its estimation in the process of model calibration. This talk will address the statistical representation of model error in computational models, and its estimation in a Bayesian inference framework. The context of physical system modeling motivates a particular strategy of embedding model error within the model, where key approximations are present. We will illustrate the construction in both homogeneous chemical ignition and large eddy simulation of turbulent flow.

Habib N. Najm  
Sandia National Laboratories  
Livermore, CA, USA  
hnnajm@sandia.gov

Khachik Sargsyan, Xun Huan  
Sandia National Laboratories  
ksargsy@sandia.gov, xhuan@sandia.gov

## MS22

### Model Discrepancy and Influenza Forecasting: A Bayesian Hierarchical Modeling Approach

Reduced order (i.e., simplified or approximate) models (ROMs) can permit a useful description of a process when complete knowledge of the system is unavailable. In some regimes, the lack of system knowledge will introduce systematic discrepancies between fitted ROM output and observations of the process. Discrepancy functions and measurement error are often introduced to explain these differences, where the ROM/observation residuals are partitioned into systematic (discrepancy) and random (measurement error) components. From a probabilistic modeling standpoint, specifying a joint distribution for the ROM, discrepancy function, and measurement error is challenging due to complex dependencies between various model components. We propose a Bayesian hierarchical model, where complex dependencies are specified through tractable conditional distributions. Information about the discrepancy conditioned on the ROM can readily be incorporated into the Bayesian hierarchical modeling framework. We illustrate these ideas in the context of flu forecasting in the United States.

David Osthus, James Gattiker  
Los Alamos National Laboratory  
dosthus@lanl.gov, gatt@lanl.gov

## MS23

### Measures of Residual Risk with Connections to Regression, Risk Tracking, Surrogate Models, and Ambiguity

Standard deviation, mean-squared error, and regression are integral parts of even the most rudimentary data analysis. Decision making based on utility theory and risk, in the Markowitzian sense of mean-plus-standard deviation, is equally widely adopted. In this presentation, we describe far-reaching extensions of these concepts that enable alternative approaches to risk mitigation and preference-driven data analysis. We give fundamental connections between regression and decision making, which lead to the construction of measures of residual risk. These measures quantify the improved situation faced by a hedging investor

compared to that of a single-asset investor, but the notion reaches further with relations to forecasting, learning, and regression. Relying on convex analysis, we establish properties of broad classes of measures of error, deviation, regret, risk, and residual risk. These measures can play central roles in the development of risk-tuned approximations of random variables, in tracking of statistics, and in estimation of the risk of conditional random variables. We illustrate the framework by develop a risk-based sparsity-inducing approach to surrogate models in high-dimensional dynamical systems as well as by learning of computationally costly simulation output from inexpensive simulations in the context of an ultra-high speed navy vessel.

Johannes O. Royset  
Operations Research Department  
Naval Postgraduate School  
joroyset@nps.edu

## MS23

### Time Consistency of Multistage Stochastic Programs

In this talk we discuss time consistency of multistage stochastic programs with relation to the modern theory of risk measures. Basically two main approaches were considered in the recent literature. In one approach optimality at every stage is defined by an objective function given by a conditional risk measure satisfying some basic assumptions. The other approach deals with optimal policies in a direct way and is closely related to Bellman's principle of optimality. One of the goals of this talk is to clarify a relation between these two approaches to time consistency.

Alexander Shapiro  
Georgia Institute of Technology  
School of Industrial and Systems Engineering  
ashapiro@isye.gatech.edu

## MS23

### Risk-Averse PDE-Constrained Optimization

The modelling of uncertainty or ambiguity in practical applications leads to models in engineering and the natural sciences that involve partial differential equations (PDEs) with uncertain parameters. Passing from modelling and simulation to optimization, we are forced to consider PDE-constrained optimization problems under uncertainty. This then requires us to employ and extend ideas from stochastic programming and risk management to handle optimization problems with distributed parameters and decision variables. In particular, we employ risk measures to handle the stochasticity of the objective functional and control-to-state mapping. This leads to non-smooth, stochastic PDE-constrained optimization problems. After establishing reasonable conditions that allow us prove the existence of a solution and derive first-order optimality conditions, we discuss smoothing techniques (including their asymptotic behavior). We then demonstrate the effects of various risk measures on the overall cost and design of the decision via numerical examples.

Thomas M. Surowiec  
Department of Mathematics  
Humboldt University of Berlin  
thomas.surowiec@gmail.com

Drew P. Kouri  
Optimization and Uncertainty Quantification

Sandia National Laboratories  
dpkouri@sandia.gov

### MS23

#### Surrogate Models For Characterizing Tail Statistics Of Systems Governed By PDEs

Performance of physical and engineering systems is often described by statistics of distribution tails such as quantile and probability of exceedance. However, these functions have poor mathematical properties for distributions described by sample statistics. Therefore, it is very difficult to optimize such quantities and build systems with the optimal characteristics. The recent trend is to use alternative characteristics, serving similar purposes: Conditional Value-at-Risk (CVaR) and Buffered Probability of Exceedance (bPOE). CVaR and bPOE can be optimized with convex and linear programming algorithms. We build surrogates for CVaR and bPOE of systems described by Partial Differential Equations (PDEs) with stochastic parameters. The surrogates are estimated with linear regression algorithms.

Stan Uryasev, Giorgi Pertaia  
University of Florida  
uryasev@ufl.edu, gpertaia@ufl.edu

### MS24

#### Derivative-Free Constrained Stochastic Optimization of a Scramjet, Using SNOWPAC

We present the derivative-free optimization tool (S)NOWPAC for derivative-free nonlinearly constrained stochastic optimization. Our target applications are computationally expensive robust optimization tasks where sample approximations for risk and deviation measures define the objective function as well as the constraints. Our proposed optimization procedure combines fully linear local surrogate models with global Gaussian processes, to mitigate the noise, within a trust region framework. Using the CUTEst optimization benchmark problem set, we demonstrate that (S)NOWPAC exhibits both fast pre-asymptotic descent and highly accurate results as the number of optimization steps increases. The fast initial descent and the noise reduction due to the Gaussian process models allow for an efficient optimization of otherwise computationally expensive supersonic propulsion system simulations. We close the talk with results on the optimization of air-fuel mixing within a scramjet under uncertain inflow conditions and nonlinear constraints on scalar dissipation to achieve uniform combustion.

Friedrich Menhorn  
Department of Informatics, Technische Universität München  
menhorn@in.tum.de

Florian Augustin, Youssef M. Marzouk  
Massachusetts Institute of Technology  
fmaugust@mit.edu, ymarz@mit.edu

### MS24

#### Numerical Simulation of Stochastic Navier-Stokes Equations Under High-Dimensional Random Forcing

In this numerical study the effect of high-dimensional stochastic forcing in time-dependent flows is investigated.

To efficiently quantify the evolution of stochasticity in such a system, the dynamically orthogonal method is used. In this methodology, the solution is approximated by a *generalized* Karhunen-Loeve (KL) expansion in the form of  $\mathbf{u}(\mathbf{x}, t; \omega) = \bar{\mathbf{u}}(\mathbf{x}, t) + \sum_{i=1}^N \mathbf{y}_i(t; \omega) \mathbf{u}_i(\mathbf{x}, t)$ , in which  $\bar{\mathbf{u}}(\mathbf{x}, t)$  is the stochastic mean, the set of  $\mathbf{u}_i(\mathbf{x}, t)$ 's is a deterministic orthogonal basis and  $\mathbf{y}_i(t; \omega)$ 's are the stochastic coefficients. Explicit evolution equations for  $\bar{\mathbf{u}}$ ,  $\mathbf{u}_i$  and  $\mathbf{y}_i$  are formulated. The elements of the basis  $\mathbf{u}_i(\mathbf{x}, t)$ 's remain orthogonal for all times and they evolve according to the system dynamics to capture the energetically dominant stochastic subspace. We consider two classical fluid dynamics problems: (1) flow over a cylinder, and (2) flow over an airfoil under up to one-hundred dimensional random forcing. We explore the interaction of intrinsic with extrinsic stochasticity in these flows. The energy cascades and correlation between stochastic energy levels in the statistical sense are also analyzed.

Hessam Babaei  
MIT  
babaei@mit.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

### MS24

#### From Data to Optimization Under Uncertainty: A Scalable Framework for Bayesian Inversion and Optimal Control of Random PDEs with Application to Turbulent Flows

We develop a framework for quantifying uncertainties across the pipeline of modeling—from observational data, to Bayesian inversion, to optimal control—for forward models given by PDEs augmented with inadequacy models. We instantiate this framework in the context of modeling incompressible turbulent flow via an algebraic turbulence model. Inadequacy of the turbulence model is represented by an inadequacy field within the turbulent viscosity term that is governed by an advection-diffusion-reaction PDE with random diffusion field. The inverse problem is to infer parameters in the turbulence model (including mean and variance of the random diffusion field) by calibrating the turbulence model with DNS data. The optimal control problem is to determine values of control variables that minimize the expected value of a quantity of interest, for random parameters and random diffusion field with uncertainty given by Bayesian calibration. Executing this pipeline with conventional UQ methods is intractable due to the infinite dimensional (large scale after discretization) random field that characterizes the inadequacy model, and the expense of solving the (highly nonlinear) forward turbulence model. To make the pipeline tractable, we invoke quadratic approximations of the diffusion field-to-observable and diffusion field-to-control objective maps, in conjunction with fast randomized linear algebra methods. We present results for a turbulent jet problem.

Peng Chen  
UT Austin  
peng@ices.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

Robert D. Moser  
University of Texas at Austin  
rmoser@ices.utexas.edu

Todd A. Oliver  
PECOS/ICES, The University of Texas at Austin  
oliver@ices.utexas.edu

Umberto Villa  
University of Texas at Austin  
uvilla@ices.utexas.edu

## MS24

### Model Error Quantification in Turbulent Combustion Computations

Quantification of model error, i.e. uncertainty associated with modeling assumptions, remains one of the most challenging aspects of uncertainty quantification and predictive simulation. We develop a novel strategy for model error quantification by directly embedding a discrepancy representation into the parameters of the model-to-be-calibrated. The embedded structure is particularly advantageous for predictive science and engineering applications: it enables physically meaningful predictions of quantities of interest (QoIs) by automatically inheriting physical laws and constraints imposed from the model, and provides an intuitive platform for “extrapolating” the model error for predicting QoIs outside those used for calibration. We characterize and propagate the model error and parameter uncertainties under a Bayesian framework. The calibration problem is solved using approximate likelihood constructions, adaptive Markov chain Monte Carlo, and polynomial chaos expansions. The overall method is demonstrated on large eddy simulation (LES) computations of turbulent flow in a Scramjet engine.

Khachik Sargsyan, Xun Huan, Zachary Vane, Guilhem Lacaze, Joseph C. Oefelein  
Sandia National Laboratories  
ksargsy@sandia.gov, xhuan@sandia.gov,  
zvane@sandia.gov, gnlacaz@sandia.gov,  
oefelei@sandia.gov

Habib N. Najm  
Sandia National Laboratories  
Livermore, CA, USA  
hnnajm@sandia.gov

## MS25

### Using Data Assimilation to Reconstruct Cardiac Electrical Dynamics

Numerical techniques have predicted that reentrant electrical scroll waves underlie many cardiac arrhythmias, but experimental limitations have hampered a detailed understanding of the specific mechanisms responsible for reentrant wave formation and breakup. To further this effort, we recently have begun to apply the technique of data assimilation, widely used in weather forecasting, to reconstruct time series in cardiac tissue. Here we use model-generated surrogate observations from a numerical experiment to evaluate the performance of the ensemble Kalman filter in reconstructing such time series for a discordant alternans state in one spatial dimension and for scroll waves in three dimensions. We show that our approach is able to recover time series of both observed and unobserved variables that match the truth. Where nearby observations are

available, the error is reduced below the synthetic observation error, with a smaller reduction with increased distance from observations. Using one-dimensional cases, we provide a deeper analysis showing that limitations in model formulation, including incorrect parameter values and undescribed spatial heterogeneity, can be managed appropriately and that some parameter values can be estimated directly as part of the data assimilation process. Our findings demonstrate that state reconstruction for spatiotemporally complex cardiac electrical dynamics is possible and has the potential for successful application to real experimental data.

Elizabeth M. Cherry  
Rochester Institute of Technology  
School of Mathematical Sciences  
excsm@rit.edu

Matthew J. Hoffman  
Rochester Institute of Technology  
mjhsma@rit.edu

Nicholas LaVigne  
Cornell University  
nsl42@cornell.edu

Nathan Holt, Darby Cairns  
Rochester Institute of Technology  
nxh7119@mail.rit.edu, dic4597@g.rit.edu

## MS25

### GPU Simulations of Heart Dynamics Close to Real Time and Over the Internet

Modeling cardiac dynamics and arrhythmias is a challenging problem. In particular simulations of complex mathematical models in 2 and 3 dimensions required until now, supercomputing power. However recent advances in graphic cards and its programming, allows for the first time to study the dynamics of complex models not only in 2D tissues but in realistic 3D heart structures at close to real time and using a single desktop computer. In this talk we show how electrical heart dynamics can be simulated and visualized using graphic cards, via a web-browser in 2 and 3D heart structures, with out the need of compilation and independent of the machine and operating system. This opens the door to the possibility of high performance computing of patient specific cardiac modeling without the need of big computer clusters. We will demonstrate real time simulations of the two most recent and complex cell models (TenTusscher et al and Ohara et al) in tissue and simulations on 3D ventricles and atria structures from a laptop computer.

Flavio H. Fenton, Abouzar Kaboudian  
Georgia Institute of Technology  
flavio.fenton@physics.gatech.edu,  
abouzar.kaboudian@gmail.com

## MS25

### Cardiac Electrophysiology: The Clinical Perspective (The Greatest Gaps in Heart Rhythm Analysis: A Clinical Cardiac Electrophysiologists Perspective on Atrial Fibrillation)

Since the application of Hodgkin and Huxleys model to heart rhythm conduction, invasive treatment of arrhythmias has grown dramatically. Atrial fibrillation (AF), the irregular contraction of the top chambers of the heart, is



the best example in our field of a treatment for an arrhythmia being designed and employed long before its mechanism is understood. Unfortunately, this treatment is with limited efficacy and poor efficiency. Our presentation will develop the following three areas of greatest need for clinicians to treat AF: 1. A large gap exists between mathematical models of AF and customized models for each patient. Several easily obtainable clinical variables, including anatomy, dominant frequency, and refractory periods may be used to help individualize these models and enhance their utility. 2. Better ways of interpreting frequency and power spectra in hopes of prognosticating ablative success are needed for AF. 3. Generating reliable mathematical models of AF that would allow for feedback control algorithms to design pace-termination or minimal lesion set strategies. Ablation of AF involves destruction of atrial tissue in hopes of creating lines of electrical block to eliminate propagation of fibrillatory wave fronts. As a result, lengthy procedures that destroy large areas of cardiac tissue are used to eliminate the arrhythmia. These three areas of need will aid clinicians in refining the treatment of this most common arrhythmia.

Michael Lloyd  
Emory University Hospital  
Atlanta, GA  
mlloyd2@emory.edu

Iravanian Shahriar, Jonathan Langberg  
Emory University Hospital  
irvanian@yahoo.com, jlangbe@emory.edu

#### MS25

##### **An Approach to Estimating the Probability of Calcium-Induced Arrhythmias in Cardiac Cells and Tissues and Their Dependence on System Parameters**

Abstract not available

Raimond Winslow  
Johns Hopkins University  
rwinslow@bme.jhu.edu

#### MS26

##### **Fast and Adaptive Partitioned Methods for Unsteady Thermal Fluid Structure Interaction**

We consider unsteady interaction between compressible fluids and structures, where heat is exchanged across a non-moving interface. Our prime examples are gas quenching in steel forging and the launch phase in a rocket nozzle. For these, we describe the models needed and challenges put on numerical methods. As a basic paradigm, partitioned methods are employed that reuse existing solvers and let these exchange boundary data at the coupling interface. In this framework, we present a fast partitioned time adaptive second order diagonally implicit Runge-Kutta method. We use a Dirichlet-Neumann method for the coupling and a linear extrapolation in time to get good starting values. This method is able to solve real life problems for engineering accuracies using only dozens of coupling iterations. This is possible due to the time adaptive method needing only few time steps, as well as the extremely fast convergence speed of the Dirichlet-Neumann iteration. Therefore, we present an analysis of the convergence speed of the Dirichlet-Neumann iteration for a fully discretized model problem, namely two coupled linear heat equations. As discretizations, we consider combinations of Finite Element

and Finite Volume methods for the subproblems. In 1D, we obtain exact formulas for the convergence rates and in 2D, good approximations.

Philipp Birken  
Centre for the Mathematical Sciences  
Lund University, Sweden  
philipp.birken@na.lu.se

Azahar Monge Sanchez  
Lund University  
Centre for the Mathematical Sciences  
azahar.monge@na.lu.se

#### MS26

##### **Coupling Requirements for Multi-Physics Problems**

The first requirement for obtaining a reliable solution to multi-physics problems is well-posedness. We will investigate well-posedness of the coupling of an incompletely parabolic system with a hyperbolic system with different size in one spatial dimension. The number and type of coupling conditions will be determined by using the so called energy-method, where one multiply the governing partial differential equations by the solution and integrate by parts. Once the coupling conditions are known for the continuous problem, we will discretize using high order finite differences on summation-by-parts form and include the coupling conditions weakly. Next, the adjoint problem will be derived using the energy-method. By using the same technique as in the primal problem, the adjoint problem will be discretized, which will lead to superconvergent functionals. Finally, by considering a physical example, it will be shown that the mathematical interface conditions contain the physically correct interface conditions. Numerical experiments corroborate the theoretical analysis.

Fatemeh Ghasemi  
Linköping University  
fatemeh.ghasemi@liu.se

Jan Nordström  
Department of Mathematics  
Linköping University  
jan.nordstrom@liu.se

#### MS26

##### **Added-Mass Partitioned (AMP) Algorithm for the Deforming Beam and Fluid Interaction**

A new partitioned algorithm for coupling incompressible flows with elastic beams is described that overcomes the added-mass instability for light solids. The algorithm requires no expensive sub-iterations and is fully second-order accurate. The new scheme is shown to be stable, even for very light beams, through the analysis of a model problem. The approach is then applied to the simulation of FSI problems involving beams undergoing large deformations using deforming composite grids.

Longfei Li  
Department of Mathematical Sciences  
Rensselaer Polytechnic Institute  
lil19@rpi.edu

#### MS26

##### **A High-Resolution Method for a Multi-Scale**

### Model of Heterogeneous Explosives

In heterogeneous explosives, sites of high temperature and pressure, called hot spots, play a significant role in detonation initiation. This study seeks to model hot spots generated by the compaction of a granular explosive. Two scales are considered, a macro scale modeled as compressible reactive flow with compaction coupled to a grain scale represented by a reaction-diffusion system. A numerical method and well-resolved results are presented for initiation and propagation of detonation waves.

Donald W. Schwendeman

Rensselaer Polytechnic Institute  
Department of Mathematical Sciences  
schwed@rpi.edu

Ashwani K. Kapila  
Rensselaer Polytechnic Inst.  
kapila@rpi.edu

### MS27

#### Modeling Leader and Follower in Cancer Invasion

As a solid tumor grows, growth conditions deteriorate inside the main cell mass. Pioneer cells then spread out from the tumor, into tissue or via the vascular and lymphatic systems, to other places in the body. These pioneer cells form the basis for new cancer cell colonies, which grow into tumors over time. Recent experiments suggest that two cellular phenotypes enhance tissue invasion through a symbiotic relationship: the highly migratory, non-proliferative, leader cells lead the invasion dragging follower cells in a cellular stream; the follower cells, non-migratory and highly proliferative, benefit from richer nutrients outside the cell pack, while the leader cells benefit by supportive chemicals secreted by the follower. We developed a mathematical model of tumor invasion based on the cellular Potts model. We discovered that a surprisingly small number of features, including biased migration, adhesion, and proliferation, must be present to generate the leader/follower stream invasion from a tumor. We showed that, agreeing with empirical observations: leader cell migration outward from the tumor leading a stream of follower cells; each stream has one leader; each leader would lead a stream; if the leader is disconnected, the stream would retract.

Yi Jiang, Sean Keeler

Georgia State University  
yjjiang12@gsu.edu, seanlkeeler@gmail.com

Jessica Konen, Adam Marcus  
Emory University  
jtape@emory.edu, aimarcus@emory.edu

### MS27

#### Stochastic Dynamics of 3D Vesicles in Viscous Flows

We study the nonlinear, nonlocal dynamics of 3D vesicles in a time-dependent, incompressible viscous flow at finite temperature by using a 3D stochastic immersed boundary method. We first validate our model by comparing with the quasi-spherical theory of a fluctuating vesicle. We then investigate a transient instability that can be observed when the direction of applied flow is suddenly reversed, which induces compressive forces on the vesicle interface, and the development of small-scale interface perturbations known as wrinkles. Our 3D numerical results can present many

more details of the vesicle dynamics, which can not be realized in 2D simulations, or even by the experiments.

Kai Liu

Mathematics  
University of California, Irvine  
liuk10@uci.edu

Shuwang Li

Department of Applied Mathematics  
Illinois Institute of Technology  
sli@math.iit.edu

John Lowengrub

Department of Mathematics  
University of California at Irvine  
lowengrb@math.uci.edu

### MS27

#### Multi-Scale Modeling and Simulation of the Growth of Bacterial Colony with Cell-Cell Mechanical Interactions

The growth of bacterial colony exhibits striking patterns that are determined by the interactions among individual, growing and dividing bacterial cells, and that between cells and the surrounding nutrient and waste. Understanding the principles that underlie such growth has far-reaching consequences in biological and health sciences. In this work, we construct a multi-scale model of the growth of *E. coli* cells on agar surface. Our model consists of detailed, microscopic descriptions of the cell growth, cell division with fluctuations, and cell movement due to the cell-cell and cell-environment mechanical interactions, and macroscopic diffusion equations for the nutrient and waste. Our large-scale simulations reproduce experimentally observed growth scaling laws, strip patterns, and many other features of an *E. coli* colony. This work is the first step toward detailed multi-scale computational modeling of three-dimensional bacterial growth with mechanical and chemical interactions. This is joint work with Dr. Mya Warren, Ms. Yue Yan, Dr. Bo Li, and Dr. Terry Hwa.

Hui Sun

University of California, San Diego  
hsun@mednet.ucla.edu

### MS27

#### Crawling and Turning of Cells in a Minimal Reaction-Diffusion Model

We study a minimal model of a crawling eukaryotic cell with a chemical polarity controlled by a reaction-diffusion mechanism. The size, shape, and speed of the cell emerge from the combination of the chemical polarity, which controls the locations where actin polymerization occurs, and the physical properties of the cell, including its membrane tension. While our model primarily describes a cell that crawls in a straight trajectory, we show that the cell may also turn, and transition to a circular trajectory. We discuss the controlling variables for this turning instability. We argue that the turning arises from a coupling between the reaction-diffusion mechanism and a shape of the cell, and discuss the possible generality of this mechanism.

Yanxiang Zhao

Department of Mathematics  
George Washington University  
yxzhao@email.gwu.edu

Brian Camley  
University of California, San Diego  
bcamley@gmail.com

Bo Li  
Department of Mathematics, UC San Diego  
bli@math.ucsd.edu

Herbert Levine  
Rice University  
herbert.levine@rice.edu

Wouter-Jan Rappel  
Department of Physics  
University of California, San Diego  
rappel@physics.ucsd.edu

## MS28

### Tree-Structured Grids Approaches for Storage, Data Access, and Discretization

The potential of tree-structured grids has been shown long ago for cheap sorting algorithms, efficient geometry representation/computer graphics, and fast summation methods for N-body problems. In the last decade, also more and more partial differential equation solvers used this type of grids due to its advantages in terms of the inherent multi-level hierarchy and the low memory requirements compared to unstructured grids. This talk gives a short overview on sophisticated approaches for data storage, data access, and discretization, in particular

1. the connection between tree-structured grids and space-filling curves,
2. basic requirements of standard finite-element methods, discontinuous Galerkin and N-body problems/fast multipole in terms of data storage, data access, and communication,
3. stream and stack concepts for finite-element discretizations of partial differential equations,
4. flexible data structures with random access to parts of the grid such as ghost layers or cell neighbors,
5. realization choices for dynamical grid adaptivity in terms of data structures (re-allocation versus scattered storage),
6. approaches for the efficient representation of complex geometries in the Cartesian structure of a tree-structured mesh.

Carsten Burstedde  
Universität Bonn  
burstedde@ins.uni-bonn.de

Miriam Mehl  
Universität Stuttgart  
miriam.mehl@ipvs.uni-stuttgart.de

## MS28

### Parallel Forest-of-Trees Algorithms for Tetrahedral and Hexahedral AMR

One motivation of adaptive mesh refinement (AMR) is to increase the accuracy of a numerical solver, while retaining or even decreasing the overall runtime, compared to uniform-mesh simulations. Especially for adaptive meshes that are refined and coarsened frequently during a single

simulation, i.e. every few time steps, partitioning according to space-filling curves (SFC) has been established as a fast and reliable method. Its favorable properties result from an efficient implicit encoding of curve indices and low memory requirements of shared metadata. Recently, SFC have been used most prominently for hexahedral meshes. We introduce a SFC for adaptive triangular and tetrahedral meshes that can be computed using bitwise interleaving operations similar to the well-known Morton curve for cubical meshes. Using this space-filling curve we have established constant-time algorithms (i.e., independent of the refinement level) to compute the parent, children and face-neighbors of a given mesh element, as well as predecessor and successor elements with respect to the SFC. Presently we are working on the AMR library t8code using a forest of octrees approach, implementing parallel algorithms for the efficient handling of the coarse mesh of trees, adaptation, and partitioning. We present our latest results, demonstrating parallel scalability with several hundred thousand MPI ranks.

Johannes Holke, Carsten Burstedde  
Universität Bonn  
holke@ins.uni-bonn.de, burstedde@ins.uni-bonn.de

## MS28

### Fast Integral Equation Solver for Elliptic PDEs in Complex Geometries

We present a new HPC framework for solving variable coefficient elliptic PDEs such as Poisson, Stokes and low-frequency Helmholtz problems on non-regular domains in three-dimensions. This framework is built on top of our Parallel Volume Fast Multipole Method (PVFMM) library. The library implements one of the fastest methods for solving constant coefficient elliptic PDEs on cubic geometries. We will present new numerical algorithms to extend this scheme to variable coefficient problems on complex geometries. Our method uses high-order piecewise polynomial representation on adaptive octrees to discretize the problem. We use Morton ordering to achieve efficient load balancing and scalable performance on several thousand MPI processes. We will discuss applications of our method to acoustic scattering problems and simulation of complex fluids. In particular, we will discuss the simulation of deformable vesicles in a Stokesian fluid. Our method allows us to compute flows in complex geometries, such as those found in microfluidic devices and blood capillaries.

Dhairya Malhotra  
Institute of Computational Engineering and Sciences  
The University of Texas at Austin  
dmalhotra@ices.utexas.edu

George Biros  
The Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
biros@ices.utexas.edu

## MS28

### Using Voxel Meshes on Octrees for High-Fidelity Simulations with Complex Geometries

Meshing complex geometries is a critical challenge for detailed simulations. Instead of employing unstructured meshes with polyhedra, we look at octree meshes with voxels to represent computational domains. Voxelization offers a robust and efficient method to discretize arbitrary geometries. However, voxels only yield a first order repre-

sentation of the surfaces. This is not sufficient for higher order simulations, but with appropriate methods, it is possible to detach the boundary description from the element shape. Obviously, any construction of additional information should maintain the robustness of the voxelization. We look into two numerical methods that fit the voxel mesh approach. The Lattice-Boltzmann method utilizes discrete velocity directions and boundary conditions only in the considered directions. The required line-surface intersections can be computed robustly. As a second method we consider a penalized high-order DG scheme. Here, the geometry is not prescribed as a surface but represented by an embedded object. This embedded boundary information can also be generated accurately by an octree refinement towards the surfaces, combined with a projection. A representation of the geometry within elements, suitable for the numerical method, can then be derived from that voxelization. Thus, the robustness in the mesh generation relying on voxels is maintained even for high-order methods.

Sabine P. Roller, Harald Klimach  
Universitaet Siegen  
sabine.roller@uni-siegen.de, harald.klimach@uni-siegen.de

## MS29

### Flow Transport for Bayesian Inference

In this work, we consider the construction of transport maps between two probability measures using flows. In the Bayesian formalism, this ODE approach is natural when one introduces a curve of probability measures that connects the prior to posterior by tempering the likelihood. We present a novel approximation of the resulting PDE which yields an ODE whose drift is a function of the full conditional distributions of the posterior.

Jeremy Heng  
St. Cross College  
University of Oxford  
jeremy.heng@stx.ox.ac.uk

## MS29

### Warp Bridge Sampling for Faster MCMC and for Multi-Modal Distributions

Warping bridge sampling (Meng and Schilling, 2002, *Journal of Computational and Graphical Statistics*) applies deterministic or stochastic transformations to multiple sets of Monte Carlo draws from different distributions in such a way that the transformed draws have similar distributional shapes between the sets, yet the transformations do not alter the normalizing constants of the original unnormalized densities. This permits more efficient Monte Carlo estimates of the ratios of normalizing constants than the original bridge sampling (Meng and Wong, 1996, *Statistics Sinica*). Our most recent work (Wang and Meng, 2016) proposed a class of Warp-U transformations intended to warp multi-modal distributions into unimodal ones, and we found that this warping strategy can also be effective for generating Monte Carlo draws themselves. This strategy creates effective paths among modes by first applying a stochastic transformation to move multiple modes into a common location and then stochastically re-creates the modes with their intended masses via its inverse transformation. We provide preliminary theoretical and empirical results to demonstrate this new potential of warp bridge sampling. (This is joint work with Lazhi Wang.)

XiaoLi Meng

Harvard University  
meng@stat.harvard.edu

## MS29

### Discrete Monge-Kantorovich Approach for Large-Scale PDE-Constrained Bayesian Inverse Problems

We present an ensemble-based method for mapping prior samples to posterior ones. This method avoids Markov chain simulation and hence discards expensive work when a sample is rejected. The underlying concept is to cast the problem of finding posterior samples into a large-scale linear programming problem for which efficient and scalable solver can be developed. Large-scale numerical results will be presented to demonstrate the capability of the method

Aaron Myers  
University of Texas at Austin - ICES  
aaron@ices.utexas.edu

## MS29

### Bayesian Filtering and Smoothing Via Measure Transport

A recent approach to non-Gaussian Bayesian inference seeks a deterministic transport map that pushes forward a reference density to the posterior. In this talk, we present an approach to nonlinear Bayesian filtering and smoothing that leverages the typical Markov structure of the data assimilation problem to compute transport maps in high-dimensions. In particular, we show how the relevant transport maps admit low-dimensional parameterizations in terms of sparsity and decomposability. We consider applications to chaotic dynamical systems.

Alessio Spantini, Daniele Bigoni, Youssef M. Marzouk  
Massachusetts Institute of Technology  
spantini@mit.edu, dabi@mit.edu, ymarz@mit.edu

## MS31

### Faster Kernel Ridge Regression Using Sketching and Preconditioning

Random feature maps, such as random Fourier features, have recently emerged as a powerful technique for speeding up and scaling the training of kernel-based methods such as kernel ridge regression. However, random feature maps only provide crude approximations to the kernel function, so delivering state-of-the-art results requires the number of random features to be very large. Nevertheless, in some cases, even when the number of random features is driven to be as large as the training size, full recovery of the performance of the exact kernel method is not attained. In order to address this issue, we propose to use random feature maps to form preconditioners to be used in solving kernel ridge regression to high accuracy. We provide theoretical conditions on when this yields an effective preconditioner, and empirically evaluate our method and show it is highly effective for datasets of up to one million training examples.

Haim Avron  
Tel Aviv University  
haimav@post.tau.ac.il

## MS31

### Generalized Hybrid Iterative Methods for Large-

### Scale Bayesian Inverse Problems

We develop a generalized hybrid iterative approach for computing solutions to large-scale Bayesian inverse problems. We consider a hybrid algorithm based on the generalized Golub-Kahan bidiagonalization for computing Tikhonov regularized solutions to problems where explicit computation of the square root and inverse of the covariance kernel for the prior covariance matrix is not feasible. This is useful for large-scale problems where covariance kernels are defined on irregular grids or are only available via matrix-vector multiplication, e.g., those from the Matérn class. We show that iterates are equivalent to LSQR iterates applied to a directly regularized Tikhonov problem, after a transformation of variables, and we provide connections to a generalized singular value decomposition filtered solution. Our approach shares many benefits of standard hybrid methods such as avoiding semi-convergence and automatically estimating the regularization parameter. Numerical examples from image processing demonstrate the effectiveness of the described approaches.

Julianne Chung  
Department of Mathematics  
Virginia Tech  
jmchung@vt.edu

Arvind Saibaba  
North Carolina State University  
asaibab@ncsu.edu

### MS31

#### Randomized Matrix-Free Trace and Log-Determinant Estimators

We present randomized algorithms for estimating the trace and determinant of Hermitian positive semi-definite matrices. The algorithms are based on subspace iteration, and access the matrix only through matrix vector products. We analyse the error due to randomization, for starting guesses whose elements are Gaussian or Rademacher random variables. The analysis is cleanly separated into a structural (deterministic) part followed by a probabilistic part. Our absolute bounds for the expectation and concentration of the estimators are non-asymptotic and informative even for matrices of low dimension. For the trace estimators, we also present asymptotic bounds on the number of samples (columns of the starting guess) required to achieve a user-specified relative error. Numerical experiments illustrate the performance of the estimators and the tightness of the bounds on low-dimensional matrices; and on a challenging application in uncertainty quantification arising from Bayesian optimal experimental design.

Ilse Ipsen  
North Carolina State University  
Department of Mathematics  
ipsen@ncsu.edu

Arvind Saibaba  
North Carolina State University  
asaibab@ncsu.edu

Alen Alexanderian  
University of Texas at Austin  
alen@ices.utexas.edu

### MS31

#### Subspace Driven Data Reduction Strategies for

### Linear Bayesian Inverse Problems

Many inverse problems involve a very large number of observations. These data often are not equally informative, however, and solving the inverse problem with all of the data can be computationally prohibitive. We introduce strategies for selecting subsets of the data that yield accurate approximations of the inverse solution, i.e., the posterior distribution. Our strategies exploit the structure of inverse problems in the Bayesian setting, and are inspired by related work on optimal low-rank approximations to the posterior distribution—extended now to the observation space and coupled with heuristics derived from rigorous error estimates.

Jayanth Jagalur Mohan, Youssef M. Marzouk  
Massachusetts Institute of Technology  
jagalur@mit.edu, ymarz@mit.edu

### MS32

#### Ensemble Grouping Strategies for Embedded Stochastic Collocation Methods

In this work we propose an efficient method for the solution of a stochastic elliptic partial differential equation (PDE) with an uncertain diffusion parameter featuring high anisotropy and irregular behavior. In the context of stochastic collocation finite element methods we use a local hierarchical polynomial approximation with respect to the uncertain parameters in a high-dimensional sample space. The deterministic PDEs are solved by using a recently developed technique that allows us to solve groups (ensembles) of samples very efficiently in an embedded fashion. Since the grouping strongly affects the spectral properties of the ensemble finite element matrices and may compromise the efficiency of the algorithm, we explore different sample-grouping techniques, based on properties of the uncertain parameter and on a surrogate model for the number of linear solver iterations. We compare their performance on a 3-dimensional anisotropic diffusion problem for different dimensions of the parameter space.

Marta D'Elia, Mohamed S. Ebeida  
Sandia National Laboratories  
mdelia@sandia.gov, msebeid@sandia.gov

Eric Phipps  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

Ahmad A. Rushdi  
University of California, Davis  
aarushdi@ucdavis.edu

### MS32

#### Techniques for Reducing Computational Complexity of Sparse Grid Stochastic Collocation Methods

Sparse grid stochastic collocation (SC) methods are a valuable tool for solving problems in uncertainty quantification, yet they suffer from a dramatic increase in costs in high-dimensions. In this talk, we apply SC methods to solve partial differential equations (PDEs) with random coefficients, and exploit multilevel and hierarchical structure in the spatial and stochastic approximation schemes to drastically improve the computation efficiency of the method. We provide a thorough analysis of the savings, and present numerical examples of our methods for both linear and

non-linear random PDEs.

Peter Jantsch  
University of Tennessee  
jantsch@math.utk.edu

Clayton G. Webster  
Oak Ridge National Laboratory  
webstercg@ornl.gov

### MS32

#### Computing Gradients in Local, Adaptive Voronoi Piecewise Surrogates

The Voronoi Piecewise Surrogates (VPS) approach is a powerful method for building surrogate models in uncertainty quantification calculations. It naturally incorporates derivative information supplied by the forward simulation within the surrogate construction, allowing accurate surrogate models to be computed using fewer samples. In this talk we discuss computing gradient information in large-scale models for use within VPS. We show how to compute these derivatives efficiently using techniques based on automatic differentiation. In particular, we focus on derivative evaluation in PDE codes implemented on emerging manycore architectures. We then compare the VPS approach with gradients to other sampling-based approaches on representative PDE problems implemented on emerging computational architectures.

Eric Phipps  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

Marta D'Elia, Mohamed S. Ebeida, H. Carter Edwards  
Sandia National Laboratories  
mdelia@sandia.gov, msebeid@sandia.gov,  
hcedwar@sandia.gov

Ahmad A. Rushdi  
University of California, Davis  
aarushdi@ucdavis.edu

### MS32

#### An Ensemble Grouping Framework Based on Voronoi Piecewise Surrogate Models

Quantifying uncertainties in computational simulations is often done in ensembles (groups) of samples in order to reduce the overall computational cost. However, grouping samples is not trivial and can be assisted by predictive global surrogate models. In this work, we introduce an adaptive ensemble grouping framework that uses different surrogate models to guide its adaptation. We employ Voronoi Piecewise Surrogates (VPS) models which does not dictate where to sample, adding a few design degrees of freedom. To complement accuracy and efficiency, we construct three surrogate models at each iteration: two for the quantity of interest (QoI) and one for the number of iterations needed by the PDE solver. Our adaptive sampling strategies aim to improve accuracy, maintain efficiency, and guarantee space filling when possible. In specific, we MC sample the two QoI surrogates to identify islands with the largest differences, select some samples within these islands that are close in their predicted number of solver iterations using the surrogate of iterations, and finally pick from an ensemble that is well-spaced while satisfying the accuracy improvement and iterations consistency requirements. We

demonstrate the application of our algorithm to different analytic smooth and discontinuous functions as well as real PDE problems.

Ahmad A. Rushdi  
University of California, Davis  
aarushdi@ucdavis.edu

Mohamed S. Ebeida, Marta D'Elia  
Sandia National Laboratories  
msebeid@sandia.gov, mdelia@sandia.gov

Eric Phipps  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

### MS33

#### Preconditioning Linear Least-Squares Problems by Identifying a Basis Matrix

We show how to identify a good basis matrix of the rectangular matrix  $A$  from the linear least-squares problem

$$\|b - Ax\|_2$$

by using an LU factorization of  $A$  followed by an iterative procedure to improve the basis. We use a separate LU factorization of this basis matrix to precondition an iterative scheme for the solution of the augmented system

$$\begin{array}{cc} I & A \\ A^T & 0 \end{array}$$

. We also discuss the case when the matrix  $A$  does not have full column rank.

Mario Arioli  
University of Toulouse  
mario.arioli@gmail.com

Iain Duff  
Science & Technology Facilities Council, UK  
and CERFACS, Toulouse, France  
iain.duff@stfc.ac.uk

### MS33

#### Using LU Factorization to Accelerate Convergence of Iterative Least Squares Solvers

We study how to use an LU factorization as a right-preconditioner of standard iterative methods (e.g., lsqr or lsqr) for solving overdetermined sparse least squares problems. Usually  $L$  is much better conditioned than  $A$  so iterating with lower trapezoidal  $L$  instead of  $A$  gives a faster convergence. In previous work, we showed that when  $L$  is not sufficiently well-conditioned, we could use a partial orthogonalization of  $L$  to accelerate convergence. In this talk, we illustrate that, alternatively to partial orthogonalization, convergence can be improved by iterating with  $LL_1^{-1}$  where  $L_1$  is the square upper triangular part of  $L$ , resulting in a cheaper algorithm. Mixed precision techniques can reduce execution time and required storage.

Gary W. Howell  
North Carolina State University  
gary\_howell@ncsu.edu

Marc Baboulin

University of Paris-Sud  
marc.baboulin@lri.fr

### MS33

#### Inner-Iteration Preconditioning for the Minimum-Norm Solutions to Least Squares Problem

Consider solving linear least squares problems for the minimum-norm solution (pseudoinverse solution). Standard iterative solutions of the problems are the CGLS, LSQR, and LSMR methods with no preconditioning. If combined with preconditioning, these methods may determine a wrong solution, i.e., the minimum-norm solution weighted with the preconditioning matrix. This can be overcome by using a two-step procedure: the first step is to solve a least squares problem; the second step is to solve a linear system of equations. For the solution of the first and second steps, we propose using the left-preconditioned and right-preconditioned generalized minimal residual (GMRES) method for least squares problems (AB- and BA-GMRES). They are combined with inner-iteration preconditioning. We give convergence conditions for proposed methods. Numerical experiments on benchmark problems show that the proposed method is more robust than previous methods.

Keiichi Morikuni

Department of Computer Science  
University of Tsukuba  
morikuni@cs.tsukuba.ac.jp

### MS33

#### LU Preconditioning for Singular Sparse Least Squares

For overdetermined systems  $Ax \approx b$  with full column rank, LU factors of  $A$  with  $L$  well-conditioned give  $U$  as a reasonable right-preconditioner, as implied by Peters and Wilkinson. When  $A$  has low column rank,  $U$  is singular. We show that it still provides a helpful preconditioner for computing the minimum-length solution  $x$  for the original system. We experiment with LUSOL on a range of sparse singular problems.

Michael A. Saunders

Systems Optimization Laboratory (SOL)  
Dept of Management Sci and Eng, Stanford  
saunders@stanford.edu

Jennifer Scott  
Rutherford Appleton Laboratory  
England  
J.A.Scott@rl.ac.uk

Nick Gould  
Numerical Analysis Group  
Rutherford Appleton Laboratory  
nick.gould@stfc.ac.uk

### MS34

#### A New Iterative Scheme Based on DG Methods for (Parallel) Time Integration

We present a new class of iterative methods for solving initial value problems (IVP) based on discontinuous Galerkin (DG) methods. We start from the weak DG formulation arising from non-linear equations and derive the new iterative methods by using different iterative schemes. Based

on our new approach, we can systematically construct explicit, implicit and semi-implicit schemes with arbitrary order of accuracy. We also show that the same schemes can be constructed by solving a series of correction equations based on the DG weak formulation. The accuracy of the schemes is proven to be  $\min\{2p + 1, K\}$  with  $p$  the degree of the DG polynomial basis and  $K$  the number of iterations. The stability has been explored numerically, showing that the implicit schemes are  $A$ -stable and the explicit schemes are competitive with existing methods. Furthermore, we combine this technique with a multi-level strategy to accelerate its convergence speed. Also, the new multi-level scheme is intended to provide a flexible framework for high order space-time discretizations and to be coupled with space-time multi-grid techniques for solving partial differential equations (PDEs). Besides its standard applications as time integrator, the newly proposed method, due to its structure, is a competitive and promising candidate for parallel in time algorithms such as Parareal, PFASST, multigrid in time, etc.

Xiaozhou Li  
Institute of Computational Science  
USI, Lugano  
xiaozhou.li@usi.ch

Pietro Benedusi, Rolf Krause  
Institute of Computational Science  
USI Lugano  
pietro.benedusi@usi.ch, rolf.krause@usi.ch

### MS34

#### Nonlinear Domain Decomposition Methods for Large Scale Problems

Abstract not available

Rolf Krause  
Institute of Computational Science  
USI Lugano  
rolf.krause@usi.ch

### MS34

#### Exascale Solvers in UG - Multigrid for HPC

Abstract not available

Gabriel Wittum  
G-CSC  
University of Frankfurt, Germany  
wittum@techsim.org

### MS35

#### How Do Electrons Move in Space? Flux Discretizations for Non-Boltzmann Statistics

When modelling semiconductor devices via the van Roosbroeck system one often uses statistical functions to describe the correspondence between carrier densities and chemical potentials. For 3D bulk semiconductors the most general choice is given by the Fermi-Dirac integral of order  $1/2$ . However, how to numerically solve the van Roosbroeck in this general (non-Boltzmann) case is still an open problem. We will present and compare several flux discretization schemes which generalize the well-known Scharfetter-Gummel scheme. Our main goal is to discretely preserve important properties from the continuous system such as existence and uniqueness of the solution, consistency with

the thermodynamical equilibrium as well as unconditional stability. We also show how these new numerical schemes can be efficiently implemented for 2D and 3D applications.

Patricio Farrell, Thomas Koprucki, J Fuhrmann  
Weierstrass Institute for Applied Analysis and Stochastics  
farrell@wias-berlin.de, koprucki@wias-berlin.de,  
fuhrmann@wias-berlin.de

### MS35

#### **NEMO5: A Parallelized Multi-Scale and Multi-Physics Nanodevices Simulation Software**

NEMO5 is a multi-scale and multi-physics nanodevice simulation software developed in the group of Prof. Klimeck and used by more than 400 groups in academia and industries, including Intel, TSMC, Samsung, etc. It successfully bridges both challenges discussed in this minisymposium. Quantum transport simulations are essential for the design of state-of-the-art nanodevices. The non-equilibrium Greens function (NEGF) method is among the most widely employed methods to describe carrier dynamics in open quantum systems. Unfortunately, the basic NEGF equations are complex, mathematically cumbersome, and their numerical solution is extremely computationally demanding. A well-known method to ease the numerical burden is the recursive Greens function method (RGF) that allows to limit the calculation and storage of the retarded Greens function to specific matrix blocks. NEMO5 uses PETSc to solve linear and non-linear systems of equations in parallel. Moreover, NEMO5 includes a general parallelization class that handles the setup of multilevel parallelization hierarchies and the distribution of problems onto the available processes. The energy  $E$  and momentum  $k$  parameters of the quantum transport equations are parallelized with MPI. NEMO5 includes both blocking and non-blocking communication methods to optimize non-trivial MPI communication in quantum transport, thus to ensure NEMO5s scalability to hundreds of thousands CPU cores.

Xinchen Guo, Daniel Lemus, Daniel Mejia  
Purdue University  
xg@purdue.edu, dlemus@purdue.edu, denphi@denphi.com

Jim Fonseca  
Network for Computational Nanotechnology  
Purdue University  
jfonseca@gmail.com

Gerhard Klimeck  
Purdue University  
gekco@purdue.edu

Tillmann Kubis  
Network for Computational Nanotechnology  
Purdue University  
tkubis@purdue.edu

### MS35

#### **Numerical Algorithms Based on Galerkin Methods for the Modeling of Reactive Interfaces in Photoelectrochemical Solar Cells**

This work concerns the numerical solution of a coupled system of self-consistent reaction-drift-diffusion-Poisson equations that describes the macroscopic dynamics of charge transport in photoelectrochemical (PEC) solar cells with reactive semiconductor and electrolyte interfaces. We

present numerical algorithms, mainly based on a mixed finite element and a local discontinuous Galerkin method for spatial discretization and three implicit-explicit time stepping techniques, for solving the time-dependent nonlinear systems of partial differential equations. We perform computational simulations under various model parameters to demonstrate the performance of the proposed numerical algorithms as well as the impact of these parameters on the solution to the model.

Michael D. Harmon  
Institute for Computational and Engineering Sciences  
mharmon@ices.utexas.edu

Kui Ren  
University of Texas at Austin  
ren@math.utexas.edu

Irene M. Gamba  
Department of Mathematics and ICES  
University of Texas  
gamba@math.utexas.edu

### MS35

#### **Modeling and Simulations Using Time-Dependent Density Functional Theory**

Time-dependent density functional theory can be used to model electron dynamics in molecules and nanostructures from first principles. We discuss our work developing a highly parallel all-electron finite element code for performing real-time electronic spectroscopy simulations without the use of a pseudopotential. Parallelism is introduced through domain-decomposition techniques and interfacing with a distributed-memory version of the FEAST eigenvalue solver and a custom Schur complement solver. Results will be presented for carbon nanotubes molecules comprised of a few hundred atoms.

James Kestyn  
ECE Department  
University of Massachusetts, Amherst, USA  
jkestyn@umass.edu

Eric Polizzi  
University of Massachusetts, Amherst, USA  
polizzi@ecs.umass.edu

### MS36

#### **Abnormal Synchrony in Evolving Brain Networks**

This talk addresses a fundamental question of how pathological synchronized rhythms, associated with epileptic seizures and Parkinsons tremors, appear in brain networks as a result of complex spatial and temporal dynamics of the brain network. There is experimental evidence showing that evolving brain networks change their functional structure during epileptic seizures from a more regular to a more random structure. In this talk, we will discuss the role of evolving network structure and the switching between healthy and abnormal rhythms that is accompanied by a change in intrinsic dynamics of neurons.

Igor Belykh  
Department of Mathematics and Statistics  
Georgia State University  
ibelykh@gsu.edu

Reimbay Reimbayev, Kevin Daley



Department of Mathematics, Georgia State University  
khorezmian@gmail.com, kmd4669201609@gmail.com

### MS36

#### Real-Time In-Situ Seismic Imaging with Sensor Networks

The seismic imaging process today involves massive seismic data collection from hundreds and thousands of seismic sensors to a central place for post computing. The whole process is expensive and often takes days even months to complete. There is great demand for real-time as it would reduce the costs and risks of E&P and mitigate the environment concerns. This talk presents an innovative Real-time In-situ Seismic Imaging (RISI) system that computes and visualizes the 3D subsurface imaging in seconds. The RISI system is a mesh network that sense and process seismic signals, and compute 3D subsurface image in-situ in real-time. Instead of data collection then post processing, the mesh network performs the distributed data processing and inversion computing under the severe bandwidth and resource constraints, and generates an evolving 3D subsurface image as more events arrive. Several innovative distributed seismic imaging algorithms have been successfully developed and validated using both synthetic and real-world seismic data set. The hardware prototype system has also been implemented and can be extended as a general field instrumentation platform, to incorporate new geophysical data processing and computing algorithms, beyond seismic.

WenZhan Song  
University of Georgia, USA  
wsong@uga.edu

### MS36

#### Enhanced Community Detection in Multilayer and Temporal Networks through Layer Aggregation

Inspired by real-world networks consisting of layers that encode different types of connections, such as a social network at different instances in time, we study community structure in multilayer networks. We analyze fundamental limitations on the detectability of communities by developing random matrix theory for the dominant eigenvectors of modularity matrices that encode an aggregation of network layers. Aggregation is often beneficial when the layers are correlated, and it represents a crucial step for the discretization of time-varying network data, whereby layers are binned into time windows. We explore two methods for aggregation: summing the layers adjacency matrices as well as thresholding this summation at some value. We analyze detectability phase transitions that are onset by varying either the density of within-community edges or community size. We identify layer-aggregation strategies that are optimal in that they minimize the detectability limit. Our results indicate good practices in the context of community detection for how to aggregate network layers, threshold pairwise-interaction data matrices, and discretize time-varying network data. We apply these results to synthetic and empirical networks, including a study of anomaly detection for the Enron email corpus.

Dane Taylor  
Department of Mathematics  
University of North Carolina at Chapel Hill  
dane.r.taylor@gmail.com

Rajmonda Caceres

Lincoln Laboratory  
Massachusetts Institute of Technology  
rajmonda.caceres@ll.mit.edu

Peter J. Mucha  
University of North Carolina Chapel Hill  
much@unc.edu

### MS37

#### Non-Smoothness in Space and Time Within Model Reduction

It is well-known that model reduction techniques for parameterized systems perform particularly well for problems whose solution depends smoothly on the involved parameters. In this talk, we consider the Reduced Basis Method (RBM) for instationary problems involving various sources of non-smoothness in space, time and parameter. In particular, we are concerned with the Hamilton Jacobi Bellman (HJB) equation, the wave equation and time-dependent obstacle problems. In all these cases, non-smooth effects may vary (travel) in space over time and the evolution is typically unknown, but of great importance. For the corresponding analysis, we use the space-time variational formulation, in which time is treated as an additional variable within the variational formulation of the problem. In special cases, certain discretizations of such space-time problems yield time-stepping schemes. Otherwise, recent tensor product solvers may be used. We discuss recent results in this framework concerning stability, approximation and reduced basis methods including efficient and reliable a posteriori error estimates.

Silke Glas  
University of Ulm, Germany  
silke.glas@uni-ulm.de

Karsten Urban  
Institute of Numerical Mathematics, University of Ulm  
karsten.urban@uni-ulm.de

### MS37

#### About Empirical Interpolation Methods and Applications

Abstract not available

Yvon Maday  
Universite Pierre et Marie Curie  
and Brown university  
maday@ann.jussieu.fr

### MS37

#### Error Control As a Key Ingredient Towards Optimal Numerical Methods

A mathematical key ingredient to achieve "optimal" numerical methods is error control via rigorous a posteriori error estimates. Such error estimates can not only be used to certify approximate solutions, but rather are the essential building block in the construction of problem adapted optimal solution spaces and related adaptive numerical methods. Examples of such optimal methods are e.g. particularly tuned mesh-adaptive finite element schemes for the approximation of PDEs or reduced basis methods (weak greedy algorithms) for the approximation of parameterized PDEs. In this talk we will address true error control for localized reduced basis methods [Ohlberger, M., Schindler,

F.: Error control for the localized reduced basis multiscale method with adaptive on-line enrichment. SIAM Journal on Scientific Computing **37**(6), A2865–A2895 (2015). ,Ohlberger, M., Rave, S., Schindler, F.: True error control for the localized reduced basis method for parabolic problems. Submitted to: Proceedings of MoRePaS III (2015)] and its usage for the construction of efficient numerical schemes for parameterized single and multiscale PDEs. Thereby we overcome the classical paradigm of offline-online splitting by allowing for local basis enrichment in the so called online-phase. Several numerical examples and applications will be shown to demonstrate the efficiency of the resulting adaptive approaches.

Mario Ohlberger  
University of Muenster  
Applied Mathematics Muenster  
mario.ohlberger@uni-muenster.de

Stephan Rave  
University of Muenster  
stephan.rave@uni-muenster.de

Felix Schindler  
University of Muenster  
Applied Mathematics Muenster  
felix.schindler@wwu.de

### MS37

#### Certified Reduced Basis Methods for 4D-Var Data Assimilation

In this talk, we present a certified reduced basis (RB) approach to four dimensional variational data assimilation (4D-VAR). We consider the particular case in which the behaviour of the system is modelled by a parametrised parabolic partial differential equation where the initial condition and model parameters (e.g., material or geometric properties) are unknown, and where the model itself may be imperfect. We consider (i) the standard strong-constraint 4D-Var approach, which uses the given observational data to estimate the unknown initial condition of the model, and (ii) the weak-constraint 4D-Var formulation, which additionally provides an estimate for the model error, and thus can deal with imperfect models. Since the model error is a distributed function in both space and time, the 4D-Var formulation leads to a large-scale optimization problem that must be solved for every given parameter instance. We introduce reduced basis spaces for the state, adjoint, initial condition, and model error. We then build upon recent results on RB methods for optimal control problems in order to derive a posteriori error estimates for RB approximations to solutions of the 4D-VAR problem. Numerical tests are conducted to verify the validity of the proposed approach.

Karen Veroy-Grepl  
Graduate School AICES  
RWTH Aachen University  
veroy@ices.rwth-aachen.de

Martin Grepl, Mark Kaercher  
RWTH Aachen University  
grepl@igpm.rwth-aachen.de, kaercher@ices.rwth-aachen.de

### MS38

#### Fast Low-Rank Approximation of a Matrix and Ex-

### tensions

Low-rank approximations is among most fundamental subjects of numerical linear algebra, with a wide range of applications to a variety of highly important areas of modern computing, which range from the machine learning theory and neural networks to data mining and analysis. We first dramatically accelerate computation of such approximations for the average input matrix, then show some narrow classes of hard inputs for our algorithms, and finally narrow such classes further by means of non-costly preprocessing with structured quasi Gaussian multipliers. Our extensive numerical tests with benchmark input matrices representing discretized PDEs consistently produce reasonably close low-rank approximations at a low computational cost. We conclude with novel extension of our acceleration to the Conjugate Gradient Algorithms.

Victor Pan  
CUNY  
Lehman  
victor.pan@lehman.cuny.edu

### MS38

#### Effective and Robust Preconditioning of General SPD Matrices with Rank Structures

We show how to use rank structured matrices to precondition general symmetric positive definite matrices with guaranteed effectiveness and robustness. A fundamental idea on the construction of the preconditioner is given. We show how to properly drop certain off-diagonal singular values to yield superior effectiveness. This is then generalized to a practical multilevel scheme. The accuracy and condition number control are given.

Jianlin Xia  
Purdue University  
xiaj@math.purdue.edu

Zixing Xin  
Department of Mathematics  
Purdue University  
zxin@purdue.edu

### MS38

#### A Fast Contour-Integral Eigensolver and the Approximation Accuracy

We discuss fast eigenvalue solutions of rank structured matrices via an accelerated contour-integral method. We justify the feasibility of fast eigenvalue count via low-accuracy matrix approximations. Linear systems with multiple shifts and multiple right-hand sides are solved quickly with a shifted factorization update strategy. We then study how to approximate the eigenvalues of some structured matrices as well as more general cases. The approximation accuracy is closely related to the rank structured compression, which gives us flexibility to control the efficiency and effectiveness. In particular, we also show that a low-accuracy structured approximation can already give satisfactory estimation of the eigenvalues.

Xin Ye, Jianlin Xia  
Purdue University  
ye83@purdue.edu, xiaj@math.purdue.edu

Raymond H. Chan  
The Chinese Univ of Hong Kong

Department of Mathematics  
rchan@math.cuhk.edu.hk

### MS38

#### Compute-Memory Tradeoff in Hierarchical Low-Rank Approximation Methods

Hierarchical low-rank approximation methods consist of a wide range of methods, from the purely algebraic  $H/H^2/HSS/HODLR$  matrices to the purely analytical fast multipole methods. Between these two extremes there exist a variety of methods that trade off computation with memory by taking semi-analytical forms, exploiting translational/rotational symmetry, using sampling and randomization. The analytical kernels have high arithmetic intensity and can extract the full potential of GPUs and Xeon Phi. Since the cost of data movement is increasing faster than arithmetic operations on future architectures, the methods that compute more to store/move less will become advantageous. Therefore, it is important to consider the whole spectrum of hierarchical low-rank approximation methods, and choose the appropriate method for a given pair of application and architecture. In this study, we take a performance oriented view of hierarchical low-rank approximation methods, and quantify the trade-offs between the algebraic and analytical variants on the latest hardware.

Rio Yokota  
Tokyo Institute of Technology  
rioyokota@gsc.titech.ac.jp

David E. Keyes  
KAUST  
david.keyes@kaust.edu.sa

### MS39

#### The Multi-Level Monte Carlo Method for Simulations of Turbulent Flows

In this paper the application of the multi-level Monte Carlo (MLMC) method on numerical simulations of turbulent flows with uncertain parameters is investigated. Several strategies for setting up the MLMC method are presented, and the advantages and disadvantages of each strategy are also discussed. A numerical experiment is carried out using the Antarctic Circumpolar Current (ACC) with uncertain, small-scale bottom topographic features. It is demonstrated that, unlike the pointwise solutions, the averaged volume transports are correlated across grid resolutions, and the MLMC method could increase simulation efficiency without losing accuracy in uncertainty assessment.

Qingshan Chen  
Clemson University  
qsc@clemson.edu

Ming Ju  
Beijing Computational Science Research Center  
jming@csr.ac.cn

### MS39

#### Multifidelity Monte Carlo Methods with Optimally-Adapted Surrogate Models

The multifidelity Monte Carlo method combines a high-fidelity model with multiple cheap surrogate models to accelerate the estimation of statistics of the outputs of the

high-fidelity model. The core of the multifidelity Monte Carlo method is an optimal model management that determines how often each of the models has to be evaluated to minimize the mean-squared error of the estimator for a given computational budget. Whereas the multifidelity Monte Carlo method considers the surrogate models as given, which cannot be changed, we allow here the adaptation of the surrogate models. We present a model management that optimally distributes the computational budget among adapting the surrogate models and evaluating the surrogate and the high-fidelity models, such that the mean-squared error of the estimator is minimized. Numerical results confirm that our approach with adaptive surrogate models leads to estimators with lower mean-squared errors than the multifidelity Monte Carlo method with static surrogate models.

Benjamin Peherstorfer  
ACDL, Department of Aeronautics & Astronautics  
Massachusetts Institute of Technology  
pehersto@mit.edu

Karen E. Willcox  
Massachusetts Institute of Technology  
kwillcox@MIT.EDU

### MS39

#### Acceleration of the Multilevel Monte Carlo Method for Certain Classes of Differential Systems

Multilevel Monte Carlo (MLMC) method is a variance reduction technique for the efficient estimation of expected values. It combines ideas of multigrid discretization and Monte Carlo sampling to achieve the optimal complexity of the estimator. We elaborate on this idea and study acceleration of the MLMC method in application to two classes of differential systems. Firstly, we consider systems of stiff jump-diffusion stochastic differential equations (SDEs). For a given combination of pathwise solvers, we present the a posteriori technique for the optimal selection of discretization levels with the optimal assignment of the solvers to each level. We also present the family of implicit two-stage split-step methods which allows exploiting all the levels of the MLMC method without the need to explicitly resolve the fastest scale of the system. Secondly, we study linear partial differential equations in domains with boundaries comprised of both deterministic and random parts and apply the method of modified potentials with kernels given by the Greens functions defined on the deterministic part of the domain. Within this approach, the size of the original differential problem is reduced by reformulating it as the boundary integral equation posed on the random part of the boundary only. MLMC method is then applied to this modified integral equation. We perform complexity analysis of the proposed acceleration techniques and supplement the obtained results with numerical examples.

Viktor Reshniak, Abdul Khaliq, Yuri Melnikov  
Middle Tennessee State University  
vr2m@mtmail.mtsu.edu, abdul.khaliq@mtsu.edu,  
yuri.melnikov@mtsu.edu

### MS39

#### Reduced Basis Methods and Their Application to Ensemble Methods for the Navier Stokes Equations

The definition of partial differential equation (PDE) models usually involves a set of parameters whose values may

vary over a wide range. The solution of even a single set of parameter values may be quite expensive. In many cases, e.g., optimization, control, uncertainty quantification, and other settings, solutions are needed for many sets of parameter values. We consider the case of the time-dependent Navier-Stokes equations for which a recently developed ensemble-based method allows for the efficient determination of the multiple solutions corresponding to many parameter sets. The method uses the average of the multiple solutions at any time step to define a linear set of equations that determines the solutions at the next time step. To significantly further reduce the costs of determining multiple solutions of the Navier-Stokes equations, we incorporate a proper orthogonal decomposition (POD) reduced-order model into the ensemble-based method.

Michael Schneier  
Florida State University  
mhs13c@my.fsu.edu

**MS40**  
**Global Stability of Fully Coupled Capsule Flow Systems**

We develop a formulation to assess the stability of regular capsule-trains advecting steadily in confined environments, such as red blood cells flowing in small capillaries. Empirical evidence shows that when less confined, this regular state breaks down into a more complex and irregular flow. This break-up is studied directly for a model system that includes full coupling between the viscous fluid flow and elastic capsule membranes. It is constructed via an orthogonal set of small disturbances to the system based on a boundary integral formulation. A set of linearly amplifying disturbances are identified and analyzed, including transient and asymptotic long-time behaviors. These include shape distortions and rigid-body-like translations and tilts of the capsules for the set of characteristic base states considered. Further, it is shown that transient disturbances can precipitate a nonlinear mechanism and bypass slower asymptotic growth. The observed apparent stability of highly confined flows is shown to arise from nonlinear interactions. Direct numerical simulations are used to verify the construction of the governing linear system and track the transition into an apparently chaotic nonlinear regime.

Spencer H. Bryngelson  
University of Illinois  
bryngel2@illinois.edu

Jonathan B. Freund  
University of Illinois at Urbana-Champaign  
jbfreund@illinois.edu

**MS40**  
**Calculation of Apparent Viscosity Using Boundary Integral Equations**

The rheological properties of particle suspension are of interest in a wide range of fields, from biology to material science. Boundary integral equations (BIE) are considered an efficient means of modelling the movement of particles in a Stokesian fluid. In this talk we present a technique to predict the apparent viscosity of a planar suspension using BIE. This method results in a well conditioned linear system and lends itself nicely to simple preconditioners. We will present results from a simulation of a Couette apparatus and provide sensitivity studies on the particle size and

volume fraction.

Lukas Bystricky  
Florida State University  
lb13f@my.fsu.edu

Bryan D. Quaife  
Scientific Computing  
Florida State University  
bquaife@fsu.edu

**MS40**  
**Hydrodynamic and Macromolecules Induced Clusters of Red Blood Cells in Microcapillary Flow**

Abstract not available

Chaouqi Misbah  
National Center for Scientific Research  
Universite Joseph Fourier Grenoble I  
chaouqi.misbah@ujf-grenoble.fr

**MS40**  
**A Second-Kind Traction Integral Equation for Wall-Bounded Viscous Flows, with Applications to Sedimentation and Viscous Erosion**

Motivated by problems in sedimentation and erosion, we derive a new boundary integral formulation for computing surface tractions in several Stokes flow problems, in particular for problems with a non-trivial background flow and/or a no-slip plane wall. The integral operator enjoys the conditioning advantages of second kind integral equations while avoiding the traditional obstacles of hypersingularity and rank deficiency. The formulation is used to illuminate the fine details of particle-wall interactions in sedimentation which lead either to glancing or to reversing trajectories, and to explore viscous erosion of bodies in a selection of fundamental background flows, where we observe the emergence of distinct limiting body shapes involving sharp corners and ridges.

William Mitchell  
University of Wisconsin  
whmitchell@math.wisc.edu.

Saverio E. Spagnolie  
University of Wisconsin  
Department of Mathematics  
Spagnolie@math.wisc.edu

**MS41**  
**A Bi-Fidelity, Low-Rank Approximation Technique for Uncertainty Quantification**

The use of model reduction has become widespread as a means to reduce computational cost for uncertainty quantification of PDE systems. In this work we present a model reduction technique that exploits the low-rank structure of the solution of interest, when it exists, for fast propagation of high-dimensional uncertainties. To construct this low-rank approximation, the proposed method utilizes models with lower fidelities (hence cheaper to simulate) than the intended high-fidelity model. After obtaining realizations to the lower fidelity models, a reduced basis and an interpolation rule are identified and applied to a small set of high-fidelity realizations to obtain this low-rank, bi-fidelity approximation. In addition to the construction of this

bi-fidelity approximation, we present convergence analysis and numerical results.

Hillary Fairbanks

Department of Applied Mathematics  
University of Colorado, Boulder  
hillary.fairbanks@colorado.edu

Jerrad Hampton

University of Colorado, Boulder  
jerrad.hampton@colorado.edu

Alireza Doostan

Department of Aerospace Engineering Sciences  
University of Colorado, Boulder  
Alireza.Doostan@Colorado.EDU

Akil Narayan

University of Utah  
akil@sci.utah.edu

#### MS41

##### **Prediction Based on the Kennedy-OHagan Calibration Model: Asymptotic Consistency and Other Properties**

Kennedy and OHagan (2001) propose a model for calibrating some unknown parameters in a computer model and estimating the discrepancy between the computer output and physical response. This model is known to have certain identifiability issue. Tuo and Wu (2016) shows that there are examples for which the Kennedy-OHagans method renders unreasonable results in calibration. In spite of its unstable performance in calibration, the Kennedy-OHagans approach has a more robust behavior in predicting the physical response. In this work, we present some theoretical analysis to show the consistency of predictor based on the Kennedy-OHagans calibration model in the context of radial basis functions context.

Rui Tuo

Oak Ridge National Laboratory  
tuorui@amss.ac.cn

Jeff Wu

Georgia Institute of Technology  
jeffwu@isye.gatech.edu

#### MS41

##### **Multifidelity Monte Carlo Methods for Uncertainty Quantification**

Multifidelity uncertainty quantification methods leverage low-fidelity models to obtain computational speedups in solving uncertainty quantification tasks. We use a control variate formulation to create the multifidelity Monte Carlo (MFMC) method. We consider general low-fidelity models with properties that cannot necessarily be well described by a priori rates. The MFMC method has an optimal strategy for allocating evaluations among multiple models, so as to minimize the variance of the control variate estimator.

Karen E. Willcox

Massachusetts Institute of Technology  
kwillcox@MIT.EDU

Benjamin Peherstorfer

ACDL, Department of Aeronautics & Astronautics

Massachusetts Institute of Technology  
pehersto@mit.edu

#### MS41

##### **Stochastic Collocation with Multi-Fidelity Models**

We shall discuss a numerical approach for the stochastic collocation method with multifidelity simulation models. The method combines the computational efficiency of low-fidelity models with the high accuracy of high-fidelity models. We shall illustrate the advantages of the method via a set of more comprehensive benchmark examples including several two-dimensional stochastic PDEs with high-dimensional random parameters. Finally, We suggest that tri-fidelity simulations with a low-fidelity, a medium-fidelity, and a high-fidelity model would be sufficient for most practical problems.

Xueyu Zhu

Scientific Computing and Imaging Institute  
University of Utah  
xzhu@sci.utah.edu

Akil Narayan, Linebarger Erin

University of Utah  
akil@sci.utah.edu, erinline@sci.utah.edu

Dongbin Xiu

Ohio State University  
xiu.16@osu.edu

#### MS42

##### **Accelerated Gradient Methods for Computing the Stationary States of Gross-Pitaevskii Equations**

The aim of this talk is to propose new fast converging algorithms based on preconditioned nonlinear conjugate gradient methods discretized by pseudo-spectral approximation schemes for computing the stationary states of Gross-Pitaevskii equations. Compared with standard methods, the proposed one is faster and is easy to implement. After presenting the scheme, we will discuss its links with some other standard methods. A thorough numerical study will be given to illustrate the behaviour of the method for 1D, 2D and 3D problems involving the rotating Gross-Pitaevskii equation with large nonlinearities and fast rotations, for strongly confining potentials.

Xavier L. Antoine

Institut Elie Cartan Nancy (IECN)  
Université de Lorraine  
xavier.antoine@univ-lorraine.fr

Qinglin Tang

University of Lorraine  
qinglin\_tang@163.com

Antoine Levitt

Inria and Ecole des Ponts ParisTech  
antoine.levitt@inria.fr

#### MS42

##### **High-Order Numerical Schemes for Computing the Dynamics of Nonlinear Schrödinger Equation**

The aim of this talk is to present high-order numerical schemes to simulate the dynamics of systems of nonlinear Schrödinger/Gross-Pitaevskii equations. The usual high

order numerical schemes based on time splitting technique can not be adapted for the computation of solutions to such equations in some situations. We introduce Runge-Kutta type schemes to obtain high order numerical schemes which can be used both for scalar and system of equations. We will consider exponential integrators such as exponential Runge-Kutta methods and Lawson methods, and high-order IMplicit-EXplicit (IMEX). We will show the properties of these numerical schemes and develop a complete numerical study to investigate the properties of the schemes.

Christophe Besse  
University of Toulouse III  
christophe.besse@math.univ-toulouse.fr

#### MS42

##### **Efficient Spectral Methods for Electronic Schrödinger Equation**

Two efficient spectral-element methods, based on Legendre and Laguerre polynomials respectively, are derived for direct approximation of the electronic Schrödinger equation in one spatial dimension. Compared to existing literatures, a spectral-element approach is used to treat the singularity in nucleus-electron Coulomb potential, and with the help of Slater determinant, special basis functions are constructed to obey the antisymmetric property of the fermionic wavefunctions. Numerical tests are presented to show the efficiency and accuracy of the proposed methods.

Yingwei Wang  
Purdue University  
wywshtj@gmail.com

Jie Shen  
Purdue University  
Department of Mathematics  
shen7@purdue.edu

Haijun Yu  
LSEC  
Academy of Mathematics and Systems Science, China  
hyu@lsec.cc.ac.cn

#### MS42

##### **Accurate and Efficient Computation of Nonlocal Potentials Based on Gaussian-Sum Approximation**

We introduce an accurate and efficient method for a class of nonlocal potential evaluations with free boundary condition, including the 3D/2D Coulomb, 2D Poisson and 3D dipolar potentials. Our method is based on a Gaussian-sum approximation of the singular convolution kernel and Taylor expansion of the density. Starting from the convolution formulation, for smooth and fast decaying densities, we make a full use of the Fourier pseudospectral (plane wave) approximation of the density and a separable Gaussian-sum approximation of the kernel in an interval where the singularity (the origin) is excluded. Hence, the potential is separated into a regular integral and a near-field singular correction integral, where the first integral is computed with the Fourier pseudospectral method and the latter singular one can be well resolved utilizing a low-order Taylor expansion of the density. Both evaluations can be accelerated by fast Fourier transforms (FFT). The new method is accurate (14-16 digits), efficient ( $O(N \log N)$  complexity), low in storage, easily adaptable to other different kernels,

applicable for anisotropic densities and highly parallelable.

Yong Zhang  
Wolfgang Pauli Institute, Wien, Austria  
sunny5zhang@gmail.com

#### MS43

##### **Discontinuous Galerkin Methods for Neutrino Radiation-Hydrodynamics**

Neutrino transport is an essential ingredient in models of core-collapse supernovae — the explosive evolutionary end stage of massive stars. Most of the gravitational energy liberated during iron core collapse is balanced by neutrino radiation, and neutrinos play a major role in driving the explosion through neutrino-matter coupling. As neutrinos are semi-transparent to the stellar fluid, neutrino-matter interactions occur under non-equilibrium conditions, and a description rooted in kinetic theory is warranted. Here we discuss our ongoing work to develop neutrino transport methods based on angular moments of the Boltzmann equation. To this end we consider the use of the Discontinuous Galerkin (DG) method, which offers some distinct advantages over more traditional finite-difference and finite-volume methods: (i) high order accuracy on a compact stencil; (ii) correct asymptotic behavior in the diffusion limit; and (iii) provide a natural framework for consistent lepton and four-momentum exchange with the stellar fluid. In this talk, we discuss a semi-implicit DG method for neutrino radiation-hydrodynamics which treats phase-space advection and neutrino-matter interactions with explicit and implicit time stepping methods, respectively.

Eirik Endeve  
Oak Ridge National Laboratory  
endevee@ornl.gov

Anthony Mezzacappa  
University of Tennessee  
mezz@utk.edu

#### MS43

##### **Neutrino Transport in Neutron Star Merger Simulations**

The recent detection by the LIGO collaboration of gravitational waves emitted by merging black holes opened a new way to observe the universe. In the coming years, gravitational wave detectors are expected to also observe the first binary neutron star and black hole-neutron star mergers. In the presence of at least one neutron star, bright electromagnetic signals can provide additional information about the properties and environment of the mergers. Neutron star mergers can also help us understand nuclear interactions, and may be the main production site of many heavy elements. The result of nucleosynthesis in neutron star mergers and the properties of the associated electromagnetic transients are strongly affected by neutrino emission and absorption. In this talk, I will review recent efforts to implement approximate neutrino transport algorithms in the general relativistic hydrodynamics simulations used to model these merger events. I will then discuss in more detail the use of an energy-integrated two-moment scheme in neutron star merger simulations, and show the impact of the choice of neutrino transport algorithm on the observable signals powered by the merger.

Francois V. Foucart  
Lawrence Berkeley National Laboratory

francois.foucart@gmail.com

### MS43

#### Neutrino Radiation-Hydrodynamics Simulations of Neutron Star Mergers

Neutron star mergers result in the dynamical ejection of up to a few percent of a solar mass of neutron rich material, because of tidal interactions and relativistic shocks. As these outflows expand and cool, they realize the conditions necessary for the activation of the so-called *rapid neutron capture process*, or *r-process*, by which heavy nuclei are synthesized. The radioactive decay of by-products of the *r-process* also powers emissions in the optical or near-infrared bands that could be detected in coincidence with the gravitational waves emitted by the binary. However, predicting nucleosynthetic yields and electromagnetic emission properties requires accurate models of the outflows. In particular their composition, which is set by complex neutrino radiation-hydrodynamic processes, is a crucial quantity to study. In this talk I will review the current state of the art in the modeling of the dynamical ejection of matter during neutron star mergers. I will show the results from preliminary studies, with simplified neutrino-transport models, which show the importance of neutrino irradiation of the outflows. Finally, I will describe our work in progress towards the first spectral (multi-group, multi-flavor) general-relativistic neutrino radiation-hydrodynamics simulations of neutron star mergers.

David Radice

Princeton University  
Institute for Advanced Studies  
dradice@ias.edu

### MS43

#### Neutrino Transport in Core Collapse Supernovae

The neutrino mechanism for powering core collapse supernovae (CCSN), where neutrinos emitted from deep in the core of the collapsed star deposit a fraction of their energy behind the stalled CCSN accretion shock, is thought to be the most likely means of producing observed Type II supernovae. Still, there is significant uncertainty whether this mechanism can actually reinvigorate the shock and if it can explain the energetics of observed CCSNe. These events appear to be marginal phenomena, where differences in hydrodynamics, radiation transport methods, initial conditions, resolution, and imposed symmetries can be the difference between a successful and a failed (simulated) explosion. With this in mind, I will discuss general relativistic moment based methods for radiative transport and their application to neutrino transport in supernovae. I will then report on a set of long-term general-relativistic three-dimensional multi-group (energy-dependent) neutrino-radiation hydrodynamics simulations of CCSNe

Luke Roberts

Michigan State University  
robertsl@nscl.msu.edu

### MS44

#### Performance of the Non-Hydrostatic Unified Model of the Atmosphere on the Intel Knights Landing

We present the progress of porting the Non-hydrostatic

Unified Model of the Atmosphere (NUMA) on the newly released Knights Landing (KNL) architecture. NUMA uses high-order continuous and discontinuous Galerkin spatial discretization with high arithmetic intensity per-degree of freedom that makes them suitable for porting to accelerators. We use a new programming language OCCA to get best performance on several many-core accelerators of different architecture. In our previous work, we have shown that NUMA achieves of upto 15x speedup on a K20X GPU relative to a 16-core AMD CPU; using the same code compiled in OpenMP mode we are targeting to achieve similar level performance on the KNL as well. Our goals for the Knights Landing architecture are: a) allowing NUMA to use OpenMP in addition to MPI b) exploit vectorization through judicious use of vector classes (double4 and double8) throughout the code. Using intrinsics for this vector class on the 512-bit wide SIMD units on the KNL we hope to get further speedup. c) take advantage of the high-bandwidth RAM on the Xeon Phi. So far we have good results on the first two aspects, while work is in progress to see benefit from the MCDRAM. In this talk we shall report our findings on the work done and results obtained so far.

Daniel Abdi, Francis X. Giraldo  
Naval Postgraduate School  
dsabdi@nps.edu, fxgiraldo@nps.edu

Lucas Wilcox  
Department of Applied Mathematics  
Naval Postgraduate School  
lwilcox@nps.edu

Tim Warburton  
Virginia Tech  
tcw@vt.edu

David Medina  
Rice University  
dmed256@gmail.com

Andreas Mueller  
Naval Postgraduate School  
amueller@nps.edu

### MS44

#### Performance Optimizations for the SU2 Higher-Order DG-FEM Fluid Solver on the Intel Xeon Phi

In this presentation we will describe a community effort to develop a new higher-order solver for the SU2 framework using a Discontinuous Galerkin Finite Element Method approach (DG-FEM). The solver is generic and enables a variety of discretizations and the solution of different PDEs, but we focus our efforts on a DG-FEM formulation for the solution of the compressible Navier-Stokes Equations (with several variants including the Euler, Reynolds-Averaged NS, Large-Eddy Simulation equations, etc.) on arbitrary unstructured meshes that result in a capability to resolve complex geometries and tackle industrial-strength problems. We also describe our efforts to optimize the performance of this new solver on the Intel Xeon Phi architecture including efforts for element and face loops, time stepping via ADER-style methods, and memory layout, optimization, and communication. The result is a high-performance solver that is ready for use by the open-source community. Our intent is to leverage this new capability for shape op-

timization problems prevalent in the aerospace industry.

Edwin van der Weide  
University of Twente  
e.t.a.vanderweide@utwente.nl

Juan J. Alonso, Thomas Economon  
Department of Aeronautics and Astronautics  
Stanford University  
jjalonso@stanford.edu, economon@stanford.edu

#### MS44

##### **Vectorization and Locality Optimizations for Seismic Imaging Methods Through Automated Code Generation**

The Full Waveform Inversion (FWI) is nowadays the most powerful numerical method for 3D subsurface image reconstruction. Petabytes of seismic data collected by suitable receivers is inverted to create an accurate, high-resolution representation of the subsurface. This technique has primary applications in geophysics, particularly for oil and gas exploration. One of the major problems of FWI is its computational cost, which grows rapidly with the complexity of the physical model. Because of the strict time windows that often characterize subsurface exploration, advancements in inversion algorithms, numerical methods, software and computer architectures are fundamental. In this talk, I will discuss our approach to the portable optimization of finite difference computations, with particular emphasis on FWI. Our framework consists of a stack of domain specific languages and optimising compilers. The mathematical specification of a finite difference method is translated by a compiler, Devito, into C code, applying a sophisticated sequence of transformations. These include standard loop transformations, such as blocking and vectorization, as well as the composition of rewrite operators to reduce the arithmetic intensity of the complex expressions produced by FWI. I will show the impact of these transformations on standard Intel Xeon architectures and the new Intel Knights Landing. I will conclude discussing future challenges and goals of our work.

Fabio Luporini  
Department of Computing  
Imperial College London  
f.luporini12@imperial.ac.uk

Gerard J Gorman  
Department of Earth Science and Engineering  
Imperial College London  
g.gorman@imperial.ac.uk

Paul Kelly  
Imperial College London  
p.kelly@imperial.ac.uk

Michael Lange  
Department of Earth Science and Engineering  
Imperial College London  
michael.lange@imperial.ac.uk

#### MS44

##### **Large-Scale Dynamic Rupture Simulations with SeisSol on KNL Platforms**

SeisSol is a software package for simulating earthquakes, supporting wave propagation simulations in heterogeneous

media as well as kinematic and dynamic rupture source models. We use a code generation approach for tuning the software for different hardware architectures; specifically, we use libxsmm as a back-end in order to generate high performance dense and sparse matrix multiplication kernels. Libxsmm was recently tuned for KNL, but also other changes were necessary such as prefetching and careful data-placement (in order to overcome the 16GB size limitation of MCDRAM). We are going to present node-level performance of SeisSol on KNL in elastic and viscoelastic media as well as the results of scalability studies.

Carsten Uphoff  
Technische Universität München  
uphoff@in.tum.de

Sebastian Rettenberger  
Technische Universität München  
Department of Informatics, Scientific Computing  
rettenbs@in.tum.de

Michael Bader  
Technical University of Munich  
bader@in.tum.de

Alice Gabriel  
Ludwig-Maximilians-Universität München  
alice-agnes.gabriel@geophysik.uni-muenchen.de

#### MS45

##### **Automatic Adjoints of Multimesh Finite Element Discretisations**

Many interesting physics-driven optimisation problems include dynamic domains (i.e. domains that change in time). One can think of the fluid-fluid interaction in a mixing process or the fluid-structure interaction of a rotating turbine or propeller. In this talk, we present an approach for solving such dynamic optimisation problems based on a multi-mesh idea: we allow the computational domain to consist of multiple meshes which can independently move. These meshes are then coupled in the variational formulation, typically through Nitsche terms. This approach is particularly promising in a PDE-constrained optimisation setting, since it allows even large movements of the domain without non-differentiable and computationally expensive remeshing steps. We implemented this approach within the FEniCS and dolfin-adjoint projects and demonstrate that it can be combined with a high-degree of automation and code generation. The user specifies the variational formulation including the coupling terms in the domain-specific language UFL, from which assembly code is automatically generated via the FEniCS form compiler. The associated adjoint equations are derived automatically from the UFL description, and solved for a specific forward state and mesh-movements. We give examples of PDE-constrained optimisation problems with dynamic domains that can now be solved in a few dozen lines of Python code.

Simon W. Funke  
Center for Biomedical Computing  
Simula Research Laboratory  
simon@simula.no

Jørgen Dokken, August Johannson  
Simula Research Laboratory



dokken@simula.no, august@simula.no

#### MS45

##### Automating Optimal Code Generation for High Order Finite Element Methods

Code generation has emerged in the last decade as a key technology in enabling the productive exploitation of advanced numerical methods for complex systems of equations. However, no existing high level code generation systems facilitate the generation of optimal implementations of high order finite element methods. The optimal algorithms for high level finite elements rely on the exploitation of the tensor product structure of the elements to reorder the assembly loop nests and avoid redundant computations. This reordering is known as sum factorisation. Here we present the extension of the Firedrake automated code generation system with the FInAT element evaluation tool to enable the automatic generation of efficient sum factored algorithms.

David Ham

Department of Mathematics  
Imperial College London  
david.ham@imperial.ac.uk

#### MS45

##### Exporting ADOL-C API for Scripting Languages Using SWIG

Scripting languages are gaining acceptance in the scientific community for purposes of prototyping and due to ease of use. In the context of algorithmic differentiation, the development effort continues to be concentrated on traditional compiled languages like Fortran and C/C++, be it source transformation or operator overloading. Although scripting languages may not be quite suited for high performance computing, providing ways to compute derivatives efficiently in these languages is a worthwhile goal. ADOL-C is an operator overloading based C++ library that provides accurate first and higher order derivatives for applications in C++. In recent years an automatic interface generator called SWIG has been developed that uses the C/C++ header files to wrap the API of a library into various scripting languages like Python, R, C-sharp, Perl, Javascript, Go, Octave etc. Although every language has its caveats, the overall process of making the C/C++ API available via SWIG is the same for all scripting languages. After the initial effort required per language, the effort required to maintain the scripting interface in sync with upstream developments in the C/C++ library is minimal. Initial studies using python and R have shown great promise encouraging us to work further and on more language interfaces, e.g. Octave. This talk will give an overview of the interface generation process, the caveats we encountered for various scripting languages, and some numerical results.

Kshitij Kulshreshtha  
Universität Paderborn  
kshitij@math.upb.de

Sri Hari Krishn Narayanan  
Argonne National Laboratory  
snarayan@mcs.anl.gov

#### MS45

##### Algorithmic Differentiation of a CAD System for

##### Shape Optimization

For the gradient-based optimization of a parametric CAD model, shape sensitivities, i.e. the derivatives of surface points with respect to the design parameters, are required. In this talk we present and analyse the algorithmic differentiation of the CAD kernel within OpenCASCADE Technology using the AD tool ADOL-C. First numerical results for the optimization of a U-bend pipe are shown including also a verification of the computed derivatives.

Andrea Walther

Universität Paderborn  
andrea.walther@uni-paderborn.de

#### MS46

##### Marginal then Conditional Sampling for Hierarchical Models

Independent sampling in the linear-Gaussian inverse problem may be performed for less computational cost than regularized inversion, when selection of the regularizing parameter is considered. We demonstrate a sequence of increasingly efficient sampling algorithms: block Gibbs, one-block, and the ‘marginal-then-conditional’ sampler with a fancy method for evaluating the ratio of determinants required for MCMC. This last method scales well with problem and data size, allowing inference over function space models with no approximation (!).

Colin Fox

University of Otago, New Zealand  
fox@physics.otago.ac.nz

Richard A. Norton  
University of Otago  
Department of Physics  
richard.norton@otago.ac.nz

J. Andrés Christen  
CIMAT, Mexico  
jac@cimat.mx

#### MS46

##### A Gaussian Process Prior Approach to Estimating Parameters in Multi-level Stochastic Differential Equations

We study the problem of estimating functional parameters describing multivariate stochastic differential equations (drift and diffusion coefficients). In our Bayesian approach we focus on a general class of Gaussian Process (GP) prior models for functional unknowns, as such priors appear naturally in certain cases. The work is motivated by the need for improved statistical inference tools for such models, that avoid (to the extent possible) discretizations. We illustrate our method with applications from oceanography.

Peter Craigmile  
The Ohio State University  
pfc@stat.osu.edu

Radu Herbei  
Ohio State

herbei@stat.osu.edu

#### MS46

##### **Fast Experimental Designs for LARGE Linear Processes: How to Image Biofilms Using a 3-D Confocal Scanning Laser Microscope**

Much progress has been made in recent years to apply mature results from numerical analysis in order to accelerate iterative samplers that solve common Bayesian linear inverse problems. Explicit sampling algorithms are now available that efficiently solve LARGE problems that have Gaussian noise by utilizing Lanczos, the method of preconditioned conjugate gradients, and Chebyshev polynomials. A combination of these approaches has recently been used to estimate the location and abundances of pathogenic biofilm bacteria attached to surfaces (e.g., medical devices), before and after antimicrobial treatments, from the 3-D movies generated by a confocal scanning laser microscope (CSLM). Because increased spatial resolution in each image of the movie (set by the microscopist) decreases temporal resolution (i.e., it takes the CSLM more time to capture more pixel data), microscopists want to know: How many pixels should be used in each planar xy-dimension? How many z-slices of planar images should be collected in the vertical dimension? How many different fields of view should be considered? How many biofilm samples from the same experiment should be imaged? How many experiments should be run? These questions regarding CSLM experimental design are addressed via accelerated iterative samplers.

Albert E. Parker

Montana State University  
Center for Biofilm Engineering  
parker@math.montana.edu

#### MS46

##### **Regularized Estimation of Likelihood-Informed Subspaces from Samples**

A large scale Bayesian inverse problem has a low effective dimension when the data are informative only over a low-dimensional subspace of the input parameter space. This subspace, called the likelihood-informed subspace (LIS), is an essential ingredient for reducing the complexity of the problem. The efficient detection of the LIS remains computationally challenging since it requires expensive evaluations of the gradient of the forward model. In this talk, we propose a sample-based (black-box) algorithm for the LIS detection which relies on a regularized estimate of the posterior covariance matrix. Numerical examples on linear-Gaussian and non-linear/non-Gaussian problems show the performance of the method.

Olivier Zahm

MIT  
zahmo@mit.edu

Youssef M. Marzouk

Massachusetts Institute of Technology  
ymarz@mit.edu

#### MS47

##### **Accelerating Tensor Contractions by Runtime Code Generation of Small Matrix Multiplications**

#### **on Intel Architecture**

Tensor contractions can be boiled down to small matrix multiplications as their main computational engine. Math libraries or compilers are unlikely to provide the best possible kernel performance for these kernels. To address this issue, we present a library which provides high performance small matrix multiplications targeting all recent x86 vector instruction set extensions up to Intel AVX-512. Our evaluation proves that large speed-ups are possible depending on the CPU and application. These speed-ups are achieved by a combination of several novel technologies. We use a code generator which has a built-in architectural model to create code which runs well without requiring an auto-tuning phase. Since such code is very specialized we leverage just-in-time compilation to only build the required kernel variant at runtime. Besides pure matrix multiplication performance we will also discuss performance impacts if matrix operands should be implicitly transposed.

Alexander Heinecke

Parallel Computing Laboratory  
Intel Corporation, Santa Clara, CA, USA  
alexander.heinecke@intel.com

Greg Henry, Hans Pabst

Intel Corporation  
greg.henry@intel.com, hans.pabst@intel.com

#### MS47

##### **A Distributed Memory Library for Sparse Tensor Functions and Contractions**

Cyclops Tensor Framework (CTF) provides a high-level abstraction for algebra with distributed sparse multidimensional arrays. CTF enables summation and contraction of tensors with user-defined element types and algebraic structures. Tensors may be sparse and/or contain symmetries. The library is general, supporting applications ranging from computational quantum chemistry to graph analysis. CTF automatically parallelizes the tensor data distribution and computation. A cyclic data layout preserves load balance in the presence of symmetry and sparsity. The contraction algorithms are selected dynamically from a large space of variants, including ones that tradeoff memory usage to minimize communication cost. The talk will focus on recent developments in CTF, including sparse matrix multiplication algorithms that use asymptotically less communication than previous approaches.

Edgar Solomonik

ETH Zurich  
solomon2@illinois.edu

#### MS47

##### **Accelerating Tensor Contractions in High-Order FEM with MAGMA Batched**

High-order finite element methods (FEM), among many other computational methods and applications, can use tensor contractions as a fundamental tool to easily express their computations. Consequently, it is possible to standardize and implement what is needed in numerical libraries. Currently, high-performance is difficult to obtain for tensor contractions using existing libraries. However, using that the contractions are many, small, and independent, along with some application specifics, we developed a MAGMA framework to group (batch) the computations and to obtain close to peak performance results. In par-

ticular, we represent contractions as tensor index reordering plus batched matrix-matrix multiplications (GEMMs). This is a key factor to achieve algorithmically many-fold acceleration (vs. not using it) due to possible reuse of data loaded in fast memory. We will present the results and the techniques leading to obtaining high-performance on multicore, GPUs, and Xeon Phi architectures.

Stanimire Tomov

Computer Science Department  
University of Tennessee  
tomov@icl.utk.edu

Azzam Haidar

Department of Electrical Engineering and Computer Science  
University of Tennessee, Knoxville  
haidar@icl.utk.edu

Ahmad Abdelfattah

University of Tennessee, Knoxville  
ahmad@icl.utk.edu

Veselin Dobrev, Ian Karlin

Lawrence Livermore National Laboratory  
dobrev1@llnl.gov, karlin1@llnl.gov

Tzanio V. Kolev

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

Jack J. Dongarra

University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

**MS47**

**Strongly-Scalable Block and Rank-Sparse Tensor Algebra**

The emerging reduced-scaling quantum mechanical methods replace high dimensional tensors with their data-compressed counterparts to reduce storage and computational costs from high-order polynomial to low-order or even linear. Such data compressed representations take many forms, from simple (element or block) sparsity to global rank compressed forms (HOSVD) to some combinations thereof. Here we will highlight our use of compressed tensor representations suitable for large-scale electronic structure computations on molecules with hundreds of atoms, and will discuss the TiledArray software framework that makes scalable computations with such data structures possible.

Edward F. Valeev

Virginia Tech  
efv@vt.edu

**MS48**

**An Overview of the GRINS Multiphysics Framework**

In this talk, we describe the GRINS multiphysics framework. GRINS is built on the libMesh finite element library, particularly the FEMSystem framework. The goal of GRINS is to provide reusable kernels, selectable at runtime, for various aspects of modeling and simulation of

complex multiphysics systems while at the same time providing user extensibility to the problem at hand. One of the salient features of GRINS is the ability to handle automatically the solution of the discrete adjoint problem. We describe the overall structure of the framework and show several examples of the classes of problems currently being studied using GRINS. Finally, we highlight planned developments in facilitating runtime selection of parameters for the use in bayesian inference and uncertainty quantification.

Paul Bauman

Mechanical and Aerospace Engineering  
University at Buffalo, State University of New York  
pbauman@buffalo.edu

**MS48**

**In Situ Data Steering with Provenance Data**

libMesh-sedimentation is an application built upon libMesh to simulate turbidity currents typically found in geological processes. In large-scale scenarios, users commonly need to track the dataflow generation and to analyze strategic simulation data, such as residuals and errors estimates, to monitor and steer simulation runs as much as possible. Considering existing solutions, such as FastBit and SDS framework, they only perform tracking after the execution of the computer simulation. This work presents DfAnalyzer, an instantiation of ARMFUL framework, which enables the Analysis of Raw data from Multiple Files. DfAnalyzer is a solution based on provenance data to extract and relate strategic simulation data for on-line queries. Moreover, we integrate libMesh-sedimentation and ParaView Catalyst to perform in situ data analysis and visualization. Our experiment used 1,040 cores, where users monitor the sediments appearance at runtime and steer simulation based on the solver convergence (i.e., residuals and errors) and visual information on the deposits.

Vitor Silva

Computer Science  
COPPE/Federal University of Rio de Janeiro  
vitor.silva.sousa@gmail.com

José Camata

NACAD - High Performance Computing Center  
COPPE/Federal University of Rio de Janeiro  
camata@nacad.ufrj.br

Marta Mattoso

Federal University of Rio de Janeiro  
marta@cos.ufrj.br

Patrick Valduriez

INRIA  
patrick.valduriez@inria.fr

Alvaro Coutinho

The Federal University of Rio de Janeiro  
Brazil  
alvaro@coc.ufrj.br

**MS48**

**An Overview of the Moose Framework and Physics Modules**

The Multiphysics Object Oriented Simulation Environment (MOOSE) framework is Idaho National Laboratory's

(INL) premier open source modeling and simulation tool. Under development since 2008, and available on GitHub since 2014, the framework provides a flexible, powerful, and analyst-friendly computational environment for the numerical solution of partial differential equations via the finite element method. The MOOSE framework makes extensive use of the facilities provided by the libMesh finite element library, including mesh management, partitioning, and generation, 1, 2, and 3D geometric element types, finite element basis functions, assembly routines, and numerical linear algebra interfaces. While there are no specific physics equations present in the base framework, the MOOSE repository also includes a diverse set of “modules” for various applications, including tensor mechanics, mechanical contact, heat transfer, chemical reactions, phase field, fluid dynamics, porous media flow, and XFEM. These modules, whose development is driven by both INL and the open science community’s simulation requirements, provide a set of valuable building blocks to researchers developing new analysis tools and conducting original research. This talk will briefly summarize and describe the physics modules which are currently available and under development, the basic equations they are designed to solve, and some of the larger-scale applications where they are currently employed.

John W. Peterson, Derek R. Gaston  
Idaho National Laboratory  
jwpeterson@gmail.com, friedmud@gmail.com

Cody J. Permann  
Center for Advanced Modeling and Simulation  
Idaho National Laboratory  
cody.permann@inl.gov

David Andrs, Andrew Slaughter, Brian Alger  
Idaho National Laboratory  
david.andrs@inl.gov, andrew.slaughter@inl.gov,  
brian.alger@inl.gov

Fande Kong  
Department of Computer Science  
University of Colorado Boulder  
fande.kong@colorado.edu

Richard Martineau  
Idaho National Laboratory  
richard.martineau@inl.gov

#### MS48

##### **Rattlesnake: A Moose/libmesh-Based Multiscale Neutronics Application**

Radiation transport studies the motion of neutrons and photons in a stationary medium by solving the linear Boltzmann transport equation (BTE). The BTE differs from standard PDEs by the number of independent variables: three spatial variables, three radiation momentum variables (two angular directions and energy), and time, leading to a large number of DOFs. Rattlesnake is the radiation transport application built on Moose designed for solving multiphysics radiation transport problems. As a Moose application Rattlesnake inherits many of its features from the libmesh FEM library making it markedly different from other transport codes and facilitating its application in multiphysics environments. Rattlesnake features a number of different angular discretization options: SN (collocation), PN (spectral expansion), and diffusion, and uses either continuous (CFEM) or discontinuous FEM (DGFEM)

for the discretization of the spatial variable. For efficient solution of the SN equations, the nonlinear diffusion acceleration (NDA) method is used. NDA uses a modified low order problem, typically diffusion, to accelerate the SN solution. This talk will focus on three features of Rattlesnake show-casing its unique design and capabilities: (1) nonlinear diffusion acceleration of CFEM and DGFEM transport schemes, (2) an inexact transport sweeper on second order meshes, (3) a model refinement study of the C5G7 problem featuring in excess of 40 billion DoF.

Sebastian Schunert  
Idaho National Lab  
sebastian.schunert@inl.gov

Yaqi Wang, Mark DeHart  
Idaho National Laboratory  
yaqi.wang@inl.gov, mark.dehart@inl.gov

#### MS49

##### **Soft Error in PCG: Sensitivity, Numerical Detections and Possible Recoveries**

Abstract not available

Emmanuel Agullo  
INRIA  
emmanuel.agullo@inria.fr

Luc Giraud  
Inria  
luc.giraud@inria.fr

#### MS49

##### **Integrating Algorithm Based Fault Tolerance into a Parallel Multigrid Library**

The vast size of forthcoming HPC systems, coupled with potential technology changes have increased concern about silent data corruption and have led to a corresponding interest in methods to provide resilience to these errors. To demonstrate the feasibility and performance of algorithm-based fault tolerance (ABFT) for complex applications, we have developed a fault-tolerant parallel multi-grid algorithm. We used the Containment Domains runtime to provide a recovery mechanism for our fault-tolerant parallel sparse matrix-matrix and matrix-vector multiplication routines. The interfaces for these routines were designed to integrate easily with the HYPRE and MFEM libraries so that resulting FT variants of these libraries could add resilience to the Smoothed Aggregation Spectral Element AMG (SAAMGE) program with minimal changes to the parent code. Our FT-AMG code requires roughly twice as much memory capacity as the non-resilient code and incurs only small increases in runtime.

Brian Austin  
NERSC  
Lawrence Berkeley National Laboratory  
baustin@lbl.gov

Xiaoye S. Li  
Lawrence Berkeley National Laboratory

xqli@lbl.gov

#### MS49

##### Resilient Dense/Sparse Linear Algebra

Abstract not available

Zizhong Chen

Department of Computer Science and Engineering  
University of California, Riverside  
chen@cs.ucr.edu

#### MS49

##### Programming Constructs for Transparent Silent-Error Mitigation in PDE Solvers

We discuss practical techniques for implementing parallel PDE solvers that can detect and correct silent data corruption while keeping the solver code understandable and maintainable. Silent data corruption is expected to be an increasing concern at large computational scales. Our focus is on previously developed mitigation approaches that use spatial interpolation to replace outlier values resulting from memory bit flips. These approaches can provide resilience to orders of magnitude higher bit-flip rates than standard solvers, with modest runtime overhead. However, it is also important to minimize the additional programmer effort needed to implement and maintain such a robust solver versus a non-robust one, particularly for solver-specific operations such as stencils. We present example implementations that remain as close as possible to the implementation of a standard solver but support swapping-in reusable components that incorporate robustness in memory accesses.

Maher Salloum

Sandia National Laboratories  
Quantum and Digital Systems Research Department  
mnsallo@sandia.gov

Jackson Mayo, Rob Armstrong

Sandia National Laboratories  
jmayo@sandia.gov, rob@sandia.gov

#### MS50

##### Engineering Sciences Workflows for V&V Applications

The Sandia Analysis Workbench (SAW) is a family of desktop applications that was developed to provide an integrated interface to many of these codes, and to improve their ease of use by providing contextual information and tools that simplify and streamline common operations. Using the Workbench, users build models, submit and manage HPC jobs, visualize results, share and store models and results in context with versioning and dependency tracking, and more. Because a diverse array of in-house and commercial tools must be integrated using limited resources, we use a component-based, data-driven approach in which tools are treated generically and customization is handled as much as possible by configuration files. While many of the codes integrated into the Workbench have scripting capabilities, overall workflow has until now been handled manually. Recently, we have recast our architectural components as nodes in an explicit workflow graph. The addition of a robust, configurable workflow engine supports V&V/UQ activities by allowing combinations of Workbench components to be used for repeatable, automated executions of parameterized models, enabling sensitivity analysis, opti-

mization, and other compound operations with minimal user effort. In this talk, we will discuss the SAW architecture and the data-driven, component-based strategies we use for integrating tools, and the challenges we are facing as we scale the system up for exascale problems.

Robert Clay, Ernest Friedman-Hill, Edward Hoffman

Sandia National Laboratories  
robert.the.mad.scientist@gmail.com, ejfried@sandia.gov,  
elhoffm@sandia.gov

#### MS50

##### Workflow Modernization Efforts at Lawrence Livermore National Laboratory

I will describe the Workflow project at LLNL, which is broadly focused on bringing workflow technologies and tools to four areas: problem setup, data management, simulation/run management, and post-processing & analysis. Our environment presents several challenges, including multiple authentication realms, a variety of security and access control requirements, and a wide variety of existing workflows based on command line tools and scripts with manual data and run management.

Daniel Laney

Lawrence Livermore National Laboratory  
laney1@llnl.gov

#### MS50

##### HPC Workflow Taxonomies: A New Methodology for Hardware/Software CoDesign

The exascale target has introduced transitional stages across the High Performance Computing (HPC) ecosystem. Understanding workflows support this transition. Differing perspectives of a workflow led us to toward developing a Workflow Taxonomy that could be the basis for characterization. The approach taken was to focus on the logic flow of simulation studies and pipeline runs, the applications with their internal computational flow, and related dataflows and develop use cases as currently deployed. A resulting study done at LANL helped in the development of a workflow white paper used as part of the procurement effort for the planned Crossroads machine for 2020. In this talk we will describe our initial efforts and present examples of our HPC Workflow Taxonomy. Its target uses span a wide range, including: Providing computational and data use workflows to industry partners working toward developing exascale architecture plans; providing a taxonomy for code development teams. In addition, we will discuss the challenges in capturing data continuously to both validate the workflow characterization and support application transition and simulation efforts.

David Montoya

LANL  
dmont@lanl.gov

#### MS50

##### Swift: Implicitly Parallel Workflow for Science and Engineering Applications

Swift (<http://swift-lang.org>) is a little language for composing parallel scripts of applications or functions. Swift is implicitly and pervasively parallel, driven by dataflow. It can run parallel or sequential tasks, orchestrating the execution of functions passing in-memory or file-based data.

Swift is extremely fast and portable. It can run a million programs, thousands at a time, hundreds per second, and execute the same script on multi-core nodes, clusters, clouds, and supercomputers.

Michael Wilde

Mathematics and Computer Science Division  
Argonne National Laboratory  
wilde@mcs.anl.gov

Justin Wozniak

Argonne National Laboratory and The University of Chicago  
wozniak@mcs.anl.gov

**MS51**

**Model Discrepancy Formulations in Dakota**

Discrepancies between experimental data and model predictions are formally attributed to measurement error and so-called model form error, in which the model is considered to be inherently inadequate. New tools in Dakota which address the long-standing issue of characterizing and mitigating model form error are presented in the context of Bayesian model calibration. Potential representations of model form error, along with estimation procedures, are considered.

Kathryn Maupin

Sandia National Laboratories  
kmaupin@sandia.gov

Laura Swiler

Sandia National Laboratories  
Albuquerque, New Mexico 87185  
lpswile@sandia.gov

**MS51**

**Bayesian Calibration of Engineering Hydraulic Fracturing Models in the Presence of Parametric Uncertainties and Model Errors**

The exploitation of oil and gas can be stimulated through hydraulic fractures (HF) which are discontinuities in the rock formation induced by the injection of high pressurized viscous fluids. This enhancement methods mimic natural phenomena and are highly complex to be reproduced through laboratory experiments. Therefore, numerical simulation provides viable and effective tools for designing and optimizing the process. As there exists great variability in geologic formations, the computational models, and consequently the predictions drawn from simulations, might lead to misleading conclusions, despite the use of efficient and robust numerical schemes. A crucial aspect is taking into consideration the unavoidable uncertainties in those computational models. On one side, one has limitations on obtaining accurate information about the rock physics, which are included in the modeling through considering parametric uncertainties. Very often, simplified engineering models are used for planning and optimizing field processes, what generates another source of uncertainties. In that case, one deals with model errors. Here, within a Bayesian framework, we explore different ways of introducing model discrepancy terms regarding well established engineering models for hydraulic fracturing trying to assess the final impact of model errors and enhancing the predictive capabilities of those models through calibration based upon more sophisticated ones, from the physical

point of view.

Fernando A. Rochinha

COPPE - Federal University of Rio de Janeiro  
Rua Nascimento Silva 100 401  
faro@mecanica.coppe.ufrj.br

Souleymane Zio

Federal University of Rio de Janeiro  
zio@ufrj.br

**MS51**

**Bayesian Selection of Model-Model Connections in Multiscale Hierarchies with Uncertainty in Calibration**

Scientific models and simulations generally have structure that is bound to hypotheses. The degree to which these hypotheses are testable is an important concern in the endeavor to improve upon existing models. A description of the testability of model structure is approached by examining the process of model calibration. A Crystal Plasticity model is used to simulate the temperature dependent single crystal yielding of bcc Fe. The model is calibrated with experimental data incrementally. Separately, the model is calibrated with pseudodata generated from one calibration of the same model. In both calibrations, the model form uncertainty is described using surrogate models of likelihood functions across calibration parameter space. The likelihood functions are used to inform decisions on the data requirements of the model. Applications to hierarchical multiscale model to model connections are discussed.

Aaron E. Tallman

Georgia Institute of Technology  
atallman3@gatech.edu

Laura Swiler

Sandia National Laboratories  
Albuquerque, New Mexico 87185  
lpswile@sandia.gov

Yan Wang, David McDowell

Georgia Institute of Technology  
yan.wang@me.gatech.edu,  
david.mcdowell@me.gatech.edu

**MS52**

**Optimal Control of Fractional Power of Elliptic Operators**

In this talk we will study the existence, regularity, and approximation of solution to optimal control problems with fractional elliptic PDEs as constraints. We realize the fractional operator as a Dirichlet-to-Neumann map via the Caffarelli-Silvestre extension. We introduce a first-degree tensor product finite element space to approximate the truncated problem and derive a priori and a posteriori error estimates. We conclude by incorporating uncertainty in the PDE models.

Harbir Antil

George Mason University  
Fairfax, VA

hantil@gmu.edu

## MS52

### An Adaptive Approach for Solving Stochastic Partial Differential Equations using Reduced Local Bases

We present in this work an adaptive approach for solving partial differential equations (PDE) with random inputs based on local, reduced bases approximations. The main goal of this research is to develop an accurate and scalable approach to solve SPDEs. To this end, we have developed a strategy based on two building blocks: 1) constructing efficient and accurate local approximations over Voronoi cells, and 2) an adaptive strategy for constructing the Voronoi partitioning of the sample space. The first pillar of our approach is a local basis defined over each Voronoi-cell constructed from local solutions at a fixed number of the closest samples to the center of the cell as well as the gradient at the center of the cell. The second ingredient of our approach is an adaptive approach to sequentially select the Voronoi-cell centers. To this end, we adopt a greedy sampling using a rigorous a-posteriori error estimator, which in turn is based on an efficient approximation of the stability constant of the discrete operator. We demonstrate the performance of our proposed approach through several numerical examples with up to thirty (30) stochastic dimensions. When compared to existing methods such as Monte Carlo and Adaptive Sparse Grid, our results indicate that our proposed approach can significantly reduce computational cost for the same accuracy.

Wilkins Aquino

Associate Professor

Dept. of Civil and Environmental Engineering, Duke Univ.

wa20@duke.edu

Zilong Zhou

Duke University

zilong.zhou@duke.edu

## MS52

### Low-Rank Tensor Methods for Constrained Optimization Problems with PDEs under Uncertainty

We develop low-rank tensor methods for solving problems that arise from discretizations of inequality-constrained optimization problems with PDEs involving uncertain parameters. Especially, we consider the optimal control of PDEs under uncertainty and variational inequalities (VIs) of obstacle type with uncertainties. The proposed second-order methods perform all computations within the low-rank tensor format (Hierarchical Tucker or Tensor Train) and by this can overcome the curse of dimensionality. If time permits, we also plan to touch on some aspects of adaptive inexactness control and of possible extensions towards coherent risk measures (especially CVaR) and state constraints.

Michael Ulbrich

Technische Universitaet Muenchen

Chair of Mathematical Optimization

mulbrich@ma.tum.de

Sebastian Garreis

TU Muenchen, Germany

garreis@ma.tum.de

## MS52

### Risk Averse Optimization for Material Science

Additive manufacturing (AM) is capable of creating compelling designs at a much faster rate than standard methods. However, achieving consistent material properties at various length scales remains a challenge. In this work, we present risk-averse optimization to produce robust designs by tightly controlling certain features of the dynamics while addressing underlying uncertainties. Aspects of the AM process are emulated with PDEs to mimic the behavior of the material during the processing. The goal is to control different aspects of the dynamics, such as source terms and boundary conditions, to achieve design targets and simultaneously accommodate uncertainties in different parts of the underlying dynamics. PDE-constrained optimization methods serve as the foundation for this work with finite element discretizations, adjoint-based sensitivities, trust-region methods, and Newton-Krylov solvers. Our final AM produced parts must achieve tight tolerances for a range of different material properties. Accordingly, a range of risk measures are considered with a specific emphasis on reliability. A numerical example demonstrates that the use of risk measures as part of the objective function results in optimal solutions and ensures that worse case scenarios are avoided.

Bart G. Van Bloemen Waanders

Sandia National Laboratories

bartv@sandia.gov

## MS53

### Adaptive Sampling for Risk-Averse PDE-Constrained Optimization

Uncertainty is ubiquitous in virtually all engineering applications and, for such problems, it is inadequate to simulate the underlying physics without quantifying the uncertainty in unknown model inputs, boundary and initial conditions, and modeling assumptions. Moreover, when certifying reliability of an engineered system, it is often not sufficient to solely approximate moments of the target quantity of interest, but rather one must accurately evaluate the risk associated with, for example, tail distribution statistics. In this presentation, we introduce an adaptive approach for approximating a certain class of risk functionals. Our target class includes many popular risk measures such as the conditional value-at-risk. To guide adaptivity, we develop local error indicators based on Voronoi tessellation and the mathematical properties of this class of risk functionals. We conclude our presentation with numerical examples demonstrating the efficiency of our approximation scheme.

Drew P. Kouri

Optimization and Uncertainty Quantification

Sandia National Laboratories

dpkouri@sandia.gov

## MS53

### Universal Scalable Robust Solvers from Computational Information Games

Can we design a universal solver? Can we, via a game/decision theoretic approach to numerical analysis and algorithm design, find a scalable algorithm that could

be applied to nearly all linear operators? One motivation for asking this question is the observation that there are (nearly) as many linear solvers as linear systems. One motivation for not asking this question is Sard's affirmation "Of course no one method of approximation of a 'linear operator' can be universal." We have asked this question anyway and obtained a positive answer via a game theoretic approach under two minor conditions: the continuity of the linear operator and the existence of a compact embedding on its image space.

Houman Owhadi  
Applied Mathematics  
Caltech  
owhadi@caltech.edu

Clint Scovel  
Caltech  
clintscovel@gmail.com

### MS53

#### Scalable Methods for Rare Event Simulation in Rotorcraft Systems

We focus on the problem of predicting rare events in rotorcraft systems. Helicopter rotors are typically designed for efficient performance in cruise conditions and in hover. At the same time, structural components and blades are sized based on desired capabilities such as maximum thrust, maximum speed, and maneuverability. At the limits of these capabilities, the rotor begins to stall, which results in loss of lift and increased power and vibratory loads. These circumstances reduce the life of the rotorcraft and may lead to accidents. For autonomous rotorcraft, one aims to predict such events in real time to prevent accidents. We compute the probabilities of extreme events in a simple rotorcraft model. We focus on dynamic importance sampling algorithms based on Hamilton-Jacobi equations. Accelerating these methods using transport maps will also be discussed.

Benjamin Zhang  
MIT  
bjz@mit.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

Tuhin Sahai  
United Technologies  
sahait@utrc.utc.com

### MS53

#### UQ Prediction Under Limited Resources

One of the primary goals of uncertainty quantification (UQ) is to understand and predict system behavior in the presence of uncertainty parameters. This introduces another layer of complexity to the underlying problem. When the baseline deterministic system is already of high complexity and requires large amount of resources to sample, UQ becomes exceedingly difficult. In this talk, we will review of the newest development in UQ algorithm research. In particular, we focus on the algorithms that deliver mathematically optimal prediction results using limited and fixed number of deterministic simulation samples.

Dongbin Xiu  
University of Utah  
dongbin.xiu@utah.edu

### MS54

#### A Luenberger Observer for Reaction-Diffusion Models with Front Position Data

Cardiac electrophysiology describes and models chemical and electrical phenomena taking place in the cardiac tissue. Given the large number of related pathologies, there is an important need for understanding these phenomena. The electric wave propagating in the cardiac tissue can be represented by a nonlinear reaction-diffusion partial differential equation called the bidomain model. The complex bidomain model must be adapted to each individual case in order to produce predictive simulations for a given patient. In this context, we can use the abundant available medical data, especially the patient electrical activation maps - which correspond to the location of the front over time. In this talk, we present a Luenberger observer for reaction-diffusion models with propagating front features, and for data associated with the location of the front over time. We start by proposing an observer for the eikonal equation that can be derived from the reaction-diffusion model by an asymptotic expansion, drawing our inspiration from image processing methods. We then carry over this observer to the underlying reaction-diffusion equation by an "inverse asymptotic analysis". We also discuss the extension to joint state-parameter estimation by using the earlier-proposed ROUKF strategy. We then illustrate and assess our proposed observer strategy with test problems and also with atrial real data. Our numerical trials show that joint state-parameter estimation is directly very effective.

Annabelle Collin  
Univ. Bordeaux and Bordeaux INP, IMB, Talence, France  
INRIA Bordeaux-Sud-Ouest, Talence, France  
annabelle.collin@inria.fr

Dominique Chapelle, Philippe Moireau  
INRIA  
Dominique.Chapelle@inria.fr, philippe.moireau@inria.fr

### MS54

#### A Computational Model for Endocardial Radiofrequency Ablation (RFA) with Open-Irrigated Electrodes

Radiofrequency ablation (RFA) is a common technique in cardiac catheterization for the treatment of arrhythmias. Although globally a pretty safe procedure, it may present some risk. Thrombus formation can occur during RFA at the electrode-tissue interface when the temperature exceeds 80°C. Open-irrigated electrodes have been developed to reduce the risk of thrombus formation by cooling the electrode-tissue interface, allowing higher RF power delivery and the creation of larger lesions. On the other hand, higher RF power delivery increases the risk of steam pops occurrence, a rather serious complication. Steam pops are caused by tissue overheating above 89°C, and may trigger explosive rupture of myocardium. If the steam pop occurs sufficiently deep in the tissue, or if the RFA is performed on atria, whose walls are thinner than the ones of the ventricles, such explosive rupture may actually result in a perforation of the cardiac chamber wall, and in dra-



matic hemorrhagic events. As of today, it is still very complicated to predict the occurrence and location of steam pops into the tissue during RFA. We present here a three-dimensional model for catheter RFA with open-irrigated electrode, its validation against in vitro experiments, and show its potential for accurate spatio-temporal prediction of steam pop occurrence.

Luca Gerardo Giorda

BCAM - Basque Center for Applied Mathematics  
lgerardo@bcmath.org

Ana Gonzalez Suarez

BCAM - Basque Center for Applied Mathematics  
Bilbao, Spain  
agonzalezs@bcmath.org

Jose M Guerra Ramos

Hospital de la Santa Creu i Sant Pau  
Barcelona, Spain  
jguerra@santpau.cat

#### MS54

##### **Influence of Cardiac Mechano-Electric Feedback on Reentry Dynamics: A Simulation Study**

In this work, we investigate the influence of the cardiac tissue deformation on re-entrant wave dynamics, using a 3D strongly coupled electro-mechanical Bidomain model posed on an ideal monoventricular geometry, including fiber direction anisotropy and stretch-activated currents. The 3D electromechanical model considered consists of four coupled components (see Colli Franzone et al., *Math. Mod. Meth. Appl. Sci.*, 26: 27 -57, 2016 and *Math. Biosci.* 280: 71 - 86, 2016): a) the quasi-static transversely isotropic finite elasticity equations for the deformation of the cardiac tissue; b) the active tension model for the intracellular calcium dynamics and cross-bridge binding; c) the anisotropic Bidomain model for the electrical current flow through the deforming cardiac tissue; d) the membrane model of ventricular myocytes, including stretch-activated channels. The numerical simulations are based on our finite element parallel solver, which employs Multilevel Additive Schwarz preconditioners for the solution of the discretized Bidomain equations and Newton-Krylov methods for the solution of the discretized non-linear finite elasticity equations. We present the results of several parallel simulations of re-entrant waves in presence of different mechano-electrical feedbacks, showing that scroll wave stability seems to be more influenced by stretch-activated currents than by the geometric influence of the deformation gradient on the electrical diffusion coefficients.

Piero Colli Franzone

University of Pavia  
colli@imati.cnr.it

Luca F. Pavarino

Department of Mathematics  
University of Milan, Italy  
luca.pavarino@unipv.it

Simone Scacchi

Universita' degli Studi di Milano  
simone.scacchi@unimi.it

Stefano Zampini

King Abdullah University of Science and Technology

stefano.zampini@kaust.edu.sa

#### MS54

##### **Numerical Modeling of the Electrical Activity in the Ventricles in Presence of Detailed Purkinje Fibers**

In this talk we aim at describing the electrical propagation in the human ventricle. To this aim, we consider a detailed Purkinje network, which is known to provide the source term for the myocardial activation. However, in some pathological cases, two electrical fronts propagate, one coming from the network as in the normal case, and one coming from the myocardium. Thus, suitable coupling methods between the network and the myocardium should be addressed. We discuss the coupling between Purkinje network and myocardium where Monodomain and Eikonal subproblems are considered. In particular, we discuss the well-posedness of the coupled problems and we compare for an idealized geometry the results. We also consider the Wolff-Parkinson-White syndrome, characterized by an anomalous myocardial source of activation which requires the solution of an implicit coupling. Our results were able to recover two important features of the electrical activation in the ventricle: the so called push-and-pull effect characterizing the propagation through the Purkinje-Muscle junctions (PMJ), and the PMJ delay when passing from the network to the myocardium and viceversa. Finally, we consider an application to a realistic case of a human ventricle and we address preliminary results about the cardiac resynchronization therapy, by searching for optimal locations of the stimulation points.

Christian Vergara

MOX, Dept. Mathematics  
Politecnico di Milano, Italy  
christian.vergara@polimi.it

#### MS55

##### **A Stable and Efficient Partitioned Algorithm for Conjugate Heat Transfer**

We describe a new partitioned approach for solving conjugate heat transfer (CHT) problems where the governing temperature equations in different material domains are time-stepped in a implicit manner, but where the interface coupling is explicit. The new approach, called the CHAMP scheme (Conjugate Heat transfer Advanced Multi-domain Partitioned), is based on a discretization of the interface coupling conditions using a generalized Robin (mixed) condition with weights determined using detailed analysis of the coupling scheme. The interface treatment combines ideas from optimized-Schwartz methods for domain-decomposition problems together with the interface jump conditions and additional compatibility jump conditions derived from the governing equations. For many problems (i.e. for a wide range of material properties, grid-spacings and time-steps) the CHAMP algorithm is stable and second-order accurate using no sub-time-step iterations. In extreme cases (e.g. very fine grids with very large time-steps) it may be necessary to perform one or more sub-iterations. A comparison is made to the classical Dirichlet-Neumann (DN) method and, where applicable, to the optimized-Schwartz (OS) domain-decomposition method and outperforms both schemes. The CHAMP scheme is also developed for general curvilinear grids and CHT examples are presented using composite overset grids that confirm the theory and demonstrate the effectiveness

of the approach.

Fanlong Meng, Jeff Banks, William Henshaw  
Rensselaer Polytechnic Institute  
mengf5@rpi.edu, banksj3@rpi.edu, henshw@rpi.edu

Donald W. Schwendeman  
Rensselaer Polytechnic Institute  
Department of Mathematical Sciences  
schwed@rpi.edu

### MS55

#### Multi-Implicit Discontinuous Galerkin Method for Low Mach Number Combustion

We present an arbitrarily high-order (in space and time) multi-implicit discontinuous Galerkin (DG) method for low Mach number combustion with complex chemistry. We discretize in space by means of the compact DG method. We use the iterative spectral deferred correction method to discretize in time with arbitrarily high formal order of accuracy. We solve the advection-diffusion-reaction equations by iterating on the weakly coupled system of equations, treating advection explicitly, and diffusion and reaction implicitly. Using this method, we are able to simulate flames with complex chemistry and a large number of chemical species, and accurately resolve the stiff transients that arise.

Will Pazner  
Division of Applied Mathematics  
Brown University  
will\_pazner@brown.edu

Per-Olof Persson  
Dept. of Mathematics  
University of California, Berkeley  
persson@berkeley.edu

### MS55

#### A Stable Algorithm for Incompressible Flows and Rigid Bodies Based on Potentials

A stable FSI algorithm is developed for the problems involving interaction between viscous incompressible flow and rigid bodies. The new algorithm is developed based on added-mass potentials and added-damping potentials. The underlying dependence between the incompressible flows and rigid bodies are exposed through this approach. The resulting scheme will be compared with our previous added-mass/added-damping partitioned algorithm, in both heavy and light rigid bodies cases.

Qi Tang  
Department of Mathematical Sciences  
Rensselaer Polytechnic Institute  
tangq3@rpi.edu

### MS55

#### A Numerical Study for the Arterial Intimal Thickening Based on the Reference Map Technique

Intimal thickening (IT) is a common condition in humans which might cause atherosclerotic lesions. In [Fok, Pak-Wing, and Rebecca Sanft. "A biochemical and mechanical model of injury-induced intimal thickening." *Mathematical Medicine and Biology* (2015): dqv040], a theoretical 1D model is proposed, to describe the growth of the ar-

terial intima due to cell proliferation. This model couples the intima growth rate with both the local stress imposed by the blood flow and the local concentration of the release of a cytokine such as Platelet-Derived Growth Factor (PDGF). However, because of the 1D assumption, this theoretical model fails to simulate the uneven distribution of the buildup on the cross section of the vessel. In this work, we numerically investigate a generalized 2D model for the arterial intimal thickening, employing the reference map technique which leads to a second-order finite-difference method for solids undergoing finite deformations. Example simulations on rabbit and rat models of atherosclerosis are presented, and the prediction results are compared with experiments.

Yue Yu  
Department of Mathematics  
Lehigh University  
yueyucarol@gmail.com

Thomas Fai  
Harvard University  
tfai@seas.harvard.edu

Pak-Wing Fok  
University of Delaware  
pakwing@udel.edu

Chris H. Rycroft  
Harvard University  
chr@seas.harvard.edu

### MS56

#### Boundary Integral Methods for Computing Forces on Particles in Unsteady Stokes and Linear Viscoelastic Fluids

Accurately characterizing the forces acting on particles in fluids is of fundamental importance for understanding particle dynamics and binding kinetics. Conventional asymptotic solutions may lead to poor accuracy for neighboring particles. We present an accurate boundary integral method to calculate forces exerted on particles for a given velocity field. The idea is to exploit a correspondence principle between the unsteady Stokes and linear viscoelasticity in the Fourier domain such that a unifying boundary integral formulation can be established for the resulting Brinkman equation. Comparison with known analytic solutions and existing asymptotic solutions confirms the accuracy in space of our numerical scheme.

Xiaofan Li, Hualong Feng  
Illinois Institute of Technology  
lix@iit.edu, hfeng8@iit.edu

Francisco Hernandez  
Illinois Inst of Tech  
fhernan4@iit.edu

Shuwang Li  
Department of Applied Mathematics  
Illinois Institute of Technology  
sli@math.iit.edu

### MS56

#### Thermodynamically Consistent Sharp Interface Model and Fictitious Domain Method for Interac-

**tion between Fluid and Interface with Mass**

lowengrb@math.uci.edu

Abstract not available

Zhiliang Xu  
 University of Notre Dame  
 zhiliangxu@nd.edu

**MS56****Krylov Integration Factor Method on Sparse Grids for High Spatial Dimension Convection-Diffusion-Reaction Equations**

Krylov implicit integration factor (IIF) methods were developed in [Chen and Zhang, *Journal of Computational Physics*, 230 (2011) 4336-4352] for solving stiff reaction-diffusion equations on high dimensional unstructured meshes. The methods were further extended to solve stiff advection-diffusion-reaction equations in [Jiang and Zhang, *Journal of Computational Physics*, 253 (2013) 368-388]. Recently we studied the computational power of Krylov subspace approximations on dealing with high dimensional problems. It was shown that the Krylov integration factor methods have linear computational complexity and are especially efficient for high dimensional convection-diffusion-reaction problems with anisotropic diffusions. In this paper, we combine the Krylov integration factor methods with sparse grid combination techniques and solve high spatial dimension convection-diffusion equations such as Fokker-Planck equations on sparse grids. Numerical examples are presented to show that significant computational times are saved by applying the Krylov integration factor methods on sparse grids.

Dong Lu  
 University of Notre Dame  
 dlv1@nd.edu

Yong-Tao Zhang  
 Department of Applied and Computational Math and Statistics  
 University of Notre Dame  
 yzhang10@nd.edu

**MS56****An Efficient Adaptive Rescaling Scheme for Interface Problems**

We present an efficient rescaling scheme for computing the long-time dynamics of expanding interfaces. The idea is to design an adaptive time-space mapping such that in the new time scale, the interfaces evolve logarithmically fast at early growth stage and exponentially fast at later times. Compared with the original rescaling method in [J. Comput. Phys. 225(1) (2007) 554-567], the adaptive scheme dramatically speed up the evolution when the size of the interface is small.

Meng Zhao, Shuwang Li  
 Illinois Institute of Technology  
 mzhao8@hawk.iit.edu, sli@math.iit.edu

John Lowengrub  
 Department of Mathematics  
 University of California at Irvine

**MS57****CFD Applications Using Adaptive Mesh and Space-Filling Curves on a GPU Supercomputer**

The mesh adaptation suitable for flow profiles locally refined by tree-data algorithms improves the computational efficiency drastically. Instead of unstructured grids with inefficient memory access, we study Cartesian-grid based and particle based CFD (Computational Fluid Dynamics) applications with adaptive mesh refinement. Space-filling curves are used to measure the computational load of inhomogeneous particle distributions and mesh structures and divides the computational domain to keep an equal load balance. For GPU implementations, efficient access to the device memory should be taken into consideration and the leaf size and the neighbor leaf access are optimized. Two kinds of practical CFD applications of SPH (Smoothed Particle Hydrodynamics) method and Lattice Boltzmann method are shown. The scalabilities on a GPU supercomputer will be also discussed.

Takayuki Aoki  
 Tokyo Institute of Technology  
 taoki@gsic.titech.ac.jp

**MS57****Minimally-Invasive Integration of P4est in Espresso for Adaptive Lattice-Boltzmann**

Dynamically-adaptive, tree-structured grids are a well-known way of solving partial differential equations with less computational effort than regular Cartesian grids. We focus on a minimally-invasive integration of such grids in existing applications. We base our work on p4est, a well-known octree mesh framework, which we integrated as a new mesh in ESPResSo, a large MD-simulation package. ESPResSo uses a thermalized D3Q19-version of the lattice-Boltzmann method (LBM) for simulating background flows. For realizing spatial adaptivity, we chose a volumetric formulation of the LBM which relies on subcycling to synchronize the evolution of time on coarse and fine grid cells, because this leads to a straight-forward generalization of the LBM to tree-structured grids. We integrated virtual quadrants at refinement boundaries into p4est while preserving p4est's good scalability and most of the existing LBM implementation in ESPResSo. To preserve most of the code, we realized random-access to neighboring cells and a grid-traversal based on p4est-mesh (a set of lookup tables encoding several properties of process-local quadrants). On top of that, we added an iterator for traversing subsets of the grid, e.g. cells of a specific level, without skipping cells explicitly. We present implementational details and performance data for various background flow scenarios using different refinement criteria, including porting of ESPResSo's GPU implementation for regular grids to p4est.

Michael Lahnert  
 University of Stuttgart  
 michael.lahnert@ipvs.uni-stuttgart.de

Miriam Mehl  
 Universität Stuttgart  
 miriam.mehl@ipvs.uni-stuttgart.de

Carsten Burstedde  
 Universität Bonn

burstedde@ins.uni-bonn.de

### MS57

#### Daino: A High-level Framework for Parallel and Efficient AMR on GPUs

Adaptive Mesh Refinement methods reduce computational requirements of problems by increasing resolution for only areas of interest. However, in practice, efficient AMR implementations are difficult considering that the mesh hierarchy management must be optimized for the underlying hardware. Architecture complexity of GPUs can render efficient AMR to be particularly challenging in GPU-accelerated supercomputers. This talk presents a compiler-based high-level framework that can automatically transform serial uniform mesh code annotated by the user into parallel adaptive mesh code optimized for GPU-accelerated supercomputers. We show experimental results on three production applications. The speedups of code generated by our framework are comparable to hand-written AMR code while achieving good strong and weak scaling up to 4000 GPUs.

Mohamed Wahib

RIKEN, Kobe, Japan  
mohamed.attia@riken.jp

Naoya Maruyama

RIKEN Advanced Institute for Computational Science  
nmaruyama@riken.jp

Takayuki Aoki

Tokyo Institute of Technology  
taoki@gsic.titech.ac.jp

### MS58

#### Transport Maps for Efficient Conditional Sampling

We present a class of nonlinear random variable transformations, in the form of block-triangular transport maps, that can be used to efficiently characterize non-Gaussian conditional probability distributions. Given two jointly distributed random variables  $\theta$  and  $d$ , our interest is in sampling the conditional distribution  $\pi(\theta|d)$ . To do this however, our map-based approach only requires samples of the joint distribution  $\pi(d, \theta)$ , which can be obtained through either model-generated data or real experimental data. Using only joint samples allows the map to be constructed offline so that  $\pi(\theta|d)$  can be efficiently sampled for *any* value of  $d$ . We discuss the details of parameterizing and constructing transport maps from samples, with a special emphasis on the use of offline computation to support near real-time applications. We also outline interesting relationships between our method and other approaches like variational autoencoders and Riemannian manifold Markov chain Monte Carlo. Several examples will be provided with a focus on Bayesian inference and the construction of non-Gaussian Markov random fields.

Matthew Parno, Youssef M. Marzouk

Massachusetts Institute of Technology  
Matthew.D.Parno@usace.army.mil, ymarz@mit.edu

### MS58

#### Dimensionality Reduction in the Wasserstein Space

When the space of probability measures  $P(X)$  is equipped with a ground metric on the ground space  $X$ , distances

based on Optimal Transport, a.k.a Wasserstein distances, provide a powerful metric on  $P(X)$  which have been shown to be efficient at comparing probability measures in numerous concrete applications appearing in such fields as Machine Learning. When the ground space  $X$  is furthermore a Hilbert space,  $P(X)$  also carries the additional structure of an infinite dimensional Riemannian manifold. Relying on the concepts of geodesics and tangent vectors, we show how this structure allows us to generalize popular dimensionality reduction algorithms such as PCA or Dictionary Learning. We present a convenient parametrization of geodesics to formulate the problem of Principal Geodesic Analysis, and propose an algorithm to minimize the resulting objective function through a maximization-minimization procedure. Experiments on toy and real data highlight the relevance and usefulness of our approach.

Vivien Seguy

Kyoto University  
vivien.seguy@iip.ist.i.kyoto-u.ac.jp

### MS58

#### Warp-Bridge Sampling to Improve Efficiency of Free Energy Calculations in Physical Chemistry

Abstract not available

Michael Shirts

University of Colorado Boulder  
michael.shirts@colorado.edu

### MS58

#### Explanation of Variability Through Optimal Transport

A methodology based on the theory of optimal transport is developed to attribute variability in data sets to known and unknown factors and to remove such attributable components of the variability from the data. Denoting by  $x$  the quantities of interest and by  $z$  the explanatory factors, the procedure transforms  $x$  into filtered variables  $y$  through a  $z$ -dependent map, so that the conditional probability distributions  $\rho(x|z)$  are pushed forward into a target distribution  $\mu(y)$ , independent of  $z$ . Among all maps and target distributions that achieve this goal, the procedure selects the one that minimally distorts the original data: the barycenter of the  $\rho(x|z)$ . Connections are found to the methodology are shown to be equivalent to  $k$ -means and principal component analysis– and to fundamental problems in statistics such as conditional density estimation and sampling. An application is shown to a time-series of ground temperature hourly data across the United States.

Esteban G. Tabak

Courant Institute  
New York University  
tabak@cims.nyu.edu

### MS60

#### Inversion Strategies for Large Scale Data and Large Scale Models

in this talk we discuss the problem of estimating a parameter function from a differential equation with multiple right hand sides. A key concept in our approach is mesh decoupling in which the forward mesh is different from the inversion mesh. While this concept allow us to deal with

millions of right hand sides in parallel it rises an important question. How should we interpolate the parameter between meshes? We show that a natural approach to this question involves multiscale techniques. We also show how to compute derivatives of the multiscaled approximation and solve the problem.

Eldad Haber

University of British Columbia, Vancouver, Canada  
ehaber@eos.ubc.ca

## MS60

### Computational Approaches for Multi-Frame Blind Deconvolution

In this talk we consider multi-frame blind deconvolution (MFBD). There are many optimization approaches that can be applied to these problems, but they are generally limited in the number of frames that can be processed. In this talk we consider efficient computational approaches to reduce a massive number of data frames to a computationally manageable set, while still preserving all information for MFBD algorithms.

James G. Nagy

Emory University  
Department of Math and Computer Science  
nagy@mathcs.emory.edu

## MS60

### High Performance Solvers for Linear Systems in Graph Laplacians

Over the past two decades there has been much theoretical progress on efficient solvers for Symmetric Diagonally Dominant matrices, which are equivalent to Laplacian matrices of undirected graphs. This talk will overview recent efforts at fast and robust implementations geared towards systems with sizes that are a factor of 100 to 1000 times larger than what's now possible. Of particular importance to this project are application-oriented benchmarks that fully test the robustness of solvers in the presence of widely varying edge weights. Resolving this issue is crucial for bridging the theory and practice of incorporating these solvers in optimization algorithms such as iterative least squares, mirror descent, and interior point methods. Works in this direction have motivated the study of a variety of theoretical algorithmic tools from the perspective of high performance computing. Crucial components in the solver algorithms include tree data structures, parallel tree contraction, and recursive invocations of iterative methods,

Richard Peng

Georgia Tech  
rpeng@cc.gatech.edu

Kevin Deweese

UC Santa Barbara  
kdeweese@cs.ucsb.edu

John R. Gilbert

Dept of Computer Science  
University of California, Santa Barbara  
gilbert@cs.ucsb.edu

Gary Miller

Carnegie Mellon University  
gmliller@cs.cmu.edu

Serban Stan

Yale University  
serban.stan@yale.edu

Haoran Xu

MIT  
haoranxu510@gmail.com

ShenChen Xu

Carnegie Mellon University  
shenchex@cs.cmu.edu

## MS60

### Efficient MCMC Samplers for Bayesian Inverse Problems with Many Measurements

Markov Chain Monte Carlo (MCMC) methods are the prevalent approaches for sampling posterior distributions corresponding to a large-dimensional parameter space. For Bayesian inverse problems with many measurements, the likelihood evaluations become prohibitively expensive because several large-scale systems need to be solved. Our approach is to use randomized matrix methods to reduce computational costs corresponding to the likelihood evaluations. We demonstrate the performance of our algorithms on 2D model problems from Diffuse Optical Tomography.

Arvind Saibaba

North Carolina State University  
asaibab@ncsu.edu

Misha E. Kilmer

Mathematics Department  
Tufts University  
misha.kilmer@tufts.edu

Eric De Sturler

Virginia Tech  
sturler@vt.edu

Eric L. Miller

Tufts University  
elmiller@ece.tufts.edu

## MS61

### Global Reconstruction of Solutions to Parametric PDEs Via Compressed Sensing

We present a novel theoretical framework for solving parametric PDEs via compressed sensing over tensor-products of Hilbert spaces. This work builds on the existing theory for the recovery of compressible solutions via  $\ell^1$ -minimization, and guarantees convergence in terms of the errors of the best  $s$ -term approximation and the residual in a given polynomial subspace. Compared to other approaches that only recover a functional of the solution, e.g. evaluation at a single point, our approach recovers the solution globally over the physical domain. We also provide extensions to the fixed point continuation algorithms and the Bregman iterations for solving the basis pursuit problem in this context. We conclude with numerical results demonstrating the efficacy of our approach in high-dimensions and comparisons with sparse grids and stochastic Galerkin approximations.

Nick Dexter

University of Tennessee  
ndexter@utk.edu

Hoang A. Tran  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
tranha@ornl.gov

Clayton G. Webster  
Oak Ridge National Laboratory  
webstercg@ornl.gov

### MS61

#### A Resilient Solver for 2D Uncertain Elliptic PDEs Via Fault-Tolerant MPI Server-Client-Based Implementation

We discuss an approach to solve 2D elliptic PDEs with uncertain coefficients that incorporates resiliency to silent data corruption (SDC) and failing MPI ranks. Resiliency to SDC is obtained at the algorithm level by recasting the PDE as a sampling problem followed by a resilient manipulation of the data to obtain a new solution. Resiliency to MPI rank failures is achieved at the implementation level by leveraging a fault-tolerant MPI implementation (ULFM-MPI) within a task-based server-client model.

Karla Morris, Francesco Rizzi  
Sandia National Laboratories  
knmorri@sandia.gov, fnrizzi@sandia.gov

Paul Mycek  
Duke University  
paul.mycek@duke.edu

Olivier Le Maitre  
Laboratoire d'Informatique pour la Mécanique et les  
Science  
olm@limsi.fr

Omar M. Knio, Andres Contreras  
Duke University  
omar.knio@duke.edu, andres.contreras@duke.edu

Khachik Sargsyan, Cosmin Safta  
Sandia National Laboratories  
ksargsy@sandia.gov, csafta@sandia.gov

Bert J. Debusschere  
Energy Transportation Center  
Sandia National Laboratories, Livermore CA  
bjdebus@sandia.gov

### MS61

#### Highly Scalable Hierarchical Sampling Algorithms for Spatially Correlated Random Fields

The ability to generate samples of random fields with prescribed statistical properties is a key ingredient for many uncertainty quantification algorithms. In this talk, we will present a highly scalable multilevel hierarchical sampling technique based on solving a stochastic partial differential equation that has linear complexity. Numerical experiments demonstrate the scalability of the sampling method applied to multilevel Monte Carlo simulations for a subsurface porous media flow application.

Sarah Osborn  
Lawrence Livermore National Lab  
osborn9@llnl.gov

Panayot Vassilevski  
Lawrence Livermore National Laboratory  
vassilevski1@llnl.gov

### MS61

#### Tackling UQ in Darma, a New Task-Based Programming Model for Extreme-Scale Computing

This talk focuses on the advantages of task-based implementations of UQ problems. Specifically, we demonstrate how to tackle UQ using DARMA (Distributed Asynchronous Resilient Models and Applications), a co-design research vehicle for asynchronous many-task (AMT) programming models. Two main features are highlighted: (a) the ability to express concurrency and parallelism among independent UQ tasks, and (b) an analysis of the benefits of reusing, within a pool of samples, tasks results to accelerate the execution time of other independent tasks. We demonstrate first a simple set of basic UQ examples written in the DARMA, and then focus on a Multi-Level Monte Carlo test case.

Francesco Rizzi  
Sandia National Laboratories  
fnrizzi@sandia.gov

Eric Phipps  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

David Hollman  
Sandia National Labs  
dshollm@sandia.gov

Jeremiah Wilke  
Sandia National Laboratories  
jjwilke@sandia.gov

Jonathan Lifflander  
Sandia National Labs  
jlifflla@sandia.gov

Hemanth Kolla  
Sandia National Laboratories  
hnkolla@sandia.gov

Nicole Slattengren  
Sandia National Labs  
nlslatt@sandia.gov

Keita Teranishi, Janine C. Bennett  
Sandia National Laboratories  
knieran@sandia.gov, jcbenne@sandia.gov

### MS62

#### Immersed Boundary Peridynamics Method (IB/PD) for Modeling Aortic Dissection

Abstract not available

Amneet P.S. Bhalla  
University of North Carolina at Chapel Hill  
mail2amneet@gmail.com

### MS62

#### Grain Growth Modeling and Novel Data Analytics

**for Weld Process**

In this talk, a model which simulates grain growth during welding will be presented; following that, the focus of the talk will be on novel data analytics for analysis of material microstructures predicted by the grain growth model. The approach uses calculation of two-point statistics for many model generated microstructure realizations followed by principal component analysis for construction of a surrogate model using important weld process parameters such as weld speed, weld pool geometry, and temperature gradients. The weld model is a novel variation on the Potts model for grain growth and is part of the SPPARKS kinetic Monte Carlo simulation framework.

John A. Mitchell  
Sandia National Laboratories  
jamitch@sandia.gov

**MS62****Concurrent Multiscale Coupling Methods for the Discretized Peridynamic Theory and the Finite Element Method**

We consider the peridynamic and the classical elastic theories in the concurrent multiscale simulation of engineering problems. In the region where material failure is expected to initiate, a discretized particle solver based on the peridynamic theory is employed. In the rest of the problem domain, the material is modeled by the classical elastic theory, and numerically discretized with the finite element method. For quasi-static problems, we employ iterations for each step in the concurrent coupling scheme, and investigate different coupling strategies as well as interface boundary conditions. The numerical experiments suggest the Robin boundary condition together with the Aitken algorithm for its robustness and efficiency. Finally, we consider a two-dimensional plane stress problem in a plate with a hole, and show that the strain concentration creates a small region where high strains cause cracks to initiate, and grow toward the outer free edges.

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

Fabiano Bargos  
University of São Paulo (USP)  
fabianobargos@usp.br

Marco Bittencourt  
Departamento de mecanico  
Universidade de Campinas  
mlb@fem.unicamp.br

Michael L. Parks  
Sandia National Laboratories  
mlparks@sandia.gov

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

**MS63****Tensor Completion with Applications to Multidimensional Inverse Problems**

We study the problem of reconstructing an approximately

low rank tensor from a small number of noisy measurements. New recovery guarantees, numerical algorithms, non-uniform sampling strategies, and parameter selection algorithms are developed. We derive a fixed point continuation algorithm for tensor completion and prove its convergence. Recovery guarantees are obtained for sub-Gaussian measurement operators and for measurements obtained by non-uniform sampling from a Parseval tight frame. We show how tensor completion can be used to solve multidimensional inverse problems arising, e.g., in the context of NMR relaxometry. Presented methods are validated on, both, simulated and experimental data. This is a joint work with Ariel Hafftko and Richard Spencer.

Wojciech Czaja  
University of Maryland, College Park  
czaja@umd.edu

**MS63****An Efficient Method for Statistical Learning by Means of Tensor Format Representations**

Abstract not available

Mike Espig  
RWTH Aachen University  
mike.espig@alopax.de

**MS63****Tensor Completion in Bayesian Inverse Problems**

We solve PDEs with uncertain parameters. The Bayesian update method is used to compute posterior probability density functions for these uncertain parameters. After applying the stochastic Galerkin approach, after KL and PC Expansions, we obtain a linear system, where all ingredients are represented as high-order tensors. Typically we are able to measure not the solution  $U$  itself, but some functional  $F(U)$  or  $U$  just in very few locations. The solution  $U$  (and possibly the functional  $F(U)$ ) of our model is a high-order tensor, where few dimensions come from the spatial part and may be many hundreds from the stochastic part. Available measurements of  $F(U)$  are only a small part of total data, therefore, we have a lot of missing values in the tensor  $F(U)$  and it is hard to use it straightforward in the Bayesian update framework. Our idea is to use tensor completion techniques to recover the missing values in  $F(U)$  (this is a preprocessing step). After that we can compute the new likelihood function (which, as we hope, will be more informative than the old one constructed just by few data) and the Bayesian update formula. We are also interested to learn from tensor completion theory, which entries of  $F(U)$  are "mostly important". This is similar to choosing indices in the famous adaptive cross approximation algorithm.

Alexander Litvinenko  
SRI-UQ and ECRC Centers, KAUST  
alexander.litvinenko@kaust.edu.sa

Hermann Matthies  
Technische Universität Braunschweig  
h.matthies@tu-bs.de

**MS63****Optimization Methods on Low-Rank Manifolds**

**and their Efficient Implementation**

Abstract not available

Ivan Oseledets

Skolkovo Institute of Science and Technology  
ivan.oseledets@gmail.com

**MS64****Batch Linear Algebra for GPU-Accelerated High Performance Computing Environments**

There is an increasing interest, driven by many scientific applications, in solving a large number of relatively small-size problems that are independent from each other. Applications include astrophysics, quantum chemistry, sparse direct solvers, CFD, and others. Most of these problems perform dense matrix operations, and may vary in size. As the size of an individual problem does not provide sufficient parallelism for modern hardware, especially accelerators, these problems must be solved concurrently as a batch in order to saturate the hardware with enough work, hence the name batched computation. This talk presents high performance batched computation using Graphics Processing Units (GPUs), with emphasis on computational kernels that perform BLAS and LAPACK functions. The kernels discussed are part of the MAGMA library. We discuss fixed and variable size problems, and highlight the user interface, kernel design, and, optimization techniques. We show that performance gains are obtained against non-batched kernels, as well as other designs that represent state-of-the-art.

Ahmad Abdelfattah

University of Tennessee, Knoxville  
ahmad@icl.utk.edu

**MS64****Extending Batched GEMM for Tensor Contractions on GPU**

Tensor contractions constitute a key computational ingredient of numerical multi-linear algebra. However, as the order and dimension of tensors grow, the time and space complexities of tensor-based computations grow quickly. Existing approaches for tensor contractions spend large amounts of time restructuring the data which typically involves explicit copy and transpose operations. In this work, we summarize existing approaches and present library-based solutions that avoid memory movement. We propose and evaluate new BLAS-like primitives that are capable of performing a wide range of tensor contractions on CPU and GPU efficiently. Through systematic benchmarking, we demonstrate that our approach can achieve, for small and moderate tensor sizes, 10x speedup on a K40c GPU and 2x speedup on dual-socket Haswell-EP CPUs against existing CUBLAS and MKL solutions respectively. Concretely, we implement a Tucker decomposition and show that using our kernels yields at least an order of magnitude speedup as compared to state-of-the-art libraries. The Strided-BatchedGEMM solution is not capable of evaluating some tensor contractions efficiently. We classify these problems and present ongoing research in modified interfaces and implementations that are capable of computing these exceptional cases with performance on par with Level-3 BLAS.

Cris Cecka

NVIDIA Corporation

ccecka@nvidia.com

**MS64****Optimizing Batched Linear Algebra on Intel(R) Xeon Phi(TM) Processors**

Various numerical algorithms including sparse solvers and finite element methods require a large number of matrix-matrix multiplication operations that can be performed independently. As these multiplications are often performed on relatively small matrices, batching the operations allows exploiting the parallelism available on today's multi- and many-core systems, such as the Intel® Xeon Phi™ x200 processors. First an overview of the Intel® Xeon Phi™ architecture is given, with emphasis on its floating-point capabilities. We then present implementation details of the batched matrix-matrix multiplication routines in the Intel® Math Kernel Library, focusing on techniques for the Intel® Xeon Phi™ processors.

Sarah Knepper, Murat E. Guney, Kazushige Goto, Shane Story, Arthur Araujo Mitrano, Timothy B. Costa, Louise Huot

Intel Corporation

sarah.knepper@intel.com, murat.e.guney@intel.com,  
kazushige.goto@intel.com, shane.story@intel.com,  
arthur.araujo.mitrano@intel.com, timothy.b.costa@intel.com, louise.huot@intel.com

**MS64****Recent Advances in Batched Linear Algebra Computation**

Batched linear algebra, the simultaneous solution of many small BLAS problems, is becoming increasingly important within the scientific computing community. From machine learning to the solution of sparse linear systems, there are a number of applications where we require the solution of multiple small linear algebra problems, as opposed to one large problem. Currently there is no standard API for batched linear algebra operations, with a range of different ideas being proposed. In this talk we discuss the current state-of-the-art software available, before analysing the effects that different API design choices might have on performance. We find that the layout of matrices in memory significantly impacts both the API design and the performance obtained on multicore systems, GPUs, and the Xeon Phi.

Samuel Relton

University of Manchester, UK  
samuel.relton@manchester.ac.uk

Mawussi Zounon

University of Manchester  
mawussi.zounon@manchester.ac.uk

Pedro Valero-Lara

University of Manchester  
pedro.valero-lara@manchester.ac.uk

**MS65****Reduced Order Modeling for Time-Dependent Optimization Problems with Initial Value Controls**

We present a new reduced order model (ROM) Hessian approximation for linear-quadratic optimal control problems where the optimization variable is the initial value. Such



problems arise in parameter identification and also as sub-problems in multiple shooting formulations of more general optimal control problems. Our ROM Hessian approximation substantially reduces the computing time needed to solve such problems. The computation of Hessian-vector products requires the solution of the linearized state equation with initial value given by the vector to which the Hessian is applied to, followed by the solution of the second order adjoint equation. We generate projection based ROMs of these two linear differential equations to generate our Hessian approximation. The challenge is that in general no fixed ROM well-approximates the application of the Hessian to all possible vectors of initial data. Our approach first selects a fixed ROM and then for each given vector of initial data augments this fixed ROM by adding one basis vector. The resulting ROM is of fixed size, but is nonlinear in the vector of initial data. Our augmented ROM can produce substantially better approximations than a fixed ROM. We show how to use our nonlinear ROMs in a conjugate gradient (CG) method to compute optimal controls. In important special cases, the structures of the problem and of the CG method only require us to augment the original fixed ROM once, rather than in each CG iteration.

#### Doerte Jando

Interdisciplinary Center for Scientific Computing,  
Heidelberg University  
doerte.jando@iwr.uni-heidelberg.de

#### Matthias Heinkenschloss

Department of Computational and Applied Mathematics  
Rice University  
heinken@rice.edu

#### MS65

##### **Inexact Parallel-in-Time Solvers and SQP Methods for PDE-Constrained Optimization**

State-of-the-art methods for optimization with transient partial differential equation (PDE) constraints require repeated executions of the PDE simulator's forward and backward (adjoint) time integration schemes. Due to the *arrow of time* this results in a major sequential bottleneck that is difficult to overcome by additional parallelization within each time step. We study parallel-in-time solvers for optimality systems arising in nonlinear PDE-constrained optimization. The solvers are based on a fully decoupled, concurrent solution of smaller optimality systems in each time subdomain. The coupling in time is recovered gradually over optimization iterations, through an inexact iterative scheme combined with a matrix-free trust-region sequential quadratic programming (SQP) algorithm. The SQP algorithm controls the degree of inexactness based on feasibility and optimality criteria, and guarantees full time continuity at the optimal solution. We present results for diffusion-dominated and advection-dominated nonlinear optimal control problems.

#### Denis Ridzal

Sandia National Laboratories  
dridzal@sandia.gov

#### MS65

##### **Parallel in Time Solutions of Variational PDE-Based Inference**

This talk discusses a scalable parallel-in-time algorithm based on augmented Lagrangian approach to perform PDE-constrained optimization. The focus application is

variational inference, also called four dimensional variational data assimilation. The simulation time window is divided into multiple subintervals to facilitate parallel cost function and gradient computations. The scalability of the method is tested by performing data assimilation with increasing problem sizes (i.e. the assimilation window). The data assimilation is performed for Lorenz-96 and the shallow water model. We also explore the feasibility of using a hybrid-method (a combination of serial and parallel 4D-Vars) to perform the variational data assimilation.

#### Adrian Sandu

Virginia Polytechnic Institute  
and State University  
asandu7@vt.edu

#### Vishwas Rao

ICES, UT Austin  
visrao@vt.edu

#### MS65

##### **Preconditioners for Time-Dependent PDE-Constrained Optimization Based on Parareal Time-Domain Decomposition**

We consider optimization problems governed by time-dependent parabolic PDEs and discuss the construction of parallel preconditioners based on the parareal time domain decomposition method for the solution of quadratic subproblems which arise within SQP methods. In the case without control constraints, the optimality system of the subproblem is directly reduced to a symmetric PDE system, for which we propose a preconditioner that decouples into a forward and backward PDE solve. In the case of control constraints we apply a semismooth Newton method (equivalent to a primal dual active set method) and apply the preconditioner to the semismooth Newton system. We prove bounds on the condition number of the preconditioned system which shows no or only a weak dependence on the size of regularization parameters for the control. We propose to use the parareal time domain decomposition method for the forward and backward PDE solves within the PDE preconditioner to construct an efficient parallel preconditioner. Numerical results are presented.

#### Stefan Ulbrich

Technische Universitaet Darmstadt  
Fachbereich Mathematik  
ulbrich@mathematik.tu-darmstadt.de

#### MS66

##### **Random Sketching for Projection Based Model Order Reduction**

In this work we show how the cost of the offline stage of the standard projection based Model Order Reduction algorithms can be reduced by random sketching – a well-known technique for handling massive data. Random sketching can be extremely helpful for the reduction of both computational and storage costs. In particular, it can be exploited for the efficient evaluation of inner products and affine decompositions. Moreover, building the reduced model with random sketching consumes just one pass over the data, which allows to avoid possible storage burden for large-scale problems. Given a parameter-dependent preconditioner, we notice that its random sketch can be used as a preconditioner itself. The sketch preserves the quality of the preconditioner with high probability. Last but not least, we discuss how random sketching can help to build

an effective parameter-dependent preconditioner by interpolation of the operator's inverse. Such preconditioner can be used to improve the quasi-optimality constants and/or to provide cheap error estimators. All our results are supported by numerical experiments.

Oleg Balabanov  
Ecole Centrale de Nantes  
olegbalabanov@gmail.com

Anthony Nouy  
Ecole Centrale Nantes  
anthony.nouy@ec-nantes.fr

## MS66

### Low-Rank Structure in Measure Transport for Bayesian Inference

Measure transport is a valuable tool for characterizing, sampling and manipulating multivariate non-Gaussian target distributions [Moselhy, Marzouk. Bayesian inference with optimal maps. 2012] [Parno, Marzouk. Transport map accelerated Markov chain Monte Carlo. 2015] [Marzouk, Moselhy, Parno, Spantini. An introduction to sampling via measure transport. 2016]. This method has a broad range of applications – e.g., the solution of Bayesian inverse problems, as well as filtering and smoothing of dynamical systems. The transport maps framework seeks a deterministic parametric map that pushes forward a tractable reference distribution to a potentially complex target distribution of interest [Bigoni, Spantini, Marzouk. Numerical methods for monotone measure transport.]. The construction of high-dimensional maps may be challenging due to the curse of dimensionality [Spantini, Bigoni, Marzouk. Measure transport, variational inference, and low-dimensional maps.]. To tackle this, we will explore the characterization of transports with low-rank structure, which can be efficiently represented using the spectral tensor-train decomposition [Bigoni, Engsig-Karup, Marzouk. Spectral tensor-train decomposition. 2016]. This allows the representation of maps using a number of parameters which grows only linearly in the dimension of the problem. This approach will be demonstrated on Bayesian inference problems arising in spatial statistics and differential equations.

Daniele Bigoni, Alessio Spantini, Youssef M. Marzouk  
Massachusetts Institute of Technology  
dabi@mit.edu, spantini@mit.edu, ymarz@mit.edu

## MS66

### Compressed Sensing Methods for Sparse Approximations of High-Dimensional Parametric PDEs

In this talk we review some recent results in the approximation of solutions to high-dimensional parametric PDEs by considering sparse approximations of their generalized polynomial chaos expansions. We further develop numerical approximations schemes which are proven to converge pointwise for any parameter vector, under some mild ellipticity conditions. We show that the convergence can be made efficient by using tools from compressed sensing and its weighted extension. The resulting numerical approximation delivers a solution which is accurate up to a prescribed discretization grid and is asymptotically optimal for a given regularity of the solution and compressibility of its GPC. The approach yields accurate solutions for any inputs at the targeted discretization grid, while keeping the computation costs in the order of one solve at the finest

level. This is joint work with Benjamin Bykowski, Holger Rauhut, and Christoph Schwab.

Jean-Luc Bouchot  
RWTH Aachen  
bouchot@mathc.rwth-aachen.de

## MS66

### Improving Sparse Recovery Guarantee for Legendre Expansions Using Envelope Bound

The sample complexity of polynomial approximation using  $l_1$  minimization has usually been derived via the uniform bound of the underlying basis. In this work, we prove a sufficient condition for sparse Legendre expansion without using this uniform boundedness condition. Our sample complexity, independent of maximum polynomial degree, is established using the restricted eigenvalue property and the unbounded envelope of all Legendre polynomials. Our analysis also reveals some easy-to-test criteria for random sample sets under which the reconstruction error can be slightly improved.

Clayton G. Webster  
Oak Ridge National Laboratory  
webstercg@ornl.gov

Hoang A. Tran  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
tranha@ornl.gov

## MS67

### Learning Dominant Wave Directions For Plane Wave Methods For High-Frequency Helmholtz Equations

We present a ray-based finite element method (ray-FEM) by learning basis adaptive to the underlying high-frequency Helmholtz equation in smooth media. Based on the geometric optics ansatz of the wave field, we learn local dominant ray directions by probing the medium using low-frequency waves with the same source. Once local ray directions are extracted, they are incorporated into the finite element basis to solve the high-frequency Helmholtz equation. This process can be continued to further improve approximations for both local ray directions and the high frequency wave field iteratively. The method requires a fixed number of grid points per wavelength to represent the wave field and achieves an asymptotic convergence as the frequency  $\omega \rightarrow \infty$  without the pollution effect. A fast solver is developed for the resulting linear system with an empirical complexity  $\mathcal{O}(\omega^d)$  up to a poly-logarithmic factor. Numerical examples in 2D are presented to corroborate the claims.

Jun Fang  
UCI  
junf2@uci.edu

## MS67

### An Efficient High Accuracy Direct Solution Technique for High Frequency Helmholtz Problems

In this talk, we present a high-order accurate discretization technique designed for variable coefficient partial differential equations. The method is based on local spectral collocation with a hierarchical merge procedure which “glues”

the local problems together. For problems with highly oscillatory solutions such as high frequency Helmholtz problems, the complexity of constructing the direct solver is  $O(N^{3/2})$  while the cost of a solve scales  $O(N \log(N))$  with a small constant. Numerical results will illustrate the performance of the proposed method to free-scattering problems with local deviations from constant coefficient, inverse scattering applications, and the performance of the adaptive algorithm.

Adrianna Gillman

Rice University  
Department of Computational and applied mathematics  
adrianna.gillman@rice.edu

**MS67**

**A Fast and High-Order Hybridizable Discontinuous Galerkin Solver for the High-frequency Helmholtz Equation**

Solving the high-frequency Helmholtz equation in heterogeneous media is an important problem. Nevertheless, in the context of numerical analysis, it is still open in view of both accuracy and complexity. The first documented algorithm with truly scalable complexity (i.e. with a runtime sub-linear in the number of volume unknowns in a parallel environment) for the high-frequency Helmholtz equation is the method of polarized traces. However, previous versions of this method were either restricted to low order discretizations, or suffered from computationally unfavorable complexity in the  $p$ -th order case. In this talk, we rectify this issue by proposing a high-order method of polarized traces based on a Hybridizable Discontinuous Galerkin (HDG) discretization. Each stage of our algorithm either can be performed in an embarrassingly parallel fashion or has a linear computational complexity with respect to the number of degrees of freedom, independently of the order  $p$ . This is an important improvement because, in the high-frequency regime, high-order discretizations are required to attenuate the pollution error. In addition, HDG is a welcome upgrade for the method of polarized traces since it delivers accurate solutions even for non-smooth media, allows for a post-processing routine to improve on accuracy, and works with flexible meshes that align with discontinuities in the medium. We present several numerical examples to corroborate the accuracy and complexity claims.

Matthias Taus

University of Texas  
matthias.taus@gmail.com

Leonardo Zepeda-Nunez  
University of California, Irvine  
lzepeda@math.mit.edu

Laurent Demanet  
Department of Mathematics, MIT  
laurent@math.mit.edu

**MS67**

**A Simple Curved Boundary Treatment for Explicit High Order Dg Methods Solving Time-Dependent Problems**

For problems defined in a 2D domain with boundary conditions specified on a curve  $\Gamma$ , we consider discontinuous Galerkin (DG) schemes with high order polynomial basis on a geometry fitting triangular mesh. It is well known

that directly imposing the given boundary conditions on a piece-wise segment approximation boundary  $\Gamma_h$  will render any finite element method to be at most second order accurate. Unless the boundary conditions can be accurately transferred from  $\Gamma$  to  $\Gamma_h$ , in general curvilinear element method should be used to obtain high order accuracy. We discuss a simple boundary treatment, which can be implemented as a modified DG scheme defined on triangles adjacent to  $\Gamma_h$ . Even though integration along the curve is still necessary, integrals over any curved element are avoided. If the domain is convex, or if the true solutions can be smoothly extended to the exterior of a generic domain, the modified DG scheme is high order accurate. In these cases, numerical tests on first order and second order PDEs including hyperbolic systems and the scalar wave equation suggest that it is as accurate as the full curvilinear DG scheme if it is stable. Rigorous stability result is difficult to establish. Nonetheless, numerical tests suggest that the modified DG scheme is stable on a reasonably coarse mesh and finer ones.

Xiangxiong Zhang

Purdue University  
Department of Mathematics  
zhan1966@purdue.edu

**MS68**

**A Method for Data-Assimilation in Navier-Stokes with a Least-Squares Rom**

Abstract not available

Richard P. Dwight  
Delft University of Technology  
r.p.dwight@tudelft.nl

**MS68**

**UQ for Correlated Inputs**

One of the most widely used methods for UQ is stochastic collocation, which is based on Gaussian quadrature. In this presentation, we propose a new method, based on clustering, to extend stochastic collocation to the case where the inputs are correlated. In this method, the cluster centers are used as collocation nodes and the weights assigned to the nodes are proportional to the cluster sizes. It can be applied when the probability distribution of the inputs is unknown and only a sample of the input distribution is available. This is frequently the case in real-life applications where a sample of inputs is given in the form of measurement data. The performance of the new method is tested with the Genz test functions serving as a benchmark. We demonstrate that the method works even for nonlinearly correlated inputs in higher dimensions (tests with dimension up to 16 are shown). A remarkable result is that the performance improves when the correlation increases. This is due to the correlation leading to the data being effectively distributed in a lower-dimensional space. Finally, this new method is applied to a lid-driven cavity flow in which the viscosity and lid velocity are correlated. This leads to a reduction in simulation time by a factor 40, while the results are remarkably similar.

Anne Eggels  
CWI Amsterdam  
The Netherlands  
eggels@cwi.nl

Daan Crommelin

Faculty of Science  
University of Amsterdam, the Netherlands  
d.t.crommelin@uva.nl

### MS68

#### UQ for Particle-Laden Turbulent Flows in Solar Receivers

Uncertainty quantification (UQ) plays a fundamental part in building confidence in predictive science. Of particular interest is the case of modeling and simulating engineering applications where, due to the inherent complexity, many uncertainties naturally arise, e.g. domain geometry, operating conditions, errors induced by modeling assumptions, etc. In this regard, one of the pacing items, especially in high-fidelity computational fluid dynamics (CFD) simulations, is the large amount of computing resources typically required to propagate uncertainty through the models. Upcoming exascale supercomputers will significantly increase the available computational power. However, UQ approaches cannot entrust their applicability only on brute force Monte Carlo (MC) sampling; the large number of uncertainty sources and the presence of nonlinearities in the solution will make straightforward MC analysis unaffordable. The present effort relies a computing-budget constrained algorithm to optimize the execution of UQ in multi-physics CFD problems. The methodology is based on multi-level/fidelity MC approaches and takes advantage of a natural hierarchy of physical models to accelerate the estimation of the effect of uncertainties. The approach is described in detail, and its performance demonstrated on a radiated turbulent particle-laden flow case relevant to solar energy receivers (PSAAP II: Particle-laden turbulence in a radiation environment).

Lluís Jofre

Stanford University  
jofre@stanford.edu

Gianluca Geraci  
University of Stanford  
ggeraci@stanford.edu

Gianluca Iaccarino  
Stanford University  
Mechanical Engineering  
jops@stanford.edu

### MS68

#### Bayesian Calibration of Model-Form Uncertainty Applied to Wind Turbine Wake Simulation

In the field of Uncertainty Quantification two types of uncertainties can be considered: aleatoric and epistemic uncertainty. The first describes uncertainties that are inherent to the problem ('irreducible') and are often modeled using statistical distributions, such as random permeability fields or random boundary conditions. Epistemic uncertainties describe uncertainties that are inherent to the model ('reducible'), such as model errors, assumptions, and numerical errors. In this talk we take a relatively new approach in which epistemic uncertainties in model parameters are addressed through Bayesian model calibration. The resulting distributions can be propagated through the model to make predictions under uncertainty. Although supported with extensive theory, Bayesian model calibration requires many model runs, which can make it unfeasible for computationally demanding problems such as encountered in CFD. To significantly reduce the computa-

tional cost of the calibration procedure, we take a new approach of using a surrogate of the model instead of the full model. Although the surrogate cannot describe all details of the full model, it can result into good estimates of the parameter distributions under suitable conditions. The application considered here is wind turbine wake simulation. Engineering models describing the shape of the wake are calibrated using data generated from full Navier-Stokes simulations to show the potential of the procedure.

Laurent van den Bos, Benjamin Sanderse  
Centrum Wiskunde & Informatica (CWI)  
Amsterdam, the Netherlands  
l.m.m.van.den.bos@cw.nl, b.sanderse@cw.nl

### MS69

#### Block Matrix Formulations for Evolving Networks

Many types of pairwise interaction take the form of a fixed set of nodes with edges that appear and disappear over time. In the case of discrete-time evolution, the resulting evolving network may be represented by a time-ordered sequence of adjacency matrices. More precisely, let  $\{G^{[k]}\}_{k=1}^M = (V, \{E^{[k]}\}_{k=1}^M)$  be a sequence of unweighted graphs evolving in discrete time, that is, the set of nodes  $V$ , with  $|V| = n$ , is fixed and the dynamism is given by the change in the set of the arcs. With this notation, given the ordered sequence of time points  $\{t_k\}_{k=1}^M$ , the network at time  $t_k$  is represented by its  $n \times n$  adjacency matrix  $A^{[k]}$ . As usual for unweighted networks, the  $(i, j)$ th entry of  $A^{[k]}$  equals 1 if there is an edge from node  $i$  to node  $j$  at time  $t_k$ , and 0 otherwise. We consider here the issue of representing the system as a single, higher dimensional block matrix, built from the individual time-slices. We focus on the task of computing network centrality measures, and present a particular block formulation that allows us to recover dynamic centrality measures [Grindrod, Higham, Parsons, E. Estrada (2011), Grindrod, Higham (2013)] respecting times arrow.

Caterina Fenu

University of Cagliari, Italy  
kate.fenu@gmail.com

Desmond Higham  
University of Strathclyde, UK  
d.j.higham@strath.ac.uk

### MS69

#### The Radau-Lanczos Method for Matrix Functions

To compute  $f(A)b_i$ ,  $A \in \mathbb{C}^{n \times n}$ ,  $b \in \mathbb{C}^n$ ,  $i = 1, \dots, k$ , block Krylov methods can be particularly efficient. The reason is that one approximates from the sum of all individual Krylov spaces, and this richer subspace may allow to get good approximations at earlier stages. In addition, *block reduction* can occur when individual subspaces start to have non-trivial intersection. On HPC architectures, block methods have the additional advantage that  $A$  must be loaded from memory only once for an entire matrix-block vector multiplication  $AB$ ,  $B = [b_1 \mid \dots \mid b_k]$ . This can have significant impact on the wall clock time. However, block Krylov methods also come with an additional cost, typically spent in orthogonalizations and other additional arithmetic, resulting in an overhead of  $\mathcal{O}(k^2n + k^3)$ . We present a convergence analysis of block Krylov subspace methods, including restarting. The focus will be on  $f$  being a Stieltjes function and  $A$  being hermitian and positive

definite. We will consider alternatives to block methods which preserve the advantage of relying on matrix block-vector multiplications, but do not extract the individual iterates from the sum of the subspaces. This avoids the additional overhead at the expense of not taking advantage of the other Krylov subspaces. We present results of a computational study which shows that nevertheless, these methods can be competitive.

Andreas J. Frommer

Bergische Universitaet Wuppertal  
 Fachbereich Mathematik und Naturwissenschaften  
 frommer@math.uni-wuppertal.de

**MS69**

**Approximating the Spectral Sums of Large-Scale Matrices using Chebyshev Approximations**

Computation of the trace of a matrix function plays an important role in many scientific computing applications, including applications in machine learning, computational physics (e.g., lattice quantum chromodynamics), network analysis and computational biology (e.g., protein folding), just to name a few application areas. We propose a linear-time randomized algorithm for approximating the trace of matrix functions of large symmetric matrices. Our algorithm is based on coupling function approximation using Chebyshev interpolation with stochastic trace estimators (Hutchinson's method), and as such requires only implicit access to the matrix, in the form of a function that maps a vector to the product of the matrix and the vector. We provide rigorous approximation error in terms of the extremal eigenvalue of the input matrix, and the Bernstein ellipse that corresponds to the function at hand. Based on our general scheme, we provide algorithms with provable guarantees for important matrix computations, including log-determinant, trace of matrix inverse, Estrada index, Schatten  $p$ -norm, and testing positive definiteness. We experimentally evaluate our algorithm and demonstrate its effectiveness on matrices with tens of millions dimensions.

Insu Han

KAIST  
 hawki17@kaist.ac.kr

Dmitry Malioutov  
 IBM Research  
 dmalioutov@us.ibm.com

Haim Avron  
 Tel Aviv University  
 haimav@post.tau.ac.il

Jinwoo Shin  
 Dept of EE.  
 KAIST  
 jinwoos@kaist.ac.kr

**MS69**

**On the Computation of the Action of the Fréchet Derivative**

The computation of the action of the Fréchet derivative of a matrix function can be a computationally challenging problem. One important application is the estimation of the condition number of the  $f(A)b$  problem. Many numerical algorithms require the computation of the action of matrix functions. The matrix exponential, for example,

is required in exponential integrators, splitting methods, or network analysis. For such problems it is important to understand the condition number to obtain an estimate of the accuracy that can be expected. On the other hand, the estimate of the condition number of  $f(A)b$  requires an efficient computation of  $L_f(A, E)b$  for multiple  $E$ . We introduce a block Krylov method for the computation of the Fréchet derivative of a matrix function on a vector. The algorithm is most efficient for low-rank direction matrices  $E$  and our main focus lies in  $E$  of rank-1. We introduce an a priori error analysis that allows a rigorous stopping criteria for a given tolerance. The efficiency of the approach is shown in several numerical experiments.

Peter Kandolf

Universitat Innsbruck  
 peter.kandolf@uibk.ac.at

Samuel Relton  
 University of Manchester, UK  
 samuel.relton@manchester.ac.uk

**MS70**

**Computing Failure Probability Using Models of Different Fidelities**

Computing failure probability is an important task in many fields such as reliability analysis and risk management. In this talk we propose a multi-fidelity method for computing failure probability. The method is based on Monte Carlo sampling and it utilizes both high- and low-fidelity models. In most cases, the majority of the samples are evaluated using a low-fidelity model, and only a small number of samples are evaluated using the original high-fidelity model. The advantage of the method is that it can achieve a similar accuracy of brute force Monte Carlo sampling with a substantially reduced cost. The purpose of the paper is to present a general theoretical framework for the multi-fidelity method and to develop practical algorithms under the framework. Different types of low-fidelity models are considered in the examples to demonstrate the accuracy and efficiency of the method.

Yuhang Chen  
 University of Utah  
 chen.7169@osu.edu

Jing Li  
 Pacific Northwest National Lab  
 jing.li@pnl.gov

Dongbin Xiu  
 Ohio State University  
 xiu.16@osu.edu

**MS70**

**Multilevel-Multifidelity Approaches for Uncertainty Quantification and Design**

In the simulation of complex physics, multiple model forms of varying fidelity and resolution are commonly available. In computational fluid dynamics, for example, common model fidelities include potential flow, Reynolds-averaged Navier Stokes, and large eddy simulation, each potentially supporting a variety of spatio-temporal resolution/discretization settings. In this presentation, we focus on novel algorithms that simultaneously exploit multiple model forms and multiple resolutions, both for uncertainty quantification (UQ) and for optimization under

uncertainty (OUU). These hybrid methods exploit multi-fidelity methods across the model form hierarchy in combination with multilevel accelerators across an associated discretization hierarchy, manifesting as multilevel-control variate methods in the UQ case and trust-region multi-grid optimization in the OUU case. These techniques will be demonstrated for both model problems and engineered systems, including high-performance aircraft nozzles and scramjets.

Michael S. Eldred  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Dept.  
mseldre@sandia.gov

Jason Monschke  
Texas A&M  
jam4375@gmail.com

John D. Jakeman  
Sandia National Labs  
jdjakem@sandia.gov

Gianluca Geraci  
University of Stanford  
ggeraci@stanford.edu

#### MS70

##### **Numerical Strategy for Model Correction Using Physical Constraints**

we present a strategy for correcting model deficiency using observational data. We first present the model correction in a general form, involving both external correction and internal correction. The model correction problem is then parameterized and casted into an optimization problem, from which the parameters are determined. More importantly, we discuss the incorporation of physical constraints from the underlying physical problem. The proposed strategy helps to improve the model prediction. It can also be used for model calibration to learn the model parameters. Several representative examples are presented, where the physical constraints take very different forms. Numerical tests demonstrate that the physics constrained model correction is an effective way to address model-form uncertainty.

YanYan He  
Scientific Computing and Imaging Center  
University of Utah  
yhe@sci.utah.edu

Dongbin Xiu  
Ohio State University  
xiu.16@osu.edu

#### MS70

##### **Stochastic Reduced Model for Multiple Scales in Porous Media**

Within the context of multi-fidelity, we consider three alternative computational models of pore-scale simulations of flow and transport, in which the geometry at pore-scale is changing due to reactions or phase transitions. The first approach is physics based, and requires the solution of flow problems with  $O(30M)$  unknowns and  $O(100K)$  time steps. The second, based on pore-network modeling, is much less complex, but is a rather crude representation of

the real physics. The third approach which we recently developed is stochastic and it constructs corrections for the pore-network models using detailed local offline simulations which are upscaled to provide probability distributions for the pore-network data. This third approach combines the advantages of the other two approaches, and we present the path forward.

Malgorzata Peszynska  
Department of Mathematics  
Oregon State University  
mpesz@math.oregonstate.edu

Timothy B. Costa  
Intel Corporation  
timothy.b.costa@intel.com

#### MS71

##### **The Earthquake Location Methods Based on Waveform Inversion**

Earthquake location is a fundamental and essential problem in seismology. However, quantities of location methods are based on the ray theory, which has low accuracy since the wavelength is not small enough comparing to the scale of the wave propagation region. We developed the adjoint waveform inversion for the earthquake location to accurately determine the earthquake's hypocenter and origin time. Due to the limitation of the iteration method for optimization problem and the singularity property of the delta function, it is not trivial to estimate the original hypocenter and origin time which would converge to the correct result. We develop a robust numerical method to greatly expand the range of convergence. Numerical experiments are presented to demonstrate the effectiveness of our method. An alternative project is introducing the auxiliary functions, whose zero set contains the real earthquake hypocenter and origin time. The benefit of this method is that the computational cost of constructing the auxiliary functions is comparable to one-step iteration of the classical waveform based adjoint method for earthquake location. Thus, it is possible for us to determine the earthquake hypocenter and origin time extremely fast and accurately.

Jing Chen, Hao Wu  
Tsinghua University  
cjcjyu597460474@163.com, hwu@mail.tsinghua.edu.cn

Xueyuan Huang  
Beijing Technology and Business University  
huangxy.math10@gmail.com

Dinghui Yang  
Tsinghua University  
dhyang@math.tsinghua.edu.cn

#### MS71

##### **Numerical Methods for Solving the Fractional Schroedinger Equation**

The fractional Schrödinger equation is a non-local partial differential equation which can describe new phenomena that are absent from the standard Schrödinger equation. However, the nonlocality introduces considerable challenges in solving this equation both numerically and analytically. In this talk, we will present several mass and/or energy conservative methods for solving the fractional Schrödinger equation. The accuracy and the dis-

persion relation of these methods are studied analytically. Also, we numerically study the dynamics of the plane waves and solitons, the performance of these methods are compared and discussed.

Siwei Duo

Missouri University of Science and Technology  
sddy9@mst.edu

**MS71**

**Solid-State Dewetting: Equilibrium and Dynamics**

Solid-state dewetting (or called agglomeration) of thin films during annealing has been simultaneously observed in various thin film materials, and has attracted ever more attention because of its extensive technological applications and important scientific interest. In this talk, we will talk about two important issues related with theoretical studies on solid-state dewetting problems: one is how to predict the equilibrium shapes of thin island films (i.e., the equilibrium problems); the other is how to derive mathematical models for describing the morphology evolution of thin solid films (i.e., the dynamical problems).

Wei Jiang

National University of Singapore and Wuhan University  
jiangwei1007@whu.edu.cn

**MS71**

**Multiscale Methods for Highly Oscillatory Quantum and Kinetic Problems**

In this talk, we will concern some multi-scale methods for solving the temporally highly oscillatory differential equations. The method we are going to address is the recent-developed two-scaled formulation approach which separates the fast and the slow variables in the problem. The method is proposed in general context to tackle with two classes of typical oscillatory regimes: the highly oscillatory regime and the diffusion regime. Through this approach, the uniformly accurate integrators with respect to the oscillation frequency can be designed. We shall present the applications of the two-scaled method to solve the Vlasov-Poisson equation and the nonlinear Dirac equation.

Xiaofei Zhao, Mohammed Lemou

IRMAR, University of Rennes 1  
zhxfnus@gmail.com, mohammed.lemou@univ-rennes1.fr

Florian Mehats

IRMAR - Université de Rennes 1  
florian.mehats@univ-rennes1.fr

Nicolas Crouseilles

INRIA-Rennes and IPSO project  
nicolas.crouseilles@inria.fr

**MS72**

**Spectral Accuracy in Time and Space for a Semi-Implicit Approach to the Incompressible Navier-Stokes Equations in Complex Geometries**

I will report on a project which combines a discontinuous Galerkin spatial discretization framework suitable for problems in complex geometry with a semi-implicit Spectral Deferred Correction method for the temporal integration of the incompressible Navier-Stokes equations. The temporal approach utilizes a gauge or auxiliary variable formulation

of the equations of motions wherein the evolved velocity is not divergence free, but the divergence free part satisfies the Navier-Stokes equations. Several different possibilities for constructing the auxiliary variable approximations will be discussed, and particular attention will be given to the choice of boundary conditions required in the implicit treatment of diffusive terms and the discrete projection operators.

Michael Minion

Lawrence Berkeley National Lab  
Stanford University  
mlminion@lbl.gov

Robert Saye

Dept. of Mathematics  
University of California, Berkeley  
rsaye@lbl.gov

Robert Saye

Lawrence Berkeley National Laboratory  
rsaye@lbl.gov

**MS72**

**Diagonally Implicit Runge-Kutta Schemes Devoid of Order Reduction**

Order reduction is a generic phenomenon that arises for time-dependent PDE boundary value problems, as well as for stiff ODEs. One existing approach to avoid order reduction is to construct schemes with high stage order. Unfortunately, high stage order is incompatible with diagonally implicit Runge-Kutta (DIRK) schemes. We show that a relaxed property, called weak stage order, can also achieve the goal. Even more, it is compatible with DIRK schemes. We finish with a systematic discussion of DIRK schemes (up to fourth order) that are devoid of order reduction.

Benjamin Seibold

Temple University  
seibold@temple.edu

Rodolfo R. Rosales

Massachusetts Inst of Tech  
Department of Mathematics  
rrr@math.mit.edu

David Shirokoff

New Jersey Institute of Technology  
david.g.shirokoff@njit.edu

Dong Zhou

Temple University  
dong.zhou@temple.edu

**MS72**

**Unconditional Stability for Multistep Imex Schemes**

In this talk we introduce a new class of linear multistep ImEx schemes that have very good unconditional stability properties. Unconditional stability is a desirable property of a time stepping scheme, as it allows the choice of time step solely based on accuracy considerations. Of particular interest are problems for which both the implicit and explicit parts of the ImEx splitting are stiff. Such splittings can arise, for example, in variable-coefficient problems, or the incompressible Navier-Stokes equations. To charac-

terize the new ImEx schemes, we introduce an unconditional stability region, which plays a role analogous to that of the stability region in conventional multistep methods. Moreover, we will show how the new diagrams explain the fundamental stability restrictions of the well-known semi-implicit backward differentiation formulas (SBDF). We further show that the new ImEx coefficients can overcome the limitations of SBDF, and highlight their utility with several examples arising from partial differential equations: such as variable diffusion, advection diffusion and, time permitting a time dependent Stokes equation.

David Shirokoff

New Jersey Institute of Technology  
david.g.shirokoff@njit.edu

### MS72

#### Order Reduction in Implicit Runge-Kutta Time-Stepping for Initial Boundary Value Problems

When advancing a time-dependent PDE forward via Runge-Kutta methods, one may observe a convergence order that is less than the actual order of the scheme. We demonstrate that this order reduction phenomenon is in fact the norm, and not the exception. Geometrically, it stems from boundary layers, produced by the fact that the scheme is too accurate near the boundary. A modal analysis reveals under which circumstances boundary layers persist over many time steps. Moreover, a systematic derivation of modified boundary conditions is presented, which remedy the order reduction phenomenon.

Dong Zhou

Temple University  
dong.zhou@temple.edu

Rodolfo R. Rosales  
Massachusetts Inst of Tech  
Department of Mathematics  
rrr@math.mit.edu

Benjamin Seibold  
Temple University  
seibold@temple.edu

David Shirokoff  
New Jersey Institute of Technology  
david.g.shirokoff@njit.edu

### MS73

#### Unfitted Finite Element Methods for PDEs on Evolving Surfaces

In this talk we review several approaches to solving PDEs on surface. We focus on an Eulerian finite element method, known as the trace FEM, for the discretization of parabolic partial differential equations on evolving surfaces. The method uses traces of volume finite element space functions on a surface to discretize equations posed on the surface. The approach is particularly suitable for problems in which the surface is given implicitly by a level set function, may undergo topological changes, and for surface-bulk coupled problems. We discuss several known variants of the method, including DG space-time finite elements and FD in time – trace FE in space. We review available error analysis, adaptivity and algebraic properties of the method.

Maxim A. Olshanski

Department of Mathematics  
University of Houston  
molshan@math.uh.edu

### MS73

#### A High Order Trace Finite Element Method for PDEs on Level Set Surfaces

We present a new high order finite element method for the discretization of partial differential equations on stationary smooth surfaces which are implicitly described as the zero level of a level set function. The discretization is based on a trace finite element formulation. Achieving higher order accuracy with trace finite elements for these problems is challenging mainly due to two reasons. First, a sufficiently accurate geometry approximation on which integrals can be evaluated accurately is required. Secondly, the irregular shape of the implicit domains requires special attention in order to control the conditioning of arising linear systems. We present and analyse a method which addresses both issues. Crucial ingredients are the mesh transformation introduced in [C. Lehrenfeld, *High order unfitted finite element methods on level set domains using isoparametric mappings*, Comp. Meth. Appl. Mech. Engrg. 2016] and a volume stabilization with an anisotropic diffusion.

Arnold Reusken

Numerical Mathematics  
RWTH Aachen University, Aachen, Germany  
reusken@igpm.rwth-aachen.de

Joerg Grande  
RWTH Aachen University, Aachen, Germany  
grande@igpm.rwth-aachen.de

Christoph Lehrenfeld  
University of Muenster  
christoph.lehrenfeld@gmail.com

### MS73

#### An Embedding Method for PDEs on Moving Surfaces

Partial differential equations (PDEs) on surfaces arise in a wide range of applications. Embedding methods are a class of methods that solve such problems in a neighborhood of the surface using standard Cartesian grid methods. The closest point method [Ruuth and Merriman, *J. Comput. Phys.* 227(3):1943-1961, 2008] is an embedding method that has been used to solve a variety of PDEs on smooth surfaces using a closest point representation of the surface. The original closest point method (CPM) was designed for problems posed on static surfaces, however the solution of PDEs on moving surfaces is of considerable interest as well. Here we propose solving PDEs on moving surfaces using a combination of the CPM and a modification of the grid based particle method [Leung and Zhao, *J. Comput. Phys.* 228(8):2993-3024, 2009]. The grid based particle method (GBPM) represents and tracks surfaces using meshless particles and an Eulerian reference grid. A modification is introduced into the reconstruction step of the GBPM to ensure that all the grid points within the computational tube surrounding the surface are active. Both finite difference and radial basis function discretizations are introduced and compared. A number of examples are presented to illustrate the numerical convergence properties of our methods. These include experiments for advection-diffusion and higher-order equations that are strongly cou-



pled to the velocity of the surface.

Argyrios Petras, [Steven Ruuth](#)  
Mathematics  
Simon Fraser University  
apetras@sfu.ca, sruuth@sfu.ca

#### MS74

##### **SENSEI: A Lightweight In Situ Interface for Contemporary Infrastructure Tools and Architectures**

SENSEI ([www.sensei-insitu.org/](http://www.sensei-insitu.org/)) is a new open source *in situ* infrastructure that allows simulations to easily choose between the popular existing open source *in situ* infrastructures ParaView Catalyst, VisIt Libsim and ADIOS. Additional *in situ* infrastructures can easily be added as well. The SENSEI interface relies on the VTK ([www.vtk.org](http://www.vtk.org)) data model to provide consistency and interoperability between other infrastructures allowing a more general and efficient workflow for data scientists. Additionally, the infrastructure can be configured to be extremely lightweight through small library sizes and memory reuse to ensure the most efficient use of HPC resources. This infrastructure has been used to generate *in situ* output on a 1 million MPI rank run on Argonnes Mira BG/Q.

[Andrew C. Bauer](#)  
Kitware Inc.  
andy.bauer@kitware.com

Patrick O'Leary  
Kitware, Inc., USA  
patrick.oleary@kitware.com

Utkarsh Ayachit  
Kitware Inc  
utkarsh.ayachit@kitware.com

#### MS74

##### **In Situ Processing Overview and Relevance to the HPC Community**

In response to a widening gap between our ability to compute and save data for subsequent use, the idea of doing analysis and other processing while data is still resident in memory has emerged as a viable alternative. Known under the umbrella term *In Situ* processing, the idea is to do as much with the data possible before writing any of it, or any derived products, to persistent storage. There is a rich history of work in this space going back several decades, with evolving R&D tracking changes in technology and needs of the computational science community. Recent work in this space includes production-quality *in situ* software infrastructures, methods for computing reduced-sized extracts, efforts to maximize software reusability across applications, and examples of application of these methods to different science problems at scale.

[E. Wes Bethel](#)  
Lawrence Berkeley National Laboratory  
ewbethel@lbl.gov

#### MS74

##### **Master of Puppets: Cooperative Multitasking for In Situ Processing**

Modern scientific and engineering simulations track the time evolution of billions of elements. For such large runs,

storing most time steps for later analysis is not a viable strategy. It is far more efficient to analyze the simulation data while it is still in memory. We present a novel design for running multiple codes in situ: using coroutines and position-independent executables we enable cooperative multitasking between simulation and analysis, allowing the same executables to post-process simulation output, as well as to process it on the fly, both in situ and in transit. We present Henson, an implementation of our design, and illustrate its versatility by tackling analysis tasks with different computational requirements, including an iterative workflow driven by a surrogate model. Our design differs significantly from the existing frameworks and offers an efficient and robust approach to integrating multiple codes on modern supercomputers. The presented techniques can also be integrated into other in situ frameworks.

[Erich Lohrmann](#)  
Georgia Tech  
erichlohmann@gmail.com

Zarija Lukic  
Lawrence Berkeley National Laboratory  
zarija@lbl.gov

Dmitriy Morozov  
Lawrence Berkeley National Lab  
dmitriy@mrzv.org

Juliane Mueller  
Lawrence Berkeley National Laboratory  
julianemueller@lbl.gov

#### MS74

##### **Extract Based In Situ Visualization and Analysis Methods and Application to Engineering Simulations**

In situ based data extracts are a subset of the volume data consisting of surfaces (slices, boundary or iso-values), lines, points or sub-volumes that have been created as the solver executes. Extracts can be orders of magnitude smaller and consequently faster to load and explore via ad hoc post-processing methods. This paper will focus upon two case studies (a helicopter landing on a ship and a flame undergoing extinguish/re-ignition) where extract based in situ visualization and analysis was used to accelerate the overall workflow. The helicopter landing problem used the Kestrel CFD code using the Libsim library for creating surface extracts that were subsequently used to create movies of the complex flowfield. The second case used the AVF-LESLIE code using the Sensei in situ infrastructure coupled to Libsim to generate surface extracts. The paper highlights how in situ processing benefited these engineering applications.

Earl P. N. Duque  
Intelligent Light  
epd@ilight.com

[Brad Whitlock](#)  
Lawrence Livermore National Laboratory  
whitlock2@llnl.gov

#### MS75

##### **Data-Driven Methods for Nonlinear Model Reduc-**

## tion in Fluid Dynamics

This talk will explore data-driven methods to extract nonlinear models and features in fluid dynamics. We will show that the elements of an identified finite dimensional approximation to the Koopman operator, obtained using extended dynamic mode decomposition, can be utilized for the construction of accurate nonlinear reduced-order models. The performance of models identified using our proposed data-driven method are compared to those found by performing a Galerkin projection of the governing equations onto proper orthogonal decomposition modes. We demonstrate that identifying nonlinear models using this method is particularly advantageous when the dataset is noisy or is confined within limited regions of space or time, or when additional properties of the governing equations are known in advance. While the focus will be on application in fluid dynamics, the techniques are sufficiently general to apply to a range of other applications.

Scott Dawson

Department of Mechanical and Aerospace Engineering  
Princeton University  
stdawson@princeton.edu

Clarence Rowley

Princeton University  
Department of Mechanical and Aerospace Engineering  
cwrrowley@princeton.edu

## MS75

### Regularization Techniques for Biochemical Reaction Networks

Models of biochemical reaction networks often contain a large number of parameters that cannot all be estimated from the limited amount of noisy experimental data usually available. Therefore, many of the nominal parameter values must be obtained from the open scientific literature. Differences in cell type, organism, etc. add large amounts of uncertainty to these literature-sourced parameters for use in the particular system to be modeled. There clearly is a need to estimate at least some of the parameter values from experimental data, however, the small amount of available data and the large number of parameters commonly found in these types of models, require the use of regularization techniques to avoid over-fitting. Parameter set selection comprises a group of related methods that seek to choose only a small subset of model parameters for estimation that will best represent the current data set and predict future ones. Parameter set selection achieves model reduction by retaining the less important parameters at their nominal values, thereby reducing the number of variable parameters in the model. A parameter sensitivity technique employing local sensitivity analysis and hierarchical clustering will be presented, along with more advanced extensions to explicitly incorporate parameter uncertainties. These techniques will then be evaluated on two biochemical reaction networks: a signal transduction example and a pharmacokinetic/pharmacodynamic type 1 diabetes example.

Daniel Howsmon, Juergen Hahn  
Rensselaer Polytechnic Institute  
Hahn Research Group  
howsm@rpi.edu, hahnj@rpi.edu

## MS75

### Model Reduction Via Spectral Analysis of the

## Koopman Operator

Associated with a nonlinear dynamical system  $T : X \rightarrow X$  is a linear infinite-dimensional operator, the Koopman operator, which evolves (vector-valued) observables  $f : X \rightarrow C^n$  on the state space via the composition  $f \mapsto f \circ T$ . As in the case of a finite-dimensional linear operator, the vector-valued observable can be decomposed into its so-called Koopman modes, which are the vector-valued coefficients of a single eigenfunction, plus the contribution due to the continuous part of the spectrum, which is particular to the infinite-dimensional setting. These modes can be computed directly from data, without direct knowledge of the form of  $T$ , using either Fourier averages or one of the variants of the Dynamic Mode Decomposition (DMD) algorithm. Truncating these modes leads to a reduced order model, similar to what is done with Proper Orthogonal Decomposition (POD). The Koopman modes, however, have well defined temporal behavior since they are associated with a single eigenvalue and thus are more dynamically relevant than POD modes.

Ryan Mohr

University of California, Santa Barbara  
mohrrm@engineering.ucsb.edu

## MS75

### Nonlinear Model Order Reduction with pyMOR

pyMOR [Milk, Rave, Schindler 2016] is a modern open source software library of model order reduction algorithms written in the Python programming language. Through its interface-based design, pyMOR combines rapid prototyping of reduction algorithms in a dynamic programming language with easy integration of the very same algorithms with external high-performance PDE solver codes. In this talk I will give a short overview on pyMOR's design and then go into more detail on how the treatment of nonlinearities with the empirical interpolation method is accomplished within this design approach.

Stephan Rave

University of Muenster  
stephan.rave@uni-muenster.de

## MS76

### The Unified Form Language and Key Points on its Translation

The Unified Form Language (UFL) is an embedded domain specific language for specification of variational formulations of partial differential equations [Alnæs et al 2014]. Key features include finite element concepts such as function spaces and test functions, as well as more generic concepts such as differential operators, tensor valued expressions, and index notation. Implemented as a Python module, UFL also includes symbolic algorithms including algorithmic differentiation of expressions and functionals, symbolic application of geometry mappings and Piola mappings, and propagation of restriction operators to expression terminals. The first part of this talk is a brief introduction to key concepts in UFL. UFL is used in finite element frameworks to achieve a high level problem specification, but to achieve high performance finite element matrix assembly a form compiler is needed. Several form compilers exist that can produce efficient low level source code for computing e.g. the element matrix from a UFL form. In the translation from UFL to low level code, a form compiler must choose strategies both to map domain specific

concepts from UFL to constructs supported by the target language, and strategies to achieve good performance. In the second part of this talk, how these topics are handled in the uflacs form compiler via value numbering of scalar subexpressions will be discussed.

Martin Alnæs  
Simula Research Lab  
martinal@simula.no

#### MS76

##### The Two-Stage Form Compiler, a Mechanism for Structure-Preserving Code Generation

TSFC is a new form compiler in use in Firedrake, an automated system for the solution of PDEs using the finite element method. The high-level interface of Firedrake uses the Unified Form Language (UFL) of the FEniCS project to allow expression of complex variational forms in a language that is close to that of the underlying mathematics. A core component of the software is the form compiler, that takes this high-level description and produces the low-level code for finite element assembly. The core concept of TSFC is to maintain the structure of the input expression as long as possible. This facilitates the application of optimisations at the highest possible level of abstraction. TSFC uses a tensor algebra language as an intermediate representation which cleanly separates the lowering of finite element constructs to tensor algebra on the one hand, from the scheduling of those tensor operations on the other. This facilitates the optimisation of this schedule either in situ or by a downstream tool, such as COFFEE, an optimizer for finite element kernels in Firedrake. We also plan to exploit structure in the finite elements themselves, such as the sum factorisation of tensor product elements. This structure will be provided by FInAT, an abstract, smarter library of finite elements.

Miklós Homolya, Lawrence Mitchell, Fabio Luporini  
Department of Computing  
Imperial College London  
miklos.homolya14@imperial.ac.uk,  
lawrence.mitchell@imperial.ac.uk,  
f.luporini12@imperial.ac.uk

David Ham  
Department of Mathematics  
Imperial College London  
david.ham@imperial.ac.uk

#### MS76

##### Generating Performance-Optimized Fem Assembly Kernels for Dune-Pdelab

dune-pdelab is a well-established discretization framework for the numerical solution of PDEs based on the Dune project. Optimizing the performance of finite element assembly for modern HPC systems is currently investigated in the EXADUNE project. While the project shows promising results, a necessity for hardware-dependent low-level programming has become apparent. Code generation techniques are a possible solution to achieve performance portability and make those results available to a broader audience. We base our approach on the Unified Form Language (UFL), developed by the FEniCS project. A special focus of our approach is the choice of an intermediate representation, which can both reflect structural information of the FEM problem and be subject to hardware-dependent transformation. This intermediate representation is based

on loopy [Klckner, transformation-based code generation for GPUs and CPUs]. We present the current state of code generation for dune-pdelab, showing examples covering a wide range of applications.

Dominic Kempf, René He  
IWR, Heidelberg University  
dominic.kempf@iwr.uni-heidelberg.de,  
rene.hess@iwr.uni-heidelberg.de

Steffen Müthing  
Heidelberg University  
steffen.muething@iwr.uni-heidelberg.de

Peter Bastian  
Interdisciplinary Center for Scientific Computing  
University of Heidelberg  
peter.bastian@iwr.uni-heidelberg.de

#### MS76

##### Using Code Generation to Improved Performance of a Multipurpose Discontinuous Galerkin Implementation Based on Dune-Fem

Abstract not available

Robert Klöforn  
International Research Institute of Stavanger  
robert.kloeforn@iris.no

#### MS77

##### Preconditioning on Parallel and Hybrid Architectures

The inherently sequential nature of sparse triangular solves can make exact ILU preconditioning unattractive on highly parallel systems like GPUs or Intel's KNL. We consider the approximate solution of the triangular systems by replacing the substitutions with the multiplication with an incomplete sparse approximate inverse (ISAI) for the triangular factors. Beside the impact of the diminished preconditioner accuracy, we in particular focus on the efficient generation of the incomplete sparse approximate inverses via batched routines.

Hartwig Anzt  
Innovate Computing Lab  
University of Tennessee  
hantz@icl.utk.edu

Thomas K. Huckle  
Institut fuer Informatik  
Technische Universitaet Muenchen  
huckle@in.tum.de

Jack J. Dongarra  
University of Tennessee, Oak Ridge National Laboratory,  
USA  
dongarra@icl.utk.edu

#### MS77

##### Efficient MPI-Openmp Programming for Scientific Applications

High-performance computing systems are using an increasing number of cores per node, which impacts both current MPI programming paradigms and motivates developing new methods for achieving improved performance using hy-

brid MPI+X approaches. In this talk we give an overview of the performance expectations in a multi-core setting, showing how to more accurately model inter-node communication in the case of multiple communicating cores. In addition, we show progress in using MPI+OpenMP effectively and highlight the challenges facing current MPI shared memory implementations.

William D. Gropp  
University of Illinois at Urbana-Champaign  
Dept of Computer Science  
wgropp@illinois.edu

Luke Olson, Philipp Samfass  
Department of Computer Science  
University of Illinois at Urbana-Champaign  
lukeo@illinois.edu, samfass2@illinois.edu

### MS77

#### Extending Kokkos for Task-Parallel and Structured Grid Algorithms with Asynchronous Parallelism

New advances in the Kokkos library extend support for a wider variety of parallel algorithms and runtime systems. In addition to the traditional one-level parallel for, reduce, and scan patterns, Kokkos now enables hierarchical data parallelism in these patterns through a ‘thread team’ interface. Kokkos also has fine-grained task parallelism that has high performance on many different architectures including NVIDIA GPUs. Furthermore, the new multidimensional range policy makes it easier to iterate over structured grids in a cache-efficient way. Lastly, support for asynchronously executing parallel algorithms on a user-specified subset of the hardware resources is available through the new ‘execution space instance’ API. These advances make it possible for asynchronous many-task codes to use Kokkos more easily.

Daniel Sunderland  
Sandia National Laboratories  
dsunder@sandia.gov

### MS77

#### OCCA : Open Concurrent Compute Abstraction

Most gains in microprocessor performance over the last decade have stemmed from increased on-chip parallelism. Modern CPUs from AMD and Intel have multiple compute cores each equipped with wide vector floating point units. High-end discrete graphics processing units (GPUs) from AMD and NVIDIA have in excess of 2000 processing elements. Differences in the non-uniform memory hierarchies and core architectures of CPUs and GPUs have resulted in the development of several different programming models like OpenMP, OpenCL, NVIDIA’s CUDA, Intel’s Thread Building Blocks, Cilk, amongst others. The OCCA project is an effort to build a unified interface that enables a moderately capable programmer to write portable code that can execute in native mode on most many-core microprocessors and potentially attain high computational efficiency. Examples applications involving the numerical solution of PDEs will illustrate the potential for developing portable and performance simulation tools using OCCA.

Tim Warburton  
Virginia Tech  
tcw@vt.edu

David Medina

Rice University  
dmed256@gmail.com

### MS78

#### Recent Development of the Parallel Selected Inversion Method on Many-Core Architecture

The parallel selected inversion method (PSelInv) has been demonstrated to be scalable to thousands of CPU cores in the MPI environment. In this talk, we will discuss the recent development of the PSelInv method on the Intel Knights Landing (KNL) many core architecture in the hybrid MPI+OpenMP environment. We find that it is crucial to use the OpenMP task feature to distribute the memory access and workload between KNL tiles. We also find that in order to obtain high parallel scalability, it is important to control the granularity of the computational tasks.

Weile Jia  
University of California at Berkeley  
wjia@lbl.gov

### MS78

#### A Hybrid Method for Computing Selected Entries of Matrix Inverse

Many problems require computing a subset of the inverse of sparse matrix. All efficient direct algorithms fall into two categories: those based on both Gaussian elimination and back substitution (two-way algorithms), and those solely based on repeated Gaussian elimination (one-way algorithms). The proposed hybrid algorithm combines one-way algorithms and two-way algorithms, to take advantages of the algorithms in both categories and avoid their disadvantages. This is achieved by executing its upward pass in one-way mode, while its downward pass in a mixture of one-way and two-way modes, followed by some special treatment at the end of the entire process. The exact way of mixture can be tailored for specific application to achieve the best performance

Song Li  
Stanford University  
lisong.pine@gmail.com

### MS78

#### Selected Inversion in Quantitative Genetics

Selected inverse of the coefficient matrix of linear equations is typically required in a quasi-Newton procedure for restricted maximum likelihood (REML) in quantitative genetics. The diagonal elements of the inverse are used by geneticists and animal breeders to evaluate the error in predicted genetic merit of economic characters in farm animals. Sparse structure of the matrix depends on a statistical model defined by a user. Quantitative geneticists often assume complicated models with several biological traits and genetic relationships among a large number of individuals. In this case, the matrix contains small but many dense blocks that often form a large, single dense block in the Cholesky factor. The classical Takahashi algorithm is inefficient in this situation, but it is still used in statistical software for REML. In this talk, we present a case study to accelerate the computations in selected inversion with a supernodal approach with an extension to handle a matrix with dense blocks. The package, Yet Another MME Solver (YAMS), implements this approach and successfully reduces the running time in REML computations for var-

ious statistical models. We also discuss additional issues of the selected inverse found in quantitative genetics and animal breeding applications.

Yataka Masuda  
University of Georgia  
yutaka@uga.edu

#### MS78

##### Scalable Selective Inversion and Stochastic Estimation of Inverse Matrices

The modern mathematical models designed in many fields of study underline growing accuracy and complexity and demand therefore an increasing need for fast and distributed computation of the inverse of sparse matrices. Several disciplines, such as Greens functions formalism, maximum likelihood estimation, genomic prediction, and covariance estimation, need just a particular subset of the entries of the inverse, so that the use of the standard inversion process is discouraged in favor of a selective inversion algorithms. One interesting approach to the selective inversion is the stochastic estimation of the diagonal of the inverse, widely used in economics, finance, interferometry, and many others. We present algorithms that aim at the approximation of selected entries of the diagonal of the inverse of sparse matrices, based on stochastic estimators for the diagonal, incomplete (block) LU factorization and eigenvector deflation techniques, with particular attention to the applications. The performance and accuracy of the mentioned methods will be presented as evaluation of the provided results in the different scenarios and in relation to the performance of direct solvers (e.g., Pardiso [M. Luisier, O. Schenk et al., Fast Methods for Computing Selected Elements of the Green's Function in Massively Parallel Nanoelectronic Device Simulations, Euro-Par 2013]). Such algorithms find the best implementation in C++, leaving wide space to parallelization.

Fabio Verbosio  
USI - Universita della Svizzera italiana  
Institute of Computational Science  
fabio.verbosio@usi.ch

Matthias Bollhoefer  
TU Braunschweig  
m.bollhoefer@tu-bs.de

Olaf Schenk  
USI - Università della Svizzera italiana  
Institute of Computational Science  
olaf.schenk@usi.ch

#### MS79

##### Reproducibility of Stochastic Methods for Computing Material Properties

Abstract not available

Derek Juba  
NIST  
derek.juba@nist.gov

#### MS79

##### An Overview of Lessons Learned at the Numerical

##### Reproducibility at Exascale Workshops

Abstract not available

Walid Keyrouz  
National Institute of Standards and Technology  
Software and Systems Division  
walid.keyrouz@nist.gov

#### MS79

##### The Forensic Reproducibility of Pseudorandom Number Generation on Modern Architectures

Abstract not available

Michael Mascagni  
Computer Science/Computational Science  
Florida State University, Tallahassee  
mascagni@fsu.edu

#### MS79

##### The Numerical Reproducibility Fair Trade: Facing the Concurrency Challenges at the Extreme Scale

Abstract not available

Michela Taufer  
University of Delaware  
taufer@udel.edu

#### MS80

##### Parameter Identification of a Counter Flow Heat Exchanger with Physics Based Model-Form Discrepancy

In this paper we consider the problems of parameter estimation and model-form discrepancy for physics based modeling of heat exchangers. Typically, one assumes that convection dominates so that diffusion can be ignored and the resulting model is defined by a hyperbolic partial differential equations. We employ a full-flux physics model, which means the model is defined by convection-diffusion equations and the hyperbolic equations are viewed as one particular form of the model. We consider the case where one assumes this parameterized model is a source of model discrepancy and measurement errors are a source of noise. Physics-based hierarchical modeling is employed develop prior knowledge about model discrepancy. We show that this approach can greatly improve the model's predictive usefulness. A simple example is used to illustrate the ideas and then we apply the results to a parameter identification problem for a distributed parameter model of a counter flow heat exchanger.

John A. Burns, Eugene Cliff  
Virginia Tech  
Interdisciplinary Center for Applied Mathematics  
jaburns@vt.edu, ecliff@vt.edu

#### MS80

##### Determination of Model Form Uncertainties Using Field Inference and Machine Learning

We propose a modeling paradigm, based on field inversion and machine learning (FIML), that harnesses available data to aid the quantification of model discrepancy in a predictive computational physics setting. This work uses inverse modeling to obtain corrective, spatio-temporally

distributed functional terms, offering a route to directly address model-form errors. Once the inference has been performed over a number of problems that are representative of the deficient physics, machine learning techniques are used to reconstruct the model corrections in terms of features that appear in the closure model. These reconstructed functional forms are then used to augment the closure model in a predictive setting. The entire framework is cast in a consistent Bayesian setup, allowing for consistent uncertainty quantification. Results are presented in fluid flow and materials problems.

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

### MS80

#### **A Stochastic Operator Approach to Model Inadequacy with Applications to Contaminant Transport**

We present recent developments on the uncertainty quantification of models. Models are imperfect representations of complex physical processes, hence exploring representations of the model inadequacies is crucial. We introduce an inadequacy model in the form of a linear operator acting on the model solution and explore methods for incorporating knowledge of model shortcomings and relevant physics. This representation is developed in the context of scalar dispersion in porous media, but the methods presented are applicable for other models.

Teresa Portone  
ICES  
UT Austin  
portone@ices.utexas.edu

Damon McDougall  
Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
damon@ices.utexas.edu

Robert D. Moser  
University of Texas at Austin  
rmoser@ices.utexas.edu

Todd A. Oliver  
PECOS/ICES, The University of Texas at Austin  
oliver@ices.utexas.edu

### MS80

#### **Reducing Model Discrepancy in Turbulent Flow Simulations: A Physics-Informed Machine Learning Approach**

Reynolds-averaged Navier–Stokes (RANS) equations are widely used in the simulations of industrial turbulent flows. The closure term of RANS equations, the Reynolds stress term, needs to be modeled and is well known as an important source of model-form uncertainties/errors. Recently, several data-driven methods have been proposed, aiming to reduce the model-form uncertainties/errors in the Reynolds stress term. In this work, we propose a data-driven, physics-informed machine learning approach that predicts the discrepancies between RANS modeled Reynolds stresses and the ones obtained from DNS. Specifically, the Reynolds stress discrepancy are parameterized and treated as functions of selected mean flow features. By using the random forests regression, statistical model of the parameterized Reynolds stress discrepancy is first trained

on the flows for which DNS data are available. The obtained statistical model is then used to predict Reynolds stresses discrepancy in new flows. The flow in a square duct at a different Reynolds number is predicted by training on the flows in a square duct at several Reynolds numbers. In addition, the flow over periodic hills is predicted by training on the flows with several different hill profiles. The results demonstrate that the predicted Reynolds stresses of both cases have a better agreement with the DNS/LES data, indicating a reduction of model-form discrepancy in Reynolds stress term by using machine learning techniques.

Jinlong Wu, Jianxun Wang, Heng Xiao  
Dept. of Aerospace and Ocean Engineering, Virginia Tech  
jinlong@vt.edu, vtwxj@vt.edu, hengxiao@vt.edu

### MS81

#### **Point Spread Function Estimation Using Radially Symmetric Prior Modeling and Enhanced Gibbs Sampling**

The point spread function (PSF) of a translation invariant imaging system is its impulse response, which cannot always be measured directly. This is the case in high energy X-ray radiography, where the PSF is estimated indirectly using calibration imagery and solving a corresponding inverse problem. When the PSF is assumed to have radial symmetry, it can be estimated from an image of an opaque straight edge. We use a non-parametric Bayesian approach, where the prior probability density for the PSF is modeled as a Gaussian Markov random field and radially symmetry is incorporated in a novel way. Monte Carlo posterior estimation is carried out by adapting a recently developed improvement to the Gibbs sampling algorithm, referred to as partially collapsed Gibbs sampling (PCG). We will present estimation results on radiographic data obtained from a high-energy X-ray diagnostic system at the Nevada National Security Site. The algorithm we present is proven to satisfy invariance with respect to the target density. Moreover, new developments in radially symmetric prior modeling and potential further applications will also be presented.

Kevin Joyce  
University of Montana  
Department of Mathematical Sciences  
joycekt@nv.doe.gov

Johnathan M. Bardsley  
University of Montana  
bardsleyj@mso.umt.edu

Aaron B. Luttman  
National Security Technologies, LLC  
luttmaab@nv.doe.gov

### MS81

#### **Multilevel Monte Carlo Algorithms for Inference**

Recent years have seen an increasing interest in uncertainty quantification (UQ). Another recent trend is the explosion of available data. Bayesian inference provides a principled and well-defined approach to the integration of data into an a priori known distribution. The posterior distribution, however, is known only point-wise (possibly with an intractable likelihood) and up to a normalizing constant. Monte Carlo methods have been designed to sample such distributions, such as Markov chain Monte Carlo (MCMC) and sequential Monte Carlo (SMC) samplers. Recently,

the multilevel Monte Carlo (MLMC) framework has been extended to some of these cases, so that approximation error can be optimally balanced with statistical sampling error, and ultimately the Bayesian inverse problem can be solved for the same asymptotic cost as solving the deterministic forward problem. This talk will concern the recent development of multilevel SMC (MLSMC) samplers [Beskos et al, 2016] and the resulting estimators for standard quantities of interest as well as normalizing constants [Del Moral, et al, 2016]. The methods have been applied successfully to nonlocal equations [Jasra et al, 2016], which are used to model anomalous diffusion and fractures in materials. MLMC data assimilation methods have also been developed, which combine dynamical systems with data in an online fashion. Examples are ML particle filters [Jasra et al, 2016] and ensemble Kalman filters [Hoel et al, 2016ab].

Kody Law  
Oak Ridge National Laboratory  
kodylaw@gmail.com

### MS81

#### Wavelet Estimators of Multivariable Nonparametric Regression Functions with Long Memory Data

In this talk, we propose wavelet estimators for multivariable nonparametric regression functions with long memory data and investigate the asymptotic rates of convergence of wavelet estimators based on level dependent thresholding. We further provide an asymptotic expansion for the mean integrated squared error (MISE) of the estimators. These results extend the corresponding univariate results in the literature to the multivariate setting.

Dongsheng Wu, Yunzhu He  
University of Alabama in Huntsville  
Department of Mathematical Sciences  
dongsheng.wu@uah.edu, yh0004@uah.edu

### MS82

#### An Adaptive Fast Multipole Accelerated Poisson Solver for Complex Geometries

We present a fast, direct and adaptive Poisson solver for complex two-dimensional geometries based on potential theory and fast multipole acceleration. The solver relies on the standard decomposition of the solution as the sum of a volume integral to account for the source distribution and a layer potential to enforce the desired boundary condition. The volume integral is computed by applying the FMM on a square box that encloses the domain of interest. For the sake of efficiency and convergence acceleration, we first extend the source distribution (the right-hand side in the Poisson equation) to the enclosing box as a  $C^0$  function using a fast, boundary integral-based method. We demonstrate with two Poisson problems on multiply connected domains with irregular boundaries that this continuous extension leads to high accuracy without excessive adaptive refinement near the boundary and, as a result, to an efficient “black box” fast solver.

Antoine Cerfon  
NYU  
cerfon@cims.nyu.edu

Travis Askham  
University of Washington

askham@uw.edu

### MS82

#### Smooth Function Extension for the Solution of Inhomogeneous Equations in Complex Geometry

In many applications, it is desirable to solve partial differential equations on irregular domains with high accuracy and speed. Integral equation based methods offer this ability for homogeneous elliptic PDEs. For problems with inhomogeneous source terms, a split of the problem allows for the evaluation of a volume potential over a simple domain. One key ingredient for the success of such an approach is a technique to efficiently compute a high-regularity extension of a function outside the domain where it is given. In this work, function extension is done by radial basis functions combined with a partition of unity technique, where the underlying matrix structure is exploited to efficiently solve each decoupled interpolation problem.

Anna-Karin Tornberg  
KTH  
akto@kth.se

### MS82

#### Integral Equation Methods for Nonlinear Fluid Models

We present a new integral equation method for simulating particulate flows suspended in fluids governed either by Navier-Stokes equations or other nonlinear models. Following existing methods in the literature, we split the problem into a boundary-only part and a volume part. We use a spectrally-accurate boundary integral method to solve the former and a regular-mesh based spectral method for the latter. We show that moving geometries can be handled seamlessly and present results that show scaling to more than a thousand soft particles hydrodynamically interacting with each other.

Shravan Veerapaneni  
Department of Mathematics  
University of Michigan  
shravan@umich.edu

### MS82

#### Fast Methods for the Evaluation of Heat Potentials in Moving Geometry

We present a hybrid asymptotic numerical method for the accurate evaluation of single and double layer heat potentials in two dimensions. These are non-local integral operators in space-time and require fast algorithms to handle the history part efficiently. In order to design high order accurate rules, however, it is the local-in-time component of the heat potential that gives rise to numerical difficulties. Our scheme requires only a local asymptotic calculation plus several boundary integrals with a Gaussian kernel, which can be accelerated by a new version of the fast Gauss transform.

Jun Wang  
Courant Institute  
New York University  
junwang@cims.nyu.edu

### MS83

#### Modeling Flow Rate–Pressure Drop Relations of

### Soft Micro-Fluidic Devices

Rectangular microfluidic channels with deformable walls are one of the simplest and most extensively studied lab-on-a-chip devices, primarily due applicability as model systems in a variety of fields like biomedical diagnostics, nanotechnology, etc. Experimentally, these devices are found to deform into a non-rectangular cross-section due to fluid–structure interactions occurring at the channel walls. These deformations significantly affect the flow profile, which, in turn, results in a non-linear relationship between the volumetric flow rate and the pressure drop. Recent perturbative calculations by the first author show that this flow rate–pressure drop relation can be obtained by analyzing a coupled system of Stokes ( $Re = 0$ ) flow in a three-dimensional (3D) rectangular channel with a top wall that is linearly elastic, specifically a Kirchhoff–Love plate. We benchmark and verify the theoretical predictions by 3D numerical simulations. The response of several shallow microfluidic channels ( $\ell \gg w \gg h$ ) with rigid base and side walls and a soft top wall (e.g., PDMS in experiments) is simulated under laminar flow conditions using the commercial software suite ANSYS. The simulations are calibrated using experimental pressure drop–flow rate data from the literature and then compared to the newly developed perturbative analytical expressions. Favorable agreement is obtained without any fitting parameters.

Ivan C. Christov, Tanmay Shidhore  
Purdue University  
christov@purdue.edu, tshidhor@purdue.edu

### MS83

#### Growth of Clogs in Microchannels: Formation of the Filter Cake

Porous membranes are used to remove contaminants from polluted water. The efficiency of the filtration process is a critical issue as worldwide water demands keep increasing. A typical membrane allows the fluid to pass through but traps contaminants, which can be modeled by colloids or micrometer-scale particles. Once a clog is formed in a pore, incoming particles aggregate upstream. This aggregate grows over time, which leads to a dramatic reduction of the flow rate. In this talk, we will consider a model that predicts the growth of the colloidal aggregate formed upon clogging of a microchannel. We will present an analytical description to capture the time-evolution of the volume of the aggregate. We show that the growth of the aggregate increases the hydraulic resistance in the channel and leads to a drop in flow rate of the filtered suspensions. The analytical description will then focus of multiple parallel channels. In this situation, the growth dynamics of the aggregates are intrinsically coupled. The results of this modeling will be compared with experimental data. Our work illustrates the critical influence of clogging events on the flow rate of porous membranes used in practical applications such as the filtration of polluted water.

Emilie Dressaire  
New York University  
Tandon School of Engineering  
dressaire@nyu.edu

### MS83

#### Molecular Dynamics Simulation Study of the Diffusion of a Janus Nano-Particle in an Explicit Solvent

Molecular dynamics simulations are used to study the diffu-

sion of a single Janus particle immersed in a Lennard-Jones fluid. In particular we consider a rod-shaped and sphere shaped particle with different surface wettability on each half-side of the particle. We analyzed the time dependence of the orientation tensor, particle displacement, and translational and rotational velocity autocorrelation functions. It is found that diffusion is enhanced when the wettability contrast is high and the local slip length on the nonwetting side is relatively large. We also observed that in contrast to homogeneous particles, the nonwetting hemisphere of the Janus particle tends to rotate in the direction of the displacement vector during the rotational relaxation time.

Ali Kharazmi  
Michigan State University  
kharazm1@msu.edu

Nikolai Priezjev  
wright state university  
nikolai.priezjev@wright.edu

### MS83

#### Mathematical Modeling of Pleated Membrane Filters

Pleated membrane filters are widely used in many applications, and offer significantly better surface area to volume ratios than equal area unpleated membrane filters. However, their filtration characteristics are markedly inferior to those of equivalent unpleated membrane filters in dead-end filtration. While several hypotheses have been advanced for this, one possibility is that the flow field induced by the pleating leads to spatially nonuniform fouling of the filter, which in turn degrades performance. We investigate this hypothesis by developing a simplified model for the flow and fouling within a pleated membrane filter. Our model accounts for the pleated membrane geometry (which affects the flow), for porous support layers surrounding the membrane, and for two membrane fouling mechanisms: (i) adsorption of very small particles within membrane pores; and (ii) blocking of entire pores by large particles. We use asymptotic techniques based on the small pleat aspect ratio to solve the model, and we compare solutions to those for the closest-equivalent unpleated filter.

Pejman Sanaei  
New Jersey Institute of Technology  
ps468@njit.edu

Giles Richardson  
University of Southampton  
g.richardson@soton.ac.uk

Thomas P. Witelski  
Duke University  
Department of Mathematics  
witelski@math.duke.edu

Linda Cummings  
New Jersey Institute of Technology  
linda.cummings@njit.edu

### MS84

#### Symbolic Math for Automated Fast Finite Differ-



**ence Computations**

Abstract not available

Navjot Kukreja  
Machine Learning and HPC researcher  
SENAI CIMATEC  
navjotk@gmail.com

**MS84****Multi-Physics Geophysical Exploration: Raising the Abstraction with Separation of Concerns**

Abstract not available

Mathias Louboutin  
PhD student at the University of British Columbia  
mloubout@eos.ubc.ca

**MS84****Large-Scale Workflows for Wave-Equation Based Inversion in Julia**

Abstract not available

Philipp A. Witte  
PhD student at the University of British Columbia  
pwitte@eos.ubc.ca

**MS84****Considerations and Techniques for Optimizing HPC Stencil Performance on the Intel Xeon Phi Processor**

Abstract not available

Charles Yount  
Principal Engineer Intel  
charles.yount@intel.com

**MS85****A New PoissonNernstPlanck Model with IonWater Interactions for Charge Transport in Ion Channels**

In this work, we propose a new PoissonNernstPlanck (PNP) model with ionwater interactions for biological charge transport in ion channels. Due to narrow geometries of these membrane proteins, ionwater interaction is critical for both dielectric property of water molecules in channel pore and transport dynamics of mobile ions. We model the ionwater interaction energy based on realistic experimental observations in an efficient mean-field approach. Variation of a total energy functional of the biological system yields a new PNP-type continuum model. Numerical simulations show that the proposed model with ionwater interaction energy has the new features that quantitatively describe dielectric properties of water molecules in narrow pores and are possible to model the selectivity of some ion channels.

Duan Chen  
Department of Mathematics and Statistics  
University of North Carolina at Charlotte  
dchen10@unc.edu

**MS85****Computational Vascular and Valvular Fluid-****Structure Interaction**

The aim of this work is to develop computational techniques for vascular and valvular fluid–structure interaction (FSI). The motivating application is the simulation of aortic valve function over the complete cardiac cycle. Due to the complex motion of the heart valve leaflets, the fluid domain undergoes large deformations, including changes of topology. We propose an immersogeometric method that directly analyzes an isogeometric surface representation of the structure by immersing it into a non-boundary-fitted discretization of the surrounding fluid domain. The variational formulation for immersogeometric FSI is derived using an augmented Lagrangian approach. We simulate the coupling of the aortic valve and the surrounding blood flow under physiological conditions, demonstrating the effectiveness of the proposed techniques in practical computations. An arbitrary Lagrangian–Eulerian (ALE)/immersogeometric hybrid methodology is also developed under the augmented Lagrangian framework for FSI. A single computation combines a boundary-fitted, deforming-mesh treatment of some fluid–structure interfaces with a non-boundary-fitted treatment of others. This approach enables us to also simulate the FSI of a bioprosthetic aortic valve implanted in a flexible artery through the entire cardiac cycle.

Ming-Chen Hsu  
Department of Mechanical Engineering  
Iowa State University  
jmchsu@iastate.edu

David Kamensky  
University of California, San Diego  
david.kamensky@gmail.com

Fei Xu  
Iowa State University  
feixu@iastate.edu

**MS85****DG-IMEX Asymptotic-Preserving Schemes and Their Stability**

For some kinetic models in diffusive regimes, we consider asymptotic -preserving methods that involve discontinuous Galerkin (DG) spatial discretizations and implicit-explicit (IMEX) Runge-Kutta temporal discretizations. A key property, uniform stability, is examined theoretically for the methods.

Fengyan Li  
Rensselaer Polytechnic Institute  
lif@rpi.edu

**MS85****Large-Scale Fracture Simulations Using Atomistic-Based Boundary Element Method**

The Atomistic-based Boundary Element Method (ABEM) is a novel technique to seek approximate solutions to the static problems of crystalline solids. Lattice Green's function is computable when the harmonic assumption is made. Two main features emerge. ABEM inherits the advantages of the classical Boundary Element Method (BEM). Another feature of ABEM is that it couples with molecular statics without ghost forces. Only truncation errors due to the harmonic assumption appear in the coupling system. By introducing several numerical techniques, this

coupling method can easily deal with a system with 1 billion atoms. These techniques include: reducing the degrees of freedom on the smooth boundaries, adaptively calculating the element-element interactions and atom-element interactions, and Neumann-to-Dirichlet iterations between the linear system from ABEM and the minimization problem from molecular statics. Numerical results of fracture problem demonstrate the effectiveness of this coupling method.

Xiaojie Wu  
Penn State University  
xxw139@psu.edu

Xiantao Li  
Department of Mathematics  
Pennsylvania State University  
xli@math.psu.edu

### MS86

#### **A Bayesian Approach for Inverse Problems in Dynamic Power Systems**

We discuss the problem of estimating the parametric uncertainty in the solution of dynamic power grid inverse problems within the framework of Bayesian inference. We investigate an adjoint-based approach to estimate the maximum a posteriori point of the parameters and their variance, which quantifies their uncertainty. Within this framework we estimate several parameters of the dynamic power system, such as generator inertias, which are not quantifiable in steady-state models. The parametric estimate is obtained through demand node perturbations and the likelihood is based on statistical skill scores. We assess the computational efficiency, and discuss the expected performance when this method is applied to large systems.

Emil M. Constantinescu  
Argonne National Laboratory  
Mathematics and Computer Science Division  
emconsta@mcs.anl.gov

Noemi Petra  
University of California, Merced  
npetra@ucmerced.edu

Cosmin G. Petra  
Lawrence Livermore  
Mathematics and Computer Science Division  
petra1@llnl.gov

Mihai Anitescu  
Argonne National Laboratory  
Mathematics and Computer Science Division  
anitescu@mcs.anl.gov

### MS86

#### **Enabling Large-Scale Power Grid Data Assimilation through Computational Enhancements**

The Power grid is an inherent dynamic system. Accurate information about its dynamic states is important for its efficient control and operation, especially with the increasing complexity resulting from uncertainties and stochastic variations introduced by intermittent renewable-energy sources, responsive loads, and new market designs. High-speed phasor deployment in the power grid over the past two decades offers an opportunity for estimating dynamic

states in real time. Methods have been proposed to assimilate the phasor measurements with a dynamic model of the power grid for such dynamic state estimation. For large-scale systems such as the west or east interconnection in North America, it is a fundamental challenge in completing all necessary computation within tens of milliseconds a required time resolution for capturing power grid dynamics and keeping up with phasor measurement cycles. A typical method for this dynamic state estimation process is based on Ensemble Kalman Filter (EnKF). Parallel computing techniques have been applied to the EnKF process which is dominated by dense-matrix operations. Data block allocation is optimized to reduce parallelization overhead. Adaptive steps in numerical integration are employed to reduce computation burden. And multi-layer implementation is experimented to improve parallelism. This talk will report the methods and accomplishments with all these computational enhancements towards real-time dynamic state estimation.

Henry Huang  
Pacific Northwest National Laboratory, USA  
Zhenyu.Huang@pnl.gov

Steve Elbert  
Pacific Northwest National Laboratory  
steve.elbert@pnl.gov

Ning Zhou  
Binghamton University  
ningzhou@binghamton.edu

Shaobu Wang  
Pacific Northwest National Laboratory  
shaobu.wang@pnl.gov

### MS86

#### **Optimization and Design of Complex Engineering Systems under Uncertainty Using High-Performance Computing**

The optimization of many engineering systems reaches extreme size and unprecedented computational complexity due to the need to consider high fidelity physics and account for uncertainties. We will present scalable optimization algorithms and solvers for solving two such optimization problems: i. material optimization for the design of stiff structures using additive manufacturing; and, ii. economic optimization and control of large scale power grid systems under stochastic loads and generation. Our computational approach is based on scalable nonlinear programming algorithms, namely interior-point methods, and is specifically aimed at matching and surpassing the parallelism of existing physics solvers (e.g., finite element solvers for the elasticity problem governing material optimization) on massively parallel computing platforms. For this we developed specialized decompositions within the optimization, namely Schur complement-based decomposition for stochastic optimization and domain decomposition equipped with low-rank representations of the Hessian for material optimization. Finally, we present large scale simulations of the two motivating applications on the U.S. Department of Energy's supercomputers and provide a detailed discussion of the effectiveness and possible further improvement of the proposed techniques.

Cosmin G. Petra  
Lawrence Livermore  
Mathematics and Computer Science Division

petra1@lnl.gov

### MS86

#### Probabilistic Density Function Method for Nonlinear Dynamical Systems Driven by Colored Noise

We present a probability density function (PDF) method for a system of nonlinear stochastic ordinary differential equations driven by colored noise. The method provides an integrodifferential equation for the temporal evolution of the joint PDF of the systems state, which we close by means of a modified large-eddy-diffusivity (LED) closure. In contrast to the classical LED closure, the proposed closure accounts for advective transport of the PDF in the approximate temporal deconvolution of the integrodifferential equation.

Alexandre M. Tartakovsky

PNNL

alexandre.tartakovsky@pnnl.gov

David A. Barajas-Solano

Pacific Northwest National Laboratory

david.barajas-solano@pnnl.gov

### MS87

#### Geometric Graph Methods for High Dimensional Data

We present new methods for segmentation of large datasets with graph based structure. The method combines ideas from classical nonlinear PDE-based image segmentation with fast and accessible linear algebra methods for computing information about the spectrum of the graph Laplacian. The goal of the algorithms is to solve semi-supervised and unsupervised graph cut optimization problems. I will present results for image processing applications such as image labelling and hyperspectral video segmentation, and results from machine learning and community detection in social networks, including modularity optimization posed as a graph total variation minimization problem.

Andrea L. Bertozzi

UCLA Department of Mathematics

bertozzi@math.ucla.edu

### MS87

#### A Max-Cut Approximation Using A Graph Based MBO Scheme

On a given graph  $G = (V, E)$  a Max-Cut is a partition of nodes  $V = V_1 \cup V_2$  such that the number of edges between the node sets  $V_1$  and  $V_2$  are maximized. This problem was proved to be NP-Complete by Richard Karp in 1972. Due to the complexity of the problem there are various types of approximation algorithms designed to obtain an accurate estimate of a graphs Max-Cut. In this talk I will discuss an alternate graph based Ginzburg-Landau functional and how the minimization of this functional relates to a graphs Max-Cut; we attempt to minimize this functional by a graph based MBO Scheme, where we use a graph heat equation with a twist, instead of diffusion being defined using a graph Laplacian we use an operator called the signless Laplacian. All of these concepts will be explained. Afterwards, I will show examples on specific networks displaying how effective the approximation

algorithm has been in testing.

Blaine Keetch

School of Mathematical Sciences

University of Nottingham

pmxbk1@nottingham.ac.uk

### MS87

#### Convex Variational Methods for Multi Class Data Segmentation on Graphs

Graph-based variational methods have recently shown to be highly competitive for various classification problems of high-dimensional data. One limitation of existing methods for problems with multiple classes is their inclination to get stuck in a local minimum due to the non-convexity of the minimization problems. This paper proposes a convex relaxation for a certain subset of graph-based multiclass data segmentation problems, and develops an efficient and robust optimization algorithm based on the corresponding dual problem. In particular, it is shown that semi-supervised classification problems of high-dimensional data can be tightly approximated by the convex relaxation. The same applies to certain unsupervised classification problems, where the usual class interphase term is appended with region terms that enforce homogeneity of the data points within each class. Specific constructions of each term are proposed for the application to segmentation of unstructured 3D point clouds. Flexible interval constraints on the class sizes can also be naturally incorporated in the convex algorithmic framework. Experiments on semi-supervised classification and unsupervised 3D point cloud segmentation indicate vanishingly small approximation errors of the relaxations, and demonstrate a greater accuracy, robustness and efficiency of the algorithms when compared to earlier methods.

Ekaterina Merkurjev

Department of Mathematics

UCSD

kmerkurjev@gmail.com

Egil Bae

University of Bergen

egil.bae@math.uib.no

### MS87

#### Primal-Dual Methods for P-Modulus on Graphs

Recently, the notion of p-Modulus has been extended to families of walks on graphs by Albin, Poggi-Corradini, et al. The p-modulus on graphs has been shown to relate to well-known graph properties such as effective resistance, shortest path, and minimal cut. Here we propose to employ a primal-dual hybrid gradients scheme (e.g., Chambolle and Pock) to the greedy algorithm for efficient p-Modulus computation. All optimization steps are simple first order. We compare against existing methods.

Dominique P. Zosso

University of California, Los Angeles

Department of Mathematics

dominique.zosso@montana.edu

### MS88

#### Building Blocks for Resilient MPI Applications

Abstract not available

George Bosilca

University of Tennessee - Knoxville  
bosilca@icl.utk.edu

### MS88

#### Partial Differential Equations Solver Resilient to Soft and Hard Faults

We present a resilient task-based domain-decomposition preconditioner for partial differential equations (PDEs) built on top of User Level Fault Mitigation MPI (MPI-ULFM). The algorithm reformulates the PDE as a sampling problem, followed by a robust regression-based solution update that is resilient to silent data corruptions (SDCs). We adopt a server-client model where all state information is held by the servers, while clients only serve as computational units. The task-based nature of the algorithm and the capabilities of ULFM together serve to support missing tasks, making the application resilient to failing clients. We present weak and strong scaling results, both for a nominal and a fault-injected case, showing excellent scalability in all cases, and study the variability in scaling due to faults. Further, we study the performance of the solver as a function of the subdomain size and the server/client configuration.

Francesco Rizzi, Karla Morris, Khachik Sargsyan  
Sandia National Laboratories  
fnrizzi@sandia.gov, knmorri@sandia.gov,  
ksargsy@sandia.gov

Paul Mycek, Andres Contreras  
Duke University  
paul.mycek@duke.edu, andres.contreras@duke.edu

Cosmin Safta  
Sandia National Laboratories  
csafta@sandia.gov

Olivier Le Maitre  
Laboratoire d'Informatique pour la Mécanique et les  
Science  
olm@limsi.fr

Omar M. Knio  
Duke University  
omar.knio@duke.edu

Bert J. Debusschere  
Energy Transportation Center  
Sandia National Laboratories, Livermore CA  
bjdebus@sandia.gov

### MS88

#### Resilience in the Parareal Method

Parallel-in-time integration is an often advocated approach for extracting parallelism in the solution of time-dependent PDE's. Due to the comparatively low parallel efficiency of parallel-in-time integration techniques, in particular

parareal, they are primarily of interest as an extension to classical approaches at parallelism such as spacial domain decomposition in the limit of poor scaling. Potential applications are expected to scale across hundreds, or possibly thousands of nodes, making algorithmic resilience towards hardware induced errors highly relevant. In this talk we present a modified algorithm that is resilient to node-losses, and discuss our experiences implementing the fault-tolerant scheme using a set of MPI interface extensions for implementing recovery strategies called ULFM. We furthermore demonstrate how the parareal algorithm may be made resilient towards Silent-Data-Corruption (SDC) errors by viewing it as a point-iterative method, locally monitoring the residual between consecutive iterations so to discard potentially corrupt iterations.

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

Allan Nielsen  
EPFL Lausanne  
allan.nielsen@epfl.ch

### MS88

#### A Soft-Fault Tolerant Advection Solver Via the Sparse Grid Combination Technique

The sparse grid combination technique (SGCT) is a cost-effective method for solving time-evolving PDEs especially for higher-dimensional problems. Instead of solving the problem on a single grid with high resolution in all dimensions, the SGCT involves solving the problem on a number of grids with higher resolution in some dimensions and lower resolution in others. Then these solutions can be combined to approximate the accuracy of the solution with the high resolution grid. The SGCT can also be easily modified to provide algorithm-based fault tolerance. This involves computing the solution a number of extra lower resolution grids for extra redundancy. Upon applying the SGCT, each different version of a solution point may be compared to detect possible corruption (soft faults). We apply this approach to a 2D advection solver, where random bit flips are applied to the intermediate solution to be detected during the SGCT. We use an error-splitting technique to derive bounds for the expected range for a solution point on a grid. If a point is outside this range, a fault-tolerant SGCT is then applied to exclude the point from the combined solution. Simulation results indicates this approach achieves similar accuracy to the no-fault case on the advection solver for up to a moderate rate of soft faults.

Peter Strazdins  
Research School of Computer Science  
Australian National University  
Peter.Strazdins@cs.anu.edu.au

### MS90

#### Perturbation-Based Reduced Bases for Parametrized Eigenvalue Problems

The aim of this talk is to present a perturbation-based reduced basis method for the approximation of parametrized eigenvalue problems. The algorithm is inspired from the ideas proposed in [Eric Cancès, Geneviève Dusson, Yvon Maday, Benjamin Stamm, Martin Vohralík, "A perturbation-method-based post-processing for the planewave discretization of Kohn-Sham models", Journal

of Computational Physics 307 (2016) 446459]. The method is illustrated on the problem of Brillouin zone integration for the computation of the electronic structure of crystalline materials, semiconductor or metals.

Virginie Ehrlicher

CERMICS - Ecole des Ponts Paristech / INRIA  
ehrlachv@cermics.enpc.fr

Eric Cancès

Ecole des Ponts Paristech  
eric.cances@enpc.fr

Damiano Lombardi

INRIA Paris-Rocquencourt  
damiano.lombardi@inria.fr

David Gontier

Ecole Normale Supérieure, Paris  
davidgontier@gmail.com

Antoine Levitt

Inria and Ecole des Ponts ParisTech  
antoine.levitt@inria.fr

## MS90

### Collocation Methods for Exploring Perturbations in Linear Stability Analysis of Dynamical Systems

Eigenvalue analysis is a well-established tool for stability analysis of dynamical systems. However, there are situations where eigenvalues miss some important features of physical models. For example, in models of incompressible fluid dynamics, there are examples where eigenvalue analysis predicts stability but transient simulations exhibit significant growth of infinitesimal perturbations. In this study, we explore an approach similar to pseudo-spectral analysis and simple to implement that can be used to predict stability. The method defines parameter-dependent perturbations of steady solutions of a dynamical system and then uses collocation to approximate the solution of accompanying parameter-dependent perturbed eigenvalue problems. For several classic benchmark problems in computational fluid dynamics, including flow over an obstacle and flow over an expanding step, we show that the resulting perturbation analysis is predictive of the behavior of transient solvers. In particular, when the perturbed eigenvalues cross the imaginary axis, indicating instability, transient solvers exhibit unsteady behavior such as vortex shedding, and when all perturbed eigenvalues have negative real part, transient solvers find steady solutions.

Howard C. Elman

University of Maryland, College Park  
elman@cs.umd.edu

David Silvester

School of Mathematics  
University of Manchester  
david.silvester@manchester.ac.uk

## MS90

### Applying Quasi-Monte Carlo Integration to a Parameterized Elliptic Eigenvalue Problem

In this talk we study an elliptic eigenproblem, with a random coefficient which can be parametrised by infinitely many stochastic parameters. The randomness in the co-

efficient also results in randomness in the eigenvalues and corresponding eigenfunctions. As such, our quantity of interest will be the expected value, with respect to the stochastic parameters, of the smallest eigenvalue which we formulate as an integral over the infinite-dimensional parameter domain. Our approximation involves three steps: truncating the stochastic dimension, discretising the spatial domain using finite elements and approximating the now finite but still high-dimensional integral. To approximate the high-dimensional integral we use quasi-Monte Carlo (QMC) methods. These are deterministic or quasi-random quadrature rules that can be proven to be very efficient for the numerical integration of certain classes of high-dimensional functions. QMC methods have previously been applied to similar elliptic source problems, however the existing framework for a rigorous analysis of the integration error does not cover the eigenvalue problem. We show that the minimal eigenvalue belongs to the spaces required for QMC theory, outline the approximation algorithm and provide numerical results.

Alexander D. Gilbert

School of Mathematics and Statistics, Uni of New South Wales  
alexander.gilbert@student.unsw.edu.au

## MS90

### Complexity Reduction for Eigenvalue Problems in Vibro-Acoustics

Parameter dependent eigenvalue problems arise in many applications in science and engineering, e.g. modal analysis of structures. The numerical approximation by standard finite element techniques of such systems is challenging due to complex geometries and possible uncertain input data. In this talk we cover several aspects of complexity reduction methods for the efficient solution of parameter dependent eigenvalue problems. We show a reduced basis method including an upper bound to build a single reduced space for the simultaneous approximation of several eigenvalues and eigenvectors of interest. In the context of the vibro-acoustic analysis of multi-story timber building structures, we apply the introduced reduced basis method for eigenvalue problems to component mode synthesis, which allows a component based decomposition and solution. Furthermore we also present results on more complex geometries with piecewise smooth but curvilinear boundaries such as a violin or a violin-bridge. Here we do not only vary the material parameters, but also the geometry. To resolve the geometry more accurately we apply isogeometric mortar techniques and combine it with empirical interpolation.

Thomas Horger

M2 - Department of Mathematics  
Technische Universität München  
horger@ma.tum.de

Barbara Wohlmuth

M2, Centre for Mathematical Sciences,  
Technische Universität München, Germany  
wohlmuth@ma.tum.de

Linus Wunderlich

M2 - Department of Mathematics  
Technische Universität München  
linus.wunderlich@ma.tum.de

## MS91

### Effective Transmission Conditions for Domain De-

### composition Methods in the Case of the Helmholtz Equation

In this talk, we review some of the transmission conditions used for non-overlapping domain decomposition methods in the case of the Helmholtz equation. Then, we propose to modify some conditions in order to improve the cost of the whole algorithm. Numerical results are presented in two and three dimensions to validate this new approach.

Yassine Boubendir

Department of Mathematical Sciences  
boubendi@njit.edu

### MS91

#### A Parareal in Time Algorithm for the Optimal Control of Evolution Equations

A parareal in time algorithm is proposed to solve the optimal control problems governed by evolution equations. This method is to solve the first order optimality system, which can be viewed as a two-point boundary value problem in time, by a time domain decomposition technique. The algorithm consists of a local parallel part and a global correction part, and it has a comparable error estimate with that of the parareal algorithm for the initial value problems. This method can also be extended to the optimal control problems with control constraints or nonlinear governing equations. A number of numerical examples are presented to show the efficiency of the proposed algorithm and to validate the convergence results.

Wei Gong

Institute of Computational Mathematics  
Chinese Academy of S  
wgong@lsec.cc.ac.cn

### MS91

#### Domain Decomposition for a Mixed Finite Element Method for Linear Elasticity with Weak Symmetry

Two non-overlapping domain decomposition methods on matching multiblock grids are presented for a mixed formulation of linear elasticity problem with weakly enforced symmetry. The proposed approaches allow to solve the problem in parallel by partitioning the computational domain into multiple subdomains, upon which local problems of lower complexity are formulated. By eliminating the interior subdomain variables, the global problem is reduced to an interface problem, which is then solved by an iterative procedure. The mixed finite element method for linear elasticity with weakly imposed symmetry of stress is used to discretize each subdomain problem. The condition numbers of the resulting algebraic interface systems are analyzed. A mixed finite element method for the problem of interest on nonmatching multiblock grids is also considered. A mortar finite element space is introduced on the nonmatching interfaces. We approximate in this mortar space the trace of the displacement, and we impose weakly a continuity of normal stress condition. The condition number of the interface system is analyzed and optimal order convergence is shown for stress, displacement and rotation. Moreover, at cell centers, superconvergence is obtained for the displacement. Computational results using an efficient parallel domain decomposition algorithm are presented in confirmation of the theory of all proposed approaches.

Eldar Khattatov

University of Pittsburgh

Department of Mathematics  
elk58@pitt.edu

Ivan Yotov

University of Pittsburgh  
Department of Mathematics  
yotov@math.pitt.edu

### MS91

#### Schur Complement Domain Decomposition Methods for the Solution of Multiple Scattering Problems

We present a Schur complement Domain Decomposition (DD) algorithm for the solution of frequency domain multiple scattering problems. Just as in the classical DD methods we (1) enclose the ensemble of scatterers in a domain bounded by an artificial boundary, (2) we subdivide this domain into a collection of nonoverlapping subdomains so that the boundaries of the subdomains do not intersect any of the scatterers, (3) we connect the solutions of the subproblems via Robin boundary conditions matching on the common interfaces between subdomains. We use subdomain Robin-to-Robin maps to recast the DD problem as a sparse linear system whose unknown consists of Robin data on the interfaces between subdomains—two unknowns per interface. The Robin-to-Robin maps are computed in terms of well-conditioned boundary integral operators. Unlike classical DD, we do not reformulate the Domain Decomposition problem in the form a fixed point iteration, but rather we solve the ensuing linear system by Gaussian elimination of the unknowns corresponding to inner interfaces between subdomains via Schur complements. Once all the unknowns corresponding to inner subdomains interfaces have been eliminated, we solve a much smaller linear system involving unknowns on the inner and outer artificial boundary. We present numerical evidence that our Schur complement DD algorithm can produce accurate solutions of very large multiple scattering problems that are out of reach for other existing approaches.

Catalin Turc

Department of Mathematical Sciences  
New Jersey Institute of Technology  
catalin.c.turc@njit.edu

### MS92

#### Efficient High-Order Finite Elements in MFEM

When utilizing high order finite elements a significant portion of the computational cost is found in computing the element matrices. For finite elements defined with an underlying tensor product structure the element matrix computations can be refactored as a series of dense tensor contractions. These contractions can be batched across elements and executed in parallel utilizing advanced architectures such as Nvidia GPUs and Intel MICs. Numerous challenges arise when attempting to write code for these architectures. We will present our exploratory work on utilizing tensor contractions for element formation in MFEM and porting these computations to Nvidia GPU architectures. Multiple abstractions to contracting the tensors and GPU porting were attempted including raw CUDA, C++ template metaprogramming with heavy inlining, OpenMP 4.5 offloading pragmas, a tensor contraction library utilizing JIT compilation, and the platform independent OKL language from the OCCA library. We will present details of the most successful approaches and many of the compli-

cations that arose in the porting process.

Aaron Fisher

Lawrence Livermore National Laboratory  
fisher47@llnl.gov

**MS92**

**Performance Portable Assembly Tools for Multi-Fluid Plasma Simulation**

Supporting scalable assembly on next generation architectures can add significantly to code complexity. This presentation will discuss the design of a general finite element assembly engine applied to multi-fluid plasma and magnetohydrodynamics simulation. Performance portability is achieved via the Kokkos programming model. The assembly library uses a directed acyclic graph (DAG) for composable physics kernels in a multiphysics setting. Assembly with hybrid data parallelism and task-based parallelism over the DAG is supported. Embedded automatic differentiation, applied via templates and operator overloading, is used for implicit and IMEX solvers. Performance results for a mixed-basis multi-fluid plasma finite element formulation will be shown for Intel Phi, NVIDIA GPU and Intel Haswell architectures.

Roger Pawlowski

Multiphysics Simulation Technologies Dept.  
Sandia National Laboratories  
rppawlo@sandia.gov

Eric C. Cyr

Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

Matthew Bettencourt

Sandia National Laboratories  
mbetten@sandia.gov

Eric Phipps

Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

Christian Trott, H. Carter Edwards

Sandia National Laboratories  
crtrott@sandia.gov, hcedwar@sandia.gov

John Shadid

Sandia National Laboratories  
Albuquerque, NM  
jshadi@sandia.gov

Richard Kramer, Allen C. Robinson

Sandia National Laboratories  
rmkrame@sandia.gov, acrobin@sandia.gov

**MS92**

**Case Studies in Using a DSL and Task Graphs for Portable Reacting Flow Simulations**

With the changing computing hardware landscape, simulation codes must be able to be flexible in targeting different architectures, such as CPUs and GPUs. Domain specific languages (DSLs) are one way to accommodate this, allowing application programmers to express their problems at a higher level while delegating the task of deploying to a

given architecture to the DSL. Task graphs, on the other hand, allow automated construction of algorithms and reasoning about the structure of a problem. By combining these, we can achieve composable problems that allow efficient utilization of modern computing architectures. Several case studies of various complexity will be shown to demonstrate the efficacy of the DSL with task graphs on CPUs and GPUs including scaling up to 18K GPUs on the breadth of Titan.

James C. Sutherland

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

Tony Saad

Institute for Clean and Secure Energy  
Department of Chemical Engineering, University of Utah  
tony.saad@utah.edu

**MS92**

**Matrix-Free Operator Evaluation for Finite Element Discretization**

There has been a long trend where the number of flops increases much faster than the amount and the bandwidth of memory. This is particularly obvious for Sunway Taihu-Light (fastest supercomputer as of June 2016) which has only 0.12 GB of memory per core. Thus, storing the matrix associated with the system can require a lot of the available memory. In many modern simulators, the use of matrices is only through sparse matrix-vector multiplications (SpMV) within Krylov solvers. However, SpMV is heavily bandwidth limited, both on CPU and GPU systems. GPUs have a much higher memory bandwidth, but so are the arithmetic capabilities. Moreover, GPUs have a lot less memory per “core” than CPUs. These considerations are in favor of stopping the separation of linear algebra from finite elements assembly routines and instead applying the vector to the operator on the fly. Previous work on matrix-free operator evaluation for finite element on CPU has shown that for high-order finite elements, using a matrix-free implementation was more advantageous. Given the success of matrix-free kernels on CPU, we extend this work to the GPU. Compared to a regular assembly, a special data structure is required and operations done during a regular assembly needs to be reordered. Our implementation is part of the deal.II finite element library and applicable to rather general differential operators. The GPU and the CPU share the same interface to allow an easy switch between the two implementations.

Bruno Turcksin

Texas A&M University  
turcksinbr@ornl.gov

Karl Ljungkvist

Uppsala University  
karl.ljungkvist@it.uu.se

Martin Kronbichler

Technische Universitat Munchen  
kronbichler@lnm.mw.tum.de

**MS93**

**Generalized Finite Element Methods Applied to Interface Problems**

The Generalized Finite Element Method (GFEM) is an ex-

tension of the standard Finite Element Method (FEM) where the trial space is obtained by augmenting the standard space of piecewise linear functions by an enrichment space. The enrichment space, consisting of non-polynomial functions, is chosen carefully such that (a) it incorporates the feature of the problem, (b) yields a sparse stiffness matrix, (c) and the GFEM yields optimal convergence, i.e.,  $O(h)$ . However, often the obvious choice of the enrichment space, satisfying (a), (b), and (c), results into *badly conditioned linear system*; in fact the conditioning of GFEM could be much worse than that of a standard FEM. Furthermore, the *conditioning is not robust* with respect to the position of the interface relative to the underlying mesh. In this talk, we show that if the enrichment space satisfies two axioms, then (i) the GFEM yields optimal order of convergence, (ii) the conditioning of GFEM is not worse than that of the standard FEM, and (iii) the conditioning of the GFEM is robust. We will also present sufficient conditions such that the enrichment space satisfies these axioms. We will also illuminate these results with computational examples.

Uday Banerjee  
Syracuse University  
banerjee@syr.edu

Ivo Babuska  
ICES, The University of Texas at Austin  
babuska@ices.utexas.edu

Qinghui Zhang  
Sun Yat-Sen University  
zhangqh6@mail.sysu.edu.cn

### MS93

#### Stable Generalized Finite Element Method for Linear Elastic Fracture Mechanics

Abstract not available

Sundararajan Natarajan  
Indian Institute of Technology Madras  
sundararajan.natarajan@gmail.com

### MS93

#### A Partition of Unity Method for Generalized Eigenvalue Problems with Application in Electronic Structure Calculations

Abstract not available

Marc Alexander Schweitzer  
INS / Meshfree Multiscale Methods  
University of Bonn / Fraunhofer SCAI  
schweitzer@ins.uni-bonn.de

### MS93

#### Exact Integration Scheme for Planewave-Enriched Partition of Unity Finite Element Method to Solve the Helmholtz Problem

In this talk, we will present an exact integration scheme to compute highly oscillatory integrals that appear in the numerical solution of the two-dimensional Helmholtz equation using the planewave-enriched partition of unity finite element method (PUFEM). In planewave-enriched PUFEM, the piecewise polynomial finite element (FE) approximation space is augmented by special, problem-

dependent, non-polynomial functions through the framework of partition-of-unity. For the Helmholtz equation in 2D, we use globally-supported planewaves as enrichment functions. The stiffness matrix that stems from the planewave-enriched PUFEM contains several *oscillatory integrals*. In the new scheme, such oscillatory integrals are computed by a recursive application of the divergence theorem, which only requires the evaluations of the corresponding oscillatory integrands at the nodes of the FE mesh. The number of such function evaluations is independent of the wave number  $k$ , which permits the scheme to be used for arbitrary high values of  $k$ . We will present numerical results for three canonical benchmark Helmholtz problems on unstructured triangular and structured rectangular meshes to demonstrate the accuracy and efficacy of the method.

N. Sukumar, Subhajit Banerjee  
University of California, Davis  
nsukumar@ucdavis.edu, subhajit.banerjee@gmail.com

### MS94

#### Toward Space-Time Parallel Climate Simulations Using Pfasst

We report on an effort to add parallelism in the time direction to the High-Order Methods Modeling Environment (HOMME) package. HOMME is a highly scalable spectral element code that provides the default dynamical core for the Community Atmospheric Model and the Community Earth System Model developed by NCAR. We will discuss the implementation of the time-parallel PFAST algorithm as the time-stepping method within HOMME, and present space-time scaling results for some standard test cases.

#### Andreas Kreienbuehl

Lawrence Berkeley National Lab  
akreienbuehl@lbl.gov

Michael Minion  
Lawrence Berkeley National Lab  
Stanford University  
mlminion@lbl.gov

Katherine J. Evans  
Oak Ridge National Laboratory  
evanskj@ornl.gov

### MS94

#### On the Time-Parallelization of the Solution of Navier-Stokes Equations Using Parareal

Unsteady turbulent flow simulations using the Navier Stokes equations require larger and larger problem sizes. Besides, new supercomputer architectures available in the next decade, will be characterized with increased computational power based on a larger number of cores rather than significantly increased CPU frequency. Hence most of the current generation CFD software will face critical efficiency issues if bounded to massive spatial parallelization ( $O(10^{7-8})$  cores), and we consider time parallelization as an attractive alternative to enhance efficiency on multi-cores architectures. Among all the algorithms developed in the last decade, Parareal remains one of the simplest solutions in the case of explicit time stepping, compressible flow solvers. A multi-grid approach for the Parareal algorithm is applied to a non-linear three-dimensional turbulent CFD problem (Decay of Homogeneous Isotropic Turbulence). The theoretical speed-up of the method is estimated, based on the actual scaling properties of an existing



massively parallel CFD solver (HYBRID) and the modelization of time communications. Influence of grid to grid operators (coarsening and refinement), and Parareal settings are studied to underline key ingredients for an efficient parallel time-integration. After assessing the performances of this approach, two enhancement strategies are proposed, based on relaxation of Parareal iteration and Krylov subspace approximation respectively.

Thibaut Lunet  
CERFACS, Toulouse  
thibaut.lunet@isae.fr

Julien Bodart  
ISAE-Supaero, Toulouse, France  
julien.bodart@isae.fr

Serge Gratton  
ENSEEIH, Toulouse, France  
serge.gratton@enseeiht.fr

Xavier Vasseur  
CERFACS  
vasseur@cerfacs.fr

#### MS94

##### The Role of Near-Resonance in Parareal Convergence

The Asymptotic Parallel-in-Time (APinT) method extends the standard Parareal method such that rapid convergence may be achieved even in the presence of oscillatory stiffness. The method employs averaging over the fast waves in the coarse solve, which provides a much more accurate approximation to the solution over long time intervals. It has been shown previously that this method converges rapidly in the asymptotic limit of scale separation. In this work, we show that the window over which averaging is performed may be chosen such that rapid convergence is possible over a much broader flow regime. A method for optimising the wave averaging window is also explored.

Adam Peddle, Beth Wingate  
University of Exeter  
adam.g.peddle@gmail.com, B.Wingate@exeter.ac.uk

Terry Haut  
Los Alamos National Laboratory  
terryhaut@lanl.gov

#### MS94

##### PFASST and Finite Elements

For time-dependent PDEs, parallel-in-time integration using the “parallel full approximation scheme in space and time” (PFASST) is a promising way to accelerate existing space-parallel approaches beyond their scaling limits. Inspired by the classical Parareal method and non-linear multigrid ideas, PFASST allows to integrate multiple time-steps simultaneously using “multi-level spectral deferred corrections” (MLSDC) with different coarsening strategies in space and time. In numerous studies, this approach has been successfully coupled to space-parallel solvers which use finite differences, spectral methods or even particles for discretization in space. In this talk, we report on using MLSDC and PFASST in time together with finite elements in space. In particular, we discuss modifications necessary to treat the mass matrix appropriately on all

levels of the space-time hierarchy. The algorithms are implemented using the Distributed and Unified Numerics Environment (DUNE) with the community-oriented PFASST++-code as well as FEniCS with the prototyping framework pySDC, allowing users to test new ideas or applications without great effort. We show first results for the advection-diffusion equation and multi-component reaction-diffusion equations such as the Gray-Scott model.

Robert Speck, Ruth Schoebel  
Juelich Supercomputing Centre  
Forschungszentrum Juelich GmbH  
r.speck@fz-juelich.de, r.schoebel@fz-juelich.de

Oliver Sander  
Institut fuer Numerische Mathematik  
TU Dresden  
oliver.sander@tu-dresden.de

#### MS95

##### Three Tricks to Tame the CFL in Discontinuous Galerkin Methods

In this talk we explore three different approaches to taming the order of accuracy dependent CFL number in discontinuous Galerkin methods for hyperbolic problems. The first approach builds on the ideas of co-volume operations appearing in the work of Hagstrom and Warburton and in papers on the so called central dG method. The second approach takes advantage of the finite speed of propagation to build local time update formulas based on a low rank approximation to a matrix exponential. The third approach introduces a new class of dG methods which we call Sobolev dG. These high order methods are built recursively from the lowest order method in such a way that the eigenvalues of the time stepping operators coincide with the low order method.

Daniel Appelo  
University of New Mexico  
appelo@math.unm.edu

Thomas M. Hagstrom  
Southern Methodist University  
Department of Mathematics  
thagstrom@smu.edu

Fengyan Li  
Rensselaer Polytechnic Institute  
lif@rpi.edu

Matt Pennybacker, Adeline Kornelus  
University of New Mexico  
pennybacker@math.unm.edu, kornelus@unm.edu

#### MS95

##### Weight-Adjusted Discontinuous Galerkin Methods for Heterogeneous Media and Curvilinear Meshes

Time-domain discontinuous Galerkin (DG) methods for wave propagation require accounting for the inversion of dense elemental mass matrices, where each mass matrix is computed with respect to a parameter-weighted  $L^2$  inner product. For curvilinear meshes and wave propagation in media with sub-element heterogeneities, these matrices are distinct over each element, necessitating additional storage. We propose weight-adjusted DG (WADG) methods which

reduce storage costs by replacing the weighted  $L^2$  inner product with a weight-adjusted inner product. This equivalent inner product results in an energy stable method, but does not increase storage costs for locally varying weights. *a priori* error estimates are derived, and numerical examples are given illustrating the application of this method to the acoustic wave equation with heterogeneous wavespeed on curved domains.

Jesse Chan  
Rice University  
jesse.chan@rice.edu

Russell Hewett  
Total  
russell.hewett@total.com

Tim Warburton  
Virginia Tech  
tcew@vt.edu

### MS95 DG Methods on Hybrid Grids for Simulating Waves

DG methods on unstructured grids have a number of desirable features for simulating waves in the time domain. In particular:

- Automatic treatment of complex geometry
- Straightforward implementation of hp-adaptivity
- Energy stability

However, for problems with large regions with smooth coefficients, they are less efficient than difference methods. Here we demonstrate that Galerkin difference methods, constructed using the same mathematical formalism of standard DG schemes but employing translation-invariant basis functions, can be used to devise arbitrary-order energy-stable methods on hybrid grids to simulate waves. The proposed methods combine the advantages of existing DG and finite difference methods.

Thomas M. Hagstrom  
Southern Methodist University  
Department of Mathematics  
thagstrom@smu.edu

Jeff Banks  
Rensselaer Polytechnic Institute  
banksj3@rpi.edu

Jeremy E. Kozdon, Lucas Wilcox  
Department of Applied Mathematics  
Naval Postgraduate School  
jekozdon@nps.edu, lwilcox@nps.edu

### MS95 On the Stability of Mesh Adaptivity for Discontinuous Galerkin Methods

In recent years, there has been a recognition that use of the skew-symmetric formulation of the discontinuous Galerkin method for linear hyperbolic problems leads to provably stable, semi-discrete schemes (continuous time, spatially discrete) for problems with variable coefficients and curvilinear elements. By working with the skew-symmetric form, it can be ensured that any quadrature and interpolation errors, which we collectively call variational errors, do

not introduce any unphysical energy growth into the solution. Here we show how the skew-symmetric form can also be used to stabilize the weak energy instability that can arise from variation errors when adaptivity in element size  $h$  and polynomial order  $p$  are allowed for, e.g., so-called *hp*-adapted methods. This weak instability for *hp*-adapted methods arises from variational errors in non-conforming surface integrals; it should be noted that this weak instability can be present even for problems with affine meshes with constant coefficients. The skew-symmetric form of the equations introduces a splitting between the volume and surface terms. This then leads to stable semi-discrete schemes if both sides of the surface integrals are discretized in the same manner. These results hold true for problems with curvilinear elements and variable coefficients.

Jeremy E. Kozdon, Lucas Wilcox  
Department of Applied Mathematics  
Naval Postgraduate School  
jekozdon@nps.edu, lwilcox@nps.edu

### MS96 Inverse Random Source Scattering Problems in Several Dimensions

This talk concerns the source scattering problems for acoustic wave propagation, which is governed by the two- or three-dimensional stochastic Helmholtz equation. As a source, the electric current density is assumed to be a random function driven by an additive colored noise. Given the random source, the direct problem is to determine the radiated random wave field. The inverse problem is to reconstruct statistical properties of the source from the boundary measurement of the radiated random wave field. We consider both the direct and inverse problems. We show that the direct problem has a unique mild solution via a constructive proof. Using the mild solution, we derive effective Fredholm integral equations for the inverse problem. A regularized Kaczmarz method is developed by adopting multi-frequency scattering data to overcome the challenges of solving the ill-posed and large scale integral equations. Numerical experiments are presented to demonstrate the efficiency of the proposed method. The framework and methodology are expected to be applicable to a wide range of stochastic inverse source problems.

Peijun Li  
Department of Mathematics  
Purdue University  
lpeijun@math.purdue.edu

### MS96 Scattering and Field Enhancement of Narrow Slits

Subwavelength apertures and holes on surfaces of noble metals (e.g., gold or silver) induce strong electromagnetic field enhancement and extraordinary optical transmission. This remarkable phenomenon can lead to potentially significant applications in biological and chemical sensing, spectroscopy, and other novel optical devices. In this talk, I will present a quantitative analysis for the field enhancement of narrow slits perforated in a slab of perfect conductor. Both the single slit and periodic slits will be discussed. We demonstrate that the enhancement of the electromagnetic field can be induced by either scattering resonances or certain non-resonant effect in the quasi-static regime. We derive the asymptotic expansions of the resonances and quantitatively analyze the field enhancement at the resonant frequencies. The field enhancement at non-resonant

frequencies in the quasi-static regime is also investigated. It is shown that the fast transition of the magnetic field in the slit induces strong electric field enhancement.

Junshan Lin  
Auburn University  
jzl0097@auburn.edu

### MS96

#### Spectral Element Simulation for Nanowire Solar Cells on HPC Platforms

We consider solving drift-diffusion systems describing the electron transport phenomena for various configurations of nanowire solar cells. Spectral element discretizations and advanced time stepping methods based on both semi-implicit explicit and characteristic schemes. Discussion includes high performance aspects on the ALCF Theta and OLCF Titan including Xeon Phi and multiGPUs.

MiSun Min  
Argonne National Laboratory  
Mathematics and Computer Science Division  
mmin@mcs.anl.gov

### MS96

#### Nonlinear Photoacoustic Tomography with Two-Photon Absorption

Two-photon photoacoustic tomography (TP-PAT) is a non-invasive optical molecular imaging modality that aims at inferring two-photon absorption property of heterogeneous media from photoacoustic measurements. In this work, we analyze an inverse problem in quantitative TP-PAT where we intend to reconstruct optical coefficients in a semilinear elliptic PDE, the mathematical model for the propagation of near infra-red photons in tissue-like optical media, from the internal absorbed energy data. We derive uniqueness and stability results on the reconstructions of single and multiple coefficients, and propose some efficient numerical algorithms for the reconstruction.

Kui Ren, Rongting Zhang  
University of Texas at Austin  
ren@math.utexas.edu, rzhang@math.utexas.edu

### MS97

#### Energy-Conservative HDG Methods for the Wave Equation

We present a class of energy-conservative hybridizable discontinuous Galerkin methods for the numerical simulation of acoustics wave equations. The methods possess several unique features which distinguish themselves from other discontinuous Galerkin methods for wave propagation problems. The first feature is that the method have significantly reduced global degrees of freedom. The second feature is that the all the approximate variables converge with the optimal order of  $k+1$  in the  $L^2$ -norm, when polynomials of degree  $k \geq 0$  are used to represent the numerical solution and when the time-stepping method is accurate with order  $k+1$ . The third feature is that the super-convergence property of the methods allows us to obtain new approximations of the displacement and velocity field which converge with order  $k+2$  (for  $k \geq 1$ ) in the  $L^2$ -norm. Finally, as for the fourth feature, the methods conserve the energy exactly. We present numerical results to illustrate

these features.

Cuong Nguyen  
Massachusetts Institute of Technology  
cuongng@mit.edu

Cristian Ciuca  
Imperial College  
cristian.ciuca13@imperial.ac.uk

Jaime Peraire  
Massachusetts Institute of Technology  
peraire@MIT.EDU

Bernardo Cockburn  
School of Mathematics  
University of Minnesota  
cockburn@math.umn.edu

### MS97

#### Stage-Parallel Implicit Runge-Kutta Time-Integration and Efficient Approximate Block Preconditioning for Discontinuous Galerkin Methods

We present some recent results in our work on efficient time-stepping and nonlinear solvers for high-order discontinuous Galerkin (DG) discretizations. Using a stage-wise ILU(0) preconditioner and a shifted block-Jacobi preconditioner between the stages, we can solve fully coupled implicit Runge-Kutta systems at a cost comparable to DIRK schemes. This also decouples the stages and allows for significant parallelism in time and better strong scaling properties. We also show a new approximate block preconditioner for DG discretizations of the compressible Navier-Stokes equations with high polynomial degrees.

Per-Olof Persson  
Dept. of Mathematics  
University of California, Berkeley  
persson@berkeley.edu

Will Pazner  
Division of Applied Mathematics  
Brown University  
will\_pazner@brown.edu

### MS97

#### Metric-Based HP-Adaptation Using a Continuous Mesh Model

Anisotropic adaptive meshing is widely recognized as an important tool in the numerical solution of convection-diffusion problems. Furthermore, high-order consistent numerical schemes using piecewise polynomial approximation call for hp-adaptation to maximize efficiency. Here we present a generalized mesh adaptation and mesh optimization method for discontinuous Galerkin Schemes using a metric-based continuous mesh model. The rationale behind the new model is to incorporate as much as possible analytic optimization techniques acting on the continuous mesh, as opposed to numerical optimization acting on the discrete mesh. In particular we discuss hp-adaptivity, as well as target-based adaptation in this context. The former requires the formulation of a suitable continuous hp-mesh, as well as a corresponding hp-optimization framework. The latter is important in a variety of applications, where one is interested in computing certain solution-dependent func-

tionals rather accurately, as opposed to minimizing the global error norm. To incorporate target-based adaptation we extend previous continuous mesh models to a weighted-norm optimization, where the weight comes from an adjoint solution providing sensitivities with respect to some relevant target functional. We present the formulation of these continuous-mesh models, as well as numerical validation for several advection-diffusion problems.

Georg May  
AICES  
RWTH Aachen  
may@aices.rwth-aachen.de

Ajay Rangarajan  
RWTH Aachen  
ajay.rangarajan@rwth-aachen.de

Vit Dolejsi  
Charles University Prague  
Faculty of Mathematics and Physics  
dolejsi@karlin.mff.cuni.cz

### MS97

#### Towards Combined CG-DG for Elliptic Problems

We combine continuous and discontinuous Galerkin methods in the setting of a model diffusion problem. Starting from a hybrid discontinuous formulation, we replace element interiors by more general subsets of the computational domain - groups of elements that support a piecewise-polynomial continuous expansion. This step allows us to identify a new weak formulation of Dirichlet boundary condition in the continuous framework. We show that the boundary condition leads to a stable discretization with a single parameter insensitive to mesh size and polynomial order of the expansion. The robustness of the discretization is demonstrated on several numerical examples. The geometric flexibility of combined CG-DG solver offers potential for improved scaling on massively parallel computers. Since the continuous patches are only weakly coupled by a hybrid variable, communication on machines with distributed memory is reduced in comparison with classical continuous Galerkin discretization. The node-local continuous expansions, on the other hand, allow for reduced memory footprint since the number of unknowns is smaller than in the discontinuous case. We present preliminary analysis of parallel communication pattern of such continuous-discontinuous solver.

Martin Vymazal  
Imperial College  
martin.vymazal@imperial.ac.uk

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

### MS98

#### On All-Regime and High-Order Lagrange-Remap Numerical Schemes for Compressible Fluids Systems

It will be the purpose of this contribution to provide an overview of recent advances in the development of high-order and all-regime Lagrange-Remap numerical schemes for compressible fluids systems with or without source terms. We will consider in particular the two cases of large

friction coefficients and low Mach numbers.

Christophe Chalons  
Université Versailles Saint-Quentin-en-Yvelines,  
christophe.chalons@uvsq.fr

### MS98

#### Conservation and Positivity Preserving Issues for Kinetic Collisional Transport Models

We'll discuss how conservation and positivity preserving properties are possible for Boltzmann type equations where the collision operator is mass preserving and dissipative. The problem for the linear case has been fully solved by means of DG schemes. However, for the case of the classical nonlinear collision operator, a Galerkin Petrov approach is more appropriate tool. This communication is from work in collaboration with from several publications whose authors include Y. Cheng, J. Proft, J Morales Escalante, E. Endeve, C. Hauck, and S. Rjasanow.

Irene M. Gamba  
Department of Mathematics and ICES  
University of Texas  
gamba@math.utexas.edu

### MS98

#### An Implicit Positivity Preserving Discontinuous Galerkin Method for Solving Conservation Laws

Explicit DG schemes for solving conservation laws suffer from restrictive CFL constraint, especially when the mesh is not uniform. In this talk, we report our recent attempt to design an implicit positivity preserving DG scheme for solving conservation laws. Primary results will be discussed and numerical results will be shown.

Tong Qin  
Brown University  
tong\_qin@brown.edu

### MS98

#### Positivity-Preserving High Order Discontinuous Galerkin Schemes for Compressible Navier-Stokes Equations

For gas dynamics equations such as compressible Euler and Navier-Stokes equations, preserving the positivity of density and pressure without losing conservation is crucial to stabilize the numerical computation. The L1-stability of mass and energy can be achieved by enforcing the positivity of density and pressure during the time evolution. However, high order schemes such as DG methods do not preserve the positivity. It is difficult to enforce the positivity without destroying the high order accuracy and the local conservation in an efficient manner for time-dependent gas dynamics equations. For compressible Euler equations, a weak positivity property holds for any high order finite volume type schemes including DG methods, which was used to design a simple positivity-preserving limiter for high order DG schemes in Zhang and Shu, JCP 2010. Generalizations to compressible Navier-Stokes equations are however nontrivial. We show that weak positivity property still holds for DG method solving compressible Navier-Stokes equations if a proper nonlinear penalty term is added to any finite volume or DG scheme. This allows us to obtain the first high order positivity-preserving schemes for

compressible Navier-Stokes equations.

Xiangxiong Zhang  
Purdue University  
zhan1966@purdue.edu

### MS99

#### Gmsfem for a Class of Nonlinear PDEs

Abstract not available

Eric Chung  
The Chinese University of Hong Kong  
Department of Mathematics  
tschung@math.cuhk.edu.hk

### MS99

#### Efficient Simulation of Asymptotically Disappearing Solutions for Wave and Maxwell's Equations

This work is on the numerical approximation of incoming solutions to the wave and equation, as well as Maxwell's equations. We consider solutions whose energy decays exponentially with time (asymptotically disappearing), meaning that the leading term of the back-scattering matrix becomes negligible. We consider a finite-element approximation, which leads to a parameter-dependent eigenvalue problem. We design an efficient method together with robust preconditioners to solve the eigenvalue problem. Numerically, we demonstrate the efficiency of the proposed approach.

James H. Adler, Xiaozhe Hu  
Tufts University  
james.adler@tufts.edu, xiaozhe.hu@tufts.edu

Vesselin Petkov  
Université Bordeaux 1  
vesselin.petkov@math.u-bordeaux1.fr

Ludmil Zikatanov  
Pennsylvania State University  
ludmil@psu.edu

### MS99

#### Finite Element Methods for the Stochastic Allen-Cahn Equation with Gradient-Type Multiplicative Noises

This talk studies the finite element approximations of the stochastic Allen-Cahn equation with gradient-type multiplicative noise that is white in time and correlated in space. The sharp interface limit as the diffuse interface thickness vanishes of the stochastic equation formally approximates a stochastic mean curvature flow which is described by a stochastically perturbed geometric law of the deterministic mean curvature flow. Two fully discrete finite element methods which are based on different time-stepping strategies for the nonlinear term are proposed. Strong convergence with sharp rates for both fully discrete finite element methods is proved. This is done with the crucial help of the Hölder continuity in time with respect to the spatial  $L^2$ -norm and  $H^1$ -seminorm for the strong solution of the stochastic Allen-Cahn equation, which are key technical lemmas. It also relies on the fact that high moments of the strong solution are bounded in various spatial and temporal norms. Finally, some numerical experiments are provided to gauge the performance of the proposed fully discrete

finite element methods and to study the interplay of the geometric evolution and gradient-type noise.

Xiaobing H. Feng  
The University of Tennessee  
xfeng@math.utk.edu

Yukun Li  
Pennsylvania State University  
Department of Mathematics  
yfl5134@psu.edu

Yi Zhang  
University of Notre Dame  
yzhang41@nd.edu

### MS99

#### An Efficient Adaptive Grid Method for the Numerical Solution of Phase-Field Models

We implement an efficient adaptive grid method for solving time dependent phase-field models. The adaptive mesh is computed based solving a parabolic Monge-Ampère equation (PMAE). An alternate procedure is used for computing the phase function and the adaptive mesh. We present several numerical experiments to demonstrate the accuracy and effectiveness of the PMAE adaptive grid method.

Mohamed Sulman  
Wright State University  
Department of Mathematics and Statistics  
mohamed.sulman@wright.edu

### MS100

#### Adaptive Quantum Chemistry Methods for Strongly Correlated Electrons with Tunable Accuracy

An outstanding challenge in the field of Quantum Chemistry is the development of new computational methods that can deal with the phenomenon of strong electron correlation. Wave functions of strongly correlated electrons have a structure that is not efficiently described by conventional wave function parameterizations. In recent years, two strategies have emerged to deal with strong electron correlation: 1) tensor factorization approaches, and 2) Monte Carlo methods. In this talk, I will discuss our efforts to develop an alternative strategy which exploits the sparsity of the FCI wave function. Our focus is on novel electronic structure theories that are adaptive and tunable. I will summarize our recent developments of the adaptive configuration interaction method and discuss its extensions to treat electronically excited states. Several strategies to generalize our adaptive configuration interaction approach will be discussed, focusing in particular on the computation of dissociation curves, conical intersections, and avoided crossings.

Francesco Evangelista, Jeffrey Schriber, Tianyuan Zhang  
Emory University  
francesco.evangelista@emory.edu,  
jeffrey.blair.schriber@emory.edu,  
tianyuan.zhang@emory.edu

### MS100

#### Methodological Developments in the Calculation of Excited-State Properties: Large Scale GW Calculations

**lations**

Although one of the main tools used in first principle simulations of materials is Density Functional Theory (DFT), several of the current approximations of exchange and correlation functionals do not provide the level of accuracy required for predictive calculations of excited state properties. The application to large systems of more accurate post-DFT approaches such as Many-Body Perturbation Theory (MBPT) – for example to heterogeneous systems, nanostructured, disordered, and defective materials – has been hindered by high computational costs. In this talk recent methodological developments in MBPT calculations will be discussed, as recently implemented in the open source code WEST [www.west-code.org], which efficiently exploits massively parallel architectures. Results using a formulation that does not require the explicit calculation of virtual states, nor the storage and inversion of large dielectric matrices will be presented; these results include quasi particle energies for systems with thousands of electrons and encompass the electronic structure of aqueous solutions, spin defects in insulators, and benchmarks for molecules and solids containing heavy elements. Simplifications of MBPT calculations based on the use of static response properties, such as dielectric-dependent hybrid functionals, will also be discussed.

Marco Govoni

Argonne National Lab and University of Chicago  
mgovoni@anl.gov

Giulia Galli

University of Chicago and ANL  
gagalli@uchicago.edu

**MS100****Recent Progress on Reducing the Cost of the Fock Exchange Operator**

The Fock exchange operator, such as in Hartree-Fock theory and Kohn-Sham density functional theory with hybrid functionals, plays a central role in characterizing ground and excited state properties in modern quantum chemistry and materials science. However, the cost associated with the Fock exchange operator is very high. We recently developed the adaptively compressed exchange operator formulation (ACE), which greatly reduces the computational cost without loss of accuracy for both insulators and metals. In particular, we can perform hybrid functional calculations in planewave basis sets with more than 4000 atoms.

Lin Lin

University of California, Berkeley  
Lawrence Berkeley National Laboratory  
linlin@math.berkeley.edu

**MS100****Electronic Density of States for Incommensurate Layers**

We prove that the electronic density of states (DOS) for 2D incommensurate layered structures, where Bloch theory does not apply, is well-defined as the thermodynamic limit of finite clusters. In addition, we obtain an explicit representation formula for the DOS as an integral over local configurations. Next, based on this representation formula, we propose a novel algorithm for computing electronic structure properties in incommensurate heterostructures, which overcomes limitations of the common approach to arti-

cially strain a large supercell and then apply Bloch theory.

Mitchell Luskin

School of Mathematics  
University of Minnesota  
luskin@math.umn.edu

Christoph Ortner

University of Warwick  
christophortner0@gmail.com

Daniel Massatt

University of Minnesota  
massa067@umn.edu

**MS101****The Design and Implementation of a Dense Linear Algebra Library for Extreme Parallel Computers**

In this talk we will present the plan for the implementation of a dense linear algebra library for highly parallel computing systems and the need for additional Batched operations.

Jack J. Dongarra

University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

Nicholas Higham

The University of Manchester  
nick.higham@manchester.ac.uk

Zounon Mawussi

INRIA  
mawussi.zounon@inria.fr

Samuel Relton

University of Manchester, UK  
samuel.relton@manchester.ac.uk

**MS101****Sparse Direct Solvers for Extreme Scale Computing**

As part of the H2020 FET-HPC Project NLA-FET, we are studying the scalability of algorithms and software for using direct methods for solving large sparse equations. Although algorithms and software for the sparse direct solution of linear equations are far more complicated than their counterparts for dense systems, we show that there are far more possibilities for exploiting parallelism when the systems are sparse. We first look at algorithms for solving sparse symmetric systems both when the system is positive-definite and when it is indefinite so that numerical pivoting is required. We study the benefits of using standard run time systems to assist us in developing codes for extreme scale computers. We show that we are competitive with prior approaches and codes but have the flexibility to move easily to new architectures. Other techniques are required when the matrix is highly unsymmetric and we discuss how methods like Markowitz pivoting can be developed to exploit parallelism.

Jain Duff

Science & Technology Facilities Council, UK  
and CERFACS, Toulouse, France

iain.duff@stfc.ac.uk

Jonathan Hogg  
STFC Rutherford Appleton Lab, UK  
jonathan.hogg@stfc.ac.uk

Florent Lopez, Stojce Nakov  
STFC-RAL  
florent.lopez@stfc.ac.uk, stojce.nakov@stfc.ac.uk

### MS101

#### Computing the Low Rank Approximation of a Sparse Matrix

In this talk we discuss an algorithm for computing a low rank approximation of a sparse matrix based on a truncated LU factorization with column and row permutations. We present various approaches for determining the column and row permutations that show a trade-off between speed versus deterministic/probabilistic accuracy. We focus then on an approach that uses tournament pivoting based on QR factorization and we compare the bounds obtained on the singular values with the ones obtained by a communication avoiding rank revealing QR factorization. We also discuss the parallel efficiency of the obtained algorithm on diverse applications going from integral equations to image processing.

Laura Grigori  
INRIA  
France  
Laura.Grigori@inria.fr

Sebastien Cayrols  
INRIA  
sebastien.cayrols@inria.fr

James W. Demmel  
University of California  
Division of Computer Science  
demmel@berkeley.edu

Alan Ayala  
Inria  
alan.ayala-obregon@inria.fr

### MS101

#### Extreme-Scale Eigenvalue Reordering in the Real Schur Form

Eigenvalue reordering in the real Schur form can be used to compute an orthonormal basis for the invariant subspace belonging to a user-selected subset of the eigenvalues. Reordering is key also in other applications such as deflations and restarts in the Jacobi-Davidson and Krylov-Schur algorithms. Algorithms for reordering rely on some method for robustly swapping adjacent blocks on the diagonal of the Schur matrix. Numerous swaps are then combined in some suitable way with the goal of making the selected blocks/eigenvalues bubble up to the top left corner. There exists several techniques for swapping adjacent blocks, but from a parallel computing perspective they are all essentially interchangeable. What matters most is not how each swap is performed but rather how the many individual swaps are ordered and grouped with the dual aims of improving the cache reuse and distributing the work effectively across the machine. As a part of the NLAFFET project funded by the EU Horizon 2020 FET-HPC pro-

gramme, we are developing new parallel formulations of time-tested numerical schemes for eigenvalue reordering. Our aim is to develop the most scalable parallel implementation to date by taking advantage of recent advances in task-based runtime systems as well as auto-tuning techniques.

Lars Karlsson, Carl Christian Kjelgaard Mikkelsen  
Umeå University, Dept. of Computing Science  
larsk@cs.umu.se, spock@cs.umu.se

Bo T. Kågström  
Umeå University  
Computing Science and HPC2N  
bokg@cs.umu.se

### MS102

#### Parallel Rational Arnoldi Algorithm and Applications

The rational Arnoldi algorithm is a popular method in scientific computing used to construct an orthonormal basis of a rational Krylov space. Each basis vector is a rational matrix function times the starting vector. Rational functions possess a partial fraction expansion which often allows to compute several basis vectors simultaneously. However, this parallelism may cause instability due to the orthogonalization of ill-conditioned bases. We present continuation strategies to minimize these effects and compare them within different application areas.

Mario Berljafa, Stefan Guettel  
The University of Manchester  
mario.berljafa@manchester.ac.uk, stefan.guettel@manchester.ac.uk

### MS102

#### The Chase Library and Its Application to Excitonic Hamiltonians

Numerically solving the Bethe-Salpeter equation for the optical polarization function is a very successful approach for describing excitonic effects in first-principles simulations of materials. While the accurate computation of exciton binding energies enable materials design, these results come at high computational cost: For modern complex materials this approach leads to large, dense matrices with sizes reaching up to  $n \sim 400k$ . Since the experimentally most relevant exciton binding energies require only few hundreds of the lowest eigenpairs, iterative schemes are a feasible alternative to prohibitively expensive direct diagonalization. The Chebyshev Accelerated Subspace iteration Eigensolver (ChASE) library, which is developed at Juelich Supercomputing Centre, is an ideal package for solving such large dense eigenvalue problems. ChASE leverages on the preponderant use of BLAS 3 subroutines to achieve close-to-peak performance. Moreover, the code is parallelized for many- and multi-core platforms. In this contribution we show its application to problems extracted from excitonic Hamiltonian and show how a new distributed CPU/GPU implementation of ChASE allows for the solution of larger eigenproblems by effectively exploiting heterogeneous multi-GPU architectures.

Edoardo A. Di Napoli  
Juelich Supercomputing Centre  
e.di.napoli@fz-juelich.de

Jan Winkelmann

AICES  
RWTH Aachen  
winkelmann@aices.rwth-aachen.de

Andre Schleife  
University of Illinois at Urbana-Champaign  
schleife@illinois.edu

### MS102

#### Applications of the Parallel Complex Moment-Based Eigensolver Package z-Pares to Large-Scale Scientific Computations

z-Pares is a software package implements the Sakurai-Sugiura method which is a contour integral eigenvalue solver using higher-order complex moments. The use of higher-order complex moments enables us to make a large number of basis for the projection with limited computational cost. z-Pares has two-level MPI parallelism which is capable of taking advantage of the hierarchical structure of recent massively parallel supercomputers. In this presentation, we show applications of z-Pares to large-scale scientific computations such as electronic structure calculations and vibration analysis on supercomputers.

Yasunori Futamura  
University of Tsukuba  
futamura@cs.tsukuba.ac.jp

Akira Imakura  
Department of Computer Science  
University of Tsukuba  
imakura@cs.tsukuba.ac.jp

Tetsuya Sakurai  
University of Tsukuba  
sakurai@cs.tsukuba.ac.jp

### MS102

#### Experiences with GPU Use in Eigenvalue Computations with SLEPc

In current supercomputers we can often find GPUs in (at least part of) the nodes. The challenge for library developers is to adapt the software to (optionally) exploit this hardware in an efficient way, without obliging the user to revamp the application code. In this talk we discuss this issue in the context of SLEPc, the Scalable Library for Eigenvalue Computations, for the computation of a few eigenpairs of a matrix with Krylov methods. We show results in two different contexts. In the first one, an application coming from molecular magnetism, the matrix is sparse and all operations (matrix-vector product and orthogonalization of vectors) are handled internally by SLEPc and carried out with MPI-GPU parallelism. In the second one, the matrix is block-tridiagonal (or banded) and interior eigenvalues are computed via shift-and-invert, hence performing linear solves at each eigensolver iteration. In this latter case, sparse matrix storage must be avoided, and we illustrate how it is possible to provide user-defined code that specifically solves block-tridiagonal linear systems with MPI-GPU parallelism. In both cases, since computing matrix coefficients is expensive, the user has to provide GPU code for that, otherwise the performance gain may be small. This is joint work with A. Lamas Daviña.

Jose E. Roman  
Universidad Politecnica de Valencia

jroman@dsic.upv.es

### MS103

#### Updating Dynamic Networks in Parallel Using Graph Sparsification

We present graph sparsification, an elegant technique for updating the properties of dynamic networks. Using graph sparsification the original network is divided into several small subgraphs. Each subgraph contains a set of specially marked edges, known as key edges, that pertain to the property to be updated. Each addition/deletion of edges is updated with respect to these key edges. In our presentation, we will show how by using graph sparsification we can develop scalable algorithms to update properties of weighted graphs such as minimum spanning tree and single source shortest paths.

Sanjukta Bhowmick  
Department of Computer Science  
University of Nebraska, Omaha  
sbhowmick@unomaha.edu

Sriram Srinivasan  
Department of Computer Science  
University of Nebraska, Omaha  
sriram882004@gmail.com

Sajal Das  
Missouri University of Science and Technology  
sdas@mst.edu

### MS103

#### Maintaining Connected Components for Infinite Graph Streams

We present an algorithm to maintain the connected components of a graph that arrives as an infinite stream of edges. Connectivity-related queries, including component spanning trees, are supported with some latency, returning the state of the graph at the time of the query. Because an infinite stream may eventually exceed the storage limits of any number of finite-memory processors, we assume an aging command or daemon where uninteresting edges are removed when the system nears capacity. Following an aging command the system will block queries until its data structures are repaired, but edges will continue to be accepted from the stream, never dropped. The algorithm will not fail unless a model-specific constant fraction of the aggregate memory across all processors is full. In normal operation, it will not fail unless aggregate memory is completely full. Unlike previous theoretical streaming models designed for finite graphs that assume a single shared memory machine or require arbitrary-size intermediate files, we distribute a graph over a ring network of finite-memory processors. We implemented our algorithm using an asynchronous message-passing system. We sketch the algorithm and give preliminary experimental results on synthetic and real graph streams.

Jonathan Berry  
Sandia National Laboratories  
jberry@sandia.gov

Matthew Oster  
Pacific Northwest National Laboratory  
matthew.oster@pnnl.gov



Cynthia Phillips, Steve Plimpton, Timothy Shead  
Sandia National Laboratories  
caphill@sandia.gov, sjplimp@sandia.gov,  
tshead@sandia.gov

**MS103****High-Performance Graph Traversal for De Bruijn Graph-Based Metagenome Assembly**

De Bruijn graph-based assembly is a popular technique for analyzing modern genomic and metagenomic DNA sequencing data. This technique is comprised of several computational building blocks, with the ultimate goal of reconstructing genomes based on overlapping shorter strings. This work describes new scalable, parallel implementations of several key compute- and memory-intensive steps of the De Bruijn graph-based assembly technique. We focus on exploiting multicore parallelism and parallel I/O capabilities available on current supercomputers, and also target memory efficiency, i.e., fully in-memory execution using as few compute nodes as possible. We demonstrate applicability of these optimized implementations by integrating them with the MEGAHIT metagenome assembler.

Vasudevan Rengasamy, Kamesh Madduri  
Pennsylvania State University  
vas.renga@gmail.com, madduri@cse.psu.edu

**MS103****Computing Graph Centrality**

Computing the centrality of a graph is a useful tool for many applications ranging from social network to traffic analysis, from drug discovery to urban planning. While many definitions of centrality exist, the most commonly used – namely betweenness and closeness centrality – are variations on the theme of all-pair shortest path. With a complexity of  $O(VE)$ , computing centrality is difficult in practice; computing centralities efficiently has attracted quite some interest in the recent years. In this talk, we will review the techniques one has to deploy to reach the highest performance when computing graph centrality. In both case of static graphs and incremental graphs, we will show that computing centrality on modern computing system involves both algorithmic and system issues.

Erik Saule  
University of North Carolina at Charlotte  
esaule@uncc.edu

A. Erdem Sariyuce  
The Ohio State University  
asariyu@sandia.gov

Kamer Kaya  
Sabanci University,  
Faculty of Engineering and Natural Sciences  
kaya@sabanciuniv.edu

Umit V. Catalyurek  
The Ohio State University  
Department of Biomedical Informatics  
umit@gatech.edu

**MS104****H Or Not to H: Hierarchical vs Dense Matrix Com-****putations on Manycore Architectures**

The talk will describe the implementation of the Cholesky factorization applied to hierarchical matrices (H-matrices) exhibiting, low rank block structure. Impacts on performance and memory footprint will be reported on Intel KNL and NVIDIA GPUs in the context of spatial statistics and computational astronomy applications. Performance analysis will be highlighted against state-of-the-art high performance dense linear algebra libraries.

Kadir Akbudak  
Bilkent University  
kadir.cs@gmail.com

Ali Charara, David E. Keyes  
KAUST  
ali.charara@kaust.edu.sa, david.keyes@kaust.edu.sa

Hatem Ltaief, Aleksandr Mikhalev  
King Abdullah University of Science & Technology (KAUST)  
hatem.Ltaief@kaust.edu.sa,  
aleksandr.mikhalev@kaust.edu.sa

George M Turkiyyah  
King Abdullah University of Science and Technology  
American University of Beirut  
gt02@aub.edu.lb

**MS104****Implementation Techniques for High Performance Blas Kernels on Modern GPU**

This talk will describe our recent experiences in developing high performance BLAS kernels on modern GPUs. We developed a model-driven method to determine the optimal thread-block size automatically on memory-bound kernels for multiple NVIDIA GPU architectures. In this talk, fast implementations of matrix-vector multiplication routines using the method will be presented. Moreover, some of our ongoing work related to the BLAS implementation will be introduced.

Daichi Mukunoki  
RIKEN, Japan  
mukunoki@hpcs.cs.tsukuba.ac.jp

**MS104****An Applications Perspective on Multi-Core, Massive Multi-Threading, and Hybrid Systems**

Multi/many-core and hybrid architectures that support massive multi-threading have raised considerable uncertainty as to what programming models might be appropriate. We will discuss applications that have been developed within the Swiss High-Performance and High-Productivity Computing platform and are exploiting emerging computer architecture quite successfully. We will see what architectural aspects are important for the various algorithmic motifs that appear in applications, and what new programming models seem to find broad acceptance among scientific programmers.

Thomas C. Schulthess, Thomas C. Schulthess  
ETH Zurich  
Swiss National Supercomputing Center (CSCS)

schulthess@cscs.ch, schulthess@cscs.ch

### MS104

#### Implementation of Parallel FFTs on Knights Landing Cluster

In this talk, we propose an implementation of parallel fast Fourier transforms (FFT) on Intel Xeon Phi clusters. To obtain peak performance of the Intel Xeon Phi processor, it must be both vectorized and multithreaded. We present an approach that makes use of the MCDRAM of the Knights Landing processor by blocking technique. We also present a computation-communication overlap method that introduces a communication thread with OpenMP. Performance results of FFTs on a Knights Landing cluster are reported.

Daisuke Takahashi

Faculty of Engineering, Information and Systems  
University of Tsukuba  
daisuke@cs.tsukuba.ac.jp

### MS105

#### Optimal Numerical Methods for Stochastic PDEs

Stochastic PDEs pose new numerical challenges because of the large number of stochastic dimensions in important applications. We consider two model problems: the numerical stochastic homogenization of elliptic problems and the stochastic drift-diffusion-Poisson system. The first model is important in materials science, while the second describes charge transport in random environments, e.g., in nanoscale semiconductor devices. In both cases, we could reduce the computational effort by orders of magnitude by devising optimal numerical algorithms based on the multi-level Monte-Carlo method and various extensions. The algorithms are optimal in the sense that the computational effort is minimal for given error bounds. Finally, we present numerical results for the model equations. The results are obtained using our codes that are mostly written in Julia.

Clemens Heitzinger

Vienna University of Technology  
clemens.heitzinger@tuwien.ac.at

Amirreza Khodadadian

Technical University Vienna (TU Vienna).  
amirreza.khodadadian@tuwien.ac.at

Gudmund Pammer

Vienna University of Technology  
gudmund.pammer@tuwien.ac.at

Stefan Rigger

TU Wien & ASU  
stefan.rigger@tuwien.ac.at

Leila Taghizadeh

Vienna University of Technology  
leila.taghizadeh@tuwien.ac.at

### MS105

#### Fast and Memory Efficient Statistical Analysis of Large Genomic Data Sets

We demonstrate how the combination of high and low-level programming language features, when made available in a single programming language, allow for fast and generic

code. Easy access to low-level data types and machine instructions makes it possible to work with large genomics data sets in compressed form without sacrificing functionality or speed. This makes problems shrink from server size to desktop size and opens the door for much larger problems than previously considered. At the same time, high-level language features make it possible to hide the low-level optimizations behind linear algebra abstractions through which a wide range of generic linear algebra routines become available. In particular, we are able to use a generic Lanczos-based SVD solver which has already proven to be superior to existing tools in the genomics software stack. More generally, as the abstractions are provided at a BLAS level, the low-level optimizations are available for a wide range of problems including many of the common statistical models for the analysis of genomics data.

Andreas Noack

Computer Science and Artificial Intelligence Laboratory  
Massachusetts Institute of Technology  
noack@csail.mit.edu

### MS105

#### Krylov Methods with a Forward Error Minimization Property and Application to Optimization

We propose an iterative method named LSLQ for solving consistent linear systems or linear least-squares problems  $Ax \approx b$  of any shape based on the Golub-Kahan process, where the dominant cost consists in products with  $A$  and  $A^T$ . In the rank deficient case, LSLQ identifies the minimum least-squares solution. LSLQ is formally equivalent to SYMMLQ (Paige and Saunders, 1975) applied to the normal equations so that the estimate norm  $\|x_k\|_2$  increases monotonically and the forward error  $\|x_k - x^*\|_2$  decreases monotonically. We provide lower and upper bound estimates on the forward error along the LSLQ iterations. The upper bound translates to an upper bound on the forward error in the Euclidean norm for LSQR, which was previously unavailable. We report numerical experiments on standard test problems and on a full-wave inversion problem arising from geophysics in which an approximate least-squares solution corresponds to an approximate gradient of a relevant penalty function that is to be minimized.

Dominique Orban

GERAD and Dept. Mathematics and Industrial Engineering  
Ecole Polytechnique de Montreal  
dominique.orban@gerad.ca

Ron Estrin

ICME  
Stanford University  
restrin@stanford.edu

Michael A. Saunders

Systems Optimization Laboratory (SOL)  
Dept of Management Sci and Eng, Stanford  
saunders@stanford.edu

### MS105

#### Shifted Laplacian Multigrid for the Elastic Helmholtz Equation

The shifted Laplacian multigrid is a well known approach for preconditioning the indefinite linear system arising from discretizing the acoustic Helmholtz equation, which is used

to model wave propagation. However, in some cases the acoustic equation is not sufficient for modeling the physics of the wave propagation, and one has to consider the elastic Helmholtz equation. Such a case arises in geophysical imaging applications, where the earth subsurface is the elastic medium. In this talk we extend the sifted Laplacian approach to the elastic Helmholtz equation, by combining the complex shift idea with approaches for linear elasticity. We show numerical experiments for problems with heterogeneous media.

Eran Treister

University of British Columbia  
erant@bgu.ac.il

Eldad Haber

University of British Columbia, Vancouver, Canada  
ehaber@eos.ubc.ca

### MS106

#### On the use of the Mori-Zwanzig Formalism to Build Closures for Reduced-Order Models

The use of reduced-order modeling (ROM) techniques in problems that lack a clear separation of scales is hampered by our inability to derive consistent closures. Accordingly, these methods are only applied successfully to problems in which interactions between the retained and discarded modes do not result in a significant loss of accuracy or stability. In this work, we use the Mori-Zwanzig (M-Z) formalism to represent the unresolved physics in coarse-grained simulations. In the M-Z approach, the unclosed terms can be formally represented as a memory integral. The general M-Z procedure to compute the memory effects is, however, intractable in real problems. We devise an alternate procedure to approximate the orthogonal dynamics and the memory kernel. This representation is used to guide the development of finite-memory models for the kernel in a ROM setting. Using this procedure, closure models are developed and demonstrated for ROMs of turbulent flow.

Ayoub Gouasmi

University of Michigan  
Department of Aerospace Engineering  
gouasmi@umich.edu

Karthik Duraisamy

University of Michigan Ann Arbor  
kdur@umich.edu

### MS106

#### A Multiscale Reduced-Order Model for NS Equations from Residual Minimization

In this contribution we explore some numerical alternatives to derive efficient and robust low-order models of the Navier-Stokes equations based on residual minimization. We start from the fact that classical Galerkin or Petrov-Galerkin approaches for ROM can be derived in the context of a residual minimization method similar to variational multi-scale modeling, VMS [Hugues TJR, Feijoo G, Mazzei L, Quincy JB. The variational multiscale method: a paradigm for computational mechanics. *CMAME* 1998; 166:324.]. Based on this, we introduce a residual minimization scheme that directly includes VMS stabilizing terms in the low-order model as proposed in [Bergmann M, Bruneau C, Iollo A. Enablers for robust pod models. *JCP* 2009; 228(2):516538.], [Weller J., Lombardi E., Bergmann, Iollo A. Numerical methods for low-order modeling of fluid flows

based on POD *Int. J. Numer. Meth. Fluids* 2009]. Here, however, the unknowns of the minimization problem are the union of the coefficients of a modal representation of the solution and of the physical unknowns at certain collocation points [Buffoni, M., Telib, H., & Iollo, A. (2009). Iterative methods for model reduction by domain decomposition. *Computers & Fluids*, 38(6), 11601167.]. The modal representation is typically based on empirical eigenfunctions obtained by proper-orthogonal decomposition, whereas the residual at collocation points are obtained by an adaptive discretization.

Andrea Ferrero

Inria Bordeaux Sud - Ouest  
andrea.ferrero@inria.fr

Michel Bergmann

INRIA Bordeaux sud ouest and IMB, Bordeaux, France  
Michel.Bergmann@inria.fr

Angelo Iollo

Institut de Mathématiques de Bordeaux  
angelo.iollo@math.u-bordeaux1.fr

### MS106

#### Strategies for Modeling Nonlinear Mechanisms from Data

The tremendous advances in recent years in computing power, new sensors and infrastructures allow the collection and real-time access of massive amounts of data. The modeling and control effort of multi-scale, high-dimensional, nonlinear dynamical systems such as turbulent fluid flows is increasingly facilitated by these developments as new insights on physical mechanisms can be discovered and later be manipulated in real-time. Although we now have access to an abundance of data, we are still facing the challenge of extracting the underlying dynamics. Recent advances in equation-free architectures connected to operator-theoretic methods are an appealing direction. Combined with compressed sensing and machine learning techniques, low-dimensional dynamics can be learned directly from data in an unsupervised manner.

Eurika Kaiser

University of Washington  
eurika@uw.edu

### MS106

#### Large Eddy Simulation Reduced Order Models

This talk proposes several large eddy simulation reduced order models (LES-ROMs) based on the proper orthogonal decomposition (POD). To develop these models, explicit POD spatial filtering is introduced. Two types of spatial filters are considered: A POD projection onto a POD subspace and a POD differential filter. These explicit POD spatial filters allow the development of two types of ROM closure models: phenomenological and approximate deconvolution. Furthermore, the explicit POD spatial filters are used to develop regularized ROMs in which various ROM terms are smoothed (regularized). The new LES-ROMs are tested in the numerical simulation of a three-dimensional flow past a circular cylinder.

Xuping Xie

Virginia Tech

xupingxy@vt.edu

### MS107

#### POD Reduced Order Modeling for Evolution Equations Utilizing Arbitrary Finite Element Discretizations

The main focus of the present work is the inclusion of spatial adaptivity for snapshot computation in the offline phase of model order reduction utilizing Proper Orthogonal Decomposition (POD-MOR). For each time level, the snapshots lie in different finite element spaces, which means in a fully discrete setting that the snapshots are vectors of different lengths. In order to overcome this obstacle, we present a discretization independent POD reduced order model, which is motivated from a continuous perspective and is set up and solved explicitly without interpolation of the snapshots. In contrary to empirical interpolation methods, we introduce a projection based approach for the treatment of nonlinear terms in order to circumvent the interpolation onto a common spatial grid. The analysis for the error between the resulting POD solution and the true solution is carried out. Finally, we present numerical examples to illustrate our approach.

Carmen Graessle

University of Hamburg  
carmen.graessle@uni-hamburg.de

Michael Hinze

Universität Hamburg  
Department Mathematik  
michael.hinze@uni-hamburg.de

### MS107

#### Reduced Order Nonlinear Fire Plume Models

The modeling of fire plumes is a complex mathematical problem, consisting of a system of partial differential equations coupling the momentum, energy, and reaction kinetics equations. Because of this complexity, simulating models of fires are typically expensive in terms of computational resources and time. Therefore, these simulation tools are used in pre- and post-fire assessments, but they are rarely used for real-time decision-making applications. Our goal is to create a reduced order model (ROM) to allow on-scene prediction capabilities during a fire. During this talk, we discuss three challenges that affect our ability to create accurate, computationally efficient ROMs along with proposed solutions. The first is the 'lifting bottleneck' associated with the nonlinear terms. These are handled differently depending on the type of nonlinearity. The next challenge is the fact that many of the coefficient terms, such as viscosity and thermal diffusivity, are temperature dependent cannot be considered constants over the temperature ranges that exist for a fire. This dependence increases the nonlinearity of the system. Finally, we must ensure that the mass fractions of fuel, air, and products remain bounded in  $[0, 1]$  and always sum to 1. For a full-order model, this is typically handled at the cell level; but for the ROM, other techniques must be used to ensure that the bounds are preserved over the entire domain.

Alan Lattimer

Jensen Hughes  
alattimer@jensenhughes.com

Jeff Borggaard  
Virginia Tech

Department of Mathematics  
jborggaard@vt.edu

Brian Lattimer

Jensen Hughes  
blattimer@jensenhughes.com

### MS107

#### Model Reduction of Multiple Transport Phenomena: The Shifted Proper Orthogonal Decomposition

Transport-dominated phenomena as shock waves and other moving fronts play an important role in many applications, e.g., in chemical engineering or aerodynamics. Efficient model reduction of these problems is still a big challenge for common methods as the classical Proper Orthogonal Decomposition (POD). In this talk, we address the optimal approximation of multiple transport phenomena with moving modes. This leads to the shifted POD which generalizes the common POD and is able to describe transport phenomena with just a few modes. The power and flexibility of this approach is demonstrated by means of numerical examples.

Philipp Schulze

Institute of Mathematics  
TU Berlin  
pschulze@math.tu-berlin.de

Julius Reiss

TU Berlin  
julius.reiss@tnt.tu-berlin.de

### MS107

#### Hybrid Discontinuous Galerkin Method for POD

Abstract not available

Yangwen Zhang

Missouri S&T  
ywzfg4@mst.edu

### MS108

#### Big Compute, Big Data, and Better Drugs

Abstract not available

Sally R. Ellingson

University of Kentucky  
sally@kcr.uky.edu

### MS108

#### Beyond Docking: Increasing the Accuracy of Virtual Screens

Abstract not available

Amir Kucharski

University of Kentucky  
sally@kcr.uky.edu

### MS108

#### Catalyzing the Normalization of Inclusion through

**Broader Engagement**

Abstract not available

Mary Ann E. Leung  
Sustainable Horizons Institute  
mleung@shinstitute.org

**MS108****A Framework for the Evaluation of Data Analyses and Visualization Tools**

Abstract not available

Samar Swaid  
Philander Smith College  
sswaid@philander.edu

**MS109****Hydromechanical Modeling Framework for Two-Scale Porous Media**

Natural geomaterials in engineering problems often exhibit porous structures at two scales due to fractures, particle aggregation, or other reasons. Coupled poromechanical processes in these two-scale geomaterials, such as consolidation in the presence of preferential flow, are beyond the modeling capabilities of classical frameworks. This presentation will describe a hydromechanical modeling framework for geomaterials with two-scale porous structures. Built on the double-porosity concept and continuum thermomechanics, the framework uniquely covers the entire ranges of fluid flow and deformations in double-porosity media—from steady-state to transient flow, and from infinitesimal to finite deformations—with thermodynamic consistency. At the core of this achievement is a novel constitutive framework that enables one to capture multiscale evolution of pore structures even under finite deformations and transient flow. Through numerical examples the frameworks superior capabilities for modeling natural geomaterials will be demonstrated, providing insights into the origin of secondary compression.

Jinhyun Choo, Ronaldo I. Borja  
Stanford University  
jinhyun.choo@alumni.stanford.edu, borja@stanford.edu

**MS109****About the Uzawa Smoother for Poroelastic Problems**

This talk deals with the efficient solution of the poroelastic equations. A multigrid method is employed with an Uzawa type iteration as smoother. The Uzawa smoother is an equation-wise procedure. It is interpreted as a combination of the symmetric Gauss-Seidel smoothing for displacements, together with a Richardson iteration for the Schur complement in the pressure field. The Richardson iteration involves a relaxation parameter which affects the convergence speed, and has to be carefully determined. The analysis of the smoother is based on the framework of local Fourier analysis and it allows us to provide an analytic bound of the smoothing factor of the Uzawa smoother as well as an optimal value of the relaxation parameter. Finally, some numerical results are presented to confirm the efficiency and robustness of this method.

Francisco José Gaspar  
University of Zaragoza

fjgaspar@unizar.es

**MS109****Computation and Joint Analysis of Coupled Flow, Geomechanics, and Geophysics in Reservoir Engineering**

Several types of data are available in geological systems, such as pressure/flow rate, surface deformation (InSAR), induced seismicity (microearthquake (MEQ)), and electromagnetic (EM) geophysical data. The geophysical data have different characteristics. For example, MEQ can detect geological failure such as fault activation, while the data could contain noise which is not distinguishable from the information of interest, particularly when the magnitude of the geological failure is not sufficiently large. InSAR data can be obtained as a series of time, while information of geological strata might be additionally needed for accurate characterization of geomechanical properties. The EM signals can be only detected after large amounts of injection. Even in the case that individual data have low quality and not sufficient for numerical inversion separately, the simulation of coupled flow-geomechanics-geophysics and the development of a novel modeling technique by integrating all the measured data will make the reservoir characterization feasible. In this study, we first investigate sensitivity of geophysical monitoring signals (i.e., MEQ, InSAR, EM) for various Earth sciences problems by using an advanced reliable forward simulator of coupled non-isothermal multiphase flow-geomechanics-geophysics. Then, we develop an inverse modeling technique (e.g., constrained optimization) for reservoir characterization, and show the numerical examples.

Jihoon Kim  
Texas A&M University  
Department of Petroleum Engineering  
jihoon.kim@tamu.edu

Evan Um  
Lawrence Berkeley National Laboratory  
evanum@gmail.com

**MS109****Preconditioning Strategies for Coupled Multiphase Poromechanics**

Geomechanical effects play an important role in the performance of many oil and gas reservoirs, particularly in unconventional settings. The mathematical framework describing fluid flow through deformable porous media consists of balance laws for the linear momentum and mass of each phase. An implicit, tightly-coupled solution of these governing equations is frequently necessary to obtain reliable model predictions, but the development of efficient and scalable algorithms is non-trivial. Here, we describe an iterative solution framework for multiphase poromechanics that scales well on large computing platforms. The linearization of the system of nonlinear algebraic equations produces a Jacobian matrix characterized by a specific block structure. Based upon an approximate block-factorization of this matrix, we propose a two-stage preconditioner. A generalized Constrained Pressure Residual approach is used to construct a reduced pressure-displacement system, involving unknowns with long range error components. In the first stage, these components are addressed with a Fixed-Stress update algorithm. Once pressure and displacement degrees of freedom have been updated, a second stage is applied to deal with the remain-

ing saturation unknowns. Numerical results are presented to illustrate performance and robustness of the proposed preconditioning approach on a variety of challenging test problems.

Joshua A. White

Lawrence Livermore National Laboratory  
white230@llnl.gov

Sergey Klevtsov  
Stanford University  
klevtsov@stanford.edu

Nicola Castelletto  
Stanford University  
Department of Energy Resources Engineering  
ncastell@stanford.edu

Hamdi Tchelepi  
Stanford University  
Energy Resources Engineering Department  
tchelepi@stanford.edu

#### MS110

##### **An Extension of the Dynamic Mode Decomposition and Applications**

Abstract not available

Travis Askham

University of Washington  
askham@uw.edu

#### MS110

##### **Network Representation and Analysis of Bluff Body Wake**

Effective control of unsteady bluff body wake requires knowledge of its nonlinear dynamics. The unsteady and nonlinear nature of the wake make it a challenging task to identify, represent and analyze the uncertain behavior of the vortical interactions involved in bluff body wake formation. Network-theoretic tools are used here to shed light on the complex vortical interactions. Two-dimensional incompressible unsteady flow over a cylinder is used as an illustrative case in the current study. The flow field data from direct numerical simulations is used to represent the vortical interactions in the wake of the cylinder through a network. The vortical elements are considered as the nodes and the induced velocity among the elements as the edges of the network. The flow field data is examined by community detection and clustering techniques from network theory to identify the coherent structures in the cylinder wake and to reduce the order of the system. The validity as well as characterization of the vortical network is analyzed through vorticity flux conservation. Ongoing developments on vortical network model will be discussed.

Muralikrishnan Gopalakrishnan Meena

Florida State University  
mg15h@my.fsu.edu

Aditya G. Nair, Kunihiko Taira  
Florida State University  
Mechanical Engineering

agn13@my.fsu.edu, ktaira@fsu.edu

#### MS110

##### **The Most Informative Data for the State, Parameters and Dynamical Model**

A theoretical basis is set for predicting the observations that are most informative for estimating the state, the parameters, and the model equations of high-dimensional dynamical systems. The methodology exploits the governing nonlinear dynamics and captures the non-Gaussian structures. Optimal observation locations are determined by maximizing the dynamic mutual information between the candidate observations and the variables of interest. The results are exemplified and the performance is quantitatively assessed using a variety of simulated fluid and ocean flows.

Pierre F. Lermusiaux, Jing Lin  
MIT

pierrel@mit.edu, linjing@mit.edu

#### MS110

##### **Sensor Placement for Multiscale Phenomena**

Multiscale processes pose challenges in determining modal decompositions with physical meaning which can render sensor placement particularly difficult. Localized features in space or time may play a crucial role for the phenomenon of interest but may be insufficiently resolved due to their low energy contribution. We consider optimal sensor placement using spatial interpolation points within the framework of data-driven modal decompositions. In recent years, the Discrete Empirical Interpolation Method or DEIM and variants like QDEIM have gained popularity for interpolating nonlinear terms arising in model reduction using Proper Orthogonal Decomposition modes. We extend this sensor placement approach to multiscale physics problems using Multi-Resolution Dynamic Mode Decomposition or mrDMD [Kutz et al., 2015], an unsupervised multi-resolution analysis in the time-frequency domain that separates flow features occurring at different timescales. The discovered sensors achieve accurate flow state reconstruction in representative multiscale examples including global ocean temperature data with an energetic El Niño mode. Interestingly, this method places sensors near coastlines without imposing additional constraints, which is beneficial from an engineering perspective.

Krithika Manohar

University of Washington, USA  
kmanohar@uw.edu

Eurika Kaiser, Steven Brunton  
University of Washington  
eurika@uw.edu, sbrunton@uw.edu

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

#### MS111

##### **Mixed-Integer PDE-Constrained Optimization for Gas Networks**

Efficiency is a growing concern in the design and control of modern energy systems. Recent decades have seen an

increased focus on optimal control of power grids and gas networks. However, due to the complexities of models including transient effects, they are rarely taken into account in existing solutions. We describe an optimization model based on a transient PDE model of a natural gas transport network. Our model includes relaxed integer controls that model compressor configurations and is the basis for relaxation approaches to an integer programming model. We focus on ensuring partial separability between network components, which allows us to distribute optimization algorithms across multiple processors, thereby ensuring scalability to larger problem instances.

Mirko Hahn

Argonne National Laboratory  
hahnmm@anl.gov

**MS111**

**Nonlinear Robust PDE Constrained Optimization Using Approximation Techniques and Model Order Reduction with Application to Electric Motor Design**

We investigate a nonlinear constrained optimization problem with uncertain parameters. By utilizing a robust worst-case formulation we obtain an optimization problem of bi-level structure. This type of problems are difficult to treat computationally and hence suitable approximations are required. We propose and investigate an approximate robust formulation that employs a quadratic approximation. The proposed framework is mixed with linear approximation techniques when appropriate. The developed method is then applied to the optimal placement of a permanent magnet in the rotor of a synchronous machine. The goal is to optimize the volume and position of the permanent magnet while maintaining a given performance level. These quantities are computed from the magnetic vector potentials given by the magnetostatic approximation of Maxwell's equation with transient movement of the rotor. Utilizing the introduced robust optimization framework we account for uncertainties in material and production precision. The problem formulation and the robustification of the optimization lead to high computational cost that requires to investigate methods for efficient realization. Since the transient movement of the rotor can be interpreted as a multi query operation, model order reduction is a promising choice. By generating reliable reduced order models with a posteriori error control the computation can be accelerated. Numerical results are presented to validate the presented approach.

Oliver Lass, Stefan Ulbrich

Technische Universitaet Darmstadt  
Fachbereich Mathematik  
lass@mathematik.tu-darmstadt.de,  
ulbrich@mathematik.tu-darmstadt.de

**MS111**

**TAO: Toolkit for Advanced Optimization**

In this talk, I will discuss TAO, the Toolkit for Advanced Optimization, that provides numerical optimization methods for high-performance computing. These methods are built upon PETSc, the Portable Extensible Toolkit for Advanced Optimization. I will discuss the overall design of the software and delve into some of the available methods.

Todd Munson

Argonne National Laboratory

Mathematics and Computer Science Division  
tmunson@mcs.anl.gov

**MS111**

**Discrete Material Optimization of Nano-Structures in Electromagnetic Applications**

A class of algorithms for the solution of discrete material optimization problems in electromagnetic applications is discussed. The algorithmic idea is based on the sequential convex programming idea, however, in each major iteration a model is established on the basis of an appropriately parametrized material tensor. The resulting nonlinear and typically nonconvex parametrization is treated on the level of the sub-problem, for which, due to block separability of the model, globally optimal solutions can be computed. Although global optimization of non-convex design problems is in general prohibitive, a smart combination of analytic solutions along with standard global optimization techniques lead to an efficient algorithm for the most relevant material parametrizations. Theoretical properties of the overall algorithm, such as global convergence as well as the avoidance of poor local optima introduced by the material parametrization are discussed. The effectiveness of the algorithm in terms of computation time as well as quality of the solution is demonstrated by a series of numerical examples ranging from the optimal design of cloaking layers for nano-particles to the identification of multiple materials with different optical properties in a matrix.

Michael Stingl

Friedrich-Alexander University of Erlangen-Nuremberg  
stingl@math.fau.de

**MS112**

**A Model Transport Equation for Epistemic UQ in RANS**

Due to their computational efficiency, the Reynolds-Averaged Navier-Stokes (RANS) turbulence models remain a vital tool for modeling turbulent flows. However, it is well known that RANS predictions are locally corrupted by epistemic model-form uncertainty. In order to quantify this uncertainty, we directly perturb the Reynolds-stress tensor at locations in the flow domain where the modeling assumptions are likely to be invalid. If this were to be done on a point-by-point basis, the resulting inference problem would be of very high dimension. To reduce the dimensionality, we propose separate model equations based on the transport of linear combinations of the eigenvalues of the anisotropy tensor. This provides us with an intuitive and low-dimensional UQ framework where the transport model decides on the magnitude and direction of the perturbations. Where the perturbations are small, the RANS result is recovered. Using traditional turbulence modeling practices we derive weak realizability constraints, and we will rely on Bayesian inference to calibrate the model on high-fidelity data. We will demonstrate our framework on a number of canonical flow problems where RANS models are prone to failure.

Wouter N. Edeling

TU Delft  
wedeling@stanford.edu

Gianluca Iaccarino  
Stanford University  
Mechanical Engineering  
jops@stanford.edu

Paola Cinnella  
 Laboratoire DynFluid, Arts et Métiers ParisTech  
 and Università del Salento  
 paola.cinnella@ensam.eu

**MS112****Data-Driven Predictive Modeling of RANS Model Discrepancies**

In this work, we develop a data-driven framework to characterize and quantify structural uncertainties in turbulence models. The first step involves the inference of the spatial distribution of modeling discrepancies. The inferred field from many different data-sets is projected onto an appropriate feature space using deep neural networks. The inference and learning steps are conditioned to respect physical and modeling constraints. Existing turbulence models are then augmented with this neural network in a predictive setting. The variability due to sparsity of data and the training of the neural network is propagated to the model outputs. The applicability of the framework is demonstrated in turbulent flows involving strong adverse pressure gradients and relaminarization.

Anand Pratap Singh  
 University of Michigan  
 anandps@umich.edu

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

**MS112****A Data-Driven, Physics-Informed Approach for Turbulence Modeling**

Turbulence modeling introduces significant uncertainties in the predictions. In light of the decades-long stagnation encountered by the traditional approach of turbulence model development, data-driven methods have been proposed as a promising alternative. In this talk, I will present a data-driven, physics-informed machine learning framework for predictive turbulence modeling based on RANS models. The framework consists of three components: (1) prediction of discrepancies in RANS modeled Reynolds stresses based on machine learning algorithms, (2) propagation of improved Reynolds stresses to quantities of interests with a modified RANS solver, and (3) quantitative, a priori assessment of predictive confidence based on distance metrics in the mean flow feature space. Merits of the proposed framework are demonstrated in a class of flows featuring massive separations. Specifically, high-fidelity simulation data from a few flows (e.g., curved backward step, channel with wavy wall) are used to train the discrepancy functions of Reynolds stress, which are subsequently used to predict flows in a new geometry (channel with periodic hill) that is not present in the training flow database. Significant improvements over the baseline RANS predictions are observed.

Heng Xiao, Jinlong Wu, Jianxun Wang  
 Dept. of Aerospace and Ocean Engineering, Virginia Tech  
 hengxiao@vt.edu, jinlong@vt.edu, vtwx@vt.edu

**MS113****A New Predictor-Corrector Method for Efficient****Modeling of Surface Effects**

The regular Cauchy–Born method is a useful and efficient tool for analyzing bulk properties of materials in the absence of defects. However, the method normally fails to capture surface effects, which are essential to determining material properties at small length scales. Here, we present a corrector method that improves upon the prediction for material behavior from the Cauchy–Born method over a small boundary layer at the surface of a material by capturing the missed surface effects. We justify the separation of the problem into a bulk response and a localized surface correction by establishing an error estimate, which vanishes in the long wavelength limit.

Andrew Binder  
 Department of Mathematics  
 University of Minnesota Twin Cities  
 bind0090@math.umn.edu

Mitchell Luskin  
 School of Mathematics  
 University of Minnesota  
 luskin@math.umn.edu

Christoph Ortner  
 University of Warwick  
 c.ortner@warwick.ac.uk

**MS113****Three-Dimensional Elasto-Plastic Simulation of Bulk Metallic Glasses**

We develop a new method for simulating three-dimensional elasto-plastic materials in the slow, quasi-static limit. The method uses a mathematical correspondence between the governing equations for these materials, and the Navier–Stokes equations for incompressible fluid flow. Using this correspondence, we adapt the well-established projection method of Chorin for incompressible fluid flow over to this new physical problem. We test the method using an elasto-plastic model of a bulk metallic glass based on the shear transformation zone (STZ) theory. The method is parallelized using the MPI library and is used to simulate the nucleation and subsequent dynamics of shear bands within the metallic glass. The three-dimensional structure of the shear bands is investigated numerically and via OpenGL visualization.

Nicholas M. Boffi, Chris H. Rycroft  
 Harvard University  
 boffi@g.harvard.edu, chr@seas.harvard.edu

**MS113****Simulations of Viscous Suspension Flows with a Meshless MLS Scheme**

This talk will focus on mesh-free methods for simulations of neutrally buoyant particles in a viscous fluid. We will focus on examples relevant to densely populated, low Reynolds number suspension flows of neutrally buoyant, non-Brownian, monodisperse particles, along with systems with bidisperse particle sizes. Numerical simulations will be done with a meshless Moving Least Squares (MLS) scheme using a polynomial reconstruction process that provides a computationally efficient method to handle general boundary conditions, while giving arbitrary order polynomial accuracy and maintaining numerical stability. We will compare the accuracy and computational efficiency with



other schemes including Incompressible Smoothed Particle Hydrodynamics (ISPH) and the Force Coupling Method.

Amanda Howard  
Brown University  
amanda\_howard@brown.edu

Nathaniel Trask  
Brown University  
Division of Applied Mathematics  
nat.trask@gmail.com

Martin Maxey  
Division of Applied Mathematics,  
Brown University  
martin\_maxey@brown.edu

#### MS114

##### **Generalized Convolutional Representation for Field Data on Graphs**

Convolutional sparse representation is an efficient tool for computing sparse representations for entire signals in terms of sums of a set of convolutions with dictionary filters. This technique has been successfully applied to natural images, video and speech in tasks as diverse as denoising, classification or superresolution. However, when the signal does not have a regular local structure but is represented by an arbitrary graph, defining equivalent convolution operations is challenging. We review some of the approaches used in the discrete signal processing and deep learning communities and discuss their merits together with some interesting applications.

Cristina Garcia-Cardona  
Los Alamos National Laboratory  
cgarcia@lanl.gov

#### MS114

##### **Consistency of Dirichlet Partitions**

I'll discuss a geometric approach to graph partitioning where the optimality criterion is given by the sum of the first Laplace-Dirichlet eigenvalues of the partition components. This eigenvalue optimization problem can be solved by a rearrangement algorithm, which we show to converge in a finite number of iterations to a local minimum of a relaxed objective. This partitioning method compares well to state-of-the-art approaches on a variety of graphs constructed from manifold discretizations, synthetic data, the MNIST handwritten digit dataset, and images. I'll present a consistency result for geometric graphs, stating convergence of graph partitions to an appropriate continuum partition

Braxton Osting, Todd Reeb  
University of Utah  
osting@math.utah.edu, reeb@utah.math.edu

#### MS114

##### **Error Estimates on Spectral Convergence of Graph Laplacian Towards the Laplace-Beltrami Operator**

A number of machine learning tasks relies on spectral properties of the graph laplacian associated to the data. Consider data points obtained as random samples of a manifold in  $R^d$ . I will discuss showing that the spectrum of the graph Laplacians on a neighborhood graph spanned by the

samples converges almost surely to the spectrum of appropriate weighted Laplace-Beltrami operators on the manifold, as the sample size increases and the neighborhood size shrinks to zero. I will present error estimates for the convergence that explicitly depend on the geometry of the manifold, the number of data points available and the size of the neighborhood used in the graph construction.

Nicolas Garcia Trillos  
Carnegie Mellon University  
nicolas\_garcia\_trillos@brown.edu

Moritz Gerlach  
Institut für Mathematik  
Universität Potsdam  
mogerlac@uni-potsdam.de

Matthias Hein  
Saarland University  
Dept. of Mathematics and Computer Science  
hein@cs.uni-saarland.de

Dejan Slepcev  
Department of Mathematical Sciences  
Carnegie Mellon University  
slepvec@math.cmu.edu

#### MS114

##### **Applications of the Graph P-Laplacian for Data Processing and Analysis**

Graph-based methods are a promising tool, which can be applied to a wide range of problems that can be modeled by a graph, e.g., network problems. One key feature of these methods is the possibility to incorporate nonlocal relationships into the analysis and processing of data rather than using only local neighborhoods. The recent trend in the literature is to translate well-studied variational problems and PDEs to the graph setting and hereby overcome drawbacks of classical approaches. In this talk we give a short introduction to the concept of partial difference equations on graphs and show that many classical numerical discretization schemes can be embedded in a graph setting and thus be interpreted as special cases in a more general framework. As particular example we discuss a family of partial difference operators on graphs known as the graph p-Laplacian and we study equations involving these operators. Finally, we demonstrate the advantages of graph-based methods for different tasks in image and point cloud processing, such as filtering, segmentation, clustering, and inpainting. Joint work with A. Elmoataz, Université de Caen, France

Daniel Tenbrinck  
Fachbereich Mathematik und Informatik  
Universität Münster  
daniel.tenbrinck@uni-muenster.de

#### MS115

##### **Acceleration of Monte Carlo Methods for Stochastic Elliptic PDEs Using Domain Decomposition and PC Approximations of Local Dirichlet Maps**

A common approach to solving Elliptic PDEs with Stochastic coefficients (SPDEs) is to resort to Monte Carlo (MC) sampling methods. We present a framework that both accelerates the resolution of SPDEs by MC methods and improves the resilience in the context of exascale computing. Our method is based on a Domain Decompo-

sition approach that uses Polynomial Chaos (PC) approximations of the (stochastic) boundary-to-boundary (BtB) operator, relating the unknown boundary values at the subdomains interfaces to one another. To do so, the BtB operator is broken into elementary contributions from each subdomain. The contributions from distinct subdomains are fully independent and can be computed in parallel. Further, the elementary SPDE problems can be parametrized using a reduced set of local random variables, enabling PC expansions with manageable dimensionality for small enough subdomains. Subsequently, the PC approximations can be jointly sampled to generate realizations of the BtB operator, without needing the resolution of additional elliptic problems, which can be inverted to obtain the corresponding sample of the boundary values. The method then allows to efficiently distribute and parallelize most of the computations. Also, given that the PC expansions are constructed independently, the method can be made robust to node failures. Finally, we validate and illustrate the efficiency of the method with extensive numerical tests on a second-order stochastic elliptic problem.

Andres Contreras, Paul Mycek  
Duke University  
andres.contreras@duke.edu, paul.mycek@duke.edu

Olivier Le Maitre  
Laboratoire d'Informatique pour la Mécanique et les  
Science  
olm@limsi.fr

Francesco Rizzi, Karla Morris, Khachik Sargsyan, Cosmin Safta  
Sandia National Laboratories  
fnrizzi@sandia.gov, knmorri@sandia.gov,  
ksargsy@sandia.gov, csafta@sandia.gov

Bert J. Debuschere  
Energy Transportation Center  
Sandia National Laboratories, Livermore CA  
bjdebus@sandia.gov

Omar M. Knio  
Duke University  
omar.knio@duke.edu

### MS115

#### Resilience Without Recovery in Algorithms for Hyperbolic Conservation Laws

Hard and soft errors are expected to occur more frequently on exascale-class supercomputers, but the use of synchronous, global checkpoint-restart will be infeasible at these scales. Generic detection and recovery mechanisms cannot take advantage of efficiencies that derive from the characteristics of the numerical algorithms. Particularly for silent data corruptions (SDCs), the algorithms may possess properties that mitigate the need for more severe intervention. For example, with minimal modifications, iterative algorithms like conjugate gradient or multigrid often can overcome isolated SDCs. Direct methods, which are the primary approach for hyperbolic systems of conservation laws, require a different mechanism by which SDCs can be overcome. For this class of problems, numerical algorithms already trade accuracy for robustness locally, and this provides a direction for the development of efficient, fault-tolerant, direct methods. Shock-capturing algorithms have several properties that can be leveraged to make such algorithms fault tolerant. We will show how the standard

flux divergence update algorithm can be decomposed at a high level into stages for which suitable fault-tolerant mechanisms can be identified. Results will demonstrate the promise of this approach. *This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.*

Jeffrey A. Hittinger, John Loffeld  
Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory  
hittinger1@llnl.gov, loffeld1@llnl.gov

### MS115

#### A Fault-Tolerant Implementation of the TeaLeaf Cg Sparse Iterative Solver

Abstract not available

Simon McIntosh-Smith  
University of Bristol  
cssnmis@bristol.ac.uk

### MS115

#### Local Recovery of Hybrid Task-Parallel Explicit PDE Solvers from Hard Failures

The existing application recovery model for MPI programs involves global termination and restart, which won't scale with the number of MPI ranks and cores. To overcome such a shortcoming, we have proposed the Local Failure Local Recovery model to keep recovery response proportional to the size of local failures, and demonstrated a scalable recovery with a prototype of fault-tolerant MPI (MPI-ULFM). This idea can be extended to emerging asynchronous task parallel programming models that exploits on-node concurrent execution of tasks to overlap recovery and non-recovery computations together. In this talk, we discuss the design of our resilient task programming models and how it is integrated with our local recovery techniques for MPI programming models. We also demonstrate a use case with 3D explicit PDE solvers from Sandia's mini-application code.

Keita Teranishi  
Sandia National Laboratories  
knteran@sandia.gov

Nicole Slattengren  
Sandia National Labs  
nslatt@sandia.gov

### MS117

#### Shape Optimization in Spectral Geometry: A Bayesian Approach

In this work, we present a novel numerical strategy to study the Szegő-Pólya conjecture on 5-gons which states that the minimizer of the first Dirichlet eigenvalue over all 5-gons of given area is the regular pentagon. This conjecture is still open. We introduce this conjecture and describe a coupled "Bayesian optimization-Finite Element" method (conforming and non-conforming piecewise linear approximations) to study this problem. Finally, we provide some numerical results.

Sebastian Dominguez  
Department of Mathematics  
Simon Fraser University

domingue@sfu.ca

Nilima Nigam  
Dept. of Mathematics  
Simon Fraser University  
nigam@math.sfu.ca

Bobak Shahriari  
University of British Columbia  
bshahr@cs.ubc.ca

#### MS117

##### Flexible Eigensolvers in Electronic Structure Calculations

Determining excited states in quantum physics or calculating the number of valence electrons in the Density Functional Theory (DFT) involve solving eigenvalue problems of very large dimensions. Moreover, very often the interesting features of these complex systems go beyond information contained in the extreme eigenpairs. For this reason, it is important to consider iterative solvers developed to compute a large amount of eigenpairs in the middle of the spectrum of large Hermitian and non-Hermitian matrices. In this talk, we present a newly developed Krylov-type methods and compare them with the well-established techniques in electronic structure calculations. We demonstrate their efficiency and robustness through various numerical examples.

Agnieszka Miedlar  
Technische Universität Berlin  
amiedlar@ku.edu

Yousef Saad  
Department of Computer Science  
University of Minnesota  
saad@cs.umn.edu

#### MS117

##### The Spectral Projector for Eigenvalue/Vector Approximation and Error Estimation

The spectral projector of a second-order differential operator, formally given as  $\mathcal{A} = -\nabla \cdot A \nabla u + cu$ , is given as a Dunford-Cauchy integral,

$$S_\gamma = \int_\gamma (zI - \mathcal{A})^{-1} dz,$$

where  $\gamma \subset \mathbb{C}$  is a contour enclosing some portion of its spectrum. We discuss how this projector can be used to approximate eigenvalues enclosed by  $\gamma$ , as well as their corresponding invariant subspaces, by using a subspace iteration procedure. We also show how a posteriori error estimates may be derived for computed approximations regardless of how they are obtained.

Jeffrey S. Ovall  
University of Kentucky, Mathematics  
jovall@pdx.edu

Jay Gopalakrishnan  
Portland state university  
gjay@pdx.edu

Luka Grubisic  
University of Zagreb

luka@math.hr

#### MS117

##### Finite Element Approximation of Eigenvalue Clusters of the Laplace-Beltrami Operator

Elliptic partial differential equations on surfaces are ubiquitous from geometry and relativity theory to applications in phase transitions, materials science, and image processing. They are typically governed by the Laplace-Beltrami operator, but more general operators arise as well. We present and analyze the approximations by Surface Finite Element Methods (SFEM) of the Laplace-Beltrami eigenvalue problem. In this context, spectral approximation is challenged by two sources of error: the geometric error due to the approximation of the surface using piecewise polynomials and the error corresponding to the finite element resolution of eigenfunctions on approximate surfaces. Our results reveal that optimal error decays for the eigenfunction approximations are guaranteed as for the "source" problem. However, our findings also indicate that this strategy is suboptimal when interested in the approximation of eigenvalues and that in this case the geometry plays a more predominant role. We provide numerical results to illustrate this apparent paradox.

Justin Owen  
Mathematics  
Texas A&M University  
jowen6@math.tamu.edu

#### MS118

##### A Domain Decomposition Based Method for the Simulation of Wind Flows Over Large Urban Areas

Accurate prediction of the wind flow field in large urban areas is the basis of many applications such as the vulnerability analysis of tall buildings and air pollution forecast. In this talk, we present a scalable domain decomposition based method for the numerical simulation of flows in urban areas with detailed geometric information of buildings. Two approaches are investigated including a 3D incompressible Navier-Stokes model and a LES with the Smagorinsky subgrid model. We report the parallel performance of the algorithm on supercomputers with a large number of processor cores.

Zhengzheng Yan  
Shenzhen Institutes of Advanced Technology  
zz.yan@siat.ac.cn

Xiao-Chuan Cai  
University of Colorado, Boulder  
Dept. of Computer Science  
cai@cs.colorado.edu

Rongliang Chen  
Chinese Academy of Sciences  
rl.chen@siat.ac.cn

#### MS118

##### Robust CFD Algorithms Through On-Line Machine-Learning: Gaussian Process Regression and Diffusion Maps

Exascale-level simulations require robust, fault-resilient, and efficient algorithm against repeated and expected software or hardware failures during computation, which may

render the simulation results unsatisfactory. If each processor can share some coarse but global information of the simulation from auxiliary data (relatively costless but of limited accuracy) we can effectively fill-in the missing data in space by machine learning techniques, e.g. Gaussian Process Regression, on the fly. We can also employ another machine learning technique, Diffusion Maps, to detect computational redundancy in time and hence accelerate the simulation by a projective time-integration, giving the overall computation a patch dynamics flavor. Learning from multifidelity data can guarantee the robustness, fault-resilience, and efficiency in large-scale exascale simulations. We will present several CFD benchmark problems to demonstrate the new capability that machine learning tools can bring to traditional scientific computing algorithms, relying on heterogeneous and multifidelity data. More broadly, in this talk we will demonstrate the symbiotic and synergistic convergence of machine learning, domain decomposition, and scientific computing in exascale simulations and other applications, e.g. multiscale simulations using black-box heterogeneous/legacy codes.

Seungjoon Lee

Division of Applied Mathematics, Brown University  
seungjoon\_lee@brown.edu

Ioannis Kevrekidis  
Princeton University  
yannis@princeton.edu

George Em Karniadakis  
Brown University  
george\_karniadakis@brown.edu

### MS118

#### Domain Decomposition in the Wave Chaos Analysis

Even though we are seeking for the highest possible fidelity, the computer representation will not be exactly the same compared to the real world. These uncertainties may arise from the imprecise knowledge of the system, small differences in the manufacturing, or numerical errors in the simulations. In most cases, those small differences can be considered as local perturbations of the entire system. Hence, the numerical solution is still a very good approximation to the exact solution of the physical problem. However, the situation can be completely different in complicated electromagnetic systems displaying wave chaos. This work concerns with a quantitative statistical analysis accounting for the uncertainty in complex wave-chaotic systems. The primary contributions are twofold: (i) a novel stochastic Greens function method for wave interaction with wave-chaotic media, which quantitatively describes the universal statistical property of chaotic systems through random matrix theory (RMT); (ii) a hybrid deterministic and stochastic formulation, in which small components (electronics, antennas, etc.) in the computational domain are modeled using first-principles and large portions (cavity enclosures, scattering environments, etc.) are modeled statistically. In this talk, we first introduce the concept of stochastic Greens function, then illustrate the domain decomposition formulation for short-wavelength wave-chaotic enclosures (wave chaos in cavities).

Zhen Peng

University of New Mexico  
Department of Electrical and Computer Engineering

pengz@unm.edu

### MS118

#### A Multi-Physics Domain Decomposition Method for Navier-Stokes-Darcy Model

In a karst aquifer, free flow and porous media flow are tightly coupled together, for which the Navier-Stokes-Darcy model has higher fidelity than either the Darcy or Navier-Stokes systems on their own. The Stokes-Darcy type model has attracted significant attention in the past ten years. However, coupling the two constituent models leads to a very complex system. This presentation discusses a multi-physics domain decomposition method for solving the Navier-Stokes-Darcy system. Computational results are presented to illustrate the features of the proposed method and the convergence analysis is demonstrated.

Changxin Qiu

Missouri University of Science and Technology  
cqrg7@mst.edu

### MS119

#### Multilevel Parallelism for SU/PG Discretization Within HPCMP Create()-AV COFFE for Compressible Rans Equations on Fully Tetrahedral Meshes

HPCMP CREATE(TM)-AV COFFE is the next-generation CFD solver within the Kestrel software suite. COFFE utilizes the SU/PG finite-element method to discretize the Navier-Stokes equations, and has been extended to higher-order. COFFE is regularly used to calculate 3-D flow-fields surrounding complex aerodynamic bodies on a supercomputer. COFFE utilizes a hierarchical parallelization strategy for the subdomain partitioned unstructured 3-D mesh, and a novel pseudo time marching/homotopy scheme is employed to march a solution to a steady-state. This solution strategy relies on the ability to solve a stiff, non-diagonally dominant parallel linear system at each pseudo step. An explanation of the discretization, parallelization strategy, and time-marching strategy will be presented as well as some applications.

Ryan S. Glasby, Jon Erwin, Stephen Wood, Douglas Stefanski

University of Tennessee  
ryan-glasby@tennessee.edu, taylor-erwin@tennessee.edu,  
swood19@vols.utk.edu, douglas.stefanski@tennessee.edu

Steve Karman  
Pointwise, Inc.  
skarman@pointwise.com

### MS119

#### Intrepid2: Performance Portable Finite Element Discretization Library

We present *Intrepid2* - an extensive tool set for finite element discretization. Departing from its previous version (*Intrepid*), the code is fully refactored to achieve portable performance on various architectures i.e., multi/many core host processors and GPUs. To do so, we use Kokkos, a C++ library that provides parallel patterns (`parallel_for/reduce/scan`) and abstractions for device specific memory and execution spaces with corresponding data layouts. In this talk, we focus on performance evalua-

tion and optimization of the code in the application context aiming at the next generation architectures.

Kyungjoo Kim  
Sandia National Laboratories  
kyukim@sandia.gov

### MS119

#### Performance Portable Software for High Fidelity Simulation of Reacting Flows Using Kokkos Array Abstractions

High fidelity turbulence combustion simulations using detailed chemical reaction mechanisms depend critically on high performance computing. This talk describes, KARFS, a new reacting flow simulation framework that is being developed to make efficient use of the hierarchical parallelism in accelerated and many-core architectures. Generic programming using templates and the array abstractions provided by the Kokkos library are utilized to achieve performance portability. Multi-level parallelism is obtained through message passing for distributed memory parallelism along with data parallel execution model for on-node shared memory parallelism. Optimized compute kernels are developed using this model for the multiphysics models needed for combustion simulations. Efficient compute kernels for thermo-chemical and molecular transport models are programmed as extensions of the Cantera library. The governing partial differential equations are solved using an explicit finite difference method. Experiences and performance results from the newest GPU accelerated and many-core systems will be presented.

Ramanan Sankaran  
Center for Computational Sciences  
Oak Ridge National Laboratory  
sankaranr@ornl.gov

Swapnil Desai  
University of Tennessee, Knoxville  
sdesai9@vols.utk.edu

Bok Jik Lee, Xiao Xu, Francisco E. Hernández Pérez,  
Hong G. Im  
King Abdullah University of Science and Technology  
bokjik.lee@kaust.edu.sa, xiao.xu@kaust.edu.sa,  
francisco.hernandezperez.1@kaust.edu.sa,  
hong.im@kaust.edu.sa

### MS119

#### Model-based Performance Optimization for GPU DG-FEM

DG-FEM offers an attractive high-order accurate and geometrically flexible method for time-domain wave problems while being highly adaptable to modern computer architectures. We present a mechanism to symbolically gather performance-relevant operation counts from DG-FEM sub-programs ('kernels'), apply these counts in a simple, linear model of kernel execution time, and demonstrate how predictions from this model can enable run-time performance optimization of a DG-FEM code. We demonstrate the utility of the method by selecting at run time a DG-FEM kernel configuration that yields the near-optimal performance. In terms of the geometric mean, our underlying simple, vendor- and GPU-type independent model achieves relative accuracy comparable to that of previously published

work using hardware specific models.

James Stevens, Andreas Kloeckner  
University of Illinois  
jdsteve2@illinois.edu, andreask@illinois.edu

### MS120

#### PFASST as a Space-Time Multigrid Method

For the numerical solution of time-dependent partial differential equations, time-parallel methods have recently shown to provide a promising way to extend prevailing strong-scaling limits of numerical codes. One of the most complex methods in this field is the "Parallel Full Approximation Scheme in Space and Time" (PFASST). PFASST already shows promising results for many use cases and many more is work in progress. However, a solid and reliable mathematical foundation is still missing. We show that under certain assumptions the PFASST algorithm can be conveniently and rigorously described as a multigrid-in-time method. Following this equivalence, first steps towards a comprehensive analysis of PFASST using block-wise local Fourier analysis are taken. The theoretical results are applied to examples of diffusive and advective type.

Matthias Bolten  
University of Kassel  
Institute of Mathematics  
bolten@mathematik.uni-kassel.de

Dieter Moser  
Juelich Supercomputing Centre  
Forschungszentrum Juelich  
d.moser@fz-juelich.de

Robert Speck  
Juelich Supercomputing Centre  
Forschungszentrum Juelich GmbH  
r.speck@fz-juelich.de

### MS120

#### MGRIT for Power-Grid Systems and BDF-K Methods

The parallel-in-time method multigrid reduction in time (MGRIT) was first designed for one-step methods. In this talk, we extend the MGRIT algorithm to multistep BDF methods for the integration of fully implicit Differential Algebraic Equations (DAE) on variable time-step grids. We will present one approach for implementing variable step-size BDF methods in a parallel-in-time context based on the XBraid software library. Results on power grid applications will also be given.

Matthieu B. Lécouvez  
Lawrence Livermore National Laboratory  
lecouvez1@llnl.gov

### MS120

#### Space-Time Multigrid Methods for Parabolic Problems

In this talk we discuss stable space-time discretization schemes, which allow the use of standard parallel finite element libraries to solve the arising space-time linear systems efficiently. Moreover we will have a look at space-time

multigrid methods which also allow parallelization with respect to space and time.

Martin Neumueller

Johannes Kepler University Linz, Austria  
neumueller@numa.uni-linz.ac.at

## MS120

### Multigrid Reduction in Time with Full Space-Time Adaptivity

The need for parallelism in the time dimension is being driven by changes in computer architectures, where performance increases are now provided through greater concurrency, not faster clock speeds. Multigrid Reduction in Time (MGRIT) is an iterative procedure that allows for temporal parallelism by utilizing multigrid reduction techniques and a multilevel hierarchy of coarse time grids. MGRIT has been shown to be effective for both linear and nonlinear problems, with speedups of up to  $50x$  seen in the linear case. Spatial and temporal adaptivity are powerful techniques used in most state of the art sequential time stepping routines. This talk will outline how those techniques can be used inside the MGRIT framework. In particular, we present a fully parallel, full multigrid (FMG) based algorithm for adapting in time and space using MGRIT, backward Euler time stepping, and a first order systems of least squares (FOSLS) spatial solver. Spatial adaptation is completed using the FOSLS error estimator and a threshold method. Temporal error estimators are calculated locally using Richardson extrapolation, but applied globally using a threshold method. Numerical results highlighting the usefulness of this approach will also be presented.

Ben O'Neill

University of Colorado, Boulder  
ben.oneill@colorado.edu

## MS121

### A Posteriori Error Analysis of the Parareal Algorithm

We derive a posteriori error estimate in a quantity-of-interest for the Parareal Algorithm. We describe a general formulation of two stage computations then perform a posteriori error analysis based on computable residuals and solution of an adjoint problem. We use the analysis to formulate novel procedures for constructing fine stage discretizations that take into account the cancellation of error. Finally, the error estimate identifies the contributions due to discretization and iteration.

Jehanzeb H. Chaudhry

University of New Mexico  
jehanzeb@unm.edu

Don Estep, Simon Tavener

Colorado State University  
estep@stat.colostate.edu, tavener@math.colostate.edu

## MS121

### Parareal for Hyperbolic Problems

Integration of hyperbolic evolution problems using Parareal has for a long time been considered impractical. Theoretical results have shown that it is not possible to give a general guarantee on the stability of the algo-

rithm when applied to this particular class of problems, and practical experiments have likewise suggested issues of slow convergence. In this talk we present our experiences using the Parareal method to accelerate a tsunami simulation tool complete with inundation modelling. The underlying PDE governing the dynamics of the model is called the shallow water wave equation. The system is in many ways typological as the SHW Eq. is a purely hyperbolic system of coupled non-linear PDEs, the solutions of which typically contain both shocks and smooth regions interacting in a non-trivial manner. Contrary to what one might expect, we find that with sufficient care in constructing the coarse operator, and a clever way of choosing time-subdomain length, we obtain parallel speed-up beyond what is possible using conventional spacial domain-decomposition techniques alone. This even for long time domains. A critical aspect in achieving the later was to decouple the time-subdomain length from the total time to be integrated, combined with theoretical considerations on how to choose the time-subdomain length optimally so to balance the communication of solution states and the convergence rate. These considerations in particular are generally applicable when using parareal.

Allan Nielsen

EPFL Lausanne  
allan.nielsen@epfl.ch

Jan S. Hesthaven

EPFL  
jan.hesthaven@epfl.ch

## MS121

### Solving Time-Periodic Navier-Stokes Problems Using Spectral Discretization in Time

We present an efficient solver for the time-periodic Navier-Stokes equations. This is done by using finite differences in space and a truncated Fourier series in time resulting in a non-linear system which is solved iteratively. This approach allows parallelization in space and time. The spacial dimensions are parallelized by domain decomposition whereas the temporal dimension is parallelized by distributing the Fourier coefficients to subgroups of processing units. Good parallel scaling has been demonstrated in previous work. The communication patterns and computational work depends on solver parameters and differs a lot between the spatial and the temporal dimensions. There is no straightforward strategy to achieve optimal load balance. In this presentation we discuss different strategies for load balancing in the context of space-time parallelism.

Daniel Hupp

ETH Zürich  
huppd@inf.ethz.ch

Peter Arbenz

ETH Zurich  
Computer Science Department  
arbenz@inf.ethz.ch

Dominik Obrist

University of Bern  
dominik.obrist@artorg.unibe.ch

## MS121

### ParaExp using Leap-Frog as Integrator for High-

## Frequency Electromagnetic Simulations

Recently, ParaExp was proposed by Gander and Gttel (2013) for the parallel-in-time resolution of linear hyperbolic problems. It splits the time domain into subdomains and computes the solution on each subdomain in parallel. The total solution is decomposed into a particular solution solved by a time integrator using zero initial values and a homogeneous solution propagated by the matrix exponential applied to the respective initial conditions. This paper deals with the application of ParaExp to electromagnetic wave problems in time-domain. These problems are typically solved by finite differences on a staggered grid pair in space and Leap-Frog (i.e., Strmer-Verlet) in time. We base ParaExp on those methods and propose the Leja method for the approximation of the matrix exponential. In Leja, errors can be controlled, only applications of the normal system matrix to vectors are needed and it has shown best performance in the examples. We will discuss academic and real world examples with respect to accuracy and efficiency. Special attention will be paid on symplecticity (energy conservation) which is regarded as important for high-frequency electromagnetic simulations.

Melina Merkel  
GSCE Technische Universität Darmstadt  
melina\_merkel@gmx.de

Innocent Niyonzima  
TU Darmstadt  
niyonzima@gsc.tu-darmstadt.de

Sebastian Schöps  
Theorie Elektromagnetischer Felder (TEMF) and GSCE  
Technische Universität Darmstadt  
schoeps@gsc.tu-darmstadt.de

## MS122

### Splitform Discontinuous Galerkin Methods for the Compressible Navier-Stokes Equations

When approximating nonlinear PDEs such as the compressible Navier-Stokes equations on discrete nodal grids, aliasing due to interpolation of nonlinear fluxes onto these grids can cause large errors which may even drive instabilities. An interesting approach to decrease the negative effect of this aliasing used primarily in the finite difference community is to reformulate the nonlinearities in different forms: Instead of using the conservative form of the PDE it is for instance possible to use the advective form by applying the chain rule. It is also possible to use an arbitrary mix between those two formulations and introduce skew-symmetric-like split formulations. By careful choice of such a reformulation it is possible to enhance the stability of the discretisations. Many questions arise, especially if the resulting method is still fully conservative, although it is based on non-conservative forms of the underlying PDEs. In this talk, we show how to incorporate the idea of different split forms of the compressible Navier-Stokes equations into the discontinuous Galerkin ansatz. We show how to recover famous splitting formulations such as the Ducros splitting or the Kennedy and Gruber splitting. We furthermore show that these novel DG schemes are still fully conservative and that with this special formulations the nonlinear stability for unresolved flows such as turbulence is highly enhanced.

Gregor Gassner  
Institute for Aerodynamics and Gasdynamics  
Universitaet Stuttgart

ggassner@math.uni-koeln.de

Andrew R. Winters  
Mathematical Institute  
University of Cologne  
awinters@math.uni-koeln.de

David A. Kopriva  
Department of Mathematics  
The Florida State University  
kopriva@math.fsu.edu

## MS122

### A Characteristic-Based CFL Condition for the Discontinuous Galerkin Method on Triangular Meshes

An exact CFL (Courant-Friedrichs-Lewy) condition is not known for the discontinuous Galerkin (DG) method applied to hyperbolic problems in two dimensions on triangular meshes. The stability condition that is most usually implemented involves scaling the time step by the radius of the smallest inscribed circle over all triangular cells. However, this is known to be not optimal, i.e. not to provide a tight bound on the largest possible time step. The main difficulty in finding a suitable CFL number is its dependence on the orientation of a mesh element with respect to the direction of the flow. Usually, the smallest CFL over all directions is taken, which is clearly not optimal. We find a natural scaling of the spectrum of the DG spatial operator by a parameter  $h_j$ , which can be seen to be the width of the cell  $\Omega_j$  along the characteristic direction of flow. We show that this parameter  $h_j$  incorporates 95% of the variation of the spectrum with respect to the orientation of  $\Omega_j$  to flow direction. We use this parameter to propose a new CFL condition. For a classical paring of the DG method with a degree  $p$  basis and an explicit Runge-Kutta scheme of order  $p+1$ , we show that a CFL number equal to  $1/(2p+1)(1+(2/(2p+2))^2)$  is within 5% accurate for  $p \leq 10$ . We show through several numerical examples that we are able to select larger stable time steps than usually obtained using the inscribed radii of the computational cells.

Lilia Krivodonova  
University of Waterloo  
lgk@math.uwaterloo.ca

## MS122

### High-Order Absorbing Boundary Conditions for Time-Domain Wave Propagation with DG Methods

DG schemes exhibit attractive features for large-scale wave-propagation simulations. For many applications, these schemes must be coupled with nonreflective boundary conditions in order to limit the size of the computational domain, which remains a challenging task. In this talk, we present a coupling of high-order absorbing boundary conditions with nodal DG schemes in three dimensions. Compatibility conditions at corners are addressed. Numerical results are proposed to validate both the accuracy and the computational efficiency of the techniques, and to illustrate the applicability of the strategy for large-scale problems.

Axel Modave  
Virginia Tech  
axel.modave@ensta-paristech.fr

Andreas Atle

TOTAL E&P  
andreas.atle@total.com

Jesse Chan  
Rice University  
jesse.chan@rice.edu

Russell Hewett  
Total  
russell.hewett@total.com

Tim Warburton  
Virginia Tech  
tcew@vt.edu

### MS122

#### Accelerating Discontinuous Galerkin Methods

In this talk we will describe the state of the art in many-core algorithms for time-dependent discontinuous Galerkin (DG) methods. We will discuss recent progress in exploiting the sparse operator structure of Bernstein-Bezier basis for tetrahedral elements on graphics processing units (GPU) [Chan & Warburton, 2016]. I will describe how speeds up result from the reduced register pressure, reduced cache usage, reduced operation count, and full unrolling of the DG volume divergence and lift operations resulting from the structure of the Bernstein-Bezier basis. An early GPU algorithm for hexahedral spectral elements was introduced in [Komatitsch et al., 2009]. Unlike tetrahedral elements, hexahedra have a natural local tensor-product, and this was exploited for degree five elements. Recently we have adapted algorithms from the finite difference GPU literature [Micikevicius, 2009] for high-order hexahedral spectral elements within a discontinuous Galerkin framework. We will discuss how we employ the excess of GPU floating point units idled during low operation count hexahedral differentiation and lift operations to process the extra operations required for the split form nodal DG scheme [Gassner et al., 2016] without incurring significant additional runtime compared to the traditional non-split flux formulation.

Tim Warburton  
Virginia Tech  
tcew@vt.edu

Jesse Chan  
Rice University  
jesse.chan@rice.edu

Gregor Gassner  
Institute for Aerodynamics and Gasdynamics  
Universitaet Stuttgart  
ggassner@math.uni-koeln.de

Axel Modave  
Virginia Tech  
axel.modave@ensta-paristech.fr

Zheng Wang  
Rice University  
zw14@rice.edu

Lucas Wilcox  
Department of Applied Mathematics  
Naval Postgraduate School

lwilcox@nps.edu

### MS123

#### Minimizing Eigenvalues for Inhomogeneous Rods and Plates

Optimizing eigenvalues of biharmonic equations appears in the frequency control based on density distribution of composite rods and thin plates with clamped or simply supported boundary conditions. In this work, we use a rearrangement algorithm to find the optimal density distribution which minimizes a specific eigenvalue. We answer the open question regarding optimal density configurations to minimize k-th eigenvalue for clamped rods and analytically show that the optimal configurations are distinct for clamped rods and simply supported rods. Many numerical simulations in both one and two dimensions demonstrate the robustness and efficiency of the proposed approach.

Weitao Chen  
University of California, Irvine  
weitaoc1@uci.edu

Ching-Shan Chou  
Department of Mathematics  
Ohio State University  
chou@math.ohio-state.edu

Chiu-Yen Kao  
Claremont Mckenna College  
Chiu-Yen.Kao@claremontmckenna.edu

### MS123

#### A Multi-Scale Model for Optical Responses of Nano-Structures

We present a semi-classical model for studying the optical responses of nano-structures, where the wave propagation is determined classically through Maxwell's equations and the motion of the nano-structure is determined quantum mechanically through Schrodinger equation. Ehrenfest molecular dynamics and Density Functional Theories are applied to deal with the many-body Schrodinger equation. The semi-classical model is numerically traceable. Linear response formulations will be derived and multi-scale solvers will be designed to solve the model. Numerical examples will be presented to demonstrate the model.

Gang Bao  
Zhejiang University  
baog@zju.edu.cn

Di Liu  
Michigan State University  
Department of Mathematics  
richardl@math.msu.edu

Songting Luo  
Department of Mathematics  
Iowa State University  
luos@iastate.edu

### MS123

#### An Asymptotic Model for Diffraction by Shallow Gratings

Thin film solar cells often use a metallic back reflector to improve efficiency. This back reflector can feature a pe-



riodic grating structure to enhance surface plasmon formation. We wish to model the effect of a shallow grating, perhaps on top of a larger amplitude grating. We introduce a simple way to apply asymptotic methods to the Helmholtz equation governing s- or p-polarized light that allow the thin grating to be replaced by an interface with non-standard transmission conditions depending on the thickness of the shallow grating. Numerical results validate our model. It is expected that this approach will allow thin gratings to be incorporated into standard finite element codes without the need for very small elements to model the grating. This is joint work with Cynthia Rivas, Manuel Solano and Rodolfo Rodriguez from the University of Concepcion, Chile together with Akhlesh Lakhtakia from Penn State University.

Peter B. Monk

Department of Mathematical Sciences  
University of Delaware  
monk@math.udel.edu

### MS123

#### Hybrid Inverse Problem for Porous Media

Abstract not available

Carlos Perez-Arancibia

Massachusetts Institute of Technology  
cperezar@mit.edu

### MS124

#### Estimation, Localization, and Modeling of Output Errors for Higher-Order Galerkin Finite Element Methods

In this presentation, we develop *a priori* estimates for the convergence of outputs, output error estimates, and localizations of output error estimates for Galerkin finite element methods. Output error estimates for order  $p$  finite element solutions are constructed using the Dual-Weighted Residual (DWR) method with a higher-order  $p' > p$  dual solution. Specifically, we analyze these DWR estimates for Continuous Galerkin (CG), Discontinuous Galerkin (DG), and Hybridized DG (HDG) methods applied to the Poisson problem. For all discretizations, as  $h \rightarrow 0$ , we prove that the output and output error estimate converge at order  $2p$  and  $2p'$  (assuming sufficient smoothness), while localizations of the output and output error estimate converge at  $2p + d$  and  $p + p' + d$ . For DG, the results use a new post processing for the error associated with the lifting operator. For HDG, these rates improve an additional order when the stabilization is based upon an  $O(1)$  length scale.

Hugh Carson

Massachusetts Institute of Technology  
hcarson@mit.edu

David I. Darmofal

Department of Aeronautics & Astronautics  
Massachusetts Institute of Technology  
darmofal@mit.edu

Steven R. Allmaras

Department of Aeronautics and Astronautics  
Massachusetts Institute of Technology  
allmaras@mit.edu

Marshall Galbraith

Massachusetts Institute of Technology  
galbramc@mit.edu

### MS124

#### An H-P Adaptive Discontinuous Galerkin Method for Unsteady Flows

While DG methods allow for varying polynomial solution approximation orders, adapting both the mesh and the order (h-p) for scalar output functionals remains a difficult task for unsteady problems. This work presents a new unsteady h-p adaptive method based on an error sampling and mesh optimization framework, in which an unsteady output adjoint drives adaptation of an anisotropic mesh and polynomial order field. The approach is demonstrated on compressible Navier-Stokes simulations.

Johann Dahm, Krzysztof Fidkowski

University of Michigan  
jdahm@umich.edu, kfid@umich.edu

### MS124

#### The Hybridized Discontinuous Galerkin Methods for Large Eddy Simulation

With the increase in computing power, Large-Eddy Simulation (LES) emerges as a promising technique to improve both knowledge of complex flow physics and reliability of flow predictions. Most LES works, however, are limited to simple geometries and low Reynolds numbers due to high computational cost. While most existing LES codes are based on 2nd-order finite volume schemes, the efficient and accurate prediction of complex turbulent flows may require a paradigm shift in computational approach.

This drives a growing interest in the development of Discontinuous Galerkin methods for LES. DG methods allow for high-order, conservative implementations on complex geometries and unstructured meshes, and offer opportunities for more accurate treatment of the sub-grid scale (SGS) turbulence. Also, they are better-suited to exploit modern HPC systems. In the spirit of making DG methods more competitive, researchers have recently developed the hybridized DG methods that result in reduced computational cost and memory footprint.

In this talk, we present an overview of hybridized DG methods for LES. Numerical accuracy, computational efficiency and SGS modeling issues are discussed in the framework of subsonic and transonic turbomachinery flows. Grid convergence studies are performed and the required resolution at different Reynolds numbers is investigated. Numerical results show rapid grid convergence and excellent agreement with experimental data at moderate computational cost.

Pablo Fernandez, Cuong Nguyen, Jaime Peraire

Massachusetts Institute of Technology  
pablof@mit.edu, cuongng@mit.edu, peraire@MIT.EDU

### MS124

#### An Adaptive Discontinuous Galerkin Reduced Basis Element Method

We present an adaptive discontinuous Galerkin reduced basis element (DG-RBE) method, which adaptively incorporates non-polynomial functions to construct approximation spaces specifically tailored to the problem of interest. The key ingredients are the following: a DG method which provides a flexible choice of approximation spaces;

dual-weighted residual (DWR) error estimates for quantities of interest; and a library of non-polynomial finite elements, called reduced basis elements (RBE), designed for specific solution features. The RBE library consists of spatially complex but parametrically simple features, such as boundary layers and trailing edge singularities. In the offline stage, these features and tailored functions are identified by solve a number of representative training cases. In the online stage, an appropriate set of library functions are incorporated to the approximation space based on the behavior of the error estimate. We assess the effectiveness of approach for compressible Navier Stokes equations over aerodynamic bodies. We demonstrate that the method on one hand extends notion of solution-based  $hp$  adaptivity by exploiting non-polynomial approximation functions; the method on the other hand extends the envelope of model reduction concepts to nonlinear equations exhibiting limited stability.

Masayuki Yano  
University of Toronto  
myano@utias.utoronto.ca

#### MS125

##### Well-Balanced Central-Upwind Schemes for the Euler Equations with Gravitation

Abstract not available

Alina Chertock  
North Carolina State University  
Department of Mathematics  
chertock@math.ncsu.edu

#### MS125

##### On Positivity of Density and Pressure for Finite Volume Methods Based on Relaxation Riemann Solvers for the Compressible Euler Equations

Abstract not available

Christian Klingenberg  
Wurzburg University, Dept. of Mathematics  
Germany  
klingen@mathematik.uni-wuerzburg.de

#### MS125

##### Well-Balanced Positivity-Preserving Central-Upwind Schemes for Shallow Water Models

I will discuss a concept of well-balanced numerical methods for the Saint-Venant system of shallow water equations and will show how one can design well-balanced central-upwind schemes, which, in addition, preserve positivity of the water depth. The performance of the designed schemes will be illustrated on a number of numerical examples

Alexander Kurganov  
Tulane University  
kurganov@math.tulane.edu

#### MS125

##### Bound Preserving Flux Limiters and Total Variation Stability for High Order Conservative Schemes

Abstract not available

Zhengfu Xu

Michigan Technological University  
Dept. of Mathematical Sciences  
zhengfux@mtu.edu

#### MS126

##### A Self-Consistent Spin-Diffusion Model for Micromagnetics: Discretization and Boundary Conditions

In recent years, spintronics has gained a lot of attention in the magnetic community. The micromagnetic model is a well established and reliable tool for the investigation of magnetization dynamics in systems subject to purely magnetic energy contributions such as exchange and magnetostatic energy. However, when it comes to spintronics, a multitude of additional physical effects has to be considered and likewise a variety of extensions to the micromagnetic model exist. The spin diffusion model introduced by Zhang, Levy and Fert is a very promising candidate and describes current and resistance in a self-consistent fashion. Moreover the spin-Hall effect and its inverse counterpart can be seamlessly integrated in the model which makes it a versatile tool for the description of a variety of effects related to spintronics. The implementation of a finite-element discretization of the spin-diffusion model coupled to the Landau-Lifshitz-Gilbert equation has to be constructed carefully due to possible discontinuities in the material parameters and solution variables. Since the self-consistent model solves simultaneously for the magnetization, the spin accumulation and the electric potential, another important issue is the choice of physically meaningful boundary conditions for these entities.

Claas Abert  
TU Wien  
claas.abert@tuwien.ac.at

#### MS126

##### Optimal Error Estimates of a Linearized Backward Euler FEM for the Landau-Lifshitz Equation

We present a fully discrete linearized backward Euler finite element method for the Landau-Lifshitz equation in which a new linearization is proposed for the gyromagnetic term. Optimal almost unconditional error estimates of the scheme in  $L^2$  and  $H^1$  norms are proved (i.e., when the stepsizes  $h$  and  $\tau$  are smaller than given constants). The analysis relies on the new linearization and an error splitting argument proposed by Li and Sun [Int. J. Numer. Anal. Model., 10 (2013), pp. 622633] and [SIAM J. Numer. Anal., 51 (2013), pp. 19591977]. Numerical results in both two- and three-dimensional spaces are provided to confirm our theoretical analysis and show the unconditional stability (convergence) of the scheme.

Huadong Gao  
School of Mathematics and Statistics  
Huazhong University of Science and Technology  
mahdgao@163.com

#### MS126

##### The Mimetic Finite Difference Method for the Landau-Lifshitz Equation

The Landau-Lifshitz equation describes the dynamics of the magnetization inside ferromagnetic materials. This equation is highly nonlinear and has a non-convex constraint that the magnitude of the magnetization is con-

stant. We present explicit and implicit mimetic finite difference schemes for the Landau-Lifshitz equation, which preserve the magnitude of the magnetization. These schemes work on general polytopal meshes, which provide enormous flexibility to model magnetic devices with various shapes. We will present rigorous convergence tests for the schemes on general meshes that includes distorted and randomized meshes. We will also present numerical simulations for the NIST standard problem #4 and the formation of the domain wall structures in a thin film. This is a joint work with K. Lipnikov.

Eugenia Kim  
University of California, Berkeley  
kim107@math.berkeley.edu

#### MS127

##### The GW Method for Computing Electronic Excited States of Molecules and Solids

The GW method is the state-of-the-art approach for computing electronic excitation energies in molecules and solids. It is in particular the most accurate method for computing band gaps in semiconductors and insulators. In this talk, I will present the mathematical formulation of the GW method and discuss its performance with respect to its main competitor, namely time-dependent density functional theory. This is a joint work with David Gontier and Gabriel Stoltz.

Eric Cancès  
Ecole des Ponts and INRIA, France  
cances@cermics.enpc.fr

#### MS127

##### Rational and Polynomial Filtering, Spectrum Slicing, and the EVSL Package

This talk will be about two different strategies for extracting extreme or interior eigenvalues of large sparse (Hermitian) matrices. The first, based on polynomial filtering, can be quite efficient in the situation where the matrix-vector product operation is inexpensive and when a large number of eigenvalues is sought, as is the case in calculations related to excited states for example. The second approach uses rational filtering and requires solving linear systems. The talk will discuss algorithmic aspects of these two approaches and will present the main features of EVSL, a library that implements filtered Lanczos and subspace iteration methods, with spectrum slicing. The EVSL package is available online at Y. Saad's web-site. This talk represents recent work with Ruipeng Li, Yuanzhe Xi, Eugene Vecharynski, and Chao Yang.

Yousef Saad  
University of Minnesota  
saad@umn.edu

#### MS127

##### Recent Progress in Numerical Methods for Electronic Excited States Calculations

Excited state properties of electrons are responsible for a variety of electronic and optical phenomena that can be harnessed to develop energy generation and conversion materials and devices. Computing the excited states of many-electron systems is often a challenging task. I will give an overview of a number of approaches that are based many-body perturbation and linear response theory, and point

out the computational bottlenecks of these approaches. I will describe numerical methods currently used to perform excited states calculations and recent efforts to reduce the computational complexity and to improve the efficiency of this type of computation on high performance computers.

Chao Yang  
Lawrence Berkeley National Lab  
cyang@lbl.gov

#### MS128

##### Revisiting Quasi-Standard Error

Abstract not available

Hongmei Chi  
Florida A&M University  
chi@cs.fsu.edu

#### MS128

##### Analyzing and Improving the Walk-on-Subdomains Algorithm in a Union of Spheres Geometry

Abstract not available

Preston Hamlin  
Department of Computer Science  
Florida State University  
hamlin@cs.fsu.edu

#### MS128

##### Stochastic Capacitance Extraction

Abstract not available

Chi-Ok Hwang  
GIST College  
Gwangju Institute of Science and Technology  
chwang@gist.ac.kr

#### MS128

##### Revisit of Block Power Method for Markov Chain Monte Carlo Applications

Abstract not available

Yaohang Li  
Department of Computer Science  
Old Dominion University  
yaohang@cs.odu.edu

#### MS129

##### Feast Eigensolver: Practices and Applications

A detailed review of FEAST eigensolver ([www.feast-solver.org](http://www.feast-solver.org)) is presented. In particular, we focus on new features including: (i) the non-Hermitian solver; (ii) the coupling with parallel distributed linear system solvers such as cluster-PARDISO and MUMPS; (iii) the new FEAST-GMRES iterative solver that does not require a preconditioner for solving very large eigenproblems.

Eric Polizzi  
University of Massachusetts, Amherst, USA  
polizzi@ecs.umass.edu

James Kestyn  
ECE Department  
University of Massachusetts, Amherst, USA  
jkestyn@umass.edu

Brendan Gavin  
ECE Department  
University of Massachusetts, Amherst  
bgavin@ecs.umass.edu

## MS129

### Variance Reduction Through Multigrid Deflation

Computing the trace of the inverse of a large matrix relies on a Monte Carlo method which requires the solution of 100s of linear systems. In our Lattice QCD application, we have used singular vectors to effectively deflate the near-null space of the matrix which when coupled with the Hierarchical Probing method deliver close to two orders of magnitude variance reduction. However, deflation has the following scalability problems:

1. the computation of  $m = 500 - 1000$  smallest singular triplets is expensive both for time and storage
2. ill-conditioning grows with problem size requiring larger  $m$
3. the time to project against this deflation space becomes significant

Effective algebraic multigrid methods (AMG) have recently been developed for Lattice QCD. In this research we use the prolongator/projection operators of an AMG method to generate a large eigenspace by computing  $m_i$  singular triplets at the  $i$ -th level of AMG. The idea is to:

1. compute and store the largest part of the space at the coarsest levels,  $m_1 \ll m_2 \ll \dots \ll m_{coarse}$
2. compute the eigenspace at the  $i$ -th level orthogonally to the spaces computed at finer levels.

We use the flexible PRIMME software to compute the singular triplets and report results on its performance as well as on the effectiveness of the resulting deflating space.

Andreas Stathopoulos  
College of William & Mary  
Department of Computer Science  
andreas@cs.wm.edu

Eloy Romero  
College of William and Mary  
eloy@cs.wm.edu

Arjun Gambhir  
Physics Department  
asgambhir@email.wm.edu

Kostas Orginos  
Physics Department  
College of William and Mary  
kostas@wm.edu

## MS129

### Recent Advancements and Future Plans for Next-Generation Eigensolvers in Anasazi

With the availability and diversity of powerful computational resources, including multi-core CPU and GPU technology, there is significant interest in numerical software libraries that allow a developer to optimize the trade-off

between effort and impact. In this talk we will discuss the current and ongoing efforts by which Trilinos is providing enabling technologies for the development of academic and industrial software targeted at next-generation architectures. The focus of this talk will be on the sparse eigensolvers (Anasazi) package in Trilinos and the scalable performance of this solver package, which fundamentally hinges upon efficient linear algebra (Tpetra/Kokkos).

Heidi K. Thornquist  
Sandia National Laboratories  
hkthorn@sandia.gov

## MS129

### Rational and Polynomial Filtering for Eigenvalue Problems

Two filtering techniques are presented for solving large hermitian eigenvalue problems by the method of spectrum slicing that consists of subdividing the spectrum in a number of subintervals and computing eigenvalues in each subinterval independently. In the first approach, the filter is a polynomial constructed as the least-squares approximation to an appropriately centered Dirac distribution. The second approach targets matrices whose spectral distribution is very irregular, as well as generalized eigenvalue problems. It is based on using a rational filter in a least-squares sense.

Yuanzhe Xi  
University of Minnesota  
yxi@cs.umn.edu

Yousef Saad  
Department of Computer Science  
University of Minnesota  
saad@cs.umn.edu

## MS130

### Balanced Multi-Criteria Graph Partitioning

Multi-criteria graph partitioning is a key component for enabling highly scalable multi-physics simulations. The main challenge is to balance all the criteria at the same time. Existing multi-criteria partitioners such as ParMetis are adapted from mono-criterion and focus their goal in minimizing communication cost. However, in practice, this approach is not robust in respect to balance of all criteria, often producing results that are not in the imbalance tolerance. We have implemented, in Scotch, specific multi-criteria partitioning algorithms, designed to meet balance constraints for all criteria. In this talk, we will present and discuss results of these new algorithms.

Cedric Chevalier, Remi Barat  
CEA/DAM  
cedric.chevalier@cea.fr, barat@ocre.cea.fr

Francois Pellegrini  
University of Bordeaux  
francois.pellegrini@labri.fr

## MS130

### Multilevel Acyclic Partitioning of Directed Acyclic Graphs for Enhancing Data Locality

In modern computational systems, the cost of data movement across nodes or within the memory hierarchies inside

a node becomes more and more significant over the cost of performing arithmetic operations. When executing a program, maximizing data reuse will reduce the data movement cost and thus, improve the general execution time. Finding a good acyclic partition of the computational directed acyclic graph (DAG) associated with the algorithm can help finding an execution improving data locality. We present a multilevel direct k-way partitioner for the acyclic partitioning of DAGs. The quality of the computed acyclic partitions are assessed at the graph level by computing the edge cut or the total volume of communication between components, and at the application level by computing the number of cache miss introduced by the partition on various algorithms.

Julien Herrmann  
Georgia Institute of Technology  
julien.herrmann@cc.gatech.edu

Aravind Sukumaran Rajam  
Ohio State University  
sukumaranrajam.1@osu.edu

Fabrice Rastello  
INRIA  
fabrice.rastello@inria.fr

P. (Saday) Sadayappan  
Ohio State University  
saday@cse.ohio-state.edu

Umit V. Catalyurek  
The Ohio State University  
Department of Biomedical Informatics  
umit@gatech.edu

### MS130

#### **Sparse Matrix-Matrix Multiplication for Modern Architectures**

Sparse matrix-matrix multiplication (SPMM) is an important kernel in high performance computing that is heavily used in the graph analytics as well as multigrid linear solvers. Because of its highly sparse structure, it is usually difficult to exploit the parallelism in the modern shared memory architectures. Although there have been various work studying shared memory parallelism of SPMM, some points are usually overlooked, such as the memory usage of the SPMM kernels. Since SPMM is a service-kernel, it is important to respect the memory usage of the calling application in order not to interfere with its execution. In this work, we study memory-efficient scalable shared memory parallel SPMM methods. We study graph compression techniques that reduce the size of the matrices, and allow faster computations. Our preliminary results show that we obtain upto 14x speedup w.r.t NVIDIAs cuSPARSE, 30% speedups w.r.t SPMM implementation provided in Intel Math Library while using 65% less memory.

Mehmet Deveci  
Sandia National Laboratories  
mndevec@sandia.gov

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

Siva Rajamanickam

Sandia National Laboratories  
srajama@sandia.gov

### MS130

#### **Partitioning Irregular Graphs at the Trillion-Edge Scale**

We introduce XtraPuLP, a new distributed-memory partitioner designed to scale to trillion-edge graphs. XtraPuLP is a significant extension from our prior shared-memory partitioner, PuLP, and is based on the label propagation community detection technique. Label propagation has been previously demonstrated to be an effective means of partitioning modern extreme-scale graph-structured datasets, including human social and interaction networks, web crawls, and brain graphs. On a collection of large sparse graphs, we show that XtraPuLP partitioning quality is comparable to other state-of-the-art partitioning methods. We also demonstrate that XtraPuLP can produce high-quality partitions of real-world graphs with billion+ vertices in minutes. Additionally, we find that using XtraPuLP partitions for distributed-memory graph analytics leads to a significant end-to-end reduction in execution time.

George M. Slota  
Rensselaer Polytechnic Institute  
slotag@rpi.edu

Siva Rajamanickam  
Sandia National Laboratories  
srajama@sandia.gov

Kamesh Madduri  
Pennsylvania State University  
madduri@cse.psu.edu

Karen D. Devine  
Sandia National Laboratories  
kddevin@sandia.gov

### MS131

#### **Exploring Efficient Block-Updates on GPU for Quantum Monte Carlo Simulation**

Evaluation of acceptance probability computed as determinant of a dense matrix of wave functions is a computational kernel in QMCPACK. This matrix undergoes a rank-one update if the event is accepted. Sherman-Morrison formula is used to update the inverse of this matrix. This effort explores the use of a delayed update algorithm where accepted events are grouped then applied together to the matrix inverse for higher efficiency on Nvidia GPU and many-core processors.

Eduardo F. D'Azevedo  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
dazevedoef@ornl.gov

Paul Kent, Ying Wai  
Oak Ridge National Laboratory  
kentpr@ornl.gov, ywli@physast.uga.edu

Tyler McDaniel  
University of North Carolina, Asheville  
bmcdanie@unca.edu

Kwai L. Wong  
 Joint Institute for Computational Science  
 University of Tennessee/ORNL  
 kwong@utk.edu

### MS131

#### Performance Evaluation of Time-Space Tiling Strategies for Iterative Stencil Computations on Multi/Many-Core CPU Systems

Iterative stencil computation appears in many scientific computations such as the electro-magnetic analysis using the FDTD method. It is well-known that the performance of iterative stencil computation is usually limited by the memory access performance of a system, which is relatively much lower than the arithmetic performance. For reducing the cost of memory access in iterative stencil computation and improving its performance, so-called time-space tiling (in other words, temporal blocking) has been studied. By using this technique, we can improve data access locality in iterative stencil computation. Several time-space tiling strategies have been already proposed, however their features on recent multi/many-core CPU systems are not fully-clarified because other factors also influence the performance; for example, sufficient parallelism is required for exploiting many computational cores on a CPU. In this study, we experimentally compare several tiling strategies as a preliminary step of developing high-performance stencil-based application programs. We evaluate the performance of typical stencil codes (e.g. 2 dimensional 5 points stencil) that employ time-space tiling on a standard Xeon processor and Xeon Phi processor systems. Throughout the performance evaluation, we discuss features of each tiling strategies on recent multi/many-core CPU systems.

Takeshi Fukaya  
 Hokkaido University, RIKEN AICS  
 fukaya@iic.hokudai.ac.jp

Takeshi Iwashita  
 Hokkaido University  
 iwashita@iic.hokudai.ac.jp

### MS131

#### GPU Computing in Iterative Solvers for Sparse Linear Systems

We present recent results in using GPU to accelerate iterative solvers based on Krylov subspace methods. First we consider sparse linear systems from computational fluid dynamics simulations. We evaluate and optimize the performance of Conjugate Gradient routines designed for GPU accelerator and compare against an industrial CPU-based implementation. Then we show how the distributed parallel Algebraic Recursive Multilevel solver (pARMS) based on MPI can benefit from heterogeneous CPU/GPU architectures. In our implementation, the preconditioning of each part of the distributed matrix (local preconditioning) is performed on a GPU and is based on the randomization of the last Schur complement system in the multilevel recursive process.

Amal Khabou, Aygul Jamal, Marc Baboulin  
 University of Paris-Sud  
 amal.khabou@lri.fr, aygul.jamal@u-psud.fr,  
 marc.baboulin@lri.fr

Masha Sosonkina

Old Dominion University  
 msosonki@odu.edu

### MS131

#### Spectral Graph Partitioning and Clustering on the GPU

The graph partitioning and clustering techniques can be used to minimize communication as well as to perform load balancing in high performance computing and to identify communities in social networks, among many other applications. The quality of the splitting is often measured by a particular metric, such as minimum balanced cut, modularity, etc. In this talk we focus on the spectral approach for finding the minimum balanced cut of a graph into multiple sub-graphs. We review the theory behind the basic algorithm and its generalization to multiple partitions. We analyze the properties of the Laplacian matrix and propose the use of new preconditioned LOBPCG eigenvalue solver, rather than the commonly used Lanczos method. Also, we discuss differences between spectral and multi-level schemes when applied to different classes of problems. Finally, we compare state-of-the-art implementation of spectral scheme in CHACO and multi-level scheme in METIS software packages on the CPU versus our spectral scheme developed on the GPU.

Maxim Naumov  
 NVIDIA  
 maxim.a.naumov@gmail.com

Timothy Moon  
 Stanford University  
 tmoon@nvidia.com

Alexandre Fender  
 NVIDIA  
 afender@nvidia.com

### MS132

#### Integral Equation Solver for Metamaterial Design

We present the implementation of an integral equation method for solving electromagnetic scattering problems with a large number of inclusions, using Julia. As the time complexity of our solver scales roughly linearly with the number of inclusions, we expect to solve large-scale problems that common alternative solvers cannot solve in reasonable time. We describe how this tool can be used in the design and analysis of two-dimensional metamaterials, such as graded photonic crystals. In addition, the possibility of automating the design process by applying optimization techniques to this solver is discussed.

Boaz Blankrot  
 Institute for Analysis and Scientific Computing  
 Technical University of Vienna  
 boaz.blankrot@tuwien.ac.at

Clemens Heitzinger  
 Vienna University of Technology  
 clemens.heitzinger@tuwien.ac.at

### MS132

#### Solving Large-Scale Electromagnetic Inverse Problems

In this talk we discuss the use of a new software package,

jInv for the parallel solution of inverse problems. We show that by using Julia it is possible to easily develop and modify inversion algorithms. We present results that show both the flexibility and robustness of the package inverting large scale electromagnetic data set.

Eldad Haber

University of British Columbia  
haber@eos.ubc.ca

### MS132

#### Lazy Array Computations with Julia

Raw data often require multiple types of transformation prior to visualization or analysis. For large data sets, the overhead of performing "preparatory" transformations can be considerable, and may hinder exploratory investigation especially when the most effective transformations are not known in advance. Julia has developed an efficient array infrastructure supporting on-the-fly transformations including elementwise computation, subregion selection, dimensionality changes, axis permutation (transposition and higher-dimensional generalizations), and index shifting. These transformations may be composed to achieve diverse transformations lazily with little or no extra computational overhead.

Timothy E. Holy

Washington University in St. Louis School of Medicine  
holy@wustl.edu

### MS132

#### Networks Analysis Using Julia

In this talk I will discuss recent work related to using Julia for analyzing large networks. These include localization analysis, link prediction, network alignment, and some more. Julia's in-place functionality, matrix representation, and function wrappers makes it an appealing programming language for networks analysis research. In this talk, I intend to focus on these strengths and showcase certain examples where Julia was used in my research.

Huda Nassar, David F. Gleich

Purdue University  
hnassar@purdue.edu, dggleich@purdue.edu

### MS133

#### Data-Driven, Physics-Constrained Model Reduction of Convection Dominated Flows

Model Order Reduction (MOR) is a promising approach for bridging the gap between high-fidelity, and time-critical applications such as Uncertainty Quantification (UQ), control design, and design optimization. However, MOR is still in its infancy particularly for complex fluid flows characterized by strongly nonlinear behavior, strong shocks, chaotic dynamics and turbulence, and combustion. In this talk we will summarize recent advances in improving the performance and stability of projection based reduced order models for nonlinear fluid flows. In particular, we will discuss methods for augmenting the standard Proper Orthogonal Decomposition (POD) algorithm with physics-based constraint equations.

Maciej Balajewicz

University of Illinois at Urbana-Champaign  
mbalajew@illinois.edu

Irina K. Tezaur

Sandia National Laboratories  
ikalash@sandia.gov

Earl Dowell

Duke University  
earl.dowell@duke.edu

### MS133

#### A Reduced-Basis Smagorinsky Turbulence Model

In this work we present a reduced basis model for the Smagorinsky turbulence model, as a first step to the construction of a reduced projection-based VMS turbulence model, which has a similar structure. This turbulence model includes a non-linear eddy diffusion term that we have to treat using the Empirical Interpolation Method, in order to obtain a linearised decomposition of the reduced basis Smagorinsky model. This model is based upon an *a posteriori* error estimation for Smagorinsky turbulence model. The theoretical development of the *a posteriori* error estimation is based on the Brezzi-Rappaz-Raviart stability theory, and adapted for the non-linear eddy diffusion term. The reduced basis Smagorinsky turbulence model is decoupled in a Online/Offline procedure. First, in the Offline stage, we construct hierarchical bases in each iteration of the Greedy algorithm, by selecting the snapshots which have the maximum *a posteriori* error estimation value. To assure the Brezzi inf-sup condition on our Reduced Basis space, we have to define a *supremizer* operator on the pressure solution, and enrich the reduced velocity space. Then, in the Online stage, we are able to compute a speedup solution of our problem, with a good accuracy. Finally we present some numerical tests, programmed in FreeFem++, in which we show the speedup the computation of a solution of a steady flow in a backward-facing step.

Enrique Delgado, Tomas Chacon, Macarena Gomez

Differential Equations and Numerical Analysis  
Universidad de Sevilla  
edelgado1@us.es, chacon@us.es, macarena@us.es

### MS133

#### An Artificial Neural Network Closure Modeling Framework for Model Order Reduction of Convective Flows

Our primary goal in developing model reduction strategies is to replace large dynamical systems with lower dimensional systems having similar range of validity, input and output characteristics. As a promising approach, POD has been successfully used to generate a representative reduced-order model (ROM) for the control, optimization, and analysis of a large number problems involving fluids. Because POD typically extracts the most energetic modes of a given system, projecting these systems and their solutions onto these low-dimensional bases (i.e., POD modes) often produces ROMs that capture the dominant characteristics of these systems. The resulting systems are low dimensional but dense and provide an efficient framework for applications where either small size or fast simulations are required. It is, however, shown that the state-of-the-art POD-ROM approaches are unlikely to enable high-fidelity predictions of massively convective flows. The novel framework we purposed in this study consists of artificial neural networks (ANNs), which provide inclusion of energy from the truncated modes to the resolved flows utilizing a training process from direct numerical simulations. Thus we chart a systematic route to expanding the applicabil-

ity of these emerging data-driven mathematical tools to put forth an equation-free ANN framework for challenging convective flow problems, with the goal of developing more accurate and robust closure and stabilization models.

Mansoor Ahmed  
Oklahoma State University  
mansoor.ahmed@okstate.edu

Omer San  
Mechanical and Aerospace Engineering  
Oklahoma State University  
osan@okstate.edu

### MS133

#### A Fast Algorithm for Ensemble-Based Reduced Order Simulations of Complex Fluid Flows

Computational efficiency is of paramount importance in many applications of complex fluid flows. Therefore, model reduction techniques have been frequently used by engineers and researchers. Among them, proper orthogonal decomposition is one of the most commonly used methods to generate reduced-order models for turbulent flows dominated by coherent structures. In this talk, we will present an ensemble-based reduced order algorithm of complex flows. It leads to a fast solver for a group of simulations simultaneously while achieving a high accuracy. Combining with nonlinear closure proper orthogonal decomposition method, this new approach provides an efficient and reliable surrogate model for simulating complex fluid flows.

Zhu Wang  
Department of Mathematics  
University of South Carolina  
wangzhu@math.sc.edu

### MS134

#### Nonlinear Model Reduction via Dynamic Mode Decomposition

We propose a new technique for obtaining reduced order models for nonlinear dynamical systems. Specifically, we advocate the use of the recently developed Dynamic Mode Decomposition (DMD), an equation-free method, to approximate the nonlinear term. DMD is a spatio-temporal decomposition of a data matrix that correlates spatial features while simultaneously associating the activity with periodic temporal behavior. With this decomposition, one can obtain a fully reduced dimensional surrogate model and avoid the evaluation of the nonlinear term in the on-line stage. This allows for a reduction in the computational cost, and, at the same time, accurate approximations of the problem. We present a suite of numerical tests to illustrate our approach and to show the effectiveness of the method in comparison to existing approaches.

Alessandro Alla  
Department of Mathematics  
University of Hamburg, Germany  
aalla@fsu.edu

Nathan Kutz  
University of Washington  
Dept of Applied Mathematics

kutz@uw.edu

### MS134

#### Iterative Rational Krylov Algorithms for Unstable Dynamical Systems and Optimality Conditions for a Finite-Time Horizon

Large-scale dynamical systems pose tremendous computational difficulties in simulation, control, design and optimization. Model reduction is one remedy to overcome these challenges. For stable linear dynamical systems, model reduction is both theoretically and computationally well established and one can provide high-fidelity, in some cases optimal, approximations using various techniques. However, most of these methods are mainly aimed at reducing stable systems and extensions to unstable systems are either not established or computationally difficult. In this talk, we investigate interpolatory reduction techniques for unstable systems. Iterative Rational Krylov Algorithm (IRKA) is a model reduction technique for reducing stable linear dynamical systems optimally in the  $\mathcal{H}_2$ -norm. While IRKA cannot provide any optimality guarantees in the case of unstable systems, there are no numerical obstacles to reducing unstable models via IRKA as long as the model has no purely imaginary eigenvalues. The numerical examples illustrate that if the system has only few unstable poles, as is the case in many prominent applications, IRKA captures these unstable poles rather accurately and produces a satisfactory reduced model. We also attempt to establish optimality conditions for model reduction of unstable systems in a finite-time horizon by following a framework based on the Lyapunov equations and finite-time Gramians.

Klajdi Sinani  
Virginia Tech  
klajdi@vt.edu

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

### MS134

#### Model Reduction, Transport Problems and Structure

Although model reduction of linear systems has reached a certain maturity over the last decades, almost every classical method fails or yields poor approximations when applied to transport problems. In the first half of the talk I use ideas from operator splitting to explain the key features preventing standard methods from succeeding and analyze state of the art methods, such as symmetry reduction and variants thereof. Second, I explain how enforcing structure in the surrogate models can offer remedies.

Benjamin Unger  
TU Berlin  
unger@math.tu-berlin.de

### MS134

#### Using Reduced Order Modeling to Solve an Assortment of Nonlocal Problems

Abstract not available

David Witman



Florida State University  
dw11d@my.fsu.edu

**MS135****Simulation of Shock Wave/Boundary Layer Interaction Using High Resolution Numerical Scheme**

Abstract not available

Ovais Khan  
Tuskegee University  
okhan@mytu.tuskegee.edu

**MS136****Fracture Propagation in Porous Media using Phase Field Approach**

The computational modeling of the formation and growth of the fluid filled fractures in poroelastic media is difficult with complex fracture topologies. Here we study the fracture propagation by approximating lower-dimensional fracture surface employing the phase field function. The major advantages of using phase-field modeling for crack propagation are i) it is a fixed-topology approach in which remeshing is avoided, ii) crack nucleation, propagation path are automatically determined based on energy minimization, that is, calculating stress intensity factors are embedded in the model. This avoids creating unstable solutions and reduces computational costs when applied to complex fracture networks and models that include fluid flow. In addition iii) joining and branching of multiple cracks also do not require any additional techniques. We develop robust and efficient numerical algorithms that can be used for three-dimensional applications and employ fixed stress iteration to solve Biot system which considers both fluid flow in the porous media and the fracture. Recently, this algorithm is extended to consider proppant transport in the fractures with locally conservative enriched Galerkin finite element approximations. Several numerical examples considering pressurized, fluid-filled and proppant filled fracture propagation in heterogeneous porous media substantiate our developments.

Sanghyun Lee  
Center for Subsurface Modeling, ICES  
UT Austin, TX, USA  
shlee@ices.utexas.edu

Andro Mikelić  
University of Lyon 1  
mikelic andro.jandro.mikelic@univ-lyon1.fr

Mary F. Wheeler  
Center for Subsurface Modeling, ICES  
University of Texas at Austin  
mfw@ices.utexas.edu

Thomas Wick  
Johann Radon Institute for Computational and Applied Math  
Austrian Academy of Sciences  
thomas.wick@ricam.oeaw.ac.at

**MS136****Coupled Modeling of Flow, Deformation and Heat Transport in Thermo-Hydro-Mechanical Systems**

Increased induced seismicity concerns at geothermal sites

around the world have motivated additional research into thermo-hydro-mechanical (THM) simulation. Concurrently, the development of automated differential equation solvers has allowed for the rapid prototyping and customization of novel solution algorithms to solve multi-physics problems such as THM. This work examines the use of an automated finite element solution framework to create a THM simulator capable of modeling evolution of the state of stress, pressure, and temperature in a poroelastic domain subjected to small perturbations representative of a geothermal operation. The simulator is benchmarked using classical problems with known analytical solutions, including Terzaghi's uniaxial compaction and Mandel's non-isothermal 3D compaction. Current research involves the incorporation of discontinuities (fractures and faults) into the THM simulator and the calculation of time, magnitude, and location of seismic activity along such discontinuities.

Saro Meguerdijian  
University of Southern California  
smeguerd@usc.edu

**MS136****A Variational Eigen-Deformation Model for Brittle Fractures in Fluid-Infiltrating Porous Media under Non-Isothermal Condition**

Many engineering applications, such as geological disposal of nuclear waste, require reliable predictions of hydro-mechanical responses of porous media exposed to extreme environments. This presentation will discuss the relevant modeling techniques designed specifically for porous media subjected to such harsh environments. In particular, we will provide an overview of the variational eigen-deformation techniques used to model brittle fractures and compaction bands. The formulation is based on applying variational principle to an energy-dissipation functional in which the saddle point leads to the governing equation. By regularizing the fracture energy and fluid dissipation functionals, the numerical scheme is able to deliver responses with exhibiting mesh bias. This formulation is particularly advantageous in the sense that (1) there is no need to introduce phase field or other field variables in the governing equations, and that (2) the resultant model is able to propagate fractures without knowing the crack path *a priori*. Numerical examples that predict branching and coalescence of fluid-driven cracks will be included.

WaiChing Sun  
Columbia University, New York  
wsun@columbia.edu

Kun Wang  
Columbia University  
kw2534@columbia.edu

**MS136****An Oscillation-Free Element-Based Finite Volume Technique for Poroelastic Problems**

One of the challenges faced by numerical techniques when solving the poromechanical equations regards to the oscillatory patterns observed in the pressure field. If the numerical formulation employs the same order of approximation for both pressure and displacements, this pathology can arise in regions experiencing undrained consolidation, which usually occurs close to permeable loaded boundaries or at the vicinity of the interface between regions

of different permeabilities. A similar pathology, known as checkerboard pressure field, is also observed when solving the Navier-Stokes equations. In this case, one of the solutions developed by the finite volume community is the Physical Influence Scheme, that tries to improve the evaluation of the velocities present in the continuity equation by obtaining an interpolation scheme based on the momentum equations. In the present work, this same technique is applied to the poroelastic equations to eliminate the pressure oscillations. The main idea is to use the stress equilibrium equations to obtain an improved interpolation function to evaluate the displacements present in mass conservation equation. The three-dimensional poroelastic equations are discretized by the Element based Finite Volume Method (EbFVM) and unstructured grids, composed by elements of different types, are employed. The numerical formulation is validated against 1D consolidation problems and some 3D realistic problems are also presented.

Herminio Tasinato Honório

Federal University of Santa Catarina, Brasil  
herminio.eng@gmail.com

Massimiliano Ferronato, Carlo Janna  
Dept. ICEA - University of Padova  
massimiliano.ferronato@unipd.it, carlo.janna@unipd.it

Clovis R. Maliska  
Mechanical Engineering Department  
Federal University of Santa Catarina  
maliska@sinmec.ufsc.br

#### MS137

### Extraction and Prediction of Coherent Patterns in Incompressible Flows Through Space-Time Koopman Analysis

Abstract not available

Dimitrios Giannakis

Courant Institute of Mathematical Sciences  
New York University  
dimitris@cims.nyu.edu

#### MS137

### Data-Driven Adaptive Control of Combustor Instabilities

A data-based adaptive strategy is developed to control the thrust produced by a Scramjet engine under normal operating conditions and near unstart. Specifically, a direct discrete time adaptive controller called the Retrospective cost adaptive control (RCAC) technique is pursued. The RCAC uses minimal modelling information and limited measurements to determine control laws. A two-dimensional computational fluid dynamic model of a Scramjet engine is used along with a heat release model to represent the dynamics of the problem. This representation involves coupled non-linear partial differential equations with  $O(10^5)$  degrees of freedom. First, the open-loop dynamic response of the model is studied to identify the dynamics of the system. Then, RCAC is used to maintain the commanded thrust in the presence of a disturbance in the inlet Mach number using the past thrust measurements and applied controls. Finally, unstart control is demonstrated using RCAC with an augmented performance criterion based on pressure measurements.

Ankit Goel

University of Michigan  
ankgoel@umich.edu

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

Dennis Bernstein  
University of Michigan  
dsbaero@umich.edu

#### MS137

### Identifying Anomalous Epidemics with Optimized Models of Sensor Networks

Disease outbreaks and epidemics may propagate geographically with anomalous, or non-diffusive, transport characteristics due to the heterogeneity of host contact networks and their behaviors. Any optimized intervention to control an epidemic should be designed with consideration for the intrinsic spatiotemporal properties of the specific network topology and disease. However, measurements of both the contact network and parameters of the possibly novel infectious agent are often incomplete or unattainable. As a practical solution for this complex problem, we suggest simplified mathematical modeling of epidemics as governed by a reaction-diffusion system to guide interventions for a generic disease in a generic population. Assuming finite and heterogeneous resources to measure epidemic progression, we investigated sensor placement strategies that optimize for minimal cost and early discrimination of possibly anomalous infection transport dynamics. We built a general model for spatial transmission of diseases using various types of non-diffusive random walks, including Lévy walks, to represent the movement of infectious agents. We then adapted principles of compressive sensing to optimize sensor coverage when including effects of intrinsic and extrinsic noise. Our results highlight where to look during the early stages of an epidemic so that interventions can be directed to maximize their effectiveness.

Kyle B. Gustafson, Joshua L. Proctor  
Institute for Disease Modeling  
kgustafson@idmod.org, JoshLProctor@gmail.com

#### MS137

### Koopman Operator Framework for Nonlinear System Identification

We present a model identification approach for input-output nonlinear systems by transforming the problem into identification of an associated bilinear albeit higher dimensional system obtained via Koopman operator spectral decomposition. We outline a systematic approach to truncate the bilinear system, and provide theoretically and computationally derived bounds on the truncation error. The proposed framework will be compared with other system identification approaches, and its usefulness illustrated through various examples.

Amit Surana

System Dynamics and Optimization  
United Technologies Research Center (UTRC)  
suranaa@utrc.utc.com

Matthew Williams  
United Technologies Research Center  
william1@utrc.utc.com

Milos Ilak  
 UTRC  
 ilakm@utrc.utc.com

**MS138****Structured Automatic Differentiation for Hedging Financial Contracts**

The hedging of large portfolios of complex financial instruments, for example credit-value-adjustment (CVA) portfolios, can be a computationally exhausting task, often requiring many hours (days) on multi-core machines. A prime contributor to this task is the calculation of sensitivities which is required for many hedging methods. The cost is due in part to the nested Monte-Carlo nature of CVA models used in practise. In this talk we expose this structure and show how sparse and structured AD methodologies can be used to significantly reduce hedging computing costs and move some problems from essentially intractable to just several hours of computing time.

Thomas F. Coleman  
 The Ophelia Lazaridis Research Chair  
 University of Waterloo  
 tfcoleman@uwaterloo.ca

**MS138****Automatic Differentiation for Machine Learning Applications**

We describe the use of automatic/algorithmic differentiation (AD) in machine learning applications. We report on the development of AD capabilities for functions implemented in python and R, describe the use of AD to provide derivatives for the stochastic gradient and stochastic quasi-Newton methods, and discuss how to exploit structure in the computation of derivatives.

Paul D. Hovland  
 Argonne National Laboratory  
 MCS Division, 221/C236  
 hovland@anl.gov

Kaitlyn MacIntyre  
 Northwestern University  
 kaitlynmacintyre2017@u.northwestern.edu

Sri Hari Krishn Narayanan  
 Argonne National Laboratory  
 snarayan@mcs.anl.gov

**MS138****Binomial Checkpointing for Arbitrary Programs with No User Annotation**

Abstract not available

Jeffrey Mark Siskind  
 Purdue University  
 qobi@purdue.edu

**MS138****Second and Higher Order Derivatives Using Reverse Mode of Automatic Differentiation**

Automatic Differentiation is a technique to augment a computer program for computing a function so that the aug-

mented program computes the derivatives as well as the function values. We revisit the Edge Pushing algorithm of Gower and Mello (2012), which evaluates the Hessian, and show that it could be interpreted as the second order reverse mode algorithm for the Hessian. The intermediate results we maintain are the first and second order derivatives of a suitably defined equivalent function  $f_k(S_k)$ , where  $S_k$  constitutes the current set of live variables (currently active variables whose values will be used in future steps in the computation). This observation can be extended to compute high order derivatives via the reverse mode. That is, the intermediate results in each step are the derivative tensors up to order  $d$  of the equivalent function  $f_k(S_k)$ . Thus a general expression for the high order chain rule for evaluating derivatives of  $f_k(S_k)$  from values at the previous step,  $f_{k+1}(S_{k+1})$ , yields an implementation of the high order reverse mode. To further improve the efficiency, we exploit the sparsity by performing updates with only the nonzero values and exploit the symmetry by keeping only the unique elements in the high order derivative tensor due to the inherently high degree of symmetry. We also show applications of our work in uncertainty quantification and quantum chemistry.

Mu Wang  
 Purdue University  
 wangmu0701@gmail.com

**MS139****Bayesian Averaging for Estimating RANS Inadequacy in Large Simulations**

Abstract not available

Richard Dwright  
 Delft University of Technology  
 r.p.dwright@tudelft.nl

**MS139****A Framework for Epistemic Uncertainty Quantification in Large-Eddy Simulation Closures**

Large-eddy simulation (LES) approaches reduce the computational cost of solving turbulent flows by removing small-scale information (and the corresponding need to use fine grids and small time integration steps) from the governing equations via low-pass filtering. However, the effect of the small-scales on the resolved flow field is not negligible and the corresponding subgrid-stresses (SGS) contributions need to be modeled. As a consequence, the assumptions introduced in the formulation of the SGS closure result in a potentially important source of structural uncertainty that propagates to the quantities of interest (QoI). The aim of this work is to develop a framework for uncertainty quantification (UQ) in LES closures that is independent of the initial SGS model formulation to overcome the limitations of traditional sensitivity studies. The framework is based on introducing perturbations to the six degrees of freedom of the modeled SGS tensor, i.e. discrepancy in magnitude, eigenvalues (shape) and eigenvectors (orientation) of the normalized turbulent stresses. The methodology is ensured to produce physically plausible states by limiting the perturbations within the realizability limits of the advection term in the filtered Navier-Stokes equations. In the mini-symposium, the strategy will be described in detail and its capabilities will be analyzed by presenting results of simulations involving turbulent flows

in different scenarios.

Stefan P. Domino  
Sandia. National Laboratories  
spdomin@sandia.gov

Lluís Jofre  
Stanford University  
jofre@stanford.edu

Gianluca Iaccarino  
Stanford University  
Mechanical Engineering  
jops@stanford.edu

### MS139

#### Deep Learning for Turbulence Modeling

Reynolds Averaged Navier Stokes (RANS) models are widely used in turbulent flow simulations because of their computational efficiency. However, they are notoriously inaccurate in flows with swirl, strong pressure gradients, separation, and stagnation. Recently, there has been increased interest in developing data-driven methods to provide improved RANS model closures. This talk will discuss a deep learning approach to data-driven turbulence modeling. A deep neural network was trained on a database of flows for which high fidelity data were available in order to predict the Reynolds stress anisotropy tensor. A specialized neural network architecture was developed to embed Galilean invariance into the model predictions. The network performance was evaluated through cross-validation on a database of flows. SAND2016-12088 A

Julia Ling  
Sandia National Laboratories  
jling@sandia.gov

Andrew Kurzawski  
UT Austin  
ajkurzawski@gmail.com

Jeremy Templeton  
Sandia National Laboratories  
jatempl@sandia.gov

### MS139

#### Estimating and Verifying K-Epsilon Model Coefficients for Jet-in-Crossflow Simulations

k-epsilon RANS models predict jet-in-crossflow interactions poorly. This is due to the use of inappropriate model parameters simply picked from the literature and the approximations inherent in RANS. In this work we calibrate RANS parameters to three different experimental datasets using a Bayesian technique. The parameters are compared with values predicted using an analytical model of the jet-in-crossflow interaction. Our preliminary results show that the bulk of the prediction error results from the use of inappropriate parameter values.

Jaideep Ray  
Sandia National Laboratories, Livermore, CA  
jairay@sandia.gov

Sophia Lefantzi, Julia Ling, Lawrence Dechant,  
Srinivasan Arunajatesan  
Sandia National Laboratories

slefant@sandia.gov, jling@sandia.gov, ljdecha@sandia.gov,  
sarunaj@sandia.gov

### MS140

#### Probabilistic, Coarse-Grained Models for PDEs with Random Coefficients

We describe a Bayesian perspective to coarse-graining (CG) of PDEs in the context of uncertainty quantification. It can also be seen as an attempt to learn reduced-order models that are physically inspired and that retain some of the structure of the full-order model. The canonical problem considered consists of a PDE where coefficients (e.g. material parameters) exhibit fine-scale, random variability. The brute-force approach for the solution of such problems requires grids that are small enough to resolve these fluctuations and ultimately yield a very large system of algebraic equations. Several physically- and mathematically-inspired techniques have been developed (e.g. classical or numerical homogenization) that generate closed systems of equations for the solution at the coarser scale while accounting for the effects of the unresolved subgrid scales. In order to make pertinent multiscale techniques feasible in the stochastic setting, it is critical that they achieve sublinear complexity i.e., the computational cost scales sublinearly with the cost of solving the fine-scale problem. We describe such a data-driven attempt where the CG model attempts to learn the features of the microscopic details that are important in terms of predicting the macroscopic response. Furthermore, we attempt to quantify the epistemic uncertainty in the CG model which arises due to the unavoidable information loss that takes place, and propagate it in the predictions produced.

Constantin Grigo  
Technische Universität München  
constantin.grigo@tum.de

Phaedon S. Koutsourelakis  
Technical University Munich  
p.s.koutsourelakis@tum.de

### MS140

#### Data and the Computational Modeling of Complex/Multiscale Systems

Obtaining predictive dynamical equations from data lies at the heart of science and engineering modeling, and is the linchpin of our technology. In mathematical modeling one typically progresses from observations of the world (and some serious thinking!) first to equations for a model, and then to the analysis of the model to make predictions. Good mathematical models give good predictions (and inaccurate ones do not) - but the computational tools for analyzing them are the same: algorithms that are typically based on closed form equations. While the skeleton of the process remains the same, today we witness the development of mathematical techniques that operate directly on observations -data-, and appear to circumvent the serious thinking that goes into selecting variables and parameters and deriving accurate equations. The process then may appear to the user a little like making predictions by "looking into a crystal ball". Yet the "serious thinking" is still there and uses the same -and some new- mathematics: it goes into building algorithms that "jump directly" from data to the analysis of the model (which is now not available in closed form) so as to make predictions. Our work here presents a couple of efforts (e.g. the reconstruction of normal forms from data) that illustrate this "new path from

data to predictions: It really is the same old path, but it is travelled by new means.

Ioannis Kevrekidis  
Princeton University  
yannis@princeton.edu

#### MS140

##### The Parallel Knowledge Gradient Method for Batch Bayesian Optimization

In this talk, we will introduce parallel knowledge gradient (pKG) algorithm for batch Bayesian optimization. The method chooses to evaluate the one-step Bayes optimal batch of points. We demonstrate empirically that the method can find global minima significantly faster than previous batch Bayesian Optimization methods on both synthetic functions and tuning hyperparameters of complex machine learning algorithms. Especially, the method provides most value in the noisy setting.

Jian Wu  
Cornell University  
jw926@cornell.edu

#### MS140

##### Deep Gaussian Processes for Uncertainty Quantification

Recently-developed deep Gaussian process (DGP) models are an attractive solution for powerful nonlinear dimensionality reduction and hierarchical feature representation as well as modeling nonstationary response surfaces. We introduce a fully Bayesian version of the DGP model for unsupervised and supervised learning tasks and demonstrate its utility in overcoming the above-mentioned challenges under limited training data while performing uncertainty quantification tasks for multiscale/multiphysics problems. First, the hierarchical structure of DGPs allows for flexible nonlinear dimensionality reduction that outperforms standard methods for the representation of multiscale/multiresolution properties. Second, we use a supervised DGP to construct a surrogate for the corresponding multi-output response surface and demonstrate its ability to learn nonstationary behavior. Finally, we examine how the Bayesian nature of our model guards against overfitting and produces estimates of the model outputs capturing the epistemic uncertainty due to limited data. Examples will be shown from UQ tasks in random heterogeneous media.

Nicholas Zabaras, Steven Atkinson  
University of Notre Dame  
nzabaras@gmail.com, satkinso@nd.edu

#### MS141

##### Title Not Available at Time of Publication

Abstract not available

Vrushali A. Bokil  
Oregon State University  
bokilv@math.oregonstate.edu

#### MS141

##### High-Order Finite-Difference Time-Domain Simulation of Electromagnetic Waves with Emphasis on

##### Interfaces Between Dispersive Optical Media

Abstract not available

Michael Jenkinson  
Rensselaer Polytechnic Institute  
jenkim2@rpi.edu

#### MS141

##### Pulse Slowdown in Two-Dimensional Active Media

Abstract not available

Gregor Kovacic  
Rensselaer Polytechnic Inst  
Dept of Mathematical Sciences  
kovacg@rpi.edu

#### MS141

##### Title Not Available at Time of Publication

Abstract not available

Jinjie Liu  
Delaware State University  
jliu@desu.edu

#### MS142

##### GPU Algorithms for Monte Carlo Particle Transport

The random nature of Monte Carlo particle transport presents significant challenges for efficient execution on GPUs. The wide variety of particle history lengths as well as types of interactions that particles can undergo are not easily adapted to the vectorized nature of GPUs and other advanced architectures such as the Intel Xeon Phi. In this talk we describe algorithmic improvements aimed at reducing thread divergence and improving efficiency for both traditional history-based transport as well as event-based transport. We primarily consider multi-group transport, but will also briefly discuss approaches for treating continuous-energy nuclear data. We compare the efficiency of GPU implementations to standard distributed memory (MPI) and shared memory (OpenMP) CPU implementations. On-node speedups of 2.5-3x relative to 8-core CPU calculations are achieved. Additionally, parallel weak scaling efficiency above 95% on more than 1000 GPUs is demonstrated on the Titan supercomputer at Oak Ridge National Laboratory.

Steven Hamilton, Thomas Evans  
Oak Ridge National Laboratory  
hamiltonsp@ornl.gov, evanstm@ornl.gov

Stuart Slattery

Computer Science and Mathematics Division  
Oak Ridge National Laboratory  
slatterysr@ornl.gov

#### MS142

##### A Fast Iterative Method for Non-Linear LTE Radiative Transfer Problems

We implement a multigrid preconditioned, asymptotic preserving, weakly penalized discontinuous Galerkin method using non-overlapping Schwarz smoothers to solve a non-linear, frequency and angle dependent radiative transfer

equation in local thermodynamic equilibrium with applications in particle transport through diffusive media. Our methodology pursues a scalable implementation in order to address large scale calculations arising from applications such as astrophysics, atmospheric radiation calculations and nuclear applications. Frequencies are addressed by a multi-group approach and angles by discrete ordinates, non-overlapping Schwarz smoothers are based on solving full local radiative transfer problems for each grid cell which are performed in parallel on a matrix-free implementation. In several tests we show the robustness of the approach for different mesh sizes, cross sections, frequency distributions and anisotropic regimes.

José Pablo Lucero Lorca

IWR, Heidelberg University, Im Neuenheimer Feld 205,  
69120 H  
pablo.lucero@iwr.uni-heidelberg.de

Guido Kanschat

Universität Heidelberg, Germany  
kanschat@uni-heidelberg.de

#### MS142

##### **Adaptive Time Step Selection in Radiative Transfer Calculations**

Time step selection in thermal radiative transfer simulations is often an afterthought when compared to the choice of spatial discretization and thermal iteration strategy. However, time step selection is a critical component of time dependent radiative transfer calculations. Taking too small of time steps inhibits productive throughput. Taking too large of time steps yields inaccurate solutions. In this presentation, we compare the results obtained using multi-stage, embedded S-stable diagonally Implicit Runge-Kutta (ESDIRK) time integrators to control time step versus more traditional, ad-hoc methods of time step control. Additionally, we compare the effects of adapting time step size within a given time step versus using adaptive time criteria to only inform subsequent time step sizes. \*This work

was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Peter G. Maginot

Lawrence Livermore National Laboratory  
maginot1@llnl.gov

#### MS142

##### **A Fully Synchronous Domain Decomposed Transport Algorithm with Splitting**

We have developed a fully synchronous domain decomposed transport algorithm that allows particle splitting for Monte Carlo radiation transport. This algorithm builds upon previously developed asynchronous domain decomposed algorithms with splitting but allows for a fully synchronous solve to avoid the race conditions that can occur with load imbalances due to splitting. Scaling studies of this algorithm on small clusters and on Titan at ORNL as applied to several radiation transport problems of interest will be performed and results will be reported. A sensitivity study to determine optimal algorithm parameters will also be performed.

Tara Pandya, Thomas Evans

Oak Ridge National Laboratory  
pandyatm@ornl.gov, evanstm@ornl.gov

J. Austin Ellis

North Carolina State University  
jaellis2@ncsu.edu

Seth Johnson

Oak Ridge National Laboratory  
johnsonsr@ornl.gov

#### MS143

##### **Adjoint Sensitivity Analysis for Scale-Resolving Turbulent Flow Solvers**

Adjoint-based sensitivity methods are powerful design tools for computational fluid dynamics. At the same time, engineers are starting to use scale-resolving simulations like Large-eddy simulations (LES) and direct numerical simulations (DNS), which resolve more scales in complex flows with unsteady separation and jets than Reynolds-averaged Navier-Stokes (RANS) methods. However, the conventional adjoint method computes large, unusable sensitivities for scale-resolving simulations, which unlike RANS simulations exhibit the chaotic dynamics inherent in turbulent flows. Sensitivity analysis based on Least-squares shadowing (LSS) avoids the issues encountered by conventional adjoint methods, but has a high computational cost even for relatively small simulations [Q. Wang, R. Hui, and P. Blonigan. Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210224, June 2014.]. The following talk discusses a new, more computationally efficient formulation of LSS and its application to turbulent flows simulated with Eddy, a discontinuous-Galerkin spectral-element-method LES/DNS solver [L. Diosady and S. Murman, Higher-order methods for compressible flows using entropy variables. *AIAA Paper 2015-0294*]. First, the new LSS formulation, called Non-intrusive LSS, is outlined, followed by a cost analysis of the method. Results are presented for benchmark turbulent flows, including flow around a low-pressure turbine blade.

Patrick J. Blonigan

Massachusetts Institute of Technology  
patrick.j.blonigan@nasa.gov

Laslo Diosady, Anirban Garai, Scott Murman

NASA Ames Research Center  
laslo.diosady@nasa.gov, anirban.garai@nasa.gov,  
scott.murman@nasa.gov

#### MS143

##### **Adjoint and Tangent Methods in a Higher-Order Space-Time Discontinuous-Galerkin Solver for Turbulent Flows**

We present a higher-order space-time discontinuous-Galerkin solver for the scale resolving solution of compressible turbulent flows. Tangent and adjoint methods are presented for sensitivity analysis. We apply this solver for the solution of compressible flows with increasing Reynolds number. We demonstrate period doubling with increasing Reynolds number and show the unbounded growth of adjoint solutions as the flow becomes chaotic. The tangent and adjoint methods described herein may be used as part of more advanced sensitivity techniques such as least-squares shadowing as will be presented in "Adjoint Sensi-

tivity Analysis for Scale-Resolving Turbulent Flow Solvers” also in this mini-symposium.

Laslo Diosady, Scott Murman  
NASA Ames Research Center  
laslo.diosady@nasa.gov, scott.murman@nasa.gov

Patrick J. Blonigan  
Massachusetts Institute of Technology  
patrick.j.blonigan@nasa.gov

Anirban Garai  
NASA Ames Research Center  
anirban.garai@nasa.gov

#### MS143

##### **Sensitivity of Wakes Via Global Stability Analysis, Covariant Lyapunov Vectors and Shadowing Method**

We apply two methods based on eigenvalue analysis to investigate the physics and sensitivity of chaotic bluff-body wakes, obtained by high-fidelity fluid dynamics computations. First, we perform a global stability analysis by calculating the eigenpairs of the Navier-Stokes operator linearized around the mean flow, from laminar to turbulent regimes. Although the analysis is mathematically restricted to laminar flows, for which fixed points of the equations are well-defined, we qualitatively predict the dominant dynamic features of the turbulent wake. The calculation of the adjoint eigenpairs enables the sensitivity calculation of control devices placed into the flow. Secondly, to tackle turbulent wakes, we extend the global stability analysis by computing the Lyapunov covariant vectors and exponents. This method does not rely on mean flow assumptions, hence, it is suitable for the calculation of chaotic flow sensitivities. We compare the Lyapunov eigenpairs with the eigenpairs calculated by global stability analysis. By using a non-intrusive method, the sensitivity of chaotic wakes is calculated via the shadowing method. The pros and cons of the two analyses are elucidated.

Luca Magri  
Stanford University  
lm547@cam.ac.uk

Georgios Rigas  
CalTech  
rigas@caltech.edu

Lucas Esclapez  
Stanford University  
lesclape@stanford.edu

Qiqi Wang  
Massachusetts Institute of Technology  
qiqi@mit.edu

#### MS143

##### **Adjoint-Based Sensitivity Analysis of Unsteady and Chaotic Flows in Fun3d**

Abstract not available

Eric Nielsen  
NASA Langley

eric.j.nielsen@nasa.gov

#### MS144

##### **Lessons from Running Open Source Projects: Building Sustainable Software and Sustainable Communities**

Computational science has given rise to quite a number of open source software projects and libraries that have several 100,000 or millions of lines of code, have been maintained over many years or decades, and are used by hundreds or thousands of users. Given academic time scales, both the user, developer, and maintainer base may be replaced several times over the lifetime of such projects. This raises the question of how these projects can be built and administered in a way that leads to sustainability both in the code base and in the community that supports the project. In this talk, we will share our experience in leading the deal.II project (see <http://www.dealii.org/>), a large C++ library that supports finite element computations to solve partial differential equations. In particular, we will discuss our approach to forging and growing a community of developers whose members are mentored to grow into roles that allow them to replace those who rotate out of the project. We will also discuss some of the technical aspects that we believe are important for long-term sustainability of open source software projects.

Wolfgang Bangerth  
Texas A&M University  
bangerth@colostate.edu

Timo Heister  
Clemson University  
Mathematical Sciences  
heister@clemson.edu

#### MS144

##### **The Research Software Engineer: An Emerging New Role in Academia in the UK**

Since 2012, in the UK, we have been defining a new role in the academic ecosystem: the research software engineer. Hypothesising that a key cause of problems with reusability, readability and reliability of scientific software has been the lack of a career structure for scientific programmers in universities, we have defined a new role and title, and have been working to normalise this with institutions and research councils. The research software engineer combines a strong scientific and mathematical background, and membership in the research community, with a focus on best practice in software engineering, and is not measured through traditional research performance metrics. Our work has resulted in a new fellowship call from the UK Engineering and Physical Sciences research council, establishing a prestigious new programme for future leaders of scientific programming. Research Software Engineering Groups, who collaborate with researchers in all fields to co-create well engineered scientific software for problems, have been established in several leading UK universities, including Cambridge, Manchester, Sheffield, Southampton, Bristol and University College London. Established scientific programming groups in UK national laboratories are also embracing the new title and professional status, helping with the establishment of a sustainable new career structure. In this talk, we will share some experiences from the four years of this campaign.

James Hetherington

University College London  
j.hetherington@ucl.ac.uk

Simon Hettrick  
University of Sheffield and Software Sustainability Institut  
s.hettrick@software.ac.uk

Rob Haines  
University of Manchester  
rhaines@manchester.ac.uk

Mike Croucher  
University of Sheffield  
m.croucher@sheffield.ac.uk

Alyss Brett  
Culham Centre for Fusion Energy  
alys.brett@ukaea.uk

#### MS144

##### **CSE Software Ecosystems: Critical Instruments for Scientific Discovery**

Software is an essential product of CSE research when complex models of reality are cast into algorithms; moreover, the development of efficient, robust, and sustainable software is at the core of CSE. While the community is beginning to embrace the fundamental role of open-source software ecosystems to support CSE collaboration and enable advances in scientific and engineering understanding, much work remains to overcome challenges in software sustainability and productivity. Difficulties arise from the confluence of disruptive changes in computing architectures, new opportunities for greatly improved simulation capabilities, and demand for greater scientific reproducibility. New architectures require fundamental algorithm and software refactoring, while at the same time enabling new multiscale and multiphysics modeling, simulation, and analysis.

This presentation will introduce activities under way throughout the community to address a variety of technical and social issues in scientific software. We will also highlight multipronged work in the IDEAS Scientific Software Productivity Project to distill and promote best practices in software engineering and productivity for CSE, develop the xSDK as a foundation of a community CSE software ecosystem, use these best practices and software to achieve science advances (e.g., in subsurface flow with hydrological and biogeochemical recycling), and engage the community in collaborative contributions.

Lois Curfman McInnes  
Mathematics and Computer Science Division  
Argonne National Laboratory  
curfman@mcs.anl.gov

Michael Heroux  
Sandia National Laboratories  
maherou@sandia.gov

David Moulton  
Los Alamos National Laboratory  
Applied Mathematics and Plasma Physics  
moulton@lanl.gov

David E. Bernholdt  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
bernholdtde@ornl.gov

Xiaoye Sherry Li  
Computational Research Division  
Lawrence Berkeley National Laboratory  
xsli@lbl.gov

Timothy D. Scheibe  
Pacific Northwest National Laboratory  
tim.scheibe@pnnl.gov

Ulrike Meier Yang  
Lawrence Livermore National Laboratory  
yang11@llnl.gov

#### MS144

##### **Data Skills and Software Training to Enable Data-Driven Discovery**

Although petabytes of data are now available, most scientific disciplines are failing to translate this sea of data into scientific advances. The missing step between data collection and research progress is a lack of training for scientists in crucial skills for effectively and reproducibly managing and analyzing large amounts of data. Already faced with a deluge of data, researchers themselves are demanding this training and need to learn while on-the-job. They require training that is immediate, accessible, appropriate for their level and relevant to their domain. This training needs to include not only technical skills, but ways of thinking about data to provide learners with the knowledge of what is possible along with the confidence to continue self-guided learning. Short, intensive, hands-on Software and Data Carpentry workshops give researchers the opportunity to engage in deliberate practice as they learn these skills, starting with strong foundational skills and receiving feedback as they learn. This model has been shown to be effective, with the vast majority (more than 90%), of learners saying that participating in the workshop was worth their time and led to improvements in their data management and data analysis skills. We have trained over 20,000 learners since 2014 on 6 continents with over 700 volunteer instructors, with the goal of providing effective training that empowers researchers to turn data into knowledge and discovery.

Tracy K. Teal  
Data Carpentry  
tkteal@datacarpentry.org

#### MS145

##### **Multi Space Reduced Basis (MSRB) Preconditioners for Large-Scale Parametrized PDEs**

In this talk we present a new 2-level preconditioner for the efficient solution of large-scale linear systems arising from the discretization of parametrized PDEs. Our preconditioner combines multiplicatively a reduced basis (RB) coarse solver and a nonsingular preconditioner, such as 1-level Additive Schwarz, Gauss-Seidel or Jacobi preconditioner. The proposed technique hinges upon the construction of a new Multi Space Reduced Basis (MSRB) method, where a RB space is built through proper orthogonal decomposition at each iteration of the Richardson or the flexible GMRES method used to solve the large-scale linear system. As a matter of fact, each reduced space is suited to solve a particular iteration and aims at fixing the scales that have not been treated by previous iterations and the fine preconditioner yet. Not only, the RB error decays exponentially fast for each space, thus yielding to very small (compared to the dimension of the original sys-



tem) reduced spaces. Since the RB accuracies obtained for each space affect the overall convergence of the iterative method in a multiplicative way, a very accurate solution of the large-scale system can be computed in very few iterations. Numerical tests have been carried out to evaluate the performance of the preconditioner in different large-scale modeling settings related to parametrized equations, up to millions of degrees of freedom, and compared with the current state-of-art algebraic multigrid preconditioners.

Niccol Dal Santo

Ecole Polytechnique Fédérale de Lausanne  
CMCS - Chair of Modeling and Scientific Computing  
niccolo.dalsanto@epfl.ch

Andrea Manzoni  
EPFL, MATHICSE-CMCS  
Switzerland  
andrea.manzoni@epfl.ch

Simone Deparis  
Ecole Polytechnique Fédérale de Lausanne  
simone.deparis@epfl.ch

Alfio Quarteroni  
Ecole Pol. Fed. de Lausanne  
Alfio.Quarteroni@epfl.ch

**MS145**

**Online Interpolation Point Refinement for Reduced Order Models**

A genetic algorithm procedure is demonstrated that refines the selection of interpolation points of the discrete empirical interpolation method (DEIM) when used for constructing reduced order models for time dependent and/or parametrized nonlinear partial differential equations (PDEs) with proper orthogonal decomposition. The method achieves nearly optimal interpolation points with only a few generations of the search, making it potentially useful for *online* refinement of the sparse sampling used to construct a projection of the nonlinear terms. With the genetic algorithm, points are optimized to jointly minimize reconstruction error and enable dynamic regime classification. The efficiency of the method is demonstrated on two canonical nonlinear PDEs: the cubic-quintic Ginzburg-Landau equation and the Navier-Stokes equation for flow around a cylinder. Using the former model, the procedure can be compared to the ground-truth optimal interpolation points, showing that the genetic algorithm quickly achieves nearly optimal performance and reduced the reconstruction error by nearly an order of magnitude.

Nathan Kutz

University of Washington  
Dept of Applied Mathematics  
kutz@uw.edu

Syuzanna Sargsyan  
University of Washington  
n/a

Krithika Manohar  
University of Washington, USA  
kmanohar@uw.edu

Steven Brunton  
Princeton University

sbrunton@princeton.edu

**MS145**

**Multifidelity Monte Carlo Estimation of Variances and Sensitivity Indices**

In variance-based global sensitivity analysis, sensitivity indices are computed which attribute the variance of the output of a forward model to the variance of its inputs, allowing the inputs to be ranked in terms of importance. In general, these sensitivity indices are unknown and must be estimated, e.g. using Monte Carlo (MC) methods. When the forward model is expensive to evaluate, the MC estimation of the sensitivity indices may become prohibitively expensive. To accelerate MC global sensitivity analysis, we have developed new multifidelity Monte Carlo (MFMC) estimators for the variance of the forward model output and the main and total effect sensitivity indices. Our MFMC approach leverages surrogate models of varying fidelities to accelerate the MC estimation. These surrogate models may be of any type and we do not require any error estimates or bounds for the surrogate predictions to be available to us. Instead, minimal high-fidelity forward model evaluations are used to maintain accuracy of the MFMC estimator. For a fixed computational cost, our MFMC estimators exhibit a mean-squared error reduction of 2-3 orders of magnitude relative to their MC counterparts. We present results for both analytical model problems as well as a groundwater flow model problem.

Elizabeth Qian

Massachusetts Institute of Technology  
elizqian@mit.edu

Dan O'Malley  
Los Alamos National Lab  
omalled@lanl.gov

Benjamin Peherstorfer  
ACDL, Department of Aeronautics & Astronautics  
Massachusetts Institute of Technology  
pehersto@mit.edu

Velimir V. Vesselinov  
Los Alamos National Lab  
vuv@lanl.gov

Karen E. Willcox  
Massachusetts Institute of Technology  
kwillcox@MIT.EDU

**MS145**

**Nonlinear Model Order Reduction for Steady Aerodynamic Design Applications**

This talk introduces a framework for building and leveraging projection-based Reduced Order Models (ROMs) for industrial-scale steady-state nonlinear Computational Fluid Dynamics (CFD) applications involving shape parameters. The proposed framework is demonstrated for the Reynolds-averaged Navier-Stokes equations past two parametrically deformed geometries: a commercial aircraft at cruise conditions, and a passenger car at highway speed. In these examples, the proposed ROM framework is shown to predict integrated forces with accuracy similar to simple data-fit approaches, and is shown to predict distributed quantities with better accuracy than simple data-fit approaches. Computational speedups are reported to be four-

to-five orders of magnitude relative to the original CFD model.

Kyle Washabaugh, Matthew J. Zahr, Charbel Farhat  
Stanford University  
washabaugh@stanford.edu, mjjzahr@lbl.gov,  
CFarhat@stanford.edu

#### MS146

##### **Regionally Implicit and Explicit Methods for the Relativistic Vlasov-Maxwell System**

In the relativistic limit, the Vlasov-Maxwell system introduces numerical difficulties as explained in (Suzuki, 2010). We develop an efficient solver for the relativistic Vlasov-Maxwell (RVM) system in order to model laser-plasma interactions; and in particular, the acceleration of electrons or ions to relativistic energies. In doing so we expand on the so called Locally Implicit Discontinuous Galerkin method (LIDG) developed in (Qiu2005,a) by defining Regionally Implicit Discontinuous Galerkin Methods. These methods are parametrized by the region parameter  $r$ . For a given cell, the region parameter determines how many neighboring cells (this collection of cells known as the region) will provide information to the prediction step of the method. We use a Rusanov Riemann solver on the interior of said region and the interior cell values on the boundary of the region. We show that these methods allow a much larger CFL number when compared to the LIDG method, and thus offer a vastly improved efficiency over the LIDG method. Here we introduce the methods for 1D and 2D problems including optimal CFL choices, convergence tests, and discussion of nonlinear problems.

Pierson Guthrey  
Iowa State University  
Department of Mathematics  
pguthrey@iastate.edu

#### MS146

##### **Vector-Potential Finite-Element Formulations for Two-Dimensional Resistive Magnetohydrodynamics**

Vector-potential formulations are attractive for electromagnetic problems in two dimensions, since they reduce both the number and complexity of equations, particularly in coupled systems, such as magnetohydrodynamics (MHD). In this talk, we consider the finite-element formulation of a vector-potential model of two-dimensional resistive MHD. Existence and uniqueness are considered separately for the continuum nonlinear equations and the discretized and linearized form that arises from Newton's method applied to a modified system. Under some conditions, we prove that the solutions of the original and modified weak forms are the same, allowing us to prove convergence of both the discretization and the nonlinear iteration.

Yunhui He  
Memorial University of Newfoundland  
yh6171@mun.ca

#### MS146

##### **A Least-Squares Approach to Two-Fluid, Electromagnetic Plasma**

A two-fluid plasma (TFP) model is presented, both as a stand-alone solver and as the preconditioner to a fully im-

PLICIT, particle-in-cell (PIC) simulation. The model couples fluid conservation equations for ions and electrons to Maxwell's equations. A Darwin approximation of Maxwell is used to eliminate spurious light waves. After scaling and modification, the TFP-Darwin model yields a nonlinear, first-order system of equations whose Frchet derivative is shown to be uniformly H1-elliptic. This system is addressed numerically by nested iteration (NI), a First-Order System Least Squares (FOSLS) discretization, adaptive local mesh refinement, and scaled AMG system solver. Numerical tests demonstrate the efficacy of this approach, yielding an approximate solution within discretization error in a relatively small number of computational work units.

Thomas Manteuffel  
University of Colorado  
tmanteuf@colorado.edu

Chris Leibs  
University of Colorado Boulder  
leibs@colorado.edu

#### MS146

##### **A Fast Linear-Time Higher-Order Unconditionally Stable 2D Implicit Wave Equation Solver**

In the current scientific computing era, computing solution to the wave equation is a basic need for several applications. Since finding analytical solutions is not an easy task for most wave propagation problems, it is essential to develop accurate and efficient numerical computation tools. Towards the end, we have implemented an implicit multi-dimensional wave equation solver which breaks the Courant-Friedrichs-Lewy (CFL) time step restriction and is highly efficient. The method is based on an A-stable (to all orders) solver for the two way wave equation. The formulation of the method is based on the Method of Lines Transpose (MOLT) - a practical approach to Rothe's method - which starts by discretizing the problem in time and then solving the resulting Helmholtz equation with a Green's function method. By leveraging an alternating direction implicit (ADI) splitting and a fast convolution method for the 1D Green's function, it can be shown that the method has  $O(N)$  computational cost. High-order accuracy is achieved through successive convolution. This existing solver has been implemented and verified for 2D constant and variable speed problems. In this talk we show the extension of the method to the problem of general scatterers with non-reflecting boundary conditions, while assessing the accuracy of our method on classical E.M. problems.

Thavappiragasam Mathialakan, Aditya Viswanathan  
Michigan State University  
mathialakan@gmail.com, aditya@math.msu.edu

Andrew J. Christlieb  
Michigan State University  
Dept. of Comp. Math., Sci & Engr.  
andrewch@math.msu.edu

#### MS147

##### **Fourth Order Time Stepping Methods for Space Fractional Nonlinear Schrodinger Equations**

A fourth-order implicit-explicit time-discretization scheme based on the exponential time differencing approach with a fourth-order compact scheme in space is proposed for space

fractional nonlinear Schrodinger equations. The stability and convergence of the compact scheme are discussed. It is shown that the compact scheme is fourth-order convergent in space and in time. Numerical experiments are performed on single and coupled systems of two and four fractional nonlinear Schrodinger equations. The results demonstrate accuracy, efficiency, and reliability of the scheme. A linearly implicit conservative method with the fourth-order compact scheme in space is also considered and used on the system of space fractional nonlinear Schrodinger equations.

Abdul M. Khaliq  
Middle Tennessee State University  
Department of Mathematical Sciences  
akhaliq@mtsu.edu

#### MS147

##### **A Unified Spectral Method for Time and Space Distributed FPDEs**

We develop a unified spectral method of Petrov-Galerkin sense for a general class of linear distributed-order fractional partial differential equations (FPDEs) in any  $(d+1)$ ,  $d \geq 1$ , spatio-temporal dimensions. To efficiently solve the corresponding linear systems, we also develop a unified fast solver with optimal complexity. The proposed scheme works with the same ease and efficiency for parabolic, elliptic, and hyperbolic problems.

Ehsan Kharazmi, Mohsen Zayernouri  
Michigan State University  
kharazm2@msu.edu, zayern@egr.msu.edu

#### MS147

##### **Sparse Approximation of Stochastic Tempered Fractional Differential Equations**

Compressive sensing has become a powerful addition to uncertainty quantification in recent years. This talk employs compressive sensing algorithms to obtain the sparse approximation of stochastic tempered fractional differential equations. In particular, to reduce the high dimensionality in stochastic space and enhance sparsity, we identify new bases for random variables through linear mappings such that the representation of the quantity of interest is sparser with new basis functions associated with the new random variables. This sparsity increases both the efficiency and accuracy of the compressive sensing-based uncertainty quantification method. Specifically, we consider rotation-based linear mappings which are determined iteratively for Hermite polynomial expansions. We demonstrate the effectiveness of the new method by solving stochastic tempered fractional partial differential equations with more than 50 random dimensions.

Guang Lin  
Purdue University  
guanglin@purdue.edu

#### MS147

##### **Numerical Solution of Fractional Partial Differential Equations Via Hybrid Functions**

A new numerical method for solving the fractional partial differential equations is presented. The method is based upon hybrid functions approximation. The Riemann-Liouville fractional integral operator for hybrid functions is given. This operator is then utilized to reduce the solution

of the fractional partial differential equations to a system of algebraic equations. Illustrative examples are included to demonstrate the validity and applicability of the technique.

Somayeh Mashayekhi  
Florida State University  
smashayekhi@fsu.edu

#### MS148

##### **Algebraic Multigrid: Theory and Practice**

This talk gives an overview of recent progress made in the design and analysis of algebraic multigrid methods. The focus is on the setup algorithm that automatically constructs the multilevel hierarchy used in the solve phase. A sharp two-grid theory is introduced and then used to derive various quality measures of the coarse spaces constructed by the setup algorithm, based on the ideas of compatible relaxation, a related identity that assumes the use of the so-called ideal interpolation operator, and an optimal form of classical algebraic multigrid interpolation that gives the best possible two-grid convergence rate. Various numerical results are presented to illustrate these theoretical results. As a test problem, we focus on a finite volume discretization of a scalar diffusion problem with highly varying (discontinuous) diffusion coefficient.

James Brannick  
Pennsylvania State University  
jjb23@psu.edu

Karsten Kahl  
Bergische Universität Wuppertal  
Department of Mathematics  
kkahl@math.uni-wuppertal.de

#### MS148

##### **Discretization-Accuracy Convergence for Full Algebraic Multigrid**

Full multigrid (FMG) is well known for converging to the level of discretization accuracy in a single cycle for a wide class of partial differential equations when the multigrid hierarchy is derived from problem geometry. When applying an FMG cycle to a hierarchy generated by algebraic multigrid (AMG), however, this scalable convergence to discretization accuracy is usually lost. This talk examines the cause of this loss and explores some improvements to standard AMG interpolation that can restore single-cycle convergence to discretization accuracy.

Wayne Mitchell  
University of Colorado at Boulder  
wayne.b.mitchell@gmail.com

Steve McCormick  
Department of Applied Math  
CU-Boulder  
stevem@colorado.edu

Robert D. Falgout  
Lawrence Livermore National Laboratory  
rfalgout@llnl.gov

Thomas Manteuffel  
University of Colorado

tmanteuf@colorado.edu

john.ruge@colorado.edu

**MS148****Improved Robustness with Root-Node Algebraic Multigrid**

Algebraic multigrid (AMG) is a successful solver for a range of problems, yet several types of applications still challenge AMG, driving up the computational cost of the method. Systems of PDEs, problems with strong anisotropy, and non-symmetric problems all require additional aspects when constructing an effective solver. In this talk, a method is presented for controlling cost, called root-node AMG, which can be interpreted as both classical and aggregation-style AMG. The basics of the root-node algorithm will be outlined along with theoretical motivation for the root-node approach. Numerical experiments will be shown that demonstrate robustness of the root-node solver for a range of problems, including systems of PDEs, non-symmetric problems, and SPD problems with strong anisotropy.

Luke Olson

Department of Computer Science  
University of Illinois at Urbana-Champaign  
lukeo@illinois.edu

Thomas Manteuffel  
University of Colorado  
tmanteuf@colorado.edu

Jacob B. Schroder  
Lawrence Livermore National Laboratory  
schroder2@llnl.gov

Ben Southworth  
University of Colorado, Boulder  
ben.southworth@colorado.edu

**MS148****Algebraic Multigrid for Non-Symmetric Linear Systems**

Algebraic multigrid (AMG) is a powerful solver for certain classes of large, sparse linear systems,  $Ax = b$ , often arising from the discretization of PDEs. However, convergence theory and most variations of AMG rely on  $A$  being symmetric positive definite. Here, new theoretical results for AMG applied to non-symmetric linear systems will be presented. These results motivate a new variation on classical AMG for non-symmetric problems and further motivate the known non-symmetric root-node AMG. Numerical results for each involving advection-diffusion and upwind discretizations of the transport equation will be presented, demonstrating their potential as robust solvers for a wide range of elliptic and hyperbolic PDEs

Ben Southworth

University of Colorado, Boulder  
ben.southworth@colorado.edu

Thomas Manteuffel  
University of Colorado  
tmanteuf@colorado.edu

John Ruge  
University of Colorado at Boulder

**MS149****A Finite Element Scheme for a Phase Field Model of Nematic Liquid Crystals**

We present a phase field model for nematic liquid crystals. Our model couples the Cahn-Hilliard equation to Ericksen's one constant model for nematic liquid crystals with variable degree of orientation. We present a special discretization of the liquid crystal energy that can handle the degenerate elliptic part without regularization. In order to develop the finite element scheme, we derive a discrete gradient flow by computing variational derivatives and setting the discrete time derivatives equal to minus the gradient. A convex splitting finite element scheme is used for the Cahn-Hilliard equation. Additional care is needed to account for the coupling of the two systems. We prove that our resulting finite element scheme is uniquely solvable and unconditionally stable. We also present numerical simulations to illustrate the method.

Amanda E. Diegel

Louisiana State University  
adiegel@lsu.edu

Shawn W. Walker  
Louisiana State University  
Department of Mathematics and CCT  
walker@math.lsu.edu

**MS149****Weak Free Surface Anchoring Model for Partially Wetting Nanoscale Nematic Liquid Films**

A weak free surface anchoring model is presented for a thin nematic liquid crystals (NLC) film. Applying the long wave approximation to the Leslie-Ericksen equations (an extension of the Navier-Stokes equations), a fourth order non-linear partial differential equation is derived for the free surface height. Unique to NLC are the dipole moments of the molecules, which induce an elastic response in the bulk of the film, and the anchoring (boundary) conditions at the free surface and substrate. In general, the anchoring conditions at the two interfaces are different. Satisfying the antagonistic conditions for very thin films incurs a large energy penalty in the bulk of the fluid. To alleviate this issue, we present a novel weak anchoring model, which dynamically relaxes the free surface anchoring to that of the substrate, as the thickness of NLC film decreases. Using linear stability analysis, we are able to draw a parallel between the stability properties of our model, and the so-called "forbidden film thicknesses" seen in experiments. In addition, large-scale simulations are carried out using a GPU, and the morphology of simulated dewetted films are compared to experimental results.

Michael A. Lam

New Jersey Institute of Technology  
Department of Mathematical Sciences  
mal37@njit.edu

Linda Cummings  
New Jersey Institute of Technology  
linda.cummings@njit.edu

Te-Sheng Lin  
National Chiao Tung University  
Department of Applied Mathematics

tslin@math.nctu.edu.tw

Lou Kondic  
Department of Mathematical Sciences, NJIT  
University Heights, Newark, NJ 07102  
kondic@njit.edu

### MS149

#### Biological Plywoods

This contribution presents recent results on the structure, self-assembly, and functionality of biological plywood materials, an ubiquitous material organization found throughout Nature, including plant cell walls, exocuticles of insects, bone, and cornea. The key objective is to demonstrate the principles used by Nature to develop multifunctional materials from simple rod-like filaments. A systematic multiscale modelling based on liquid crystal material physics is used to characterize the formation of these plywoods through chiral directed self-assembly. The role of orientational diffusivity in the temporal kinetics is established. In real system, second phases, curved confinements lead to defects which are shown to follow topological features determined by the plywoods architecture. We develop a geometric modeling methodology to characterize the order and orientation of the plywoods. The approach is validated with cornea-like tissues and with the exocuticle of a beetle. Finally we study the nano-wrinkling in surface layers of biological plywoods, which are responsible for optical functionalities and cell growth. In the former, we develop and apply a model to explain the diffraction patterns and color changes in tulips. Taken together, the results provide a quantitative understanding of the principles that govern the formation of biological plywoods for armor and strength as well as functionalities such as structural color.

Alejandro D. Rey  
Department of Chemical Engineering  
McGill University  
alejandro.rey@mcgill.ca

### MS149

#### Finite Element Approximation of Nematic Liquid Crystals with Variable Degree of Orientation

We consider the simplest one-constant model, put forward by J. Ericksen, for nematic liquid crystals with variable degree of orientation. The equilibrium state is described by a director field  $\mathbf{n}$  and its degree of orientation  $s$ , where the pair  $(s, \mathbf{n})$  minimizes a sum of Frank-like energies and a double well potential. In particular, the Euler-Lagrange equations for the minimizer contain a degenerate elliptic equation for  $\mathbf{n}$ , which allows for line and plane defects to have finite energy. We present a structure preserving discretization of the liquid crystal energy with piecewise linear finite elements that can handle the degenerate elliptic part without regularization, and show that it is consistent and stable. We prove  $\Gamma$ -convergence of discrete global minimizers to continuous ones as the mesh size goes to zero. We develop a quasi-gradient flow scheme for computing discrete equilibrium solutions and prove it has a strictly monotone energy decreasing property. We present simulations in two and three dimensions to illustrate the method's ability to handle non-trivial defects. Our results include electric and colloidal effects. This work is joint with R.H. Nochetto and S. Walker.

Wujun Zhang

Department of Mathematics  
University of Maryland  
wujunzhang@gmail.com

Ricardo Nochetto  
Department of Mathematics  
University of Maryland, College Park  
rhn@math.umd.edu

Shawn W. Walker  
Louisiana State University  
Department of Mathematics and CCT  
walker@math.lsu.edu

### MS150

#### Toward Distributed Eigenvalue and Singular Value Solver Using Data Flow System

The computation of the Eigenvalue problems and Singular Value Decomposition, has a long history, with many improvements over the years, both in implementations and algorithmically. To address cache-based memory hierarchies, it was reformulated to use Level 3 BLAS in the LAPACK library. ScaLAPACK was introduced to take advantage of distributed computing. Algorithmically, the divide and conquer algorithm reduced the number of operations. Still, methods remained memory bound, so two stage algorithms were developed recently to reduce memory operations and increase the computational intensity. We investigate the impact of these changes by testing various historical and current implementations on a common, modern multicore machine and a distributed computing platform.

Azzam Haidar  
Department of Electrical Engineering and Computer Science  
University of Tennessee, Knoxville  
haidar@icl.utk.edu

Aurelien Bouteiller  
University of Tennessee  
bouteill@eecs.utk.edu

George Bosilca  
University of Tennessee - Knoxville  
bosilca@icl.utk.edu

Stanimire Tomov  
Computer Science Department  
University of Tennessee  
tomov@icl.utk.edu

Jack J. Dongarra  
University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

### MS150

#### Communication Avoiding and Synchronous Reducing Techniques for Dense Parallel Eigenvalue Solver

Recently and in near-future, communication latency becomes a significant and severer bottleneck in parallel computing. A lot of communication optimization techniques such as communication reduction, avoidance and hiding could be applied for parallel numerical libraries as well as parallel simulation codes. In this study, we demonstrate several strategies or remarkable reduction of the amount

of communication latency focused on eigenvalue solver, EigenExa. We confirmed that the optimized routines for the Householder tridiagonalization show good performance improvement up to 20 to 50 percentage compared with the naive implementation when we use more than a thousand nodes on the K computer.

Toshiyuki Imamura  
RIKEN Advance Institute for Computational Science  
imamura.toshiyuki@riken.jp

Yusuke Hirota  
RIKEN AICS  
yusuke.hirota@riken.jp

Susumu Yamada, Masahiko Machida  
Japan Atomic Energy Agency  
yamada.susumu@jaea.go.jp,  
machida.masahiko@jaea.go.jp

### MS150

#### Computing a Partial SVD of General Matrices

It is well known that the SVD of a real matrix  $A$ ,  $A = USV^T$ , can be obtained from the eigenpairs of the matrix  $C_V = A^T A$  or  $C_U = AA^T$ . Alternatively, the SVD can be obtained from the eigenpairs of the augmented matrix  $C_{UV} = [0 \ A; A^T \ 0]$ . This formulation is particularly attractive if  $A$  is a bidiagonal matrix and only a partial SVD (a subset of singular values and vectors) is desired. This presentation focuses on  $C_{UV}$  formed with bidiagonal matrices, discusses how tridiagonal eigensolvers can be applied in that context, and significant savings that can be thus propagated to the SVD of general matrices. The presentation will also discuss opportunities for parallelism, accuracy and implementation issues.

Osni A. Marques  
Lawrence Berkeley National Laboratory  
Berkeley, CA  
oamarques@lbl.gov

### MS150

#### Towards Materials Design with Extreme-Scale Quantum Simulations

With ever improving accuracy of ab initio electronic structure methods, quantum simulations have now become a predictive tool that can be used to search for new materials with desired properties. For simulations tools to be relevant for such searchers, calculations for individual compounds have to be very reliable and optimized to minimize the machine and energy footprint, as they have to be repeated automatically tens or hundreds thousands of times. Today, this can be accomplished with clusters that have hybrid CPU-GPU nodes. We will discuss the algorithmic developments that were necessary to run modern electronic structure codes on such systems.

Raffaele Solcà  
ETH  
Eidgenössische Technische Hochschule Zürich  
rasolca@student.ethz.ch

Anton Kozhevnikov  
ETH Zürich  
Institut f. Theoretische Physik  
antonk@ethz.ch

Thomas C. Schulthess, Thomas C. Schulthess  
ETH Zurich  
Swiss National Supercomputing Center (CSCS)  
schulthess@cscs.ch, schulthess@cscs.ch

### MS151

#### Stochastic Collocation Methods Via Compressive Sensing and Its Applications in Uncertainty Quantification

Stochastic computation has received intensive attention in recent years, due to the pressing need to conduct uncertainty quantification (UQ) in practical computing. One of the most widely used techniques in UQ is generalized polynomial chaos. In this talk, we will discuss collocation method via compressive sampling for recovering sparse polynomial chaos expansions (PCE) using randomized quadratures. The framework includes both the bounded measures such as the uniform and the Chebyshev measure, and the unbounded measures which include the Gaussian measure. We provide theoretical analysis on the validity of the approach. Several numerical examples are given to confirm the theoretical results.

Ling Guo  
Department of Mathematics, Shanghai Normal  
University, China  
lguo@shnu.edu.cn

Akil Narayan, Yuhang Chen  
University of Utah  
akil@sci.utah.edu, chen.7169@osu.edu

Tao Zhou  
Institute of Computational Mathematics, AMSS  
Chinese Academy of Sciences  
tzhou@lsec.cc.ac.cn

Dongbin Xiu  
Ohio State University  
xiu.16@osu.edu

### MS151

#### Multivariate Quadrature Rules for Correlated Random Variables

Numerical quadrature is ubiquitous in computer science and engineering. A number of techniques, such as sparse grids and cubature rules, are available for computing integrals of tensor-product measures, however integrating functions parameterized by correlated variables is much more difficult. In this talk I will present a gradient-based optimization approach for computing multivariate quadrature rules for correlated variables and arbitrary shapes. The efficacy of this method will be demonstrated with various numerical examples motivated by a number of challenges faced when quantifying uncertainty.

John D. Jakeman  
Sandia National Labs  
jdjakem@sandia.gov

Akil Narayan  
University of Utah

akil@sci.utah.edu

### MS151

#### Effective Quadratures: Least Squares Polynomials for Parametric Studies

We present a new sampling strategy for constructing polynomial chaos approximations for expensive physics simulation models. The proposed approach, *effectively subsampled quadratures* involves sparse subsampling an existing tensor grid using QR column pivoting such that well-conditioned least squares estimates can be computed. We also demonstrate how to construct gradient-enhanced polynomial approximations using the same subsamples.

Pranay Seshadri  
Cambridge University  
Department of Engineering  
ps583@cam.ac.uk

Akil Narayan  
University of Utah  
akil@sci.utah.edu

### MS151

#### Gradient-Enhanced Stochastic Collocation Methods for Uncertainty Quantification

The talk is concerned with stochastic collocation methods with the gradient information, namely, we consider the case where both the function values and derivative informations are available. Particular attention will be given to discrete least-squares and the compressive sampling methods. Stability results of these approaches will be presented.

Ling Guo  
Department of Mathematics, Shanghai Normal  
University, China  
lguo@shnu.edu.cn

Akil Narayan  
University of Utah  
akil@sci.utah.edu

Dongbin Xiu  
Ohio State University  
xiu.16@osu.edu

Tao Zhou  
Institute of Computational Mathematics, AMSS  
Chinese Academy of Sciences  
tzhou@lsec.cc.ac.cn

### MS152

#### Homogenization of the Transmission Eigenvalue Problem for a Periodic Media

In this talk we consider the inverse acoustic/electromagnetic scattering problem of determining information about the macro and micro-structure of a periodic media where the period is characterized by a small parameter. To this end, we study the transmission eigenvalue problem as the small parameter tends to zero to obtain the homogenized eigenvalue problem. This is a non-linear and non-selfadjoint eigenvalue problem which makes its investigation mathematically challenging. We prove (weak) convergence of the eigenvalues/functions as well as show that the effective material properties can be

determined by the measured homogenized transmission eigenvalues.

Isaac Harris  
Texas A&M University  
Department of Mathematics  
iharris@math.tamu.edu

### MS152

#### An Introduction to Interior Transmission Eigenvalues and Non-Destructive Testing

The efficient numerical calculation of interior transmission eigenvalues is a challenging task due to the fact that the corresponding interior transmission problem is neither elliptic nor self-adjoint. It is explained what interior transmission eigenvalues are, how they can be computed to high accuracy with boundary integral equations, and how they can be used to visualize the interior of a given three-dimensional object which is the aim in non-destructive testing in order to uncover location, size, and geometry of inhomogeneities. Additionally, some open problems both from the theoretical and practical point of view are illustrated.

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

### MS152

#### Spectral Results of the Interior Transmission Eigenvalues for Maxwell's Equations

We consider the transmission eigenvalue problem for Maxwell's equations corresponding to non-magnetic inhomogeneities with contrast in electric permittivity that has fixed sign (only) in a neighborhood of the boundary. Following the analysis made by Robbiano in the scalar case we study this problem in the framework of semiclassical analysis and relate the transmission eigenvalues to the spectrum of a Hilbert-Schmidt operator. Under the additional assumption that the contrast is constant in a neighborhood of the boundary, we prove that the set of transmission eigenvalues is discrete, infinite and without finite accumulation points. A notion of generalized eigenfunctions is introduced and a denseness result is obtained in an appropriate solution space.

Shixu Meng  
University of Delaware  
sxmeng@udel.edu

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

### MS152

#### A Spectrum Projection Method for Transmission Eigenvalues

In this talk, we consider a nonlinear integral eigenvalue problem, which is a reformulation of the transmission eigenvalue problem arising in the inverse scattering theory. The boundary element method is employed for discretization, which leads to a generalized matrix eigenvalue problem. We propose a novel method based on the spectral projection. The method probes a given region on the complex plane using contour integrals and decides if the region con-

tains eigenvalue(s) or not. It is particularly suitable to test if zero is an eigenvalue of the generalized eigenvalue problem, which in turn implies that the associated wavenumber is a transmission eigenvalue. Effectiveness and efficiency of the new method are demonstrated by numerical examples.

Jiguang Sun  
Department of Mathematical Sciences  
Michigan Technological University  
jiguangs@mtu.edu

### MS153

#### Title Not Available at Time of Publication

Abstract not available

Christophe Berthon  
University of Nantes  
christophe.berthon@math.univ-nantes

### MS153

#### Numerical Methods for the Chemotaxis Models

Abstract not available

Yekaterina Epshteyn  
University Of Utah  
epshteyn@math.utah.edu

### MS153

#### Convergence Estimates for the Approximation of Scalar Conservation Equations

Abstract not available

Bojan Popov  
Department of Mathematics  
Texas A&M University  
popov@math.tamu.edu

### MS154

#### Computational Efficiency of a Parallel High-Order Spectral Method for 3D Water Wave Simulations

CFD- or finite-element-type simulations of the three-dimensional non-linear evolution of ocean and coastal waves are still widely impractical on today's high performance computers because of the complexity of the underlying equations. A common simplifying assumption is that the fluid is inviscid and incompressible, and the flow is irrotational which leads us to have to solve Laplace equation with boundary conditions at the free surface. For reasons of computational efficiency, we propose to solve the latter with a high-order spectral model (HOSM) to accurately simulate relevant processes within an acceptable amount of computing time. In fact, we are going to present a MPI parallel HOSM that is capable of simulating non-linear wave interaction and effects that originate from a non-trivial bathymetry profile. In its current optimized version, the computation of the bottom influence and the global communications needed for the fast Fourier transforms (FFTs) are two major limiting factors. We are going to discuss recent computational improvements that we have undertaken in order to reduce runtime and increase efficiency and parallel scalability on a supercomputer. This includes rearranging indices and loops, exploring the viability of using single as opposed to double precision as well

as reducing the number of FFTs involved in the computations.

Nicole Beisiegel  
University College Dublin  
School of Mathematical Science & Statistics  
nicole.beisiegel@ucd.ie

Frédéric Dias  
University College Dublin  
Ecole Normale Supérieure de Cachan  
frederic.dias@ucd.ie

### MS154

#### Requirements and Progress Towards Answering the Big Questions in Storm Surge Modelling

Coastal hazards are a growing problem worldwide due to not only the current and projected sea-level rise but also due to increasing populations and economic dependence on coastal areas. Today, coastal hazards related to strong storms are one of the most frequently recurring and wide spread hazards to coastal communities today. In particular storm surge, the rise of the sea surface in response to wind and pressure forcing from these storms, can have a devastating effect on the coastline. Furthermore, with the addition of climate change related effects, the ability to predict these events quickly and accurately is critical to the protection and sustainability of these coastal areas. Computational approaches to this problem must be able to handle its multi-scale nature while remaining computationally tractable and physically relevant. In this talk I will outline some of the approaches we are developing to address the issues faced by computational models of storm surge including the use of adaptive mesh refinement, and embedded sub-scale physics. Combining these new approaches promises to address some of the problems in current state-of-the-art models while continuing to decrease the computational overhead needed to calculate a forecast or climate scenario.

Kyle T. Mandli, Colton Conroy, Jiao Li  
Columbia University  
Applied Physics and Applied Mathematics  
kyle.mandli@columbia.edu, cjc2235@columbia.edu,  
jl4170@columbia.edu

### MS154

#### Scaling at Exascale in Blended Isogeometric, Discontinuous Galerkin, and Particle-in-Cell Approaches

Recent HPC studies performed on various architectures will be presented, including on second generation Xeon Phi Knights Landing (KNL) processors. This will include studies on a general coprocessor based code called ArcSyn3sis, that runs blended isogeometric discontinuous Galerkin methods using the Open Concurrent Compute Abstraction (OCCA) library. Next work being performed on a discontinuous Galerkin shallow water solver, DGSWEM, using an asynchronous MPI alternative called High Performance ParalleX (HPX) will be presented. And finally, we will discuss results from a pair of particle-in-cell gyrokinetic solvers XGC1 and XGCa, that utilize the Adaptable IO System (ADIOS).

Craig Michoski  
UT Austin  
michoski@ices.utexas.edu



Robert D. Moser  
University of Texas at Austin  
rmoser@ices.utexas.edu

Clint Dawson  
Institute for Computational Engineering and Sciences  
University of Texas at Austin  
clint@ices.utexas.edu

Varis Carey  
University of Colorado, Denver  
varis.carey@ucdenver.edu

Chris Simmons  
ICES, University of Texas  
csim@ices.utexas.edu

#### MS154

##### Multiscale Flood Simulations in Urban Regions

Risk assessment and prediction of floods in complex environmental settings, like urban environments, is a challenging task, especially when done in real time. While often (simplified) hydrological models coupled with meteorological models for precipitation input are sufficient, they are no longer valid when three-dimensional effects take place in the flow and considerably impact the flow dynamics. In such situations, a full three-dimensional solution is indispensable to obtain a sufficiently correct prediction of flow variables. Within our researches, floods and impact of floods on infrastructure should be predicted, reaching from the river down to the scale of the built infrastructure, such as railway, subway, tunnels, waste water channels, buildings, and building infrastructure. A real-time simulation based on an efficient and transparent 2D–3D coupling for an integrated prediction of local flows together with large scale flow dynamics should therefore support experts for a better protection against extreme flood events, allowing him or her to interactively explore and evaluate various scenarios and, thus, enabling decision makers to decide on a new level of information detail and accuracy.

Christoph M. Ertl  
Chair of Computation in Engineering, Prof. Rank  
Technische Universität München  
christoph.ertl@tum.de

Christoph M. Ertl  
Technische Universität München  
Chair for Computation in Engineering  
christoph.ertl@tum.de

Florian Mintgen  
TU München  
f.mintgen@tum.de

Nevena Perovic  
Chair for computation in engineering  
TU München  
nevena.perovic@tum.de

Hao Zeng  
TU München  
hao.zeng@tum.de

Michael Manhart  
TU Muenchen  
Fachgebiet und Laboratorium für Hydromechanik

m.manhart@bv.tum.de

Ralf-Peter Mundani  
TUM, Faculty of Civil, Geo and Environmental  
Engineering  
Chair for Computation in Engineering  
mundani@tum.de

#### MS155

##### Fast Charge Equilibration Solvers for Polarizable Force Fields

Polarizable and reactive force fields which incorporate dynamic charges into their molecular models have been shown to address the shortcomings of classical force fields that define static charges and static bonding topologies. A crucial component of polarizable and reactive force fields is a charge distribution model. Accurate solution of the charge distribution problem, as formulated by several different models proposed to date, manifests itself as large-scale sparse linear systems. Krylov subspace methods employed to solve these systems represent a significant limitation in terms of solution time and scalability in large simulations. In this talk, we will discuss our recent efforts to accelerate the convergence rate of these iterative techniques via scalable thread parallel Incomplete LU based preconditioning methods. We will present extensive numerical tests regarding the trade-offs involved in choosing a good preconditioner and its efficient implementation on shared memory computers. While our focus has been on the Charge Equilibration (QEq) model, preliminary results suggest that our techniques can be highly useful for more complex models such as SQE and ACKS2.

H. Metin Aktulga  
Lawrence Berkeley National Lab  
hma@cse.msu.edu

Kurt O'Hearn  
Michigan State University  
ohearnku@msu.edu

#### MS155

##### Greatly Scalable Multiple-Copy Enhanced Sampling Algorithm on Massively Parallel Computing

One broad strategy to improve sampling relies on multiple copies to enhance sampling efficiency of MD. Multiple copy algorithms adopt a divide-and-conquer strategy to execute the desired computation with massive parallel resources. We report a robust implementation of generalized multiple copy algorithms for molecular dynamics simulations in the greatly scalable program NAMD that is on the parallel programming system charm++. Multiple concurrent NAMD instances are launched with internal partitions of charm++ and located continuously within a single communication world. Communication between NAMD instances are executed by low-level point-to-point communication functions and further implemented in NAMDs Tcl scripting interface. The communication enabled Tcl scripting provides a sustainable application interface for end user to realize generalized multiple copy enhanced sampling without modifying source code. Novel applications of enhanced sampling algorithms with fine-grained inter-copy communication structure, including replica exchange solute tempering, global lambda exchange along an absolute alchemical reaction path, Hamiltonian exchange umbrella sampling of complex collective variables in multidimensional order parameter space and string method with swarms-of-

trajectories, were carried out on Mira, a 10 flops IBM Blue Gene/Q supercomputer, to demonstrate the versatility and massive scalability of the present implementation.

Wei Jiang  
Leadership Computing Facility, Argonne National  
Laboratory  
Lemont, IL 60439  
wjiang@alcf.anl.gov

### MS155

#### Scalable, Reproducible and Fast Quantum and Reactive Molecular Dynamics Simulations

We have developed a divide-conquer-recombine algorithmic framework for large quantum molecular dynamics (QMD) and reactive molecular dynamics (RMD) simulations. The framework is globally scalable and reproducible, as well as locally fast, on current and future computing platforms, with minimal architectural assumptions. The framework has achieved parallel efficiency over 0.98 on 786,432 IBM Blue Gene/Q processors for 39.8 trillion electronic degrees-of-freedom QMD in the framework of density functional theory and 67.6 billion-atom RMD. We will discuss several applications including (1) 16,616-atom QMD simulation of rapid hydrogen production from water using metallic alloy nanoparticles, (2) 6,400-atom nonadiabatic QMD simulation of exciton dynamics for efficient solar cells, and (3) 112 million-atom RMD simulation of metacarbon synthesis by high temperature oxidation of SiC nanoparticles. We will also discuss the merger of exaflop/s high performance computing and exabyte big data analytics for the discovery of new materials at our MAterials Genome Innovation for Computational Software (MAGICS) center.

Aiichiro Nakano, Rajiv K. Kalia  
University of Southern California  
anakano@usc.edu, rkalia@usc.edu

Fuyuki Shimojo  
Kumamoto University  
shimojo@kumamoto-u.ac.jp

Priya Vashishta  
University of Southern California  
priyav@usc.edu

### MS155

#### Performance Issues for Modeling Materials via Molecular Dynamics on Current and Future Hardware

Optimizing our LAMMPS molecular dynamics (MD) code for current and future accelerator hardware (many-core, GPU, KNL, etc) has been challenging because of the wide variety of materials models (force fields) it supports for both pairwise and manybody interactions. I'll discuss some of our successes and challenges, as well as lessons learned. This includes creation of hardware-specific code (for CPU, GPU, Phi), as well as efforts to use Kokkos, a Sandia-developed library which attempts to insulate applications like MD codes from hardware details. I'll also describe load-balancing strategies that work well for MD, especially for coarse-grained material models, and on-the-fly visualization ideas which we've found particularly useful for large-scale particle simulations, including MD.

Steve Plimpton  
Sandia National Laboratories

sjplimp@sandia.gov

Mike Brown  
Intel  
michael.w.brown@intel.com

Stan Moore, Christian Trott  
Sandia National Laboratories  
stamoor@sandia.gov, crtrott@sandia.gov

### MS156

#### Monitoring HPC Systems for Improving Efficiency

Performance variations are becoming more prominent with new generations of HPC machines. Understanding these variations and resilience to anomalous application behavior are critical challenges for extreme-scale computing. To address these challenges, data collection systems gather detailed information about system resources. This data is applicable to a wide range of tasks including identification of applications, anomalous/malicious behavior, and performance variation root causes. This talk introduces our recent work to address two key problems: 1) automatic detection of anomalous HPC application behavior and 2) automatic identification of observed performance anomaly root causes. To address the first problem, we build semi-supervised classification models that represent healthy application behavior and enable identifying deviation from the defined norm. To tackle the harder second problem, we train classification models using well-documented anomalous application behavior and match suspicious application behavior to the previously labeled anomalies. To efficiently identify anomalies, we perform statistical analysis to create definitive features. The high-dimension, noisy, and high volume of the data make analysis a challenge. Thus we apply clustering and feature extraction to reduce the dimensionality of the feature space covered by the employed machine learning algorithms. This talk presents results based on data collected on our local clusters and DOE machines.

Ayse Coskun  
Boston University  
Electrical and Computer Engineering Department  
acoskun@bu.edu

### MS156

#### Analysis of User Job Level Performance Data Using XDMoD with Application Kernels

HPC systems are a complex combination of hardware (servers, networks, storage) and software, and it is desirable to have a way to continuously ensure that the infrastructure is running with optimal efficiency as well as the ability to proactively identify underperforming hardware, system software, and user jobs. XDMoD, through integration with monitoring frameworks such as TACC-Stats, Performance Co-Pilot, and Ganglia, allow end-users and system support personnel to obtain detailed job level information for all jobs running on the HPC resource. In addition to job level performance data, XDMoD, through the implementation of benchmark application kernels, provides diagnostic information on system wide performance. While application kernel data are used by XDMoD as a quality assurance tool, they are inherently an ideal probe of the status of the system and provide primary information on all system contentions issues that a typical user would encounter. An analysis of application kernel run data will be presented

along with the implications of this analysis.

Thomas Furlani

State University of New York at Buffalo  
furlani@buffalo.edu

#### MS156

##### Network Congestion: Challenges and Approaches to Understanding Application Impact

In large-scale HPC systems, contention for shared resources is a major cause of performance variability. Analysis of continuous system-wide state information is necessary for understanding the conflicting demands for shared resources and for assessing the resultant impact on application performance. Such analysis requires obtaining right-fidelity system data and accounting for adaptive mechanisms in system software and dynamic system workloads. In this work, we present our methods for data collection, congestion analysis, and application impact assessment as they have been applied to large-scale platforms, both during production and in controlled experiments, across multiple-generations of networks. We present challenges and opportunities in enabling continuous, run-time operational analysis of network state and potential application impact and in using such data to improve system operations.

Ann Gentile, James Brandt, Anthony Agelastos,  
Benjamin Allan

Sandia National Laboratories  
gentile@sandia.gov, brandt@sandia.gov, am-  
agela@sandia.gov, baallan@ca.sandia.gov

#### MS156

##### Detecting Application Interference Through System Log Analysis with Spark

Performance variability on HPC platforms is a critical issue with serious implications on the users: irregular runtimes prevent users from correctly assessing performance and from efficiently planning allocated machine time. The hundreds of applications concurrently sharing thousands of resources escalate the complexity of identifying the causes of runtime variations. On production systems, implementing trial-and-error approaches is practically impossible. On that account, making use of existing information represents a preferable path to solution. Cray systems collect large amounts of data related to user applications, data that can be highly valuable for understanding performance variability. Novel analytics tools enable the exploring of ways to use the data for identifying and understanding performance variability. In this context, we have developed a Spark-based tool for analyzing system logs with the goal of identifying applications that show high variability (victims) and applications potentially causing the variability (aggressors). Understanding the nature of both types of applications is crucial to developing a solution to these issues. This talk goes through the different steps of the analysis (data aggregation and filtering, victim and aggressor detection, validation) and the configuration parameters that help refine the search space. The analysis was implemented using Sparks RDD and DataFrame API and was tested on data coming from production systems.

Diana Moise

Cray, Inc.

dmoise@cray.com

#### MS157

##### Quantitative Assessment of Transportation Network Vulnerability with Dynamic Traffic Simulation Methods

Safe, secure, and reliable operation of our nation's transportation networks is critical for sustaining modern society. This includes mitigating congestion and assessing impacts of natural, accidental, and intentional disruptions to such networks. Past research on transportation network vulnerability has primarily focused on static traffic models, many of which are formulated as traditional optimization problems. However, transportation networks are dynamic because their usage varies over time. As a result, a realistic characterization of network vulnerability must account for these dynamic properties. This talk describes a dynamic traffic simulation-based approach to assess the vulnerability of transportation networks to disruptions. The approach includes prioritization of critical links over time and is generalizable to the case where both link and node disruptions are of concern. Motivating case study examples are presented and insights into the time varying criticality of links are discussed.

Venkateswaran Shekar  
University of Massachusetts Dartmouth  
vshekar@umassd.edu

Samrat Chatterjee  
National Security Directorate  
Pacific Northwest National Laboratory  
samrat.chatterjee@pnnl.gov

Mahantesh Halappanavar  
Pacific Northwest National Laboratory  
mahantesh.halappanavar@pnnl.gov

Lance Fiondella  
University of Massachusetts Dartmouth  
lfiondella@umassd.edu

#### MS157

##### Dynamic Brain Networks

Current functional neuroimaging (fMRI) studies focus primarily on correlational network analyses, though some still apply classical frequency analysis methods. Few, however, consider global spatiotemporal patterns of cortical activity beyond identifying groups of voxels whose signals fluctuate together. In this talk, I will discuss a recently developed modal decomposition algorithm called Dynamic Mode Decomposition (DMD) and its applicability to study global brain activity patterns of fMRI data acquired by the Human Connectome Project (HCP). DMD was developed by Schmid [J. Fluid Mech., 2010] as an equation-free method to model the behavior of complex dynamical systems to identify spatially coherent modes in a flow field. Importantly, each mode is associated with an eigenvalue defining a unique frequency of oscillation and growth/decay rate. So far, DMD has been applied to the analysis of brain sleep spindles using EEG in a small data set by Brunton et al. [arXiv, 2014], but has not yet been applied in a large-scale population study or to high-resolution fMRI data such as that acquired by the HCP. I will discuss how we can use DMDs dynamic modes to build new graphical representations of the cortex, and also how we can use them to shed

light on the relationship that various factors like age and neurological disorders (such as Alzheimers and Schizophrenia) have with brain function.

Kristian Eschenberg  
University of Washington  
keschenb@uw.edu

Tom Grabowski  
Integrated Brain Imaging Center  
University of Washington  
tgrabow@uw.edu

David Haynor  
Dept. Radiology  
University of Washington  
haynor@uw.edu

### MS157

#### Advances in Algorithms and Applications for Dynamic Graphs

Over the last decade, dynamic graph formulations have emerged in multiple application sectors. In social networking, the webs for connectivity and social discourse are maintained as complex networks that continuously evolve at rates that make them intractable to store, let alone analyze. In life sciences, high-throughput technologies for genotyping, phenotyping and imaging are being deployed at massive scales, in order to capture dynamic data at various scales. Similar trends also persist in other domains—e.g., cybersecurity, urban networks. While the machinery to collect dynamic data at scale have made significant strides, software capabilities and in particular, algorithmic modeling and network characterizations are mostly in their nascent stages. In this introductory talk, we will provide an overview of some of the notable algorithmic advances for dynamic networks analysis, as motivated by several application use-cases in the life sciences, cybersecurity and brain imaging. Specifically, we will focus on algorithms for community detection, graph matching, and network comparison, and their related challenges. Of emphasis will be to explore the connections (or gaps thereof) between algorithmic modeling and real world application. The talk will be followed by other talks that serve to highlight a subset of the specific application use-cases in more depth. The second part of the mini-symposium will focus on algorithms and network models for dynamic graph analyses.

Ananth Kalyanaraman  
Boeing Centennial Chair, Associate Professor  
School of EECS, Washington State University  
ananth@eeecs.wsu.edu

Mahantesh Halappanavar  
Pacific Northwest National Laboratory  
hala@pnnl.gov

### MS157

#### Dynamic Networks of Microbial Biofilms

Cell-based therapeutics is a key component of precision medicine, i.e., the new paradigm for disease prevention and treatment aimed at providing customized healthcare solutions on a patient-by-patient basis. While there has been significant progress towards understanding the cellular behavior and controlling individual cells, our understanding still lacks a computational framework able to capture the population-level cell interactions and emerging behaviors

that are critical to fighting diseases like antibiotic-resistant infections or cancer. In this talk, we bring a new perspective on molecular communication and nano-networking at population-level (as opposed to single cell-level) which is critical to engineer cells behavior, reprogram the cell-cell communication, and develop new strategies to control the dynamics of population of cells. In particular, we discuss the significance of the network-based approach to explore the subtleties of bacteria inter-cellular network and its implications to biofilm dynamics. Indeed, as computational models become more and more powerful, a network-centric approach to studying cell populations can improve our understanding their social behaviors and possibly help controlling the infectious diseases they cause. This is a major step towards developing new drugs and targeted medical treatments to fight biofilm-related infections.

Radu Marculescu  
Dept. Electrical and Computer Engineering  
Carnegie Mellon University  
radum@cmu.edu

Chieh Lo  
Carnegie Mellon University  
chiehl@andrew.cmu.edu

### MS158

#### A Model Reification Approach to Fusing Information from Multifidelity Information Sources

Engineers often have a variety of models available to aid in any given task. Typically, one of these models is used based on its cost and fidelity. We present a method that utilizes fused information from multiple models that results in superior predictive performance than any of the constituent models. Our methodology first estimates the correlation between each model using a reification procedure and then fuses the available information in a Bayesian sense.

Doug Allaire  
Texas A&M  
dallaire@tamu.edu

William Thomison  
Texas A&M University  
dillon.thomison@gmail.com

### MS158

#### A Bayesian Approach to Optimal Coupling of Multidisciplinary Models

Design of complex engineering systems requires coupled analyses of multiple disciplines. In many cases, disciplines may be weakly coupled, so that some of the interaction terms can be neglected without significantly impacting the accuracy of the system output while reducing the computational cost of using the model for design optimization and/or uncertainty quantification. As a result, we propose a new Bayesian approach based on Sequential Monte Carlo that balances accuracy and complexity to find optimal system coupling.

Ricardo Baptista  
MIT  
rsb@mit.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

Benjamin Peherstorfer  
ACDL, Department of Aeronautics & Astronautics  
Massachusetts Institute of Technology  
pehersto@mit.edu

Karen E. Willcox  
Massachusetts Institute of Technology  
kwillcox@MIT.EDU

### MS158

#### Multi-Information Source Optimization with Unknown Constraints

In the *constrained multi-information source optimization* problem we study an optimization task where our goal is to optimize a design specified by multiple parameters under unknown constraints. Both the objective and the constraints are expensive-to-evaluate black-box functions. However, we also have access to a variety of information sources that offer cheaper approximations, e.g., numerical simulations that employ models of the true functions of varying complexity. These information sources are subject to *model discrepancy*, i.e. their internal model inherently deviates from reality. Note that our notion of model discrepancy goes considerably beyond typical observational noise that is common in multi-fidelity optimization: in our scenario information sources can be biased and are not required to form a hierarchy. We present a novel algorithm that is based on a rigorous mathematical treatment of the uncertainties arising from the model discrepancies. Its optimization decisions rely on a stringent Bayesian value of information analysis that trades off the predicted benefit and its cost. Moreover, we present an experimental evaluation demonstrating that our method consistently outperforms the state-of-the-art techniques: it finds designs of considerably higher objective value and additionally inflicts less cost in the exploration process.

Matthias Poloczek  
Cornell University  
poloczek@cornell.edu

Peter I. Frazier  
Dept. of Operations Research and Information  
Engineering  
Cornell University  
pf98@cornell.edu

### MS158

#### Reducing the Error of Monte Carlo Algorithms by Learning Control Variates

Control variates are a powerful technique to reduce the error of Monte Carlo, but conventionally it is necessary to know a good function approximation a priori. Stacked Monte Carlo (StackMC) is a post-processing technique that overcomes this limitation by constructing a control variate from data samples. This talk shows extensions to StackMC for use with importance sampling, Latin-hypercube sampling and quasi-Monte Carlo sampling, as well as use with multiple fitting functions and discrete input spaces.

David Wolpert  
Santa Fe Institute  
david.h.wolpert@gmail.com

Brendan Tracey  
Stanford University

tracey.brendan@gmail.com

### MS159

#### A Multiscale, Coulomb Collision Machine-Learning Approach for Particle-in-Cell Plasma Algorithms

Coulomb collisions in plasmas are governed by the Landau-Fokker-Planck equation. Collision modules in PIC simulations are typically Monte-Carlo-based. Monte Carlo is attractive for its simplicity, efficiency in multiple dimensions, and conservation properties. However, it is noisy, of low temporal order (typically  $\mathcal{O}(\sqrt{\Delta t})$ ), and has to resolve the collision frequency for accuracy [Dimitis, et al., *JCP*, **228** (2009)]. In this study, we explore a non-stochastic, multiscale alternative to Monte Carlo for PIC. The approach is based on a Green-function-based reformulation [Hu, Krommes, *PoP*, **1** (1994)] of the Vlasov-Fokker-Planck equation, which can be readily incorporated in modern multiscale collisionless implicit PIC algorithms [Chen, Chacón, and Barnes, *JCP*, **230** p.7018 (2011); Chen, Chacón, *CPC*, **185** (2014); Chen, Chacón, *CPC*, **197** (2015)]. An asymptotic-preserving operator-splitting approach allows the collisional step to be treated independently from the particles while preserving the multiscale character of the method. A significant element of novelty in our algorithm is the use of machine learning strategies to avoid a velocity space mesh for the collision step [Yoon and Chang, *PoP*, **21** (2014)]. The resulting algorithm is non-stochastic and first-order-accurate in time. We will demonstrate the method with several relaxation examples.

Luis Chacon, Guangye Chen  
Los Alamos National Laboratory  
chacon@lanl.gov, gchen@lanl.gov

### MS159

#### Spectral Methods for Multiscale Plasma Physics Simulations

The Vlasov-Maxwell (VM) equations are extremely challenging numerically because of their high dimensionality, nonlinearities and the wide disparity of spatial and temporal scales. In this work, we present a spectral method for the VM equations based on a decomposition of the plasma phase-space density in Hermite or Legendre modes. It leads to a truncated system for the expansion coefficients (i.e. moments). Its most important feature is that, with a suitable spectral basis, the low-order moments are akin to the typical moments (mass, momentum, energy) of a fluid/macroscopic description of the plasma, while the kinetic/microscopic physics can be retained by adding more moments. In addition, spectral convergence, stability and exact conservation laws in the limit of finite time step can be proven. Selected results illustrating the properties and the potential of the method will be presented. A comparison between PIC and the spectral method on standard electrostatic test problems shows that the spectral method can be orders of magnitude faster/more accurate than PIC. Some attempts to optimize the spectral decomposition in velocity space and multi-dimensional fully electromagnetic tests with efficient preconditioning techniques will also be presented. With the built-in fluid/kinetic coupling and favorable numerical properties, spectral methods might offer an optimal way to perform accurate large-scale simulations including microscopic physics.

Gian Luca Delzanno, Gianmarco Manzini  
Los Alamos National Laboratory  
delzanno@lanl.gov, gmanzini@lanl.gov

Juris Vencels  
KTH  
vencels@kth.se

Stefano Markidis  
KTH, Sweden  
markidis@kth.se

Vadim Roytershteyn  
Space Science Institute  
vroyersh@gmail.com

### MS159

#### Fluid-Kinetic Coupling with the Hermite-Fourier Spectral Method

The Hermite-Fourier spectral method for solving the Vlasov-Maxwell equations combines the fluid and kinetic approaches into one framework. In fact, the method links near-Maxwellian plasmas that can be represented with very few moments (fluid description) with plasmas where the kinetic physics drives strong non-Maxwellian behavior and therefore many moments are needed for an accurate representation of the distribution function. In this work, we show first how the Hermite-Fourier spectral method combines the fluid and kinetic plasma models and second how to transition continuously from one description to the another one by increasing/decreasing the number of Hermite functions dynamically during the simulation. Finally, the applications of such adaptive Hermite-Fourier spectral method are discussed.

Stefano Markidis  
KTH, Sweden  
markidis@kth.se

Gian Luca Delzanno, Gianmarco Manzini  
Los Alamos National Laboratory  
delzanno@lanl.gov, gmanzini@lanl.gov

Ivy Bo Peng, Erwin Laure  
KTH  
Sweden  
bopeng@kth.se, erwinl@kth.se

### MS159

#### Hybrid Method and Its Applications to Selected Problems in Space Physics

In many problems in space, astrophysical, and laboratory plasma the dominant kinetic effects are associated with ions. For such problems, the so-called hybrid class of methods provides an efficient description. Hybrid approximation retains full kinetic treatment of the ions using a Vlasov equation, but approximates electrons using a fluid model, for example as a massless fluid with a prescribed equation of state. As a result, the spatial and temporal scales associated with kinetic electron physics are removed from the problem. In this talk, I will review the method and will discuss its advantages and disadvantages using several practical examples drawn from space physics. Specifically, I will describe applications of a hybrid particle-in-cell method to global simulations of the Earth's magnetosphere and plasma turbulence.

Vadim Roytershteyn  
Space Science Institute

vroytersh@gmail.com

### MS160

#### Mathematical Models for the Correspondence Problem

We introduce the fascinating correspondence problem also known as image registration. Roughly spoken, the goal is to automatically establish correspondences between points in different projections of a scene. In particular in medical imaging, this problem is very important and used for applications such as motion correction or data fusion. Several examples displaying different facets of the problem are discussed. A mathematical framework for the correspondence problem is outlined. Starting point is a variational formulation, where a joint energy is to be minimized on an appropriate set. Modular building blocks such as distance measures and regularizers are briefly discussed and related to particular applications. Finally, a brief outlook on constrained image registration is presented. Constraints are used to improve the modelling by restricting the admissible set in a smart way.

Jan Modersitzki  
University of Lübeck  
Institute of Mathematics and Image Computing  
jan.modersitzki@mic.uni-luebeck.de

### MS160

#### A Lagrangian Solver for Diffeomorphic Image Registration

We present an efficient solver for diffeomorphic registration problems in the framework of Large Deformations Diffeomorphic Metric Mappings (LDDMM). We use an optimal control formulation, in which the (stationary or in-stationary) velocity field of a hyperbolic PDE needs to be found such that the distance between the final state of the system (the transformed/transported template image) and the observation (the reference image) is minimized. As transformation models, we consider both the transport equation (assuming intensities are preserved during the deformation) and the continuity equation (assuming mass-preservation). We consider the reduced form of the optimal control problem and solve the resulting unconstrained optimization problem using a discretize-then-optimize approach. A key idea of our method is the elimination of the PDE constraint using a Lagrangian hyperbolic PDE solver. Lagrangian methods rely on the concept of characteristic curves, which we compute approximately and in parallel using a fourth-order Runge-Kutta method. The solver can be differentiated efficiently enabling fast Gauss-Newton based optimization methods. We present quickly converging iterative linear solvers using multigrid or spectral preconditioners that render the overall optimization efficient and scalable. We demonstrate the potential of our new approach using several synthetic and real world test problems.

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

Andreas Mang  
The University of Texas at Austin  
The Institute for Computational Engineering and Sciences

andreas@ices.utexas.edu

### MS160

#### Stochastic Newton and Quasi-Newton Methods for Large-Scale Least-Squares Problems

In this work, we describe stochastic Newton and stochastic quasi-Newton approaches to efficiently solving large linear least squares problems, for applications where the size of the data exceeds the memory capabilities or for problems with time-dependent data acquisition. We present a novel framework called combined subsampling, where stochasticity is introduced to overcome computational limitations, and we describe stochastic approximation methods for computing solutions. Theoretical results for consistency of both the stochastic Newton and stochastic quasi-Newton methods are provided. Numerical results demonstrate the potential benefits of this approach.

Julianne Chung, Matthias Chung, Joseph T. Slagel  
Department of Mathematics  
Virginia Tech  
jmchung@vt.edu, mcchung@vt.edu, slagelj@vt.edu

Luis Tenorio  
Colorado School of Mines  
ltenorio@mines.edu

### MS160

#### Reduced Order Models for One-Step Multispectral Quantitative Photoacoustic Tomography

Photoacoustic tomography, a high contrast, high resolution imaging modality, is an ill-posed inverse coefficient problem for a coupled wave equation and diffusion equation pair. Standard inversion methods such as regularized quasi-Newton optimization involve a significant computational cost which increases linearly with the number of optical wavelengths. To accelerate the inversion, we use a POD-based reduced order model for the wavelength dependence. We demonstrate the computational gains on a synthetic problem motivated by neuroimaging.

Arvind Saibaba  
North Carolina State University  
asaibab@ncsu.edu

Sarah Vallelian  
Statistical and Applied Mathematical Sciences Institute  
scvallel@ncsu.edu

### MS161

#### The Effect of Snapshot Divergence Error on the Accuracy of Reduced Order Modeling

We consider using ROM bases derived from strongly divergence-free (Scott-Vogelius) and weakly divergence-free (Taylor-Hood) finite element simulations. We find that the strongly divergence-free bases provides more accurate predictions of lift and drag, in particular over longer time intervals.

Muhammad Mohebujjaman  
Clemson University  
mmohebu@clemson.edu

Leo Rebholz  
Clemson University

Department of Mathematical Sciences  
rebholz@clemson.edu

Traian Iliescu  
Department of Mathematics  
Virginia Tech  
iliescu@vt.edu

Xuping Xie  
Virginia Tech  
xupingxy@vt.edu

### MS161

#### libROM: A Distributed-Memory Adaptive Incremental Proper Orthogonal Decomposition

Model reduction is increasingly applied to large-scale dynamical systems in order to reduce the computational effort required for multi-query applications such as optimization, control, and sensitivity analysis. A bottleneck in proper orthogonal decomposition-based approaches, including those used as inputs to hyperreduction schemes, is the use of direct linear algebra on very large snapshot matrices. We present a distributed-memory algorithm to overcome these limitations using a combination of incremental algorithms for singular value decomposition and adaptive control of snapshot selection using methods inspired by adaptive time-stepping algorithms for ordinary differential equations. An implementation of our algorithm is available on GitHub in the C++ library libROM. Results are presented for example applications in transport phenomena to illustrate libROMs performance, and to suggest how libROM can be incorporated into users applications.

Geoffrey M. Oxberry, Bill Arrighi, Tanya Kostova-Vassilevska, Robert W. Anderson  
Lawrence Livermore National Laboratory  
oxberry1@llnl.gov, arrighi2@llnl.gov, vas-  
silevska1@llnl.gov, anderson110@llnl.gov

### MS161

#### Adaptive Reduced-Order Modeling for Flows under Uncertainty

Space-adaptive numerical methods have recently found their way into reduced-order modeling of parametrized PDEs, because they promise simultaneous control of the snapshot discretization error and the reduced basis approximation error. We consider the snapshot computation with adaptive finite elements for POD-Galerkin reduced-order modeling of time-dependent incompressible flows. It is convenient to represent the snapshots as members of a reference finite element space containing all snapshot spaces. In our talk we are going to address the resulting challenges with respect to Galerkin orthogonality, discrete divergence-freeness and computational cost. Our motivation for reduced-order modeling is the quantification of uncertainty for flow problems with random data. In this context, a reduced-order model can be viewed as a surrogate model. This means, for given realizations of the input data, the model provides realizations of the corresponding output data at a low numerical cost. When put in the framework of stochastic sampling, this enables the estimation of statistical quantities at essentially the cost of the snapshot computations.

Sebastian Ullmann  
TU Darmstadt  
ullmann@gsc.tu-darmstadt.de

Jens Lang  
 Technische Universität Darmstadt  
 Department of Mathematics  
 lang@mathematik.tu-darmstadt.de

### MS161

#### **Hierarchical Model Reduction Methods for Incompressible Navier-Stokes Equations in Pipes and Networks: Challenges and Perspectives**

Emerging applications of scientific computing require a progressively improved efficiency without affecting accuracy. Reduction of computational costs is pursued by either general methods (for instance based on a off-line/on-line paradigm) or by specific customization of a method for a specific problem. In this talk, we consider a dimensional model reduction called HiMod (Hierarchical Model Reduction) specifically developed for problems where the geometry induces dynamics consisting of a mainstream and transverse components. The background application is hemodynamics in network of pipes (arteries). When we consider a fluid in a pipe or in a network, we can take advantage from this particular configuration. HiMod constructs a "psychologically 1D" model by adding 2D transverse details to a 1D backbone finite element discretization of the incompressible Navier-Stokes equations. By an educated choice of the spectral basis that describes the transverse components, we can obtain a significant reduction of the size of the numerical problem with a consequent significant computational saving. In the talk, we address recent advances of the method including an extensive discussion of the choice of the modal function, the simulation of real vascular geometries and the use of isogeometric methods for the mainstream. Work in collaboration with S. Perotto, S. Guzzetti, P.J. Blanco, L. Alonso Alvarez, A. Reali. Supported by NSF DMS 1412963

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

### MS162

#### **Solution of Large-Scale Constrained Riccati Equations for Low-Order Models and Controllers of Flows**

Riccati equations play a prominent role in model reduction and controller design of dynamical systems. If flows are concerned, the solution process of Riccati equations faces two major obstacles. Firstly, models for the simulation of flows are typically of high-order which is particularly problematic for the solution of matrix equations like the Riccati equations. Secondly, the divergence free constraint has to be respected also by the Riccati solutions which calls for extension of established solution techniques. In this talk, we will present a general approach to the solution of high-dimensional constrained Riccati equations and illustrate the performance of the approach and the implementation for the design of low-dimensional compensators for the stabilization of the cylinder wake.

Jan Heiland  
 TU Berlin, Germany  
 heiland@mpi-magdeburg.mpg.de

### MS162

#### **Ghosts of Model Reduction: Past, Present and**

### **Yet-To-Come**

The emergence of model reduction is related to the rise of digital information processing: To enable the simulation of differential equation models with numerical complexities beyond computational capabilities, reduced order models are sought which approximate the original model's behavior. The evolution of this discipline at the interface between applied mathematics and scientific computing, is illustrated exemplarily by Gramian-based approaches for input-output systems. Starting with linear time-invariant systems, this class of model reduction methods developed various specializations, among others, for: nonlinear, parametric, time-varying, structured or descriptor systems. These techniques are outlined, related to current research and to set in perspective to future challenges such as controlling large-scale complex networks with high-dimensional parameter-spaces. Furthermore, issues such as comparability and metrics for efficient model reduction are discussed.

Christian Himpe  
 Computational Methods in Systems and Control Theory  
 Max Planck Institute Magdeburg  
 himpe@mpi-magdeburg.mpg.de

Sara Grundel  
 MPI Magdeburg  
 grundel@mpi-magdeburg.mpg.de

Peter Benner  
 Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany  
 benner@mpi-magdeburg.mpg.de

### MS162

#### **Reduced-Order Optimal Feedback Control of Semilinear Parabolic Dynamics**

Feedback control design plays a fundamental role in modern engineering. For an optimality-based formulation of the control problem, the Dynamic Programming Principle allows the characterization of the associated value function as the viscosity solution of a first-order, fully nonlinear Hamilton-Jacobi-Bellman equation. The equation is defined over the state-space of the controlled dynamical system and therefore, even control problems over low-dimensional dynamics lead to HJB equations of high complexity. In this talk, we present an approximation framework to compute (sub)optimal feedback controllers based on the solution of a Generalized HJB equation and a policy iteration algorithm. The controller is computed upon a state space representation of reduced dimension, obtained via spectral elements. Problems arising from the feedback control of partial differential equations illustrate the effectiveness of our approach in a high-dimensional context.

Dante Kalise  
 Radon Institute for Computational and Applied Mathematics (RICAM)  
 dante.kalise@oeaw.ac.at

Karl Kunisch  
 Karl-Franzens University Graz  
 Institute of Mathematics and Scientific Computing  
 karl.kunisch@uni-graz.at

### MS162

#### **Feedback Control for Parametric PDEs Using RB**



## Methods

Feedback control for PDEs is in general a computationally demanding procedure, since the spatial discretization leads to very large discrete systems. When parametric problems are considered, e.g. geometric variations, material coefficients or uncertainties, the procedure can easily become infeasible in contexts where in real-time new controllers are desirable, or where the computational power is limited, e.g. on microcontrollers or smartphones. We thus apply the reduced-basis method to feedback control problems, starting with the standard (and most simple) linear-quadratic regulator setup. We then extend this formulation by introducing disturbances and thus come up with  $\mathcal{H}_2$ -optimal control problems. Finally, nonlinear MPC schemes will be tackled with the RB method. The goal is always to obtain online/offline efficient schemes and to derive a-posteriori error bounds that certify the quality of the approximations.

Andreas Schmidt, Bernard Haasdonk  
University of Stuttgart  
schmidta@mathematik.uni-stuttgart.de,  
haasdonk@mathematik.uni-stuttgart.de

## MS163

### Analysis and Approximation of a Fractional Cahn-Hilliard Equation

We derive a Fractional Cahn-Hilliard Equation (FCHE) by considering a gradient flow in the negative order Sobolev space  $H^a$ ,  $a \in [0, 1]$  where the choice  $a = 1$  corresponds to the classical Cahn-Hilliard equation whilst the choice  $a = 0$  recovers the Allen-Cahn equation. It is shown that the equation preserves mass for all positive values of fractional order  $a$  and that it indeed reduces the free energy. The well-posedness of the problem is established in the sense that the  $H^1$ -norm of the solution remains uniformly bounded. We then turn to the delicate question of the  $L^8$  boundedness of the solution and establish an  $L^8$  bound for the FCHE in the case where the non-linearity is a quartic polynomial. As a consequence of the estimates, we are able to show that the Fourier-Galerkin method delivers a spectral rate of convergence for the FCHE in the case of a semi-discrete approximation scheme. Finally, we present results obtained using computational simulation of the FCHE for a variety of choices of fractional order  $a$ . It is observed that the nature of the solution of the FCHE with a general  $a \in [0, 1]$  is qualitatively (and quantitatively) closer to the behaviour of the classical Cahn-Hilliard equation than to the Allen-Cahn equation, regardless of how close to zero be the value of  $a$ .

Mark Ainsworth  
Division of Applied Mathematics  
Brown University  
Mark\_Ainsworth@brown.edu

Zhiping Mao  
Brown University, USA  
zhiping\_mao@brown.edu

## MS163

### Implicit-Explicit Difference Schemes for Fractional Differential Equations with Non-Smooth Solutions

We propose second-order implicit-explicit (IMEX) time-stepping schemes for nonlinear fractional differential equations with fractional order  $0 < \beta < 1$ . From the known

structure of the non-smooth solution and by introducing corresponding correction terms, we can obtain uniformly second-order accuracy from these schemes. We first prove the convergence and linear stability of the proposed schemes. Numerical examples illustrate the flexibility and efficiency of the IMEX schemes and show that they are effective for nonlinear and multi-rate fractional differential systems, and also multi-term fractional differential systems with non-smooth solutions.

Wanrong Cao  
Southeast University of China  
wanrongcao@gmail.com

Fanhai Zeng  
Division of Applied Mathematics  
Brown University  
fanhai\_zeng@brown.edu

Zhongqiang Zhang  
Worcester Polytechnic Institute  
zzhang7@wpi.edu

George Em Karniadakis  
Brown University  
george\_karniadakis@brown.edu

## MS163

### Thermodynamically Viable Fractional Wave Equations for Power Law Attenuation in Viscoelastic Media

Many complex media of great practical interest, such as in medical ultrasound and sediment acoustics, display an attenuation that increases with a power law as a function of frequency. Usually measurements can only be taken over a limited frequency range, while fractional wave equations model attenuation over all frequencies. There is therefore some freedom in how the models behave outside of this limited interval, and many different fractional wave equations have been proposed. In addition, it is desirable that a wave equation models physically viable media and for that two conditions have to be satisfied. The first is causality, and the second is a criterion that comes from thermodynamic considerations and implies that the relaxation modulus is a completely monotonic function. The latter implies that attenuation asymptotically cannot rise faster than frequency raised to the first power. These criteria will be explained and used to evaluate several of the fractional wave equations that exist.

Sverre Holm  
Universitetet i Oslo  
sverre@if.uio.no

## MS163

### Space-Time Duality and Anomalous Diffusion

Abstract not available

Jim Kelly  
Michigan State University  
kellyja8@stt.msu.edu

## MS164

### Incompressible Two-Phase Flows with Open/Outflow Boundaries: Boundary Condi-

## tions and Algorithm

This work focuses on the motion of a mixture of two immiscible incompressible fluids in a domain with open boundaries. The domain boundary is open in the sense that the fluids can freely leave or even enter the domain through such boundaries. In particular, we concentrate on situations where the interface formed between the two fluids passes through the open portions of the domain boundary. The problem therefore involves truly two-phase outflow/open boundaries. The challenge facing the design of effective numerical techniques for treating such problems is manifold. Some primary issues are associated with the viscosity contrasts, density contrasts, surface tension, and the presence of fluid interface, backflows and strong vortices on the open boundaries. Large density ratios and large viscosity ratios of the two fluids make two-phase outflow simulations tremendously challenging. In this talk we present a family of boundary conditions, and an associated numerical algorithm, for two-phase outflow simulations within the phase field framework. These open boundary conditions have the salient feature that they all ensure the energy stability of the two-phase system, even in situations where strong vortices, backflows, and large viscosity contrast are present at the open boundaries. Numerical experiments will be presented to demonstrate the long-term stability of the method in situations where large density contrast, large viscosity contrast?

Suchuan Dong  
Purdue University  
sdong@purdue.edu

## MS164

### Scalable Compact Localized Exponential Time Differencing Method for Simulating Coarsening Dynamics of Cahn-Hilliard Equations

Numerical simulation of microstructure coarsening is a subject of great interest in computational materials science. The coarsening dynamics in a binary mixture can be modeled by the celebrated Cahn-Hilliard equations. In this talk, we first present a stabilized compact exponential time differencing (cETD) method for solving the Cahn-Hilliard equations. The proposed method combine linear operator splittings, compact discretizations of spatial operators, exponential time integrators, multistep or Runge-Kutta approximations and fast Fourier transform, to produce efficient and accurate numerical algorithms. Then we propose a scalable implementation of the cETD method based on domain decomposition and operator localization, which is highly suitable for parallel computing. Various numerical experiments are carried out to computationally study the coarsening dynamics under different diffusion mobilities, including extreme spatial scale simulations on the Sunway TaihuLight supercomputer.

Lili Ju  
University of South Carolina  
Department of Mathematics  
ju@math.sc.edu

## MS164

### Computational Models of the Endoplasmic Reticulum, Bifurcation and Evolution

The Endoplasmic Reticulum is a continuous membrane that separates the luminal and cytosolic fluids with cells. It takes a wide variety of morphologies, from flat, cister-

nal phases, fenestrated, and reticulated tubule networks. Each of these phases is associated with an essential role in cellular functioning, serving as a host sit for protein transcription as well as the principle transport mechanism for cellular Calcium. We present a family of models that describe the interaction of surface proteins with the curvature of the underlying membrane morphology, reducing the the coupled evolution to a geometric flow.

Keith Promislov  
Michigan State University  
kpromisl@math.msu.edu

## MS164

### Linear Unconditional Energy-Stable Splitting Schemes for Mixtures of Nematic-Isotropic Flows with Anchoring Effects

The study of interfacial dynamics between two different components has become the key role to understand the behavior of many interesting systems. Indeed, two-phase flows composed of fluids exhibiting different microscopic structures are an important class of engineering materials. The dynamics of these flows are determined by the coupling among three different length scales: microscopic inside each component, mesoscopic interfacial morphology and macroscopic hydrodynamics. Moreover, in the case of complex liquids composed by the mixture between isotropic (newtonian fluid) and nematic (liquid crystal) flows, its interfaces exhibit novel dynamics due to anchoring effects of the liquid crystal molecules on the interface. In this talk I will introduce a new differential problem to model mixtures composed by isotropic fluids and nematic liquid crystals, taking into account viscous, mixing, nematic, and anchoring effects and reformulating the corresponding stress tensors in order to derive a dissipative energy law. Then, I will present two new linear unconditionally energy-stable splitting schemes that allows us to split the computation of the three pairs of unknowns (velocity-pressure, phase field-chemical potential and director vector-equilibrium) in three different steps. Finally, I will present several numerical simulations in order to show the efficiency of the proposed numerical schemes.

Giordano Tierra  
Department of Mathematics  
Temple University  
gtierra@temple.edu

## MS165

### Cluster Sampling Filters for Non-Gaussian Data Assimilation

We present a fully non-Gaussian version of the Hamiltonian Monte Carlo (HMC) sampling filter. The Gaussian prior assumption in the formulation of the original HMC sampling filter is relaxed. Specifically, a clustering step is introduced after the forecast phase of the filter, and the prior density function is estimated by fitting a Gaussian Mixture Model (GMM) to the prior ensemble. Using the data likelihood function, the posterior density is then formulated as a mixture density, and is sampled using a HMC approach (or any other scheme capable of sampling multimodal densities in high-dimensional subspaces). The main filter presented is named "cluster HMC sampling filter" (CIHMC). A multi-chain version of the CIHMC filter, namely MC-CIHMC is also proposed to guarantee that samples are taken from the vicinities of all probability modes of the

formulated posterior.

Ahmed Attia  
Virginia Tech  
attia@vt.edu

### MS165

#### Multilevel Higher Order Quasi Monte Carlo Methods for Bayesian Inverse Problems

We consider forward and inverse uncertainty quantification of partial differential equations with distributed uncertain inputs. Focusing on the Bayesian approach to inverse problems together with suitable uncertainty parametrization, the problem reduces to the computation of high-dimensional integrals. Under certain sparsity conditions on the parametric input, the resulting integrands fulfill analogous conditions, allowing recently developed higher-order QMC methods to be applied. We focus in this talk on computational aspects of this problem, including single-level and multilevel formulations of the discrete Bayesian inverse problem, which are shown to outperform conventional Monte Carlo-based approaches. Numerical results confirming the theoretical improvement of the time-to-solution for a given error level are presented. [J. Dick et al. Multilevel higher order Quasi-Monte Carlo Bayesian Estimation. Tech. rep. 2016-34. Seminar for Applied Mathematics, ETH Zurich, 2016.] We also comment on software for the application of these methods, and their pertinence in a high-performance setting. [R. N. Gantner and Ch. Schwab. Computational Higher Order Quasi-Monte Carlo Integration. In: MCQMC14, Leuven, Belgium, April 2014. Ed. by Ronald Cools and Dirk Nuyens. 2016, pp. 271-288.] This work was supported by CPU time from the CSCS under project ID d41, by the SNSF under Grant No. SNF149819, and by the Australian Research Council under project No. DP150101770.

Robert N. Gantner  
ETH Zurich  
robert.gantner@sam.math.ethz.ch

Josef Dick  
School of Mathematics and Statistics  
The University of New South Wales  
josef.dick@unsw.edu.au

Quoc T. Le Gia  
School of Mathematics and Statistics  
University of New South Wales, Australia  
qlgia@unsw.edu.au

Christoph Schwab  
ETH Zuerich  
SAM  
christoph.schwab@sam.math.ethz.ch

### MS165

#### Adaptive Dimension Reduction to Accelerate Infinite-Dimensional Geometric MCMC

Bayesian inverse problems highly rely on efficient and effective MCMC algorithms for uncertainty quantification. Infinite-dimensional MCMC algorithms, directly defined on function spaces, are robust under refinement of physical models. Recently, a class of algorithms also starts to take advantage of geometric information provided by, for example, quadratic approximation of the parameter-to-observation map so that they are capable of exploring com-

plex probability structures, as frequently arise in UQ for PDE constrained inverse problems. However, the required geometric information, in particular the Fisher information metric, is very expensive to obtain in high dimensions. The issue can be mitigated by dimension reduction techniques. By carefully splitting the unknown parameter space into a low-dimensional geometry-concentrated subspace, and an infinite-dimensional geometry-flat subspace, one may then apply geometry-informed MCMC methods to the low-dimensional subspace and simpler methods to the infinite-dimensional complement (DILI, Cui et al 2016). In this paper, we explore randomized linear algebraic algorithms (Halko et al 2011) to efficiently obtain a local low-dimensional subspace to speed up geometric infinite-dimension MCMC methods. It enables adaptive dimension reduction at low cost.

Shiwei Lan, Mark Girolami  
University of Warwick  
s.lan@caltech.edu, m.girolami@warwick.ac.uk

Andrew Stuart  
Mathematics Institute,  
University of Warwick  
a.m.stuart@warwick.ac.uk

### MS165

#### Optimization-Based Samplers in a Measure Transport Framework

In the Bayesian statistical paradigm, uncertainty in the parameters of a physical system is characterized by a probability distribution. Information from observations is incorporated by conditioning on the resulting data; this conditioning describes an “update” from the prior distribution to the posterior distribution. Quantities of interest—such as credible regions and event probabilities—can then be obtained by taking expectations with respect to this posterior. Solving a Bayesian inference problem thus amounts to characterizing the posterior, and a flexible way to do so is via sampling. Markov chain Monte Carlo (MCMC) algorithms can produce asymptotically exact samples from the posterior using only unnormalized density evaluations, but their efficiency rests on the design of effective proposal distributions. Recent optimization-based MCMC methods, for example randomize-then-optimize (RTO), repeatedly solve optimization problems to generate good proposal samples. We analyze RTO using a new interpretation that describes each optimization as a projection that yields the action of a transport map. From this analysis, several new variants of RTO—adaptive RTO, mixtures of proposals, and transformed random walks—follow naturally.

Zheng Wang  
Massachusetts Institute of Technology  
USA  
zheng\_w@mit.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

### MS166

#### Numerical Approximation of Optimal Transport Maps via Monge-Ampere Equations

We describe an almost-monotone filtered approximation of the Monge-Ampere equation, which is designed to capture weak (viscosity) solutions. The optimal transport

constraint is reformulated as a Hamilton-Jacobi equation, which can also be approximated using a monotone scheme. The combined scheme allows for the numerical approximation of optimal transport maps in a range of challenging settings. Applications to beam shaping are discussed.

Brittany Froese  
New Jersey Institute of Technology  
Department of Mathematical Sciences  
bdfroese@njit.edu

#### MS166

##### **Freeform Multifocal Lenses Designed with Supporting Quadric Method (SQM): Geometry and Diffraction Effects**

Design of freeform refractive lenses is known to be a difficult inverse problem. But solutions, if available, can be very useful, especially in devices required to redirect and reshape the radiance of the source into an output irradiance redistributed over a given target according to a prescribed pattern. In this report we present results of theoretical and numerical analysis of refractive lenses designed with the Supporting Quadric Method. It is shown that such freeform lenses have a particular simple geometry and qualitatively their diffractive properties are comparable with rotationally symmetric lenses designed with classical methods.

Vladimir Olikier  
Department of Mathematics and Computer Science  
Emory University  
oliker@mathcs.emory.edu

#### MS166

##### **Inverse Reflector Problem in Arbitrary Coordinates**

In inverse modelling of optical systems, the shape of an optical surface is governed by the Monge-Ampère equation. I shall present a coordinate-free formulation of this equation for the inverse reflector problem using tensor calculus. The solution of the Monge-Ampère equation proceeds in two stages, i.e., we first compute a mapping from source to target, and from this we compute the shape of the reflector. For the numerical computation of the mapping we employ a least-squares method. I shall demonstrate the performance of the method for a circular source employing polar coordinates.

Jan Ten Thijs Boonkamp  
Department of Mathematics and Computer Science  
Eindhoven University of Technology  
j.h.m.tenthijeboonkamp@tue.nl

#### MS166

##### **Monge-Ampère Equation in Freeform Illumination Optics**

Freeform optics design is an important technique widely used in optical systems to redistribute the spatial energy distribution of a light source. In practical applications, it is usually the case that the influence of the size or the angular extent of the actual source on the performance of an illumination system can be ignored. This allows the actual source to be considered as an ideal source (point source/collimated beam) in the design, meanwhile a good agreement between the actual performance and the nominal performance of the illumination system still can be guaranteed. Under the

assumption of an ideal source, the design of freeform illumination optics becomes essentially a mathematical problem. There are several zero-tendue methods that can be applied to the design of freeform illumination optics, such as the optimization method, the ray targeting method, the Supporting-Ellipsoid method and the Monge-Ampère (MA) equation method. Among these methods, the MA method may be the most advanced, because a good numerical solution to the prescribed illumination design can easily be obtained without cumbersome Monte Carlo ray tracing and smooth surfaces can be constructed by use of the numerical solution without step discontinuities. In this presentation, we will give a detailed derivation of the elliptic MA equation and introduce the application of MA method to the illumination design of one freeform surface and two freeform surfaces in LED lighting and laser beam shaping.

Rengmao Wu  
College of the Optical Sciences  
The University of Arizona  
wrengmao@163.com

Zhenrong Zheng  
State Key Laboratory of modern Optical Instrumentation  
Zhejiang University  
zrzr@zju.edu.cn

#### MS167

##### **Image Processing, Optimization, and Uncertainty Quantification for Large-Scale Experimental Science**

Abstract not available

Richard Archibald  
Computational Mathematics Group  
Oak Ridge National Laboratory  
archibaldrk@ornl.gov

#### MS167

##### **Using Conjugate Gradient Iterations for Sampling in Hierarchical Gaussian Inverse Problems**

In this talk, we focus on the problem of using truncated conjugate gradient (CG) iterations to approximately solve a linear system whose solution yields an independent sample from a Gaussian distribution. The question is, if we only have the ability to obtain approximate solutions of the linear system, how can we obtain a Markov chain whose stationary distribution is the Gaussian distribution of interest? To this end, we present the *gradient scan Gibbs sampler*, which recently appeared in the literature, and then discuss preconditioning techniques for accelerating its convergence properties. Such methods are motivated by applications in large-scale linear inverse problems, where linear systems are large enough that they can only be approximately solved using an iterative method, such as conjugate gradient, but where uncertainty quantification, and hence provably convergent MCMC methods, are needed.

Johnathan M. Bardsley  
University of Montana  
bardsleyj@mso.umt.edu

#### MS167

##### **Robust-Adaptive Decision-making for Financial**

## Forecasting

This work illustrates a robust-adaptive strategy applied to financial portfolio allocation. The application relies on actual financial data and, while simplified, is representative of what would need to be dealt with for real analysis. The Capital Asset Pricing Model (CAPM) selects one of nine potential stock indices for investment. The choice is repeated daily and the goal is to maximize the value of the portfolio at the end of a 16-year period. We establish, first, that a purely random selection strategy (select the stock index randomly) offers the potential to outperform selection based on CAPM forecasts (select the stock index predicted to perform best). Both purely random and CAPM-based strategies track well with the expected market growth as indicated by the S&P-500 Index. We further establish that the proposed robust-adaptive strategy grows the value of the portfolio several times more than what can be achieved based solely on CAPM forecasts. The improvement is up to a factor of forty if a wait-and-see option is exercised whereby the analyst decides to either keep investing (at the risk of losing value) or sell everything (and protect the current value of the portfolio) until the decision is made to invest again. The discussion addresses why a robust-adaptive approach to decision-making offers the potential of superior results for a wide range of applications from financial analysis to condition monitoring and big data analytics.

Francois Hemez

Los Alamos National Laboratory  
hemez@lanl.gov

## MS167

### Estimating Particle Size Distributions in Dynamic Experiments via Mie Scattering

An important problem in the study of aerosols and gases is the size distribution of particulates entrained in the gas. In particular, in dynamic material studies, it is often of interest to understand the distribution of sizes of particles released from a metal surface when it is shocked. Mie scattering theory gives a closed form solution to Maxwells equations under the assumption that the particles are perfectly spherical, so that one can, in principle, determine the size distribution of particles by measuring the scattering of incident laser light at different angles. In this work we present the mathematical model for computing size distributions from measured data using the Mie model, which results in a Volterra integral operator of the first kind. We will discuss some of the challenges associated with inverting such operators, which are ill-posed but not severely, as well as the challenges introduced by analyzing real data. We will also demonstrate a statistical formulation of the problem and show Monte Carlo results that give the sensitivity of this approach to experimental factors such as the number of light probes and the scattering angles chosen and present results from real experiments taken at the U.S. Department of Energys Special Technologies Laboratory.

Aaron B. Luttmann, Marylesa Howard

National Security Technologies, LLC  
luttmaab@nv.doe.gov, howardmm@nv.doe.gov

## MS168

### Solving Multi-Scale Stochastic Partial Differential

## Equations Using Deep Neural Networks

The applicability of traditional methods for stochastic partial differential equations (SPDEs) including Monte Carlo (MC) methods (and its variants), operator based methods, moment methods and generalized polynomial chaos (gPC) is typically restricted by the fact that they are either incapable of dealing with the curse of dimensionality or require a vast number of evaluations of the (generally expensive) forward model to obtain convergent statistics. The challenge is compounded when one is dealing with multiscale systems which require resolving small-scale effects on the large scale to get accurate predictions. We demonstrate a modification of the multiscale finite element method (MFEM) where the multiscale basis functions are predicted using a deep neural network (DNN). The DNN is trained with data obtained by solving the corresponding homogeneous PDE on coarse block elements of the finite element (FE) grid with a broad range of length scales of the random field. The full solution is obtained by coupling the locally adaptive multiscale basis functions with the global FE formulation. Our approach is validated with examples on steady flow through random porous media.

Ilias Bilonis, Rohit Tripathy

Purdue University  
ebilionis@gmail.com, rtripath@purdue.edu

## MS168

### Probabilistic Meshless Methods for Partial Differential Equations and Bayesian Inverse Problems

Partial differential equations (PDEs) are a challenging class of problems which rarely admit closed-form solutions, forcing us to resort to numerical methods which discretise the problem to construct an approximate solution. We seek to phrase solution of PDEs as a statistical inference problem, and construct probability measures which quantify the epistemic uncertainty in the solution resulting from the discretisation [Jon Cockayne, Chris Oates, Tim Sullivan, and Mark Girolami. Probabilistic Meshless Methods for Partial Differential Equations and Bayesian Inverse Problems, 2016]. We explore construction of probability measures for the strong formulation of elliptic PDEs, and the connection between this and meshless methods. We seek to apply these probability measures in Bayesian inverse problems, parameter inference problems whose dynamics are often constrained by a system of PDEs. Sampling from parameter posteriors in such problems often involves replacing an exact likelihood involving the unknown solution of the system with an approximate one, in which a numerical approximation is used. Such approximations have been shown to produce biased and overconfident posteriors when error in the forward solver is not tightly controlled. We show how the uncertainty from a probabilistic forward solver can be propagated into the parameter posteriors, thus permitting the use of coarser discretisations which still produce valid statistical inferences.

Jon Cockayne

The University of Warwick  
j.cockayne@warwick.ac.uk

## MS168

### A Fully Bayesian Solver for ODEs Based on Runge-Kutta

In an ongoing push to construct probabilistic extensions of classic ODE solvers for applications in statistics, ma-

chine learning and control engineering, Schober et al. recently showed that Runge-Kutta methods can be recast as a Bayesian (Kalman) filter with an Integrated Wiener Process prior on the solution. This filter returns a posterior measure which returns the result of a Runge-Kutta step as a posterior mean and additionally outputs the covariance from Gaussian Process (GP) regression with low computational overhead compared to classical Runge-Kutta solvers. In an attempt to amend this methods rough uncertainty calibration while sustaining its small cost overhead, Kersting and Hennig proposed a novel way to quantify the uncertainty in this filtering framework by probing the vector field using Bayesian Quadrature. By linking the free parameters of this new probabilistic ODE solver to the (local) Lipschitz constants and differentiability of the underlying vector field of the ODE, we derive a fully Bayesian ODE filter by encoding these properties in a Matrn prior on the vector field. If sufficient prior knowledge is available, this Bayesian solver returns a well-calibrated uncertainty quantification and average-case error analysis, while staying sufficiently close to Runge-Kutta solvers and inheriting their proven desirable properties.

Hans Kersting, Philipp Hennig  
Max Planck Institute for Intelligent Systems  
hans.kersting@tuebingen.mpg.de, phen-  
nig@tuebingen.mpg.de

#### MS168

##### **Risk-Averse Assessment, Learning, and Optimization Using Surrogate Models**

We formalize a unifying framework for handling parametric, surrogate, and learning uncertainty. Specifically, we develop an optimization workflow that offers a principled and risk-averse handling of uncertainties not only originating from random inputs, but also in the training of probabilistic surrogate models, and the exploration of the input/design space towards identifying feasibility and optimal configurations. The effectiveness of the proposed methodology is demonstrated through simple pedagogical examples and established benchmark problems, as well a realistic application involving the risk-averse design optimization of super-cavitating hydrofoils using multi-fidelity simulations.

Paris Perdikaris  
Massachusetts Institute of Technology  
parisp@mit.edu

Johannes O. Royset  
Operations Research Department  
Naval Postgraduate School  
joroyset@nps.edu

#### MS169

##### **Goal Oriented Adaptivity Based on the Forward-Adjoint Symmetrization of the Transport Equation**

In this talk, we will more thoroughly investigate the numerical properties of the second-order form of neutron transport equation based on combining the forward and adjoint equations. As shown in our previous exploratory work [Hanus and McClarren, JCTT Vol. 45, 2016], this formulation allows obtaining the solution of both the forward and the adjoint problem by solving only the forward problem and performing a relatively cheap post-processing step. This feature has been utilized in our previous work in the framework of adjoint-weighted goal oriented adaptive

FEM. Here we will primarily focus on the biggest shortcoming outlined in that work the oscillatory behavior caused by inadequate approximation of the gradient of the finite element solution, required to reconstruct the solution of the original transport problem. To address this issue, we will apply a locally averaging gradient recovery technique and discuss its effect on the final solution, as well as on the quality of the error indicators used in the adaptivity procedure.

Milan Hanus  
Texas A&M University  
mhanus@tamu.edu

Ryan McClarren  
Texas A&M University  
USA  
rgm@tamu.edu

#### MS169

##### **Implicit Solution of the Vlasov-Poisson System**

Most Eulerian and semi-Lagrangian codes use a splitting in time approach for implicit, charged particle transport solvers. The reason for this is advection with respect to only position or only velocity can be solved explicitly and even exactly, but advection with respect to both position and velocity requires solving a linear system whose order is the size of the phase space discretization. This becomes unacceptable for higher dimensional problems. We propose a new approach which does not involve splitting in time, but splitting the phase space into subdomains. If each subdomain is appropriately chosen, the transport problem can be solved explicitly on each subdomain. A Krylov solver is then used to solve for the boundaries between subdomains, which involves a much smaller linear system than the one required for the no splitting approach.

Cory Hauck, Charles K. Garrett  
Oak Ridge National Laboratory  
hauckc@ornl.gov, ckgarrett@lanl.gov

#### MS169

##### **Rattlesnake - MOOSE-Based Radiation Transport Application for Multi-Physics Simulations at INL**

Rattlesnake is the MOOSE based application for simulating the physics of radiation transport. Radiation transport studies the motion of particles, typically neutrons or photons, in a stationary host-medium. To this end, radiation transport codes solve the linear Boltzmann transport equation (BTE). In contrast to standard PDEs the BTE is an integro-differential equation featuring a seven-dimensional phase space: three spatial dimensions, three dimensions for the momentum of motion, and time. Rattlesnake is a MOOSE based application and has been designed from the outset to be deployed within a multi-physics framework. Temperature and densities affected by the solution of BTE have a significant influence on the material properties - the nuclear cross sections - that the BTE relies on. Hence, BTE coupled with other physics is a high-dimensional non-linear problem in the multiphysics environment. Typical radiation transport methods assume constant cross sections over each element, while Rattlesnake allows for direct injection of the feedback on quadrature points leading to a consistent treatment within the FEM framework. We focus on three Rattlesnake features: (1) multiphysics coupling with a few benchmark problems (2) multi-scale approach where multiple discretization methods are matched on various prob-

lem subdomains utilizing either mortar FEM or interface penalty method, (3) a multigroup cross section management toolkit, YAKXS, used by Rattlesnake.

Yaqi Wang  
Idaho National Laboratory  
yaqi.wang@inl.gov

Sebastian Schunert  
Idaho National Lab  
sebastian.schunert@inl.gov

Mark DeHart, Richard Martineau  
Idaho National Laboratory  
mark.dehart@inl.gov, richard.martineau@inl.gov

### MS169

#### Fmm Preconditioner for Radiative Transport Equation with Isotropic Coefficients

We propose in this work a fast numerical algorithm for solving the equation of radiative transfer (ERT) in isotropic media. The algorithm has two steps. In the first step, we derive an integral equation for the angularly averaged ERT solution by taking advantage of the isotropy of the scattering kernel, and solve the integral equation with a fast multipole method (FMM). In the second step, we solve a scattering-free transport equation to recover the original ERT solution. Numerical simulations are presented to demonstrate the performance of the algorithm for both homogeneous and inhomogeneous media. This is a joint work with Kui Ren and Rongting Zhang.

Yimin Zhong  
University of Texas at Austin  
yzhong@math.utexas.edu

### MS170

#### Exploring Model-Form Uncertainties in Boyancy-Driven Les Simulations

Abstract not available

Stefan P. Domino  
Sandia. National Laboratories  
spdomin@sandia.gov

### MS170

#### Lyapunov Stability of Scale-Resolving Turbulent Flow Simulations

With the increase in computing power, scale-resolving turbulent simulations –such as LES and DNS– emerge as promising approaches to improve both knowledge of complex flow physics and reliability of flow predictions. While these techniques have become essential tools in science, no comparable impact has been observed in engineering. The chaotic dynamics that these simulations inherit from the underlying turbulent flow are largely responsible for this lag. In particular, conventional sensitivity analysis breaks down for chaotic systems and this compromises critical tasks in engineering such as flow control, design optimization, error estimation, and uncertainty quantification.

While a number of sensitivity analysis methods have been proposed for chaotic dynamics, they all come at a high computational cost. This is ultimately related to the positive portion of the Lyapunov spectrum of the system. As such, a better understanding of the dynamic behavior of scale-

resolving simulation is necessary both to estimate the cost of these methods and to drive informed strategic decisions about what methods have the most promise for complex flow configurations.

In this talk, we investigate the effect of mesh resolution, discretization order, and flow conditions on the Lyapunov spectrum of LES and DNS simulations. From these results, we provide guidelines on how different approaches would perform and in particular which ones are likely to perform best for complex turbulent flows.

Pablo Fernandez, Nisha Chandramoorthy, Qiqi Wang  
Massachusetts Institute of Technology  
pablof@mit.edu, nishac@mit.edu, qiqi@mit.edu

### MS170

#### Lyapunov Spectrum Shift of an LES from Addition of Localized Artificial Viscosity

Large eddy simulations have been quite successful in capturing the complex nonlinear dynamics of turbulent fluid flows. A Lyapunov exponent analysis of such simulations has shown that there can be hundreds or even thousands of unstable modes which cause the tangent and adjoint solvers to diverge exponentially from almost all initial conditions. But, the number of positive Lyapunov exponents are typically much smaller than the number of dimensions of the nonlinear system which can range in 10's of millions to billions. This presentation demonstrates a method on further reducing that number. In many industrial turbulent flows it has been observed that there are only a few regions in the domain of the problem that contribute to the instability of the linearized solvers. By using energy analysis, such regions are identified, and localized artificial viscosity is added. This curbs the instability of the linearized flow solutions and reduces the number of unstable modes. Results are shown on a turbomachinery flow problem of transonic flow over a turbine vane using an adjoint solver for LES.

Chaitanya Talnikar  
MIT  
talnikar@mit.edu

Qiqi Wang  
Massachusetts Institute of Technology  
qiqi@mit.edu

### MS171

#### A Software Infrastructure for Solving Quantum Physics Problems on Extremely Parallel Systems

In the German Research Foundation (DFG) project ESSEX (Equipping Sparse Solvers for Exascale), we develop scalable sparse eigensolver libraries for large quantum physics problems. Partners in ESSEX are the Universities of Erlangen, Greifswald, Wuppertal, Tokyo and Tsukuba as well as DLR. The project pursues a coherent co-design of all software layers where a holistic performance engineering process guides code development across the classic boundaries of application, numerical method and basic kernel library. The ESSEX Sparse Solver Repository (ESSR) follows a distributed software development strategy using the distributed version control system Git and supports application driven fault tolerance. ESSR includes the kernel library GHOST (General, Hybrid, and Optimized Sparse Toolit) and the flexible software framework PHIST for implementing iterative methods on HPC systems. PHIST (Pipelined Hybrid Iterative Solver Toolkit)

has been developed containing an interface to the existing numerical software framework Trilinos. PHIST also includes adapters to basic building block libraries so that high-level algorithm developments can benefit from high-performance kernel implementations, e.g. sparse matrix-vector multiplication kernels. Moreover, PHIST provides systematic and continuous testing of all software components and allows us to develop stable implementations of innovative iterative methods in an evolving hard- and software environment.

Achim Basermann

German Aerospace Center (DLR)  
Simulation and Software Technology  
Achim.Basermann@dlr.de

Jonas Thies

German Aerospace Center (DLR)  
jonas.thies@dlr.de

Melven Roehrig-Zoellner

German Aerospace Center (DLR)  
Simulation and Software Technology  
Melven.Roehrig-Zoellner@DLR.de

**MS171**

**What We Have Learned About Using Software Engineering Practices in Scientific Software**

The increase in the importance of Scientific Software motivates the need to identify and understand which software engineering (SE) practices are appropriate. Because of the uniqueness of the scientific software domain, existing SE tools and techniques developed for the business/IT community are often not efficient or effective. Appropriate SE solutions must account for the salient characteristics of the scientific software development environment. To identify these solutions, members of the SE community must interact with members of the scientific software community. This presentation will discuss the findings from a series of case studies of scientific software projects, an ongoing workshop series, and the results of interactions between my research group and scientific software projects.

Jeffrey C. Carver

Department of Computer Science  
University of Alabama  
carver@cs.ua.edu

**MS171**

**Potential for Big Data Technologies to Radically Change the Software Engineering of HPC Visualization and Analysis Tools**

Forbes Magazine estimated the Big Data Analytics market at \$125 billion in 2015. That is 25x the entire DOE 2015 budget for HPC computing. This market uses Hadoop Map-Reduce and Apache Spark for parallel processing of vast amounts of textual data. The parallel programming paradigm is highly simplified relative to its message-passing based, HPC counterpart. Some HPC researchers have considered Big Data technologies for HPC workflows. However, for visualization tools, too much focus has been on rendering. While scalable, parallel surface and volume rendering is important, this does not represent the majority of recent software engineering investments in tools such as VisIt or ParaView. More and more, these tools represent large investments in scalable, parallel data processing algorithms involving end-to-end advancements from par-

allel I/O, to task management to computed results which are often numerical metrics instead rendered images. In addition, new algorithms can involve varying degrees of machine learning, a cornerstone of the Big Data toolbox. We will give an overview of research in the application of Big Data technologies for HPC visualization and analysis with some emphasis on approaches studied at LLNL. We will look in depth at how best to represent HPC scientific mesh and field data for Big Data tools and outline potential challenges and rewards of a new parallel programming model for writing HPC data analysis algorithms using PySpark.

Mark Miller

Lawrence Livermore National Laboratory  
miller86@llnl.gov

**MS171**

**Role of the Molecular Science Software Institute Toward Sustainable Software**

The Molecular Sciences Software Institute (MolSSI) is a new initiative funded by the National Science Foundation that will serve as a nexus for science, education, and cooperation serving the community of computational molecular scientists a broad field that includes biomolecular simulation, quantum chemistry, and materials science. Ultimately, the MolSSI will enable computational scientists to tackle science and software challenges that are orders of magnitude larger and more complex than those currently within our grasp. The MolSSI will deploy a team of Software Scientists and support Software Fellows in research groups across the U.S. to build sustainable, open-source software components for the benefit of the full spectrum of computational chemistry codes. This talk will focus on how MolSSI will catalyze significant advances in software infrastructure, education, standards, and best-practices to open new windows on the next generation of scientific Grand Challenges.

Theresa Windus

Iowa State University  
twindus@iastate.edu

Daniel Crawford

Virginia Tech  
crawdad@vt.edu

Cecilia Clementi

Rice University  
cecilia@rice.edu

Robert Harrison

Brookhaven National Laboratory and Stony Brook University  
rharrison@bnl.gov

Teresa Head-Gordon

Department of Bioengineering  
University of California, Berkeley  
TLHead-Gordon@lbl.gov

Shantenu Jha

Department of Electrical and Computer Engineering  
Rutgers University  
shantenu.jha@rutgers.edu

Anna Krylov

University of Southern California  
krylov@usc.edu



Vijay Pande  
Stanford University  
Pande Lab  
pande@stanford.edu

### MS172

#### Manifold Clustering for Nonlinear Local Reduced Order Structural Models with Contact

Building on the work presented in [Balajewicz, Projection-based model reduction for contact problems, 2015] featuring the usage of a Non-Negative Matrix Factorization (NNMF) algorithm for constructing a Reduced-Order Basis (ROB) for contact forces, a complete framework for constructing stable, robust, and accurate local Hyper Reduced Order Models (HROMs) for comprehensive nonlinear contact problems is described in this presentation. This framework incorporates the local ROB approach [Am-sallem, Nonlinear model order reduction based on local reduced-order bases, 2012] for both displacement and Lagrange multiplier degrees of freedom, when the latter are used for enforcing contact constraints. Unlike previous works, in this approach clusters are formed by directly minimizing the intrinsic dimensionality of the resulting local subspaces through a scalable manifold clustering algorithm based on  $l_1$ -minimization. This approach provides insight into the number and corresponding dimensionality of physical regimes present in a given parametric problem. Additionally, matrix completion techniques based on fast convex optimization methods are used in conjunction with NNMF in order to improve the compression of the dual snapshots and robustness of the dual ROB. The performance of this framework is assessed for several contact problems and contrasted with that of a counterpart where local hyper reduction is performed using k-means clustering and standard NNMF.

Todd Chapman  
Stanford University  
tac688@stanford.edu

### MS172

#### Lagrangian Basis Method for Dimensionality Reduction of Convection Dominated Nonlinear Flows

High-fidelity simulations of high-Reynolds numbers flows require very fine spatial and temporal discretizations to accurately resolve the multi-scale features of the flow, which makes an efficient reduced order model a very desirable and yet hard to obtain goal. It is well known that standard projection techniques using POD are not well suited for the reduction of convection dominated flows featuring moving sharp gradients and discontinuities. In this talk, we summarize foundations of a new model reduction approach specifically developed for the reduction of such flows. In the proposed new approach, global basis functions are used to approximate the flow in the Lagrangian frame of reference, that is, basis functions are used to approximate both the position and states of fluid particles. Thus, flow features requiring a large set of basis functions in the Eulerian frame of reference require relatively few in the Lagrangian frame. The proposed approach is successfully demonstrated for the reduction of several simple but representative nonlinear flows.

Rambod Mojgani, Maciej Balajewicz  
University of Illinois at Urbana-Champaign

mojgani2@illinois.edu, mbalajew@illinois.edu

### MS172

#### Modeling and Predicting Reversals of Earth's Dipole by Low-Dimensional Models

Low-dimensional models for Earth's magnetic dipole may be a powerful tool for studying large-scale dipole dynamics over geological time scales, where direct numerical simulation remains challenging. We investigate the utility of several low-dimensional models by calibrating them against the signed relative paleointensity over the past 2 million years. Model calibrations are done by "data assimilation" which allows us to incorporate nonlinearity and uncertainty into the computations. We find that the data assimilation is successful, in the sense that a relative error is below 8% for all models and data sets we consider. The successful assimilation of paleomagnetic data into low-dimensional models suggests that, on millennium time scales, the occurrence of dipole reversals mainly depends on the large-scale behavior of the dipole field, and is rather independent of the detailed morphology of the field. This, in turn, suggests that large-scale dynamics of the dipole may be predictable for much longer periods than the detailed morphology of the field, which is predictable for about one century. We explore these ideas and introduce a concept of "coarse predictions", along with a sound numerical framework for computing them, and a series of tests that can be applied to assess their quality.

Matthias Morzfeld  
University of Arizona  
Department of Mathematics  
mmo@math.arizona.edu

Alexandre Fournier, Gauthier Hulot  
Institute de Physique du Globe de Paris  
fournier@ipgp.fr, gh@ipgp.fr

### MS172

#### Error Surrogates for Reduced-Order Models Based on Machine Learning Techniques

Characterizing the epistemic uncertainty introduced by substituting a full-order model with a reduced-order model (ROM) is crucial for developing rapid yet reliable uncertainty quantification (UQ) procedures involving time-dependent partial differential equations. We propose an automatic procedure for the construction of ROM error surrogate (ROMES) based on machine learning techniques, such as random forests and high-dimensional regression. First, a large number of additional inexpensive error indicators is computed at each ROM query. Then, high-dimensional regression methods identify the most relevant error indicators by relying on automatic feature selection; finally, selected features are adaptively combined to provide unbiased, low variance error surrogates. We show how the resulting ROMES can enhance the accuracy of UQ strategies exploiting ROMs.

Stefano Pagani  
Politecnico di Milano  
stefano.pagani@polimi.it

Kevin T. Carlberg  
Sandia National Laboratories  
ktcarlb@sandia.gov

Andrea Manzoni

EPFL, MATHICSE-CMCS  
Switzerland  
andrea.manzoni@epfl.ch

**MS173****Parallel Monolithic Multigrid Preconditioners for MHD**

Magnetohydrodynamics (MHD) models describe a wide range of plasma physics applications, from thermonuclear fusion in tokamak reactors to astrophysical models. These models are characterized by a nonlinear system of partial differential equations in which the flow of the fluid strongly couples to the evolution of electromagnetic fields. As a result, the discrete linearized systems that arise in the numerical solution of these equations are generally difficult to solve, and require effective preconditioners to be developed. We investigate monolithic multigrid preconditioners for a one-fluid, viscoresistive MHD model in two dimensions that utilizes a second Lagrange multiplier added to Faraday's law to enforce the divergence-free constraint on the magnetic field. We consider the extension of a well-known relaxation scheme from the fluid dynamics literature, Vanka relaxation, to this formulation, as well as approximate block-factorization smoothers. To isolate the relaxation scheme from the rest of the multigrid method, we utilize structured grids, geometric interpolation operators, Galerkin coarse grid operators, and inf-sup stable elements for both constraints in the system. Parallel numerical results are shown for the Hartmann flow problem, a standard test problem in MHD.

Thomas Benson

Lawrence Livermore National Laboratory  
benson31@llnl.gov

Scott Maclachlan  
Department of Mathematics and Statistics  
Memorial University of Newfoundland  
smaclachlan@mun.ca

James H. Adler  
Tufts University  
james.adler@tufts.edu

Raymond S. Tuminaro, Eric Cyr  
Sandia National Laboratories  
rstumin@sandia.gov, eric.c.cyr@gmail.com

**MS173****A Generalized Approach for Scalable N-Body Calculations**

Particle-based schemes play an important role in resolving kinetics in many plasma calculations. Nonetheless, their extension to complex domains and non-standard boundary conditions can pose a challenge to traditional approaches. In this talk a finite element-based approach to particle-particle particle-mesh (P3M) methods is presented. P3M methods work by splitting a potential into a quickly decaying short-range component and a smooth long-range component that is accurately resolved on the underlying mesh. Unlike traditional P3M methods, which utilize a Gaussian screen function to perform this decomposition, the method presented herein is constructed using screen functions based on a finite element basis, thus allowing generalizations to the method. One consequence is that this approach no longer depends on the FFT, but instead

utilizes multigrid methods which exhibit excellent scalability. Additionally, the finite element formulation facilitates a more general approach to handling boundary conditions. Numerical evidence is presented in support of this approach for a range of applications.

Scott High

Department of Computer Science  
shigh2@illinois.edu

Natalie N. Beams, Jonathan B. Freund  
University of Illinois at Urbana-Champaign  
beams2@illinois.edu, jbfreund@illinois.edu

Luke Olson

Department of Computer Science  
University of Illinois at Urbana-Champaign  
lukeo@illinois.edu

**MS173****Scalable Block Preconditioners for Maxwells Equations and Continuum Plasma Physics Models**

Multiple time-scales can arise in electromagnetic systems due to source currents, boundary effects, and coupling to other physics. In such situations, time-scales associated with light waves can be many orders of magnitude faster than time-scales of interest. When efficient linear solvers are available, implicit time integration can allow simulations to take large time-steps compared to the explicit limit. In this work, we develop block preconditioners that allow integration at slow time-scales of interest. We consider a compatible discretization of the Maxwell equations, using edge elements for the electric field and face elements for the magnetic field. We pursue block preconditioners that segregate edge and face variables so that existing multilevel solvers can be employed in subblocks. We propose a new augmentation-based approach to approximating the electric field Schur complement that requires only the application of traditional algebraic multigrid. We demonstrate the robustness and algorithmic scalability of our solution method on test problems that contain multiple realistic time-scales. Additionally, we extend these techniques to a multifluid model for plasma physics. This model couples a nodal discretization of the fluid equations to the edge-face electromagnetic discretization. We demonstrate that the augmentation-based approach for the electromagnetic subsystem allows for implicit time integration at time-scales much slower than the speed of light.

Edward G. Phillips

Sandia National Laboratories  
egphill@sandia.gov

John Shadid  
Sandia National Laboratories  
Albuquerque, NM  
jnshadi@sandia.gov

Eric C. Cyr  
Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

**MS173****Mimetic Finite Difference Schemes for Maxwell**

## Equations

We are interested in the efficient numerical solution of Maxwell's equations. Mimetic finite difference (MFD) schemes are considered. They are structure-preserving discretizations which work on general unstructured and irregular grids and result in discrete operators that satisfy the exact sequence connecting grad, div and curl operators on the continuous level. By finding the relationship between such MFD schemes and the finite element method, we are able to analyze the convergence of the MFD discretizations and construct efficient multigrid methods.

Carmen Rodrigo  
University of Zaragoza  
carmenr@unizar.es

James H. Adler  
Tufts University  
james.adler@tufts.edu

Francisco José Gaspar  
University of Zaragoza  
fjgaspar@unizar.es

Xiaozhe Hu  
Tufts University  
xiaozhe.hu@tufts.edu

Ludmil Zikatanov  
Pennsylvania State University  
ludmil@psu.edu

### MS174

#### A Second Order Finite Difference Method for Time-Space Fractional Diffusion Equations with Riesz Derivative

In this paper a second order finite difference method is presented to solve time-space fractional diffusion equations. Specifically, the central difference is used to approximate the Riesz fractional derivative in space. To obtain the second-order accuracy in temporal discretization for Caputo derivative, a trapezoidal formula is used to solve a system of Volterra integral equations transformed from spatial discretization. The stability and convergence of this method are proved. Finally numerical experiments are given.

Sadia Arshad  
COMSATS Institute of Information Technology, Pakistan  
sadia@lsec.cc.ac.cn

Jianfei Huang  
College of Mathematics, Qingdao University, Qingdao  
266071,  
China  
jfhuang@lsec.cc.ac.cn

Yifa Tang  
LSEC, ICMSEC, Academy of Mathematics and Systems  
Science,  
Chinese Academy of Sciences, Beijing 100190, China  
tyf@lsec.cc.ac.cn

### MS174

#### Highly Efficient Methods for Space-Fractional

## Reaction-Diffusion Systems

High-order time stepping methods for space-fractional systems of reaction-diffusion equations have been developed. A fourth-order compact spatial approximation is used in order to achieve fourth order accuracy in time and space. The methods are based on the exponential time differencing Runge-Kutta method and the Pade approximations of the matrix exponential functions. In order to avoid the multiplications of matrices with higher powers, a partial fractional approach is used. The stability and convergence of the methods are analyzed and numerical experiments for different fractional orders are conducted on several examples. These methods are proved to maintain high order accuracy in the case of non-smooth initial data. To enhance the computational efficiency, the short memory approach is utilized.

Khaled Furati  
Department of Mathematics & Statistics  
King Fahd University of Petroleum & Minerals  
kmfurati@kfupm.edu.sa

Mohammad Yousuf  
King Fahd University of Petroleum and Minerals  
Saudi Arabia  
myousuf@kfupm.edu.sa

Harish Bhatt  
Department of Mathematics and Statistics  
Texas A&M University-Corpus Christi  
harish.bhatt@tamucc.edu

Abdul Khaliq  
Middle Tennessee State University  
abdul.khaliq@mtsu.edu

### MS174

#### Fractional-Order Modeling and Simulation of the Human Ear System

We develop a physical-mathematical model of the human ear employing fractional-order lumped elements in order to account for viscoelastic characteristics of biological tissues in the human ear. In order to fit the parameters of the proposed model and assess its performance, we experimentally obtain the impedance estimates along with gain estimates using distortion product otoacoustic emissions (DPOAEs) from normal hearing adults. Employing fractional-order elements generalizes the traditional lumped element models and yields new attractive possibilities for fine tuning of the resulting models to better capture the viscoelastic response of biomaterials, exhibiting nonlocal and history effects.

Maryam Naghibolhosseini  
Michigan State University  
naghib@msu.edu

### MS174

#### Well-Balanced Central Fractional Derivatives in Reaction-Diffusion Modeling Applications

This discussion concerns variational properties of a much natural and balanced class of fractional derivative operators based on both left-sided and right-sided  $\alpha$ -th order formulations, where  $\alpha \in (\frac{1}{2}, 1)$ . Approximations of fractional differential equations equipped with such naturally balanced fractional derivatives are constructed via Ritz-Galerkin approaches. It is found that not only the novel

fractional derivative structures preserve desirable dynamic features, but also equations equipped with them possess expected variational properties when a Ritz-Galerkin formula is applied. Several simulation examples are given to illustrate our results.

Qin Sheng

Department of Mathematics  
Baylor University  
Qin.Sheng@baylor.edu

Haiwei Sun

University of Macau  
hsun@umac.mo

Yufeng Xu

Central South University  
China  
xuyufeng@csu.edu.cn

### MS175

#### Reducing Parallel Communication Costs in Algebraic Multigrid

Algebraic multigrid (AMG) is often viewed as a scalable solver for sparse linear systems, yet the method lacks parallel scalability due to high costs associated with communication. Sparse matrix-vector multiplication (SpMV), for example, is a dominant operation in each iteration of the solve phase of AMG and is communication bound, particularly at coarser levels of the AMG hierarchy. Indeed, as the sparsity pattern of a matrix increases in density, the communication costs of each SpMV also increases. As a result, performing a SpMV on a coarse level of the hierarchy can be more costly than the same operation on the original fine level, due to the increased communication costs. The communication costs associated with coarse levels of the AMG hierarchy can be greatly reduced through changes to both the AMG algorithm and the parallel implementation. The sparsity pattern of each coarse level can be altered to decrease the number of nonzeros, increasing the number of iterations required while limiting per-iteration cost. Furthermore, the parallel strategy for communication can be rearranged to limit the costly inter-node communication. In this talk we will highlight several strategies to reduce communication costs that are applicable to a range AMG methods and other sparse computations.

Amanda Bienz

University of Illinois at Urbana-Champaign  
bienz2@illinois.edu

William D. Gropp

University of Illinois at Urbana-Champaign  
Dept of Computer Science  
wgropp@illinois.edu

Luke Olson

Department of Computer Science  
University of Illinois at Urbana-Champaign  
lukeo@illinois.edu

### MS175

#### Algebraic Multigrid for Directed Graph Laplacian Systems

Currently there is a lack of robust multigrid solvers for directed graph Laplacian linear systems including directed

graphs with scale-free degree distribution. We seek to apply similar multigrid techniques used for undirected graph for directed graphs due to the robustness of multigrid solvers for undirected graphs Laplacian systems. We use Petrov-Galerkin restriction and prolongation operators to lift the restriction of symmetry, applied Lean Algebraic Multigrid affinity measurement for coarsening and we re-define low-degree elimination to be applicable for the directed case. We have performed preliminary tests with our native prototype on a variety of real-world directed graphs and our results show that our solver has promise to be robust.

Alyson Fox

University of Colorado Boulder  
fox.alyson@gmail.com

Geoff Sanders

Lawrence Livermore National Laboratory  
sanders29@llnl.gov

Thomas Manteuffel

University of Colorado  
tmanteuf@colorado.edu

### MS175

#### Nonlinear Algebraic Multigrid for the Power Flow Equations

The power flow problem is a standard problem in power engineering. It is a networked-structured, complex quadratic system of equations that today is typically solved via Newtons method. However, the Newtons method requirement to repeatedly solve large linear systems is becoming a computational bottleneck as problem sizes grow. Here we develop an algebraic multigrid approach to address this bottleneck.

We first discuss the power flow problem itself and point out its similarity with graph Laplacian linear systems, despite the fact that it is complex and nonlinear. Based on this similarity, we apply standard algebraic multigrid ideas, but using a multiplicative coarse-grid correction rather than the typical additive one. The result is a coarse problem that has the structure of a power flow problem, allowing for multilevel recursion. We also discuss a nonlinear Gauss-Seidel smoother. Experiments demonstrate the scalable nature of our approach.

Colin Ponce, Panayot Vassilevski

Lawrence Livermore National Laboratory  
ponce11@llnl.gov, vassilevski1@llnl.gov

David Bindel

Cornell University  
bindel@cs.cornell.edu

### MS176

#### Active Liquid Crystals and their Applications to Cell Motility

I will discuss a set of models for active liquid crystals that are derived following the generalized Onsager principle. Their predictions in various confined geometries will be discussed briefly. Then, I will apply these models to study cell motility. A multiphase complex fluid model incorporating the active liquid crystal as one of the active matter layer will be introduced and discussed. Numerical simulations using the model for cell migration on patterned substrates

will be presented.

Qi Wang  
University of South Carolina  
Beijing Computational Science Research Center  
qwang@math.sc.edu

#### MS176

##### **A Direct Proof on the Eigenvalue Preservation for the Beris-Edward System Modeling Liquid Crystals**

The Beris-Edward is a hydrodynamic system modeling nematic liquid crystals in the setting of Q-tensor order parameter. Mathematically speaking it is the incompressible Navier-Stokes equations coupled with a Q tensor equation of parabolic type. In our talk we revisit the simplified Beris-Edward system modeling nematic liquid crystals and study its eigenvalue preservation property. We give an alternative but direct proof to the eigenvalue preservation of the initial data for the Q-tensor. The advantage of this proof is that such result is not only valid in the whole space case, but the bounded domain case as well.

Xiang Xu  
Old Dominion University  
x2xu@odu.edu

#### MS176

##### **Numerical Approximations to a New Phase Field Model for Two Phase Flows of Complex Fluids**

We derive a new phase field theory for immiscible mixtures of nematic liquid crystals and viscous fluids using the variational principle coupled with the generalized Onsager principle. A novel phase transition mechanism is implemented to couple the nematic liquid crystal phase with the viscous fluid phase to arrive at the dissipative hydrodynamic model for incompressible fluid mixtures. Through a delicate explicit-implicit numerical discretization, we develop a decoupled, linear scheme for a simplified version of the phase field model, as well as a coupled, nonlinear scheme for the full model. Both schemes are shown as unconditionally energy stable with consistent, discrete dissipative energy laws. Several numerical examples are presented to show the effectiveness of the new model and the new numerical schemes developed for it.

Xiaofeng Yang  
University of South Carolina  
xfyang@math.sc.edu

#### MS176

##### **On Energy Stable Schemes for a Hydrodynamic Q-Tensor Model of Liquid Crystals**

The hydrodynamic Q-tensor model has been widely used for studying flows of liquid crystals and liquid crystal polymers. It is derived from a variational approach together with the generalized Onsager principle, in which the total energy decreases in time. In this talk, I will present a novel, linear, second order semi-discrete scheme in time to solve the governing system of partial differential equations in the model. The new scheme is developed following the energy quadratization strategy so that it is linear and unconditionally energy stable at the semi-discrete level. The convergence rate in time is established in a mesh refinement test. Several numerical examples are presented to

demonstrate the effectiveness of the model and the numerical scheme in simulating defect dynamics in flows of liquid crystals.

Jia Zhao  
University of North Carolina at Chapel Hill  
zhaojia@email.unc.edu

Xiaofeng Yang  
University of South Carolina  
xfyang@math.usc.edu

Qi Wang  
University of South Carolina  
Beijing Computational Science Research Center  
qwang@math.sc.edu

#### MS177

##### **Fast Batched SVD on GPUs**

Hierarchical matrices are an efficient way for storing the dense matrices of very large dimension that appear in the discretization of integral operators associated with elliptic PDEs, in Schur complement methods exploiting dimension reduction, in spatial statistics and computational astronomy when describing pairwise relations between data points, etc. They exploit the fact that the underlying matrices, while formally dense, are data sparse. They have a structure consisting of blocks, many of which can be well-approximated by low rank factorizations, resulting in compressing the dense matrix in an accuracy-controlled manner. Singular value decomposition of these very small hierarchical blocks is a crucial operation and should be optimized by ensuring high occupancy on throughput-oriented manycore architectures. This representation can avoid superlinear growth in memory requirements to store  $n \times n$  dense matrices in a scalable manner requiring  $O(n)$  units of storage with a constant depending on the representative rank  $k$  for the low rank blocks.

Wajih Halim Boukaram, David E. Keyes  
KAUST  
wajihhalim.boukaram@kaust.edu.sa,  
david.keyes@kaust.edu.sa

Hatem Ltaief  
King Abdullah University of Science & Technology (KAUST)  
hatem.Ltaief@kaust.edu.sa

George M Turkiyyah  
King Abdullah University of Science and Technology  
American University of Beirut  
gt02@aub.edu.lb

#### MS177

##### **Some Notes on Divide-and-Conquer Eigensolvers**

Most existing divide-and-conquer eigensolvers for symmetric band matrices (BDC eigensolvers) are based on transforming rank- $k$  modification problems arising in this context into a sequence of rank-one modification problems. As a consequence, the aggregation of the eigenvector information tends to become the computational bottleneck. We will discuss advances in two aspects: On the one hand, we will summarize new insights resulting from performance and accuracy comparisons with algorithmic variants which handle the rank- $k$  modification problems as a whole. On

the other hand, we will present and compare various algorithmic modifications of the divide-and-conquer strategy for computing eigenvalues *only*.

Wilfried N. Gansterer

Department of Computer Science  
University of Vienna  
wilfried.gansterer@univie.ac.at

Michael Moldaschl, Kastor Felsner  
University of Vienna  
michael.moldaschl@univie.ac.at, kas-  
tor.felsner@univie.ac.at

**MS177**

**Accelerating the Singular Value Decomposition with a Hybrid Two-Stage Algorithm**

The classic algorithm for computing the Singular Value Decomposition (SVD) of a matrix first reduces it to bidiagonal form, then applies an iterative algorithm to reduce it to diagonal form. The performance of the one-stage reduction to bidiagonal is well known to be memory bandwidth limited, because of the use of matrix-vector products (GEMV). Due to the growing gap between memory bandwidth and computation speed, two-stage reductions were developed that first reduce to band form, then reduce from band form to bidiagonal form. These remove the bandwidth limitation by concentrating operations in matrix-matrix products (GEMM). Also, when computing singular vectors, the divide-and-conquer algorithm has improved the performance of the SVD. Here, we investigate the use of accelerators such as GPUs and the Intel Xeon Phi to speed up SVD computations. We demonstrate significant speedups compared to one-stage implementations and CPU-only two-stage implementations. All three phases of the SVD algorithm are accelerated: reduction to bidiagonal, bidiagonal SVD using divide-and-conquer, and computation of singular vectors.

Mark Gates

Computer Science Department  
University of Tennessee  
mgates3@icl.utk.edu

Azzam Haidar  
Department of Electrical Engineering and Computer Science  
University of Tennessee, Knoxville  
haidar@icl.utk.edu

Jack J. Dongarra  
University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

**MS177**

**Efficient Transformation of the General Eigenproblem with Symmetric Banded Matrices to a Banded Standard Eigenproblem**

The solution of symmetric eigenproblems plays a key role in many computational simulations. Generalized eigenproblems are transformed to a standard problem. This transformation has the drawback that for banded matrices in the generalized eigenproblem the banded structure is not preserved. The matrix of the standard eigenproblem will generally be a full matrix. We followed the ideas of the Group of Lang (University of Wuppertal) who modified

Crawford's algorithm and implemented this procedure to ELPA. By keeping the banded structure we save one reduction step on the matrix and one backtransformation step for the eigenvectors. This provides a good speedup compared to the standard transformation procedure with Cholesky factorization.

Michael Rippl

Technische Universität München  
Germany  
rippm@in.tum.de

Thomas K. Huckle  
Institut fuer Informatik  
Technische Universitaet Muenchen  
huckle@in.tum.de

Bruno Lang  
University of Wuppertal  
lang@math.uni-wuppertal.de

**MS178**

**Continuous Alternating Least Squares for Regression of Low-Rank Functions**

We present a nonlinear regression algorithm for creating low-rank functional approximations from sampled data. The resulting approximation is represented in the function-train format, a recently developed continuous analogue to the tensor-train decomposition. The representation is obtained using a rank-revealing algorithm to solve a corresponding optimization problem. In contrast to other low-rank functional decompositions, our approximation can be parameterized in a manner where the output can be nonlinear with respect to unknown parameters. The flexibility of this approach provides a greater capability to adapt to structure within the data. We investigate the effect of several sampling strategies on the resulting approximation quality, and demonstrate their performance in the context of an elliptic PDE with an uncertain permeability field.

Alex A. Gorodetsky

Massachusetts Institute of Technology  
goroda@mit.edu

John D. Jakeman  
Sandia National Labs  
jadjakem@sandia.gov

**MS178**

**Low Rank Approximation and Integration of High Dimensional Functions for Quantum Chemistry Applications**

A method is proposed for a fast evaluation of high-dimensional integrals of potential energy surfaces (PES) that arise in many areas of quantum dynamics/chemistry. Our approach decomposes the integrand into suitable low-rank tensor formats using a set of integrand evaluations. The high dimensional integration problem is then reduced to a relatively short sum of products of easy-to-evaluate low-dimensional integrals, each of which can be estimated using appropriate quadrature rules. The decomposition is achieved by suitable version of alternating least squares algorithm depending on the type of tensor format considered. This approach also eradicates a force-constant evaluation as the hotspot of many quantum dynamics simulations and also delays the curse of dimensionality. In this work, we specifically consider canonical and tensor train formats for

approximation. This general method is applied to two applications in quantum chemistry.

Prashant Rai  
Sandia National Laboratories  
Livermore, CA, USA  
pmrai@sandia.gov

Khachik Sargsyan  
Sandia National Laboratories  
ksargsy@sandia.gov

Habib N. Najm  
Sandia National Laboratories  
Livermore, CA, USA  
hnnajm@sandia.gov

Matthew Hermes, So Hirata  
University of Illinois Urbana Champagne  
morilore@gmail.com, sohirata@illinois.edu

### MS178

#### A New Sampling Method for Alternating Least-Squares Regression

One approach to regression of scattered data in high dimensions uses alternating least-squares (ALS) to construct separated representations (or canonical decompositions). Such functional representations are built via a sequence of univariate regression problems corresponding to the individual dimensions of the data. Despite recent applications of separated representations to non-intrusive algorithms for uncertainty quantification, there is currently no established theory for the selection of optimal sampling distributions. Inspired by optimal sampling of multivariate orthogonal polynomials, we propose a sampling method that improves the stability and accuracy of ALS relative to standard sampling approaches where the quantity of interest is sampled according to the probability distribution of the input variables. We present numerical examples to demonstrate the improvement achieved by this sampling method with no additional costs.

Matthew Reynolds  
University of Colorado, Boulder  
matthew.reynolds@colorado.edu

Alireza Doostan  
Department of Aerospace Engineering Sciences  
University of Colorado, Boulder  
Alireza.Doostan@Colorado.EDU

### MS179

#### Distributed-Memory Graph Ordering

Ordering vertices of a graph is key to minimize fill-in and data structure size in sparse direct solvers, maximize locality in iterative solvers, and improve performance in graph algorithms. Except for naturally parallelizable ordering methods such as nested dissection, many important ordering methods (e.g., reverse Cuthill-McKee and minimum degree) have not been efficiently mapped to distributed-memory architectures. We report on our progress in developing scalable parallel implementations of such challenging graph-ordering algorithms.

Ariful Azad, Aydin Buluc  
Lawrence Berkeley National Laboratory  
azad@lbl.gov, abuluc@lbl.gov

Mathias Jacquelin  
Lawrence Berkeley National Lab  
mjacquelin@lbl.gov

Esmond G. Ng  
Lawrence Berkeley National Laboratory  
egng@lbl.gov

### MS179

#### Impact of Blocking Strategies for Sparse Direct Solvers on Top of Generic Runtimes

Among the preprocessing steps of a sparse direct solver, reordering and block symbolic factorization are two major steps to reach a suitable granularity for BLAS kernels efficiency and runtime management. In this talk, we present a reordering strategy to increase off-diagonal block sizes. It enhances BLAS kernels and allows to handle larger tasks, reducing runtime overhead. Finally, we will comment the resulting gain in the PaStiX solver implemented over STARPU and PARSEC.

Gregoire Pichon  
INRIA  
gregoire.pichon@inria.fr

Mathieu Faverge  
Bordeaux INP - Labri - Inria  
mathieu.faverge@inria.fr

Pierre Ramet  
Bordeaux University - INRIA  
pierre.ramet@inria.fr

Jean Roman  
INRIA  
Jean.Roman@inria.fr

### MS179

#### Refined Ordering in Supernodes to Improve Block Structure and Data Locality

In this work, we propose a low complexity heuristic which aims at improving the block structure of non zeros below the diagonal blocks by bringing non-zero rows closer together. This is of particular importance to achieve higher computational efficiency on GPU platforms for instance. We compare our heuristic to the state-of-the-art heuristic based on solving the Traveling Salesman Problem in terms of reordering time, quality of the block structure, and finally impact on numerical factorization time.

Mathias Jacquelin  
Lawrence Berkeley National Lab  
mjacquelin@lbl.gov

Barry Peyton  
Dalton State College  
bpeyton@daltonstate.edu

Esmond G. Ng  
Lawrence Berkeley National Laboratory  
egng@lbl.gov

### MS179

#### Reducing the Complexity of the Minimum Local

**Fill Algorithm**

The minimum local fill (MLF) heuristic computes a fill-reducing permutation for sparse Cholesky factorization. It is generally believed that MLF is very expensive and produces marginally better permutations than the well-known minimum degree (MD) heuristic. Recently we have introduced implementations of MLF that are effective in reducing fill but are much less time-consuming than conventional implementations of the past. In this talk, we will describe further enhancements that result in more efficient implementations of the MLF heuristic. The first enhancement is a technique to compute the required local fill counts in a more indirect, and ultimately, more efficient manner. The second enhancement is a technique for limiting the number of times entire elimination cliques are searched during the computation of Reach sets.

Barry Peyton  
Dalton State College  
bpeyton@daltonstate.edu

Mathias Jacquelin  
Lawrence Berkeley National Lab  
mjacquelin@lbl.gov

Esmond G. Ng  
Lawrence Berkeley National Laboratory  
egng@lbl.gov

**MS180****Modelling Liquid-Vapor Phase Change with Metastability**

We propose a model that describes phase transition including metastable states present in the van der Waals Equation of State. From a convex optimization problem on the mixture energy, we deduce a dynamical system that is able to depict all the stable equilibria as attraction points including metastable states. The dynamical system is then used as a relaxation source term in an isothermal 44 two-phase model. Although the model turns to be conditionally hyperbolic, we manage to prove that, for smooth solutions, the hyperbolicity regions are invariant domains of the system with relaxation. This is a joint work with Hala Ghazi (Univ. of Nantes) and Francois James (Univ. of Orleans).

Helene Mathis  
University of Nantes  
helene.mathis@univ-nantes.fr

**MS180****High Order Maximum Principle Preserving Semi-Lagrangian Finite Difference Weno Schemes for the Vlasov Equation**

Abstract not available

Jingmei Qiu  
University of Houston  
jingqiu@math.uh.edu

**MS180****Low-Dissipation Centred Fluxes for Discontinuous Galerkin Finite Element Methods**

We propose a class of centred (no explicit upwinding) numerical fluxes that depend on a parameter  $\alpha$  and orig-

inate from the multidimensional extension of the FORCE scheme on unstructured meshes. It turns out that such a parameter generates an  $\alpha$  family of numerical methods that are well suited for use in DG methods, as demonstrated here. Given an order of accuracy  $N$  and a corresponding CFL number  $C_N$ , there exists a real number  $\alpha_N$  that exactly matches the stability condition of DG, for any order of accuracy  $N$ . The resulting schemes, in their first order mode, are monotone and have surprisingly small dissipation, which is particularly evident for slowly moving waves. We shall present details on the construction, analysis and assessment of the numerical methods, for linear model problems and for non-linear systems. The performance of the schemes is indeed very encouraging, showing in particular, that the performance of DG does depend on the numerical flux and that the centred fluxes proposed here are well suited to the DG approach.

Eleuterio F. Toro, Beatrice Saggiorato  
University of Trento  
eleuterio.toro@unitn.it, saggiorato.beatrice@gmail.com

Svetlana Tokareva  
University of Zurich  
svetlana.tokareva@math.uzh.ch

**MS180****Positivity-Preserving Well-Balanced Discontinuous Galerkin Method for Tidal Bores**

Abstract not available

Yulong Xing  
University of California, Riverside  
xingy@ucr.edu

**MS181****Modeling the 1976 Teton Dam Failure using GeoClaw**

Using the GeoClaw software package, we model the 1976 Teton Dam failure that led to massive flooding in eastern Idaho. The flood boundaries reported by GeoClaw show excellent agreement with the boundaries recorded by the USGS at the time of the event and with historical descriptions of the dam failure. Performance results for recent work on developing a GeoClaw extension of ForestClaw, a parallel library based on p4est for solving PDEs on adaptive Cartesian meshes will also be reported.

Donna Calhoun  
Boise State University  
donna.calhoun@boisestate.edu

Melody Shih  
Columbia University  
yoyoshih13@gmail.com

Kyle T. Mandli  
Columbia University  
Applied Physics and Applied Mathematics  
kyle.mandli@columbia.edu

**MS181****Many Waves on Many Cores**

ExaHyPE is a H2020 project where an international consortium of scientists write a simulation engine for hyper-



bolic equation system solvers based upon the ADER-DG paradigm. Two grand challenges are tackled with this engine: long-range seismic risk assessment and the search for gravitational waves emitted by binary neutron stars. The code itself is based upon a merger of flexible spacetime data structures with highly optimised compute kernels for the majority of the simulation cells. This talk is a tour de force through ExaHyPE. We start with a brief sketch of the project and the underlying workflow of the ADER-DG scheme. It becomes obvious that the scheme - if not implemented naively - suffers significantly from high memory and data movement demands on modern architectures. In the main part of the talk, we thus highlight various techniques how to reduce the number of data movements to and from the memory subsystem, how to squeeze out bits from the MPI data exchange and how to eliminate some MPI messages completely. A brief discussion of first results closes the talk.

Dominic E. Charrier, Tobias Weinzierl  
Durham University  
dominic.e.charrier@durham.ac.uk,  
tobias.weinzierl@durham.ac.uk

### MS181

#### Performance and Time to Solution of Dynamically Adaptive Tsunami Simulations

We attempt a comparison of time-to-solution for solving the shallow water equations in the context of tsunami simulation. Using several benchmark and tsunami scenarios, we discuss the performance of a Finite Volume scheme on tree-structured triangular meshes that may include regular grid patches. We will also compare the Finite Volume performance with recent results using discontinuous Galerkin methods. We consider different HPC platforms including the Salomon supercomputer, where heterogeneous compute nodes (with Xeon Phi coprocessors) pose additional challenges on the load balancing. Here, we particularly report on experiences with using the Xeon Phi coprocessors in symmetric mode.

Chaulio Ferreira, Leonhard Rannabauer, Michael Bader  
Technical University of Munich  
chaudio.ferreira@tum.de, rannabau@in.tum.de,  
bader@in.tum.de

### MS181

#### Metaprogramming for Unstructured-Mesh Applications in Ocean Modeling

Recent advances in computer hardware technology make more realistic ocean simulations possible; however, this progress has a price. On the one hand, the growing complexity of the physical and mathematical models requires the development of new and efficient numerical methods. On the other hand, the trend towards heterogeneous and highly parallel architectures increases the programming effort necessary to implement, develop, and maintain these models. We propose one possible approach both to significantly improve the computational performance of the discontinuous Galerkin finite element method and, at the same time, to decrease the effort to implement new methods and applications. The first goal will be reached by utilizing a number of techniques aiming to address the existing performance bottlenecks; those include a new approach to domain decomposition and performance oriented adaptive numerical schemes. The second goal will be achieved by resorting to modern software design strategies, including

code generation and automatic optimization techniques for the compute-intensive kernels. The performance and code design advances resulting from the work proposed here can cover a lot of ground towards making unstructured mesh models the mainstream of ocean science and, in particular, available to users with limited access to HPC systems.

Harald Koestler  
University Erlangen-Nürnberg  
harald.koestler@fau.de

Sebastian Kuckuk  
Universitaet Erlangen-Nuernberg, Germany  
sebastian.kuckuk@fau.de

### MS182

#### Vectorization of Multi-Body Potentials: Performance and Portability

As today's supercomputers become more and more powerful, simulations can cover bigger length-scales and time-scales using more accurate, but also more expensive force fields. In the materials science community, many-body potentials are widely used for their predictive power with respect to certain material properties, at the expense of higher computational cost. The challenge lies in mapping the complex calculations necessary to evaluate such potentials onto the available computing devices. Since modern architectures concentrate the computational power in wide SIMD units, and compilers commonly have trouble generating efficient code for them, a dedicated optimization effort is necessary. Special care is needed to minimize the effort required to implement a potential on a new architecture, and to allow for portability at the algorithmic level. Our research provided insights in the vectorization of the Tersoff, REBO and AIREBO potentials, widely used for semiconductor, carbon material, and carbohydrate simulations. We target a diverse set of hardware ranging from x86 CPUs (Westmere to Skylake), to Xeon Phi accelerators of both generations, and even GPUs. The improvements typically double the simulation throughput in large-scale, parallel runs, and higher speedups are possible when deploying accelerators.

Paolo Bientinesi, Markus Höhnerbach  
AICES, RWTH Aachen  
pauldj@aices.rwth-aachen.de,  
hoehnerbach@aices.rwth-aachen.de

Ahmed E. Ismail  
Department of Mechanical Engineering  
RWTH Aachen University  
ahmed.ismail@mail.wvu.edu

### MS182

#### Multilevel Summation Method for Efficient Computation of Nonbonded Forces in Molecular Dynamics Simulations

The multilevel summation method (MSM) provides an efficient algorithm for computing the long-range nonbonded electrostatic forces arising in molecular dynamics (MD) simulations. The method imposes a splitting on the  $1/r$  interaction kernel, calculating the short-range part exactly and interpolating the remaining long-range parts from a hierarchy of nested grids, where interpolation between grid levels produces an algorithm with operation count that scales linearly in the number of atoms. The methodology is quite flexible, permitting use for simulations that

employ non-periodic, semi-periodic, or periodic boundary conditions, and allowing generalization to other interaction kernels, most notably the long-range dispersion forces. The use of B-spline interpolation is shown to improve the accuracy of the approximation by an order of magnitude over previous formulations. The parallelization of MSM in the MD program NAMD makes use of a combined domain decomposition and force decomposition approach to provide scaling to large numbers of processors, together with fine-grained SIMD (single instruction, multiple data) parallelism employed to greatly enhance the performance of the localized convolution calculations.

David J. Hardy  
Theoretical Biophysics Group, Beckman Institute  
University of Illinois at Urbana-Champaign  
dhardy@illinois.edu

Robert D. Skeel  
Department of Computer Science  
Purdue University  
rskeel@purdue.edu

## MS182

### Particle-Particle Particle-Mesh (P3M) on Knights Landing Processors

Particle-particle particle-mesh methods are often used in molecular dynamics codes to approximate the effects of long-range forces between atoms where it would not be feasible to compute all pair-wise interactions. While short-range interactions are computed in a pair-wise fashion, the forces produced by long-range interactions are obtained by mapping particle charge to a grid, solving Poisson's equation in the frequency domain for the electrical potential, and then mapping the local potential back to the particles. Using the popular molecular dynamics code LAMMPS, we present vectorization and new implementations of the two mapping algorithms. We also discuss how using larger stencil sizes when mapping charges and forces better takes advantage of the Xeon Phi architecture, both by making use of its large vector registers and because a larger stencil allows a coarser grid to be used. This shifts work from the poorly-scaling FFTs used to solve Poisson's equation and to the newly-accelerated and highly parallel mapping functions. The acceleration of the PPPM method as a whole also affects the optimal input parameters in a similar fashion; using a smaller cutoff to shift work from the pair-wise short-range computation to the long-range PPPM computation saves time even while using a finer charge grid to preserve accuracy.

William J. Mcdoniel, Paolo Bientinesi  
AICES, RWTH Aachen  
mcdoniel@ices.rwth-aachen.de,  
pauldj@ices.rwth-aachen.de

Ahmed E. Ismail  
Department of Mechanical Engineering  
RWTH Aachen University  
ahmed.ismail@mail.wvu.edu

Markus Höhnerbach  
AICES, RWTH Aachen  
hoehnerbach@ices.rwth-aachen.de

Rodrigo Canales  
RWTH Aachen University

rodrigo.canales@rwth-aachen.de

## MS182

### Towards Autotuning Between OpenMP Schemes for Molecular Dynamics on Intel Xeon Phi

Shared-memory performance is a key component for obtaining good scalability on manycore architectures such as the Intel Xeon Phi. In the area of Molecular Dynamics, the performance of the non-bonded short-range force calculation kernel is crucial to the overall runtime of the simulation. The characteristics of this kernel, however, vary strongly. Systems with a small number of particles need a very fine-grained OpenMP scheme. At the same time, fine-grained schemes may be inferior to more coarse-grained ones for larger numbers of molecules. Particle density, the number of interaction sites per molecule and the type of potential also affect the required granularity. The spatial distribution, on the other hand, may impose a need for load-balancing, while the runtime environment may place restrictions on the available memory for parallelization-related buffers that some schemes require. In the present work we evaluate several OpenMP schemes for the force calculation on multiple simulation systems via ls1-mardyn [Niethammer et al., ls1-mardyn: The Massively Parallel molecular dynamics code for large systems, 2014], [Tchipev et al., Optimized Force Calculation of Molecular Dynamics Simulations for the Intel Xeon Phi, 2015] on up to 256 threads. We observe that none of the evaluated schemes can deliver the best results across all simulation parameters and outline our work on autotuning, in order to not burden the end-user with a manual choice of the best scheme.

Nikola P. Tchipev, Andrei Costinescu, Steffen Seckler  
Technical University of Munich  
Department of Informatics  
tchipev@in.tum.de, andrei.costinescu@tum.de,  
steffen.seckler@tum.de

Philipp Neumann  
Technische Universität München, Germany  
Scientific Computing in Computer Science  
neumanph@in.tum.de

Hans-Joachim Bungartz  
Technical University of Munich, Department of Informatics  
Chair of Scientific Computing in Computer Science  
bungartz@in.tum.de

## MS183

### Detecting and Attributing Application Performance Anomalies to Root Causes

Resources such as the filesystem and interconnection network are shared by all jobs running on a supercomputer and sometimes, even by multiple supercomputers at a facility. Understanding the performance variability of individual applications and attributing it to its root causes is crucial to optimize application performance and overall system throughput. Identifying the root causes can help us in developing better algorithms for resource-aware scheduling, placement and routing. In this talk, I will present recent work on monitoring, visualizing and analyzing performance variability on different platforms.

Abhinav Bhatele  
Lawrence Livermore National Laboratory  
bhatele@llnl.gov

Staci Smith, David K. Lowenthal  
University of Arizona  
smiths949@email.arizona.edu, dkl@cs.arizona.edu

Vanessa Cedeno  
Virginia Tech  
vcedeno@vbi.vt.edu

Todd Gamblin  
Lawrence Livermore National Laboratory  
tgamblin@llnl.gov

### MS183

#### **TOKIO: Using Lightweight Holistic Characterization to Understand, Model, and Improve HPC I/O Performance**

I/O performance is playing an increasingly large role in scientific application efficiency, not only because of the gap between storage and computational hardware performance, but also because scientific discovery is more data-driven now than ever before. I/O performance variance and predictability are therefore major concerns for HPC data centers. Why are some applications more efficient than others, why does performance vary from run to run, and what can we do about it? A myriad of factors including application I/O strategies, resource contention, resource allocation, and bursty I/O traffic may contribute to the problem, and focusing too deeply on any single factor might lead us to overlook true root causes and critical correlations. In this talk, we will explore this issue with a particular focus on the TOKIO (TOtal Knowledge of I/O) project, a collaboration between Lawrence Berkeley National Laboratory and Argonne National Laboratory. The goal of TOKIO is to advance I/O instrumentation techniques, data integration, and analysis to improve our understanding of I/O in the datacenter and produce actionable insight for the scientific computing community.

Philip Carns  
Argonne National Laboratory  
carns@mcs.anl.gov

### MS183

#### **Understanding and Avoiding Performance Variability in High Performance Networks**

Performance variability on large systems is manifested from multiple different sources. One important subsystem that can introduce performance variability is the interconnection network. In this talk I will discuss work covering two different areas of understanding network induced performance variability. The first area is the network fabric itself and how inter-job interference can create performance variability. Overtime, a tool for studying fabric contention that can impact performance, was developed in order to better understand the impact of other jobs usage of the network as well as understanding the impact of the most frequently used application communication patterns in the presence of interference on real large systems under study. The second area is how network traffic can cause contention for local resources at the node level, causing slowdown due to memory subsystem contention. This work concentrates on onloaded vs. offloaded networking interfaces over a wide variety of generations of x86 processor servers. Overall, it is shown that such interference can have great impact at scale even on applications known to have low performance

variability.

Ryan Grant  
Sandia National Laboratories and University of New Mexico  
regrant@sandia.gov

Taylor Groves  
University of New Mexico  
tgroves@sandia.gov

Kevin Pedretti, Ann Gentile  
Sandia National Laboratories  
ktpedre@sandia.gov, gentile@sandia.gov

Dorian Arnold  
University of New Mexico  
darnold@cs.unm.edu

### MS183

#### **Resource Management Techniques for Reducing Interference Among Message-Passing Applications**

Interference among nearby jobs has been recently identified as the dominant reason for the high performance variability of message-passing applications running on High Performance Computing (HPC) systems. Typically, HPC systems are dynamic with multiple jobs coming and leaving in an unpredictable fashion, sharing simultaneously the system interconnection network. In such environment contention for network resources is causing random stalls in the progress of application execution degrading application and system performance overall. Eliminating job interactions is the key for guaranteeing performance predictability of applications. In order to achieve this objective we propose resource management techniques at different levels of decision: 1) job scheduling policy that controls the node fragmentation in the system, and reduces the number of applications sharing the same part of the network, and 2) QoS policy that allows for fair sharing of the network link by multiple HPC applications of different bandwidth requirements using virtual channels arbitration enabled in modern interconnects (e.g. InfiniBand). Experimental results show that the proposed job scheduling policy leads to few jobs sharing network resources and thus having fewer jobs interactions while the QoS policy is able to effectively reduce the degradation from the remaining jobs interactions. These two software techniques are complementary and could be used together without additional hardware.

Ana Jokanovic, Jose Carlos Sancho  
Barcelona Supercomputing Center  
ana.jokanovic@bsc.es, jose.sancho@bsc.es

German Rodriguez  
Rockley Photonics  
german.rodriguez@rockleyphotonics.com

### MS184

#### **Massive-Scale Streaming Analytics for Dynamic Graphs**

Emerging real-world graph problems include: detecting community structure in large social networks; improving the resilience of the electric power grid; and detecting and preventing disease in human populations. Unlike traditional applications in computational science and engineering, solving these problems at scale often raises new chal-

allenges because of the sparsity and lack of locality in the data, the need for additional research on scalable algorithms and development of frameworks for solving these problems on high performance computers, and the need for improved models that also capture the noise and bias inherent in the torrential data streams. In this talk, the speaker will discuss the opportunities and challenges in massive data-intensive computing for applications in computational science and engineering.

David A. Bader

Georgia Institute of Technology  
bader@cc.gatech.edu

#### MS184

##### Models for Principled Characterization of Dynamic, Spatially Embedded, Multiscale Networks

Advances in neuroimaging techniques have made it possible to reconstruct whole-brain networks composed of structural (physical fiber tracts) and functional (statistical relationships) connections among brain regions. Analysis of these static networks has revealed a host of non-random attributes, including highly connected hubs, modular architecture, and rich clubs. In this talk I will discuss two recent advances in brain network modeling. First, I will introduce the multi-layer network model for characterizing time-varying functional brain networks. I will cover recent work showing that the brains flexibility the extent to which its multi-layer modular organization is stable across time can be used to predict an individuals learning rate, is associated with executive function and state of arousal, and is also altered in psychiatric diseases such as schizophrenia. Second, I will discuss the role that the brains intrinsic geometry plays in shaping its network architecture. I will cover two recent studies one in which we show that simple wiring rules can explain a wide variety of the brains topological features and another in which we modify classical community detection tools to uncover space-independent community structure in brain networks.

Danielle Bassett

School of Engineering and Applied Science  
University of Pennsylvania  
dsb@seas.upenn.edu

Richard betzel  
Dept. Bioengineering  
University of Pennsylvania  
rbetzel@seas.upenn.edu

#### MS184

##### Dynamic Network Analysis: From Inference to Insight

From gene interactions and brain activity to cellphone calls and zebras grazing together, large, noisy, and highly dynamic networks of interactions are everywhere. Unfortunately, in this domain, our ability to analyze data lags substantially behind our ability to collect it. Moreover, we may be collecting the wrong data for the questions we want to answer in the first place. From collecting the data and inferring the networks to producing meaningful insight at scale, challenges are there every step of the way and computational approaches have been developed to meet those challenges. We will present computational approaches that address some of the questions about dynamic interaction networks: whom should we sample? how often? what is the "right" network? what are the meaningful patterns and

trends? and how can we use the network to gain insight into other aspects of the node behavior? The methods leverage the topological graph structure of the networks and the size of the available data to, somewhat counter-intuitively, to produce more accurate results faster. We will demonstrate the scientific implications of the computational analysis on networks of zebras, baboons, and interacting brains cells.

Tanya Y. Berger-Wolf

University of Illinois at Chicago  
tanyabw@uic.edu

#### MS184

##### Scalable Algorithms for Graph Matching and Edge Cover Computations

Computing a matching and an edge cover in a graph are two basic problems in computer science, with many applications in network science, machine learning, computational science and engineering, etc. Most variants of matching and edge cover problems can be solved in polynomial time, yet the running times are high and the algorithms are sophisticated. It is even more challenging to design parallel algorithms, since many algorithms rely on searching for long paths in a graph, or implicitly communicate information along long paths, and thus have little concurrency. We design new approximation algorithms for variant matching and edge cover problems that have high concurrency. For b-Matching, we obtain a 1/2-approximation algorithm; and for b-Edge cover, we describe a 3/2-approximation algorithm. Both algorithms have been implemented on shared memory and distributed memory machines, and we report results from tens of thousands of cores. We also discuss some applications of b-Matchings and b-Edge covers.

Alex Pothén

Purdue University  
Department of Computer Science  
apothén@purdue.edu

Arif Khan  
Purdue University  
khan58@purdue.edu

Mostofa Patwary  
Northwestern University  
mostofa.ali.patwary@intel.com

Pradeep Dubey  
Intel Labs  
pradeep.dubey@intel.com

#### MS185

##### Multifidelity Moment Estimation in Coupled Multidisciplinary Systems Under Uncertainty

The objective of this work is to tackle the complexities involved with uncertainty propagation in feedback coupled multidisciplinary systems. We present a multifidelity Monte Carlo approach to estimate the first and second moments of the system output. The method employs control variates to exploit multiple surrogate models and reduce the variance in the final estimate as compared to standard Monte Carlo simulation. We also explore several ways to build different fidelities for the system output.

Anirban Chaudhuri, Karen E. Willcox  
Massachusetts Institute of Technology

anirbanc@mit.edu, kwillcox@MIT.EDU

### MS185

#### Modeling and Identification of Multi-Physics Uncertainties in the Response of Heated Structures

This investigation is focused on the response of uncertain heated structures, i.e., for which both structural displacements and temperature fields must be determined with the former strongly affected by the latter. Projection-based, nonlinear, reduced order modeling techniques have recently been devised to address this class of multi-physics problems at a much reduced cost than full finite element computations. In this light, the present effort centers on the introduction, directly in these ROMs, of uncertainty in the disciplinary properties (structural and thermal ones) but also on the coefficient of thermal expansion field connecting both disciplines. It is first recognized that the linear and nonlinear stiffness structural properties and the thermal terms in the structural equations can be regrouped in a symmetric positive definite matrix. Then, maximum entropy concepts are utilized to formulate a stochastic model of this uncertain matrix and of the corresponding conductance and capacitance matrices. Next, the identification of the hyperparameters of this stochastic model is next addressed from measurements of the coupled thermal-structural response. The application of these concepts is finally demonstrated on a representative hypersonic aircraft panel example.

Pengchao Song, X.Q. Wang  
Arizona State University  
pengchao.song@asu.edu, xiaoquan.wang.1@asu.edu

Marc P. Mignolet  
Arizona State  
marc.mignolet@asu.edu

### MS185

#### Sparsity Identification in Non-Gaussian Distributions, with Applications to Multi-Disciplinary Systems

In a Markov network, represented as an undirected probabilistic graphical model, the lack of an edge denotes conditional independence between the corresponding nodes. Sparsity in the graph is of interest as it can accelerate inference. We develop an algorithm to identify the sparsity structure of continuous high-dimensional non-Gaussian distributions. The algorithm relies on exploiting the connection between the sparsity of the graph and the sparsity of transport maps, which deterministically couple one probability measure to another. We present results for datasets drawn from multi-disciplinary engineering systems.

Rebecca Morrison, Youssef M. Marzouk  
Massachusetts Institute of Technology  
rmorriso@mit.edu, ymarz@mit.edu

### MS185

#### Lookahead in Single and Multiple Fidelity Bayesian Optimization

Bayesian optimization (using a function prior) is usually implemented using a greedy search policy. Expected improvement, for example, heuristically balances exploration and exploitation. Lookahead policies, alternately, decide

evaluations knowing future optimization behavior. This talk compares lookahead and greedy policies, explores domains where lookahead excels, and prove that lookahead can be arbitrarily better than greedy. We show exact numeric results for lookahead in standard GP optimization, and results in multifidelity optimization.

Brendan Tracey  
Stanford University  
tracey.brendan@gmail.com

David Wolpert  
Santa Fe Institute  
dhw@santafe.edu

### MS186

#### MHD with Embedded Particle-in-Cell (MHD-EPIC): Capturing Kinetic Effects in Global Simulations

A new modeling capability to embed the implicit Particle-in-Cell (PIC) model iPIC3D into the extended magneto-hydrodynamic model BATS-R-US has been developed recently in University of Michigan. The PIC domain can cover the regions where kinetic effects are most important, such as the magnetopause reconnection sites. The BATS-R-US code, which is a MHD model, with its block-adaptive grid can efficiently handle the rest of the computational domain where the MHD or Hall MHD description is sufficient. The current implementation of the MHD-EPIC model allows two-way coupled simulations in two and three dimensions with multiple embedded PIC regions. The MHD and PIC grids can have different grid resolutions and grid structures. The MHD variables and the moments of the PIC distribution functions are interpolated and message passed in an efficient manner through the Space Weather Modeling Framework (SWMF). Both BATS-R-US and iPIC3D are massively parallel codes fully integrated into, run by and coupled through the SWMF. This model has been successfully applied to the global magnetosphere simulations of Ganymede, Mercury and Earth.

Yuxi Chen  
University of Michigan  
yuxichen@umich.edu

Gabor Toth  
University of Michigan  
Space Physics Research Laboratory  
toth@umich.edu

Stefano Markidis  
KTH, Sweden  
markidis@kth.se

Ivy Bo Peng  
KTH  
Sweden  
bopeng@kth.se

### MS186

#### Method of Lines Transpose (molt), An O(N) Implicit Maxwell Solver Coupled with Particle in Cell

The Maxwell's solver proposed by E. Wolf et al. was developed based on the method of lines transpose (MOLT) framework and proven to be numerically efficient due to its A-stable property. Why the method ensures a divergence

free magnetic field, a costly elliptic divergence cleaning technique was used to enforce Gauss's law. To improve the efficiency, in this talk we present a novel hyperbolic cleaning technique for the MOLT Maxwell solver. Specifically, the Lagrangian multiplier is first written in a wave equation form and hence the same MOLT wave equation solver for Maxwell's equations can be applied. The proposed scheme is further coupled with a particle-in-cell method for the simulation of plasmas. The numerical results of several benchmark test problems will be presented to demonstrate the performance of our scheme.

Andrew J. Christlieb

Michigan State University  
Dept. of Comp. Math., Sci & Engr.  
andrewch@math.msu.edu

Yingda Cheng  
Department of Mathematics  
Michigan State University  
ycheng@math.msu.edu

Wei Guo  
Michigan State University  
wguo@math.msu.edu

#### MS186

##### **Semi-Lagrangian Solution of the Gyrokinetic Vlasov Equation: Latest Developments in the Selalib Code**

Selalib is an object-oriented Fortran library for the solution of Vlasov-like kinetic equations, by means of semi-Lagrangian or particle-in-cell methods [SELALIB: <http://selalib.gforge.inria.fr>]. The library is highly modular: the building blocks of our solvers may be changed easily, allowing one to test different algorithms. A typical semi-Lagrangian simulation in Selalib is based on a dimensionally split form of the Vlasov equation: this results in a sequence of one- or two-dimensional advection equations, to be solved by combining the method of characteristics with an appropriate interpolation strategy (local or global). For each advection equation the original domain is decomposed across processors according to a different layout; Selalib provides the tools for creating these layouts, and for performing data remapping across them. We present here how this framework is implemented in the case of the Gyrokinetic Vlasov equation, for the simulation of turbulence in magnetic fusion devices. We use field-aligned interpolation (and differentiation) to reduce the number of poloidal planes in our computational mesh, while keeping the flexibility to select different discretization strategies in the poloidal plane: we consider multi-block curvilinear and hexagonal meshes that are conformal to the magnetic flux surfaces, as well as non-conformal Cartesian meshes with embedded boundaries. Pros and cons of the various strategies are discussed.

Yaman Güçlü

Department of Mathematics  
Michigan State University  
yaman.guclu@ipp.mpg.de

Michel Mehrenberger, Laura Mendoza  
Institut de Recherche Mathématique Avancée (IRMA),  
University of Strasbourg  
mehrenbe@math.unistra.fr, mendoza@unistra.fr

Ahmed Ratnani

Max-Planck-Institut für Plasmaphysik, NMPP division  
ahmed.ratnani@ipp.mpg.de

Eric Sonnendrücker, Edoardo Zoni  
Max-Planck-Institut für Plasmaphysik, NMPP division,  
and  
Technische Universität München, Department of  
Mathematics  
eric.sonnendruecker@ipp.mpg.de,  
edoardo.zoni@ipp.mpg.de

#### MS186

##### **An Eulerian Discontinuous Galerkin Scheme for the Fully Kinetic Vlasov-Maxwell System**

Collisionless or weakly collisional plasmas are ubiquitous in the universe, necessitating a kinetic description. The Vlasov-Maxwell system of equations describes the evolution of the probability distribution function of the particles subject to self-consistently generated electromagnetic fields. Typically, the Vlasov-Maxwell system is solved in a Lagrangian framework using the particle-in-cell method; however, we focus on the development of a fully kinetic (ions and electrons) Eulerian approach, which is ideal for studying problems requiring low amplitude and high sensitivity, such as energy dissipation in turbulence and magnetic reconnection. In particular, we describe the development and performance of a high-order discontinuous Galerkin scheme in 6D phase space. The resulting system is applied to turbulence and reconnection studies.

Jason TenBarge

IREAP, University of Maryland  
jason.tenbarga@gmail.com

James Juno  
IREAP  
University of Maryland  
jjuno@umd.edu

Ammar Hakim  
Princeton Plasma Physics Laboratory  
ahakim@pppl.gov

#### MS187

##### **Scalable Solvers for Joint Inversion with Several Structural Coupling Terms**

Joint inversion is the solution of a single inverse problem involving two or more parameter fields, defined on the same physical domain. By doing so, one can reasonably hope to obtain better reconstructions for each parameter field, by combining the information coming from multiple coherent data sets, rather than by treating each ones individually. However very often the different parameter fields do not have a known, reliable relationship connecting them to each other. A joint inverse problem can nonetheless be formulated in that case by explicitly coupling the different parameters. Our work focuses on the different ways to couple these parameters, and to solve the corresponding joint inverse problems. In a first part, we review the main techniques used to combine the structure of the different parameters: cross-gradient, vectorial TV, and nuclear norm. These terms are widely used in practice, but still poorly understood in the context of inverse problems. Moreover, they are highly nonlinear, posing significant computational challenges for large-scale problems. In a second part, we study the properties of these coupling terms when inter-

preted as anisotropic diffusion operators, and use these insights to derive scalable numerical solvers for the joint inverse problems.

Benjamin Crestel  
University of Texas at Austin  
Institute for Computational Engineering and Science  
crestel@ices.utexas.edu

Tan Bui-Thanh  
The University of Texas at Austin  
tanbui@ices.utexas.edu

Georg Stadler  
Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

### MS187

#### Efficient Solvers for Coupled Imaging Problems

This talk presents some large-scale coupled optimization problems arising in imaging and efficient numerical methods for their solutions. In particular, we are interested in linear and nonlinear inverse imaging problems involving motion which are affected by noisy data. Our goal is to simultaneously reconstruct the images and recover their motion parameters. We discuss and compare decoupled, partially coupled, and fully coupled optimization techniques. For the latter, we propose a novel method that is very efficient when the motion can be represented using relatively few parameters, as is the case for rigid, affine, or smooth low-dimensional transformation models. Since the imaging problems at hand are typically ill-posed, we utilize prior knowledge of the expected solution by adding regularization and constraints. For our work, we consider simple bound constraints on both the image and motion parameters and explore the use of hybrid methods for determining regularization parameters. Our numerical experiments provide a detailed convergence analysis and outline the potential of the methods for large-scale imaging problems.

James Herring  
Emory University  
Department of Mathematics and Computer Science  
james.lincoln.herring@emory.edu

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

### MS187

#### Parallel Algorithms for PDE-Constrained Optimization Problems with Hyperbolic Constraints

We describe a distributed-memory solver for large-scale PDE-constrained optimization problems with hyperbolic constraints in three dimensions. Our contributions are the design of effective preconditioners for the Hessian, an automatic parameter selection, and the parallelization of our solver on distributed-memory architectures. Hyperbolic PDE-constrained optimization problems have numerous applications ranging from geophysical sciences, atmo-

spheric sciences, to medical imaging sciences. They are in general challenging to solve in an efficient way: They are infinite-dimensional in the continuum, high-dimensional after discretization, and ill-posed in nature resulting in large, severely ill-conditioned systems after discretization that have to be solved numerically. We solve for the first order optimality conditions using a globalized, preconditioned, inexact Gauss-Newton-Krylov method. We propose efficient multigrid and domain decomposition strategies to effectively precondition the Hessian. We will showcase convergence and scalability results of our solver for up to 10 billion unknowns.

Andreas Mang  
The University of Texas at Austin  
The Institute for Computational Engineering and Sciences  
andreas@ices.utexas.edu

Amir Gholaminejad  
PhD Student, Institute for Computational Engineering and Sciences, University of Texas at Austin  
i.amirgh@gmail.com

George Biros  
The Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
biros@ices.utexas.edu

### MS187

#### Arock: An Algorithmic Framework for Asynchronous Parallel Coordinate Updates

Finding a fixed point to a nonexpansive operator, i.e.,  $x^* = Tx^*$ , abstracts many problems in numerical linear algebra, optimization, and other areas of scientific computing. To solve fixed-point problems, we propose ARock, an algorithmic framework in which multiple agents (machines, processors, or cores) update  $x$  in an asynchronous parallel fashion. Asynchrony is crucial to parallel computing since it reduces synchronization wait, relaxes communication bottleneck, and thus speeds up computing significantly. At each step of ARock, an agent updates a randomly selected coordinate  $x_i$  based on possibly out-of-date information on  $x$ . The agents share  $x$  through either global memory or communication. If writing  $x_i$  is atomic, the agents can read and write  $x$  without memory locks. Theoretically, we show that if the nonexpansive operator  $T$  has a fixed point, then with probability one, ARock generates a sequence that converges to a fixed point of  $T$ . Our conditions on  $T$  and step sizes are weaker than comparable work. Linear convergence is also obtained. We propose special cases of ARock for linear systems, convex optimization, machine learning, as well as distributed and decentralized consensus problems. Numerical experiments of solving sparse logistic regression problems are presented.

Ming Yan  
Michigan State University  
Department of CMSE; Department of Mathematics  
yanm@math.msu.edu

Ming Yan  
Michigan State University  
Department of CMSE  
yanm@math.msu.edu

Zhimin Peng  
University of California. Los Angeles  
Department of Mathematics

zhimin.peng@math.ucla.edu

Yangyang Xu  
University of Alabama  
yangyang@ima.umn.edu

Wotao Yin  
University of California, Los Angeles  
wotaoyin@math.ucla.edu

### MS188

#### Spectral Density of the Koopman Operator: Interpretation and Approximation

Koopman operator models the time evolution of measurements taken from a dynamical system, such as pointwise velocities collected from a fluid flow. This operator is linear; its eigenvalues are associated with modes of evolution of measurements, which are known to match physical observations of fluid flows. There has been a substantial progress in approximating steady eigenvalues and modes of the Koopman operator using the Dynamic Mode Decomposition algorithm; however, the operator often has a continuous part of the spectral density, which does not correspond to eigenvalues and has attracted little attention. In this talk we discuss the challenges in detecting, approximating, and interpreting the continuous spectral density of the Koopman operator in the context of fluid kinematics.

Marko Budisic  
Department of Mathematics  
Clarkson University  
marko@clarkson.edu

Ryan Mohr  
University of California, Santa Barbara  
mohrrm@engineering.ucsb.edu

Mihai Putinar  
Professor of Mathematics  
University of California at Santa Barbara  
mputinar@math.ucsb.edu

### MS188

#### Using Derivative Snapshots in ROMs Error Estimates for Linear Systems, Error Bounds for Non-linear Systems and Numerical Examples

We consider a POD ROM method when the time derivatives are used as snapshots in addition to snapshots from the solution and compare it to the case when only solution snapshots have been used. We will comment on previously derived upper bounds for the 2-norm of the errors in both cases and will present newly derived *a priori* error estimates for linear dynamical systems. Numerical examples will be presented to demonstrate the insights gained from the new estimates.

Tanya Kostova-Vassilevska  
Lawrence Livermore National Laboratory  
vassilevska1@llnl.gov

### MS188

#### Computational Reduction Strategies for Bifurcations and Stability Analysis in Fluid-Dynamics: Applications to Coanda Effect

We focus on reducing the computational costs associated

with the hydrodynamic stability of solutions of the incompressible Navier-Stokes equations for a Newtonian and viscous fluid in contraction-expansion channels. In particular, we are interested in studying steady bifurcations, occurring when non-unique stable solutions appear as physical and/or geometric control parameters are varied. The formulation of the stability problem requires solving an eigenvalue problem for a partial differential operator. An alternative to this approach is the direct simulation of the flow to characterize the asymptotic behavior of the solution. Both approaches can be extremely expensive in terms of computational time. We propose to apply Reduced Order Modeling (ROM) techniques to reduce the demanding computational costs associated with flow stability analysis. The application that motivated the present study is the onset of asymmetries (i.e., symmetry breaking bifurcation) in blood flow through a regurgitant mitral valve, depending on the Reynolds number and the regurgitant mitral valve orifice shape.

Giuseppe Pitton  
SISSA, International School for Advanced Studies Trieste  
Italy  
gpitton@sissa.it

Annalisa Quaini  
Department of Mathematics, University of Houston  
quaini@math.uh.edu

Gianluigi Rozza  
SISSA, International School for Advanced Studies  
Trieste, Italy  
gianluigi.rozza@sissa.it

### MS188

#### Incremental POD Method of Snapshots for Finite Element Fluid Computations

We discuss an incremental algorithm for proper orthogonal decomposition (POD) computations. We modify an incremental matrix SVD algorithm of Brand to accommodate data arising from Galerkin-type simulation methods for time dependent PDEs. The algorithm initializes and efficiently updates the POD eigenvalues and modes during the time stepping in a PDE solver without storing the simulation data. We demonstrate the effectiveness of the algorithm using finite element computations for fluid flows.

Hiba Fareed  
Mathematics and Statistics  
Missouri S&T  
hf3n3@mst.edu

Jiguang Shen  
School of Mathematics  
University of Minnesota  
shenx179@umn.edu

John Singler  
Missouri S&T  
Mathematics and Statistics  
singlerj@mst.edu

Yangwen Zhang  
Missouri S&T



ywzfg4@mst.edu

### MS189

#### Reduced Order Modelling for Fluid-Structure Interaction and Optimal Control Problems

In this talk we will focus on recent advances in reduced order modelling for parametrized problems in computational fluid dynamics, with a special attention to the case of optimization based domain decomposition (opt-DD) reduced order methods and reduced order fluid-structure interaction problems (FSI). For the latter, we will propose reduced order versions of existing numerical scheme, based either on a monolithic [Ballarin, Rozza; International Journal for Numerical Methods in Fluids, 2016] or a segregated approach [Ballarin, Rozza, Maday; submitted, 2016]; advantages and drawbacks of reduced order monolithic vs reduced order partitioned approaches will be highlighted. For the former, we will propose an extension of high fidelity opt-DD methods [Gunzburger, Peterson, Kwon; Computers & Mathematics with Applications, 1999] to the reduced order framework. Test cases for opt-DD model reduction of elliptic problems will be provided. Finally, a possible opt-DD ROM for FSI problems will be introduced.

Francesco Ballarin

SISSA, International School for Advanced Studies  
francesco.ballarin@sissa.it

Gianluigi Rozza

SISSA, International School for Advanced Studies  
Trieste, Italy  
gianluigi.rozza@sissa.it

### MS189

#### Pod-Based Multiobjective Optimal Control by Use of the Reference Point Method.

A bicriterial optimal control problem governed by a parabolic partial differential equation (PDE) and bilateral control constraints is considered. For the numerical optimization the reference point method is utilized. The PDE is discretized by a Galerkin approximation utilizing the method of proper orthogonal decomposition (POD). POD is a powerful approach to derive reduced-order approximations for evolution problems. Numerical examples illustrate the efficiency of the proposed strategy.

Dennis Beermann

University of Constance  
dennis.beermann@uni-konstanz.de

### MS189

#### Accelerating Newtons Method for Large-Scale Time-Dependent Optimal Control via Reduced-Order Modeling

We use projection based reduced order models (ROMs) to substantially decrease the computational cost of Newton's method for large-scale time-dependent optimal control problems. Previous approaches of using projection based ROMs in optimization solve sequences of computationally less expensive optimization problems, which are obtained by replacing the high fidelity discretization of the governing partial differential equation (PDE) by a ROM. This may not be feasible or efficient for several nonlinear PDEs because 1) ROMs may not be well-posed because the ROM solution does not satisfy properties obeyed by the so-

lution of the high-fidelity PDE discretization, such as conservation properties, or non-negativity, 2) ROMs can become untrustworthy when the controls differ slightly from those at which the ROM was generated, or 3) validation that a ROM is still sufficiently accurate at a new control can be very expensive. We use the high-fidelity PDE discretization for objective function and gradient evaluation, and use state PDE and adjoint PDE solves to generate a ROM to compute approximate Hessian information. The ROM generation is computationally inexpensive, since it reuses information that is already computed for objective function and gradient evaluation, and the ROM-Hessian approximations are inexpensive to evaluate. The resulting method often enjoys convergence rates similar to those of Newton's method, but at the computational cost per iteration of the gradient method.

Caleb C. Magruder

Rice University, Computational and Applied Mathematics  
cm47@rice.edu

Matthias Heinkenschloss

Department of Computational and Applied Mathematics  
Rice University  
heinken@rice.edu

### MS189

#### Efficient PDE-Constrained Optimization under Uncertainty using Adaptive Model Reduction and Sparse Grids

This work introduces a framework for accelerating optimization problems governed by partial differential equations with random coefficients by leveraging adaptive sparse grids and model reduction. Adaptive sparse grids perform efficient integration approximation in a high-dimensional stochastic space and reduced-order models reduce the cost of objective function and gradient queries by decreasing the complexity of primal and adjoint PDE solves. A trust region method that allows for inexact gradient and objective evaluations manages these two sources of inexactness and ensures global convergence. Numerical results show the proposed method is up to two orders of magnitude less expensive than existing methods when applied to the model problem of stochastic 1d optimal flow control.

Matthew J. Zahr

Stanford University  
mjzahr@lbl.gov

Kevin T. Carlberg

Sandia National Laboratories  
ktcarlb@sandia.gov

Drew P. Kouri

Optimization and Uncertainty Quantification  
Sandia National Laboratories  
dpkouri@sandia.gov

### MS190

#### Title Not Available at Time of Publication

Abstract not available

Zhaopeng Hao

Mathematical Sciences of Worcester Polytechnic Institute

zhao@wpi.edu

**MS190****Fractional Modeling of Multiphase Flows**

In this work, we develop a fractional extension of a mass-conserving Allen-Cahn phase field model that describes the mixture of two incompressible fluids. The fractional order controls the sharpness of the interface, which is typically diffusive in integer-order phase-field models. The model is derived based on an energy variational formulation. The spatial discretization is based on a Petrov-Galerkin spectral method whereas the temporal discretization is based on a stabilized ADI scheme both for the phase-field equation and for the Navier-Stokes equation. We demonstrate the spectral accuracy of the method with fabricated smooth solutions and also the ability to control the interface thickness between two fluids with different viscosity and density in simulations of two-phase flow in a pipe and of a rising bubble. We also demonstrate that using a formulation with variable fractional order we can deal simultaneously with both erroneous boundary effects and sharpening of the interface at no extra resolution.

Fangying Song

Division of Applied Mathematics  
Brown University  
fangying\_song@brown.edu

Chuanju Xu

School of Mathematical Sciences  
Xiamen University  
cjxu@xmu.edu.cn

George Karnidakis

Brown University  
george\_karniadakis@brown.edu

**MS190****An Iterative Mathematical-Computational Framework for Data-Driven FPDE Modelling and Simulation**

Abstract not available

Mohsen Zayernouri

Michigan State University  
zayern@egr.msu.edu

**MS190****Backward Fractional Advection-Dispersion Equations to Predict Contaminant Source**

The forward Fractional Advection Dispersion Equation (FADE) provides a useful model for non-Fickian transport in heterogeneous porous media. This presentation introduces the corresponding backward FADE model, to identify source location and release time. The backward method is developed by combining sensitivity analysis and fractional adjoint, and the resultant backward FADE differs significantly from the traditional backward Advection Dispersion Equation (ADE) because the fractional derivative is not self-adjoint and the probability density function for backward locations is highly skewed. Finally, the method is validated using tracer data from well-known field experiments. The backward ADE cannot reliably identify the source in these applications, since the forward ADE does not provide an adequate fit to the concentration data.

Yong Zhang

University of Alabama  
yzhang264@ua.edu

Mark Meerschaert

Dept. of Statistics and Probability  
Michigan State University, USA  
mcubed@stt.msu.edu

Roseanna Neupauer

University of Colorado Boulder  
roseanna.neupauer@colorado.edu

**MS191****A Diffuse Interface Model for Two-Phase Ferrofluid Flows**

A ferrofluid is a liquid which becomes strongly magnetized in the presence of applied magnetic fields. It is a colloid made of nanoscale monodomain ferromagnetic particles suspended in a carrier fluid. These particles are suspended by Brownian motion and will not precipitate nor clump under normal conditions. Ferrofluids are dielectric and paramagnetic. There are two well established PDE models used as a mathematical description for the behavior of ferrofluids: the Rosensweig and Shliomis models. These deal with one-phase flows, which is the case of many technological applications. However, some applications arise naturally in the form of a two-phase flow: one of the phases has magnetic properties while the other one does not (magnetic manipulation of microchannel flows, microvalves, magnetically guided transport, etc.). We develop a model describing the behavior of two-phase ferrofluid flows using phase field techniques and present an energy-stable numerical scheme for it. For a simplified version of this model and the corresponding numerical scheme we prove, in addition to stability, convergence and, as a consequence, existence of solutions. With a series of numerical experiments we illustrate the potential of these simple models and their ability to capture basic phenomenological features of ferrofluids such as the Rosensweig instability

Ignacio Tomas

Texas A&M university  
itomas@tamu.edu

**MS191****Numerical Schemes for Phase Field Model for Moving Contact Line Model**

The threshold dynamics method developed by Merriman, Bence and Osher (MBO) is an efficient method for simulating the motion by mean curvature flow when the interface is away from the solid boundary. Direct generalization of the MBO type method to the wetting problems with interface intersecting the solid boundary is not easy because solving heat equation on general domain with wetting boundary condition is not as efficient as that for the original MBO method. The dynamics of the contact point also follows a different dynamic law compared to interface dynamics away from the boundary. We develop an efficient volume preserving threshold dynamics (MBO) method for drop spreading on rough surfaces. The method is based on minimization of the weighted surface area functional over an extended domain that includes the solid phase. The method is simple, stable with the complexity  $O(N \log N)$  per time step and it is not sensitive to the inhomogeneity

or roughness of the solid boundary. We also extend the idea to an efficient method for image segmentation.

Xiaoping Wang  
Hong Kong University of Science and Technology  
mawang@ust.hk

#### MS191

##### **An Energy Stable Numerical Scheme for the Cahn-Hilliard-Wilmore and Functionalized Cahn-Hilliard Models**

The numerical approximations the Cahn-Hilliard Wilmore/Functionalized Cahn-Hilliard equations are considered. The key challenge for the energy stability is associated with a non-convex, non-concave term in the energy expansion. To overcome this difficulty, we add nonlinear regularization terms, so that a convex-concave decomposition of the energy is available. In turn, the unique solvability and unconditional energy stability are established at the theoretical level, due to the convex splitting nature of the numerical scheme. Moreover, a full order convergence analysis for the fully discretized scheme is provided. Some numerical simulation results are also presented in this talk.

Cheng Wang  
University of Massachusetts  
Dartmouth, MA  
cwang1@umassd.edu

#### MS191

##### **Decoupled, Energy Stable Scheme for Continuum Hydrodynamics Allen-Cahn and Cahn-Hilliard Phase Field Model**

The continuum hydrodynamics model is a coupled system that consists of Navier-Stokes equations and phase-field equation with generalized Navier boundary conditions. we develop a stabilized, decoupled, time discretization scheme for the coupled nonlinear system. For the time discretization, we present a linear decoupled energy stable scheme to solve the phase-field model with moving contact line by adding some stabilizing terms. For the spatial discretization, the finite element method is used to solve the model. The Navier-Stokes equation is solved by the projection method, and the schemes are solved by preconditioned conjugate gradient method. The scheme is unconditionally energy stable and leads to linear and decoupled elliptic equations to be solved at each time step. Stability analysis and ample numerical simulations are presented thereafter. Numerical results show the time accuracy of the scheme and simulate the processes of wetting and dewetting phenomena.

Hui Zhang  
Beijing Normal University  
hzhang@bnu.edu.cn

#### MS192

##### **Joint Model and Parameter Reduction for Large-Scale Bayesian Inversion**

Algorithmic scalability to high dimensional parameters and computational efficiency of numerical solvers are two central challenges in solving large-scale PDE-constrained inverse problems. Here we will investigate the intrinsic dimensionality in both parameter space and model space by exploiting the interaction among various information

sources and model structures. We will also discuss various strategies for jointly identifying low-dimensional parameter and model subspaces. The resulting subspaces naturally lead to accelerated sampling methods that can overcome the above-mentioned challenges and demonstrate potential reductions for problems with high-dimensional data.

Tiangang Cui  
Monash University  
tiangang.cui@monash.edu

Youssef M. Marzouk, Karen E. Willcox  
Massachusetts Institute of Technology  
ymarz@mit.edu, kwillcox@MIT.EDU

#### MS192

##### **On Covariance Operators Derived from Elliptic PDEs**

We consider the problem of boundary effects on inverse elliptic covariance operators defined on a bounded domain. We propose two simple, easy to implement methods to ameliorate these. One is using a Robin boundary instead of a Neumann or a Dirichlet one. The other involves ensuring the pointwise variance throughout the domain is constant. We also revise a previously suggested method of extending the computational domain. All methods are presented in infinite dimensions. We present numerical studies of our methods over non-standard domains.

Yair Daon  
New York University  
New York City  
yair.daon@gmail.com

Georg Stadler  
Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

#### MS192

##### **Iterative Updating of Model Error for Bayesian Inversion**

One of the computational challenges associated with large-scale inverse problems is the cost of forward model evaluations. Often a compromise must be made between accuracy (fine models) and speed (coarse models). We outline an algorithm that iteratively estimates the distribution of the model error arising from using a coarse model, allowing for more accurate sampling of the posterior distribution. The convergence of the algorithm in the linear Gaussian case is analyzed, wherein the posterior remains Gaussian and can be characterized by the evolution of its mean and covariance. We show that these both converge exponentially fast, with the limiting covariance being non-degenerate. In the non-linear case, numerically all measures will be approximated by ensembles of particles; we show convergence of some different particle approximations in the large particle limit. Finally we present some numerical results to illustrate the behavior of the algorithm.

Matthew M. Dunlop  
Mathematics Institute  
University of Warwick  
mdunlop@caltech.edu

Andrew Stuart  
Computing + Mathematical Sciences

California Institute of Technology  
astuart@caltech.edu

Erkki Somersalo  
Case Western Reserve University  
ejs49@case.edu

Daniela Calvetti  
Case Western Reserve Univ  
Department of Mathematics, Applied Mathematics and  
Statistic  
dxc57@case.edu

### MS192

#### Bayesian Calibration of Inadequate Stochastic PDE Models

Partial differential equations (PDE) are a powerful tool to model complex physical systems across many disciplines. PDE models often depend on some uncertain input parameters, whose values cannot be directly measured. Calibration is the process of fitting the PDE model to some available, possibly noisy, experimental data. The Bayesian approach not only provides an estimate of the uncertain parameters, but also allows propagating uncertainty to the outputs of the system. For many applications of practical interests, the PDE model is often inadequate, that is it only partially captures the complex dynamics of the physical system leading to a discrepancy between its outputs and reality. We model the inadequacy by introducing a stochastic term in the governing equations. This approach leads to the solution of a Bayesian inverse problem where evaluations of the likelihood function require computing the expected value of the observables with respect to a stochastic field. This poses a substantial computational challenge for Markov chain Monte Carlo methods, since it requires many evaluation of a computationally expensive and possibly nonlinear PDE forward problem for different realization of both the calibration parameters and the stochastic field. In this talk we explore scalable algorithms for the Bayesian calibration of inadequate PDE models that exploit derivative information. An application to an algebraic closure model for turbulent flows will be presented.

Umberto Villa  
University of Texas at Austin  
uvilla@ices.utexas.edu

Todd A. Oliver  
PECOS/ICES, The University of Texas at Austin  
oliver@ices.utexas.edu

Robert D. Moser  
University of Texas at Austin  
rmoser@ices.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

### MS193

#### Eikonals, Ray Mappings and Nonimaging Optics

Rays and wavefronts are two equivalent ways to formulate geometric optics. They are both included in Hamilton's eikonal functions. I shall describe examples of lens design where eikonals, ray mapping, and wave front formulations

are alternatively applied.

Jacob Rubinstein  
Department of Mathematics  
Technion - Israel Institute of Technology  
koby@tx.technion.ac.il

### MS193

#### A Least-Squares Method for Optical Design

The optical system design problem is an *inverse problem*: "Find an optical system which contains reflectors and/or lenses that gives the desired light output at the target for a given input at the source". We show that the optical design problem of one/two surfaces gives rise to a Monge-Ampère equation with transport boundary condition and this Monge-Ampère equation is obtained through a generalized geometrical design method which is based on the principle of equal optical path length and conservation of the luminous flux. We solve this boundary value problem using a recently developed least-squares algorithm. It is one of the few numerical algorithms capable to solve these type of problems efficiently. The least-squares algorithm is an iterative minimization procedure. At each iteration, three steps are performed: two of these are nonlinear minimization steps, which can be performed pointwise, and the third step involves two Poisson problems. The least-squares method gives the optimal mapping which transforms the given energy density at the source to the desired energy density at the target, and the optical system is obtained via this mapping.

Nitin Kumar Yadav  
Department of Mathematics and Computer Science  
Eindhoven University of Technology  
n.k.yadav@tue.nl

### MS193

#### Active Flux Schemes for Liouville's Equation from Geometrical Optics

The classical problem of geometrical optics can be formulated as follows: given an optical system and an input, find the output light distribution. Using ray tracing, any such problem may be solved to arbitrary accuracy. However, the order of convergence for such methods is very slow, so that many rays are often needed. Especially for lighting optics, the resolution of the problem typically requires several million rays. We aim to improve the convergence rate, and thereby the computational efficiency, by constructing solvers based on Liouville's equation. Away from optical interfaces Liouville's equation is simply a linear advection equation, so that any high-order method may be directly applied. However, when optical transitions are present, the Hamiltonian is discontinuous and Liouville's equation is no long well-defined. In such cases, we locally apply ray tracing methods to find the physically correct solution. In selecting a method, the total light energy should be preserved so that high-order conservative methods are preferred. At the same time, we wish to avoid large stencils that would give complications near optical interfaces. Finally, the method should allow for easy implementation of local ray tracing near optical interfaces. Active flux schemes fit the bill nicely, they are a class of methods that use averages for conservation and point values that define fluxes. We apply these schemes to several optical problems, showing their strengths and advantages.

Bart van Lith

Department of Mathematics and Computer Science  
Eindhoven University of Technology  
b.s.v.lith@tue.nl

#### MS194

### Uncertainty Quantification in Reduced Order Modeling

Reduced order modeling (ROM) is often used to develop compact representations for models governed by deterministic differential equations. A classical approach in ROM is to make use of the well-known proper orthogonal decomposition (POD). In POD, training data is used to construct a reduced basis that serves as a surrogate model. As new data (test data) becomes available, the new data can be projected onto the reduced space giving an approximation of the test data. However, it is often the case that both the training data and the test data are contaminated with noise. The noise in the training data causes errors in the basis functions. By projecting noisy data onto noisy basis functions, the resulting approximation has a nonlinear dependence upon both sources of errors, even in the case where the training and the test are both assumed to follow a simple additive error model. As a result, an additive error model cannot represent the approximation of the test data in the reduced space. In order to formulate a meaningful measure of the uncertainty present in the results of the POD, the nonlinear relationship between the noise in the training data and the test data must be accounted for. In this talk we first explore this relationship, and then show that under certain conditions the test data can be approximated by a normal random variable. This information is then used to construct an asymptotic prediction envelope for repeated measurements.

Jared Catenacci

North Carolina State University  
catenajw@nv.doe.gov

Aaron B. Luttmann

National Security Technologies, LLC  
luttmaab@nv.doe.gov

#### MS194

### Improved Probabilistic Principal Component Analysis for Application to Reduced Order Modeling

A fundamental application of classical PCA is reduced order modeling (ROM), where the principal components are computed from a training data (TD) set and used as a reduced basis for the data space. In our previous work, we made a ROM for stochastic systems by projecting rapidly calculated (but noisy) data onto the basis vectors to obtain rapid, accurate predictions. All of these steps were done using techniques that do not account for noise in the data and  $L_2$  projection to create the reduced representations of the noisy data. In this work, stochastic approaches are used. The PPCA is used to generate the basis, which then allows estimating the noise in the TD. The standard PPCA has also been improved so that the variance of the basis expansion coefficients when representing the TD can be estimated. Another benefit of PPCA is to use model selection (MS) criteria to automatically truncate the basis used for the ROM. A new statistical approach for the projection step uses the standard deviation information from the PPCA to create a “smart” projection that does not deviate too much from the TD. This gives improved results when projecting noisy data. The holistic approach gives a fully stochastic method for computing a ROM from noisy

TD, determining the ideal MS, and projecting noisy test data onto the ROM. We demonstrate our framework on synthetic data and on an application to computing radiation transport with MC simulations.

Indika G. Udagedara

Clarkson University  
udagedig@clarkson.edu

#### MS194

### On the Role of Physics for Extrapolating First-Principles Models

Advancements in numerical modeling have allowed for a paradigm shift such that first-principles physics-based simulations are increasingly relied upon to predict the behavior of engineered systems before they are tested. The predictive capability of such models can be limited when they are extrapolated to settings that are different from those used for calibration. Research has focused on increasing predictive capability through different techniques including calibration of model parameters and development of bias-correcting discrepancy terms. The applicability of these methods can be limited because using the same model in the calibration and extrapolation domains fundamentally assumes that the physics are the same in both domains. Herein, we propose that reliable extrapolation depends on quantifying the extent to which physics in the calibration domain is representative of physics encountered in the extrapolation regime. This concept will be illustrated using experimental results obtained in the development of a model to predict the payload capacity of a quadcopter. We demonstrate that the model remains reliable even when extrapolated to use propeller blades larger than those used for calibration. This result suggests that considering only a mathematical description of the model to infer predictive capability is incomplete because in this application the physics remains unchanged in the extrapolation regime.

Kendra Van Buren

Los Alamos National Laboratory  
klvan@lanl.gov

#### MS195

### Comparison of Several Advanced High Order Vlasov Schemes

Vlasov schemes as alternative to particle-in-cell methods are very effective when high resolution is needed. In this talk, we compare the numerical performance of several recently developed high order Vlasov schemes including high order semi-Lagrangian (SL) schemes using various spatial discretization such as WENO, discontinuous Galerkin (DG), and convected schemes; high order Eulerian schemes such as Runge-Kutta DG methods and energy-conserving DG methods; and a high order WENO method based on the method of lines transpose approach. In particular, we test the codes for several benchmarks, such as Landau damping and two-stream instability which are widely used in the literature, and compare their numerical efficiency in terms of CPU time and the ability to resolve filamentation structures and to converse the physical invariants. Furthermore, we formulate several plasma sheath problems including the floating wall problem, the source-collector problem and a problem with boundary condition satisfying the generalized Bohm criterion. These tests are of importance in physics but they are rarely used to benchmark Vlasov schemes. The performance of the high order schemes for

simulating the sheath problems will be investigated.

Wei Guo

Michigan State University  
wguo@math.msu.edu

Andrew J. Christlieb  
Michigan State University  
Dept. of Comp. Math., Sci & Engr.  
andrewch@math.msu.edu

#### MS195

### Positivity-Preserving Hybrid Semi-Lagrangian Discontinuous Galerkin Schemes for the 2D2V Vlasov-Poisson System on Unstructured Meshes

Abstract not available

David C. Seal

United State Naval Academy  
seal@usna.edu

James A. Rossmanith  
Iowa State University  
Department of Mathematics  
rossmani@iastate.edu

#### MS195

### Solving Fluid and Continuum Kinetic Equations in Plasma Physics Using the Discontinuous Galerkin Method

The discontinuous Galerkin method is used to solve equations in plasma physics in fluid and kinetic regimes. Some challenges in modeling plasmas include the need to resolve disparate spatial and temporal scales with physically-relevant dispersive and diffusive phenomena. A multi-fluid plasma model, where ions and electrons are treated as separate fluids, using the discontinuous Galerkin method is used for applications in the fluid regime. A direct discretization of the Vlasov equation using the discontinuous Galerkin method is used to study applications in kinetic regimes. Typically kinetic simulations use particle-in-cell codes for each of the ion and electron species, but particle codes are subject to noise. Direct discretization of the Vlasov equation can provide smooth solutions of the distribution function in phase space. The high-dimensionality of the Vlasov equation has made continuum kinetic simulations inaccessible until recently. Simulation results will be presented for fluid and kinetic applications using the discontinuous Galerkin method.

Bhuvana Srinivasan

Aerospace and Ocean Engineering  
Virginia Tech  
srinbhu@vt.edu

Ammar Hakim  
Princeton Plasma Physics Laboratory  
ahakim@pppl.gov

Petr Cagas, Yang Song  
Aerospace and Ocean Engineering  
Virginia Tech  
pcagas@vt.edu, yangsong@vt.edu

#### MS195

### IFP: An Optimal, Adaptive, Fully Conserva-

### tive and Asymptotic Preserving 1D2V Vlasov-Rosenbluth-Fokker-Planck for Spherical ICF Implosion Calculation

We present a new, 1D2V Vlasov-Rosenbluth-Fokker-Planck code for simulating implosions of spherical inertial confinement fusion (ICF) capsules. The code adopts a multi-grid preconditioned Anderson Acceleration as a solver in order to step over stiff collision time-scales while ensuring exact conservation (mass, momentum, and energy) in order to ensure long time accuracy of simulation results. We optimize mesh resolution requirements by 1) adapting the velocity-space mesh based on the species' local thermal-velocity, 2) treating the cross-species collisions exhibiting disparate thermal velocities by a coarse-graining strategy, 3) ensuring an analytical Maxwellian in the limit of  $\Delta t v_{col} \rightarrow \infty$ , using a new equilibrium preserving discretization scheme, and 4) tracking capsule implosions with a moving mesh in physical space. We demonstrate the efficiency and accuracy properties of the approach with challenging numerical examples in spherical geometry.

William T. Taitano, Luis Chacon, Andrei Simakov, Brett Keenan

Los Alamos National Laboratory  
taitano@lanl.gov, chacon@lanl.gov, simakov@lanl.gov, keenan@lanl.gov

#### MS196

### High-Dimensional Approximation and Structured Sparsity

Standard compressed sensing theory is built on the concept of sparsity. However, since the early days it has been known that many problems to which compressed sensing techniques can be applied possess substantial additional structure beyond sparsity. By leveraging such structure e.g. through the design of a suitable recovery algorithm one aims to obtain a better approximation than through a sparsity-promoting algorithm alone. In this talk I will first survey some common structured sparsity notions that have been studied in compressed sensing, before focusing on structures that can be beneficial to high-dimensional approximation. In particular, I will discuss how the notion of lower sets can be viewed within the context of structured sparsity, and exhibit a number of different algorithms for promoting this type of structure. Time permitting, I will also discuss joint sparsity and so-called sparsity in levels within the function approximation context.

Ben Adcock, Yi Sui, Casie Bao

Simon Fraser University  
ben\_adcock@sfu.ca, ysui@sfu.ca, cbao@sfu.ca

#### MS196

### Model Order Reduction and Sparse Approximations for Dynamical Systems with Random Parameters

We consider dynamical systems of ordinary differential equations or differential algebraic equations including many physical parameters. A quantity of interest (QoI) is defined as the output of a system. The parameters are replaced by independent random variables to achieve an uncertainty quantification. Due to the concept of the polynomial chaos, the QoI is expanded into a series with orthogonal basis polynomials depending on the random parameters. Using all polynomials up to a moderate total degree, a huge number of basis functions occurs, because the num-

ber of random parameters is relatively large. Our aim is to identify a sparse approximation for the QoI, where only a few basis polynomials are required for a sufficiently accurate representation. We apply methods of model order reduction (MOR) to construct a sparse approximation. In the case of linear dynamical systems, the input-output behavior is described by a transfer function in the frequency domain. The approximation error of the sparse representation can be estimated by Hardy norms of the transfer functions. Typical MOR techniques for linear systems are Krylov subspace methods or balanced truncation, for example. We also investigate the more difficult case of nonlinear dynamical systems. Now a transfer function is not available in general. For an MOR, we apply proper orthogonal decomposition. Furthermore, numerical results are shown for test examples, which represent mathematical models of electric circuits.

Roland Pulch  
University of Greifswald  
pulchr@uni-greifswald.de

#### MS196

##### **Error Correction for Multivariate Function Approximation**

We discuss the problem of constructing an accurate function approximation when data are corrupted by unexpected errors. The unexpected corruption errors are different from the standard observational noise, in the sense that they can have much larger magnitude and in most cases are sparse. By focusing on overdetermined case, we prove that the sparse corruption errors can be effectively eliminated by using  $\ell_1$ -minimization, also known as least absolute deviations method. We establish probabilistic error bounds of the  $\ell_1$ -minimization solution with the corrupted data. Several numerical examples are presented to verify the theoretical finding.

Yeonjong Shin, Dongbin Xiu  
Ohio State University  
shin.481@osu.edu, xiu.16@osu.edu

#### MS196

##### **New Sufficient Conditions for Sparse Recovery via Nonconvex Minimizations**

In this talk, we present a new theory for compressed sensing that reveals that nonconvex minimizations are at least as good as  $\ell_1$  minimization in exact recovery of sparse signals. Our theoretical recovery guarantees are developed through a unified null space property based-condition that encompasses all currently proposed nonconvex functionals in literature. Several nonconvex functionals will be explored and the specific conditions in order to guarantee improved recovery will be given. Numerical examples, related to polynomial approximation of complex-valued functions in high dimensions, will be provided to support the new theory and demonstrate the computational efficiency of nonconvex minimizations.

Clayton G. Webster  
Oak Ridge National Laboratory  
webstercg@ornl.gov

Hoang A. Tran  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division

tranha@ornl.gov

#### MS197

##### **A Model for the Transport of Miscible Fluids in the Presence of Anomalous Diffusion**

The modeling of fluid mixing during transport in porous media is critical to engineering applications such as enhanced recovery in oil and gas operations and ensuring environmental safety with respect to contaminant transport in aquifers. Recently, a strongly nonlocal model was developed that utilizes the mathematical tools and terminology from peridynamic mechanics to model fluid transport in porous media. This model recovers the classical theory as a special case as well as has the capability to model anomalous diffusive transport. Here we present an extension to this model for miscible fluids that takes into account convective processes. We present the model and a meshfree numerical technique to study the effect of small-scale heterogeneities in permeability and viscosities on the macroscopic mixing of two fluids.

John Foster  
The University of Texas at Austin  
jfoster@austin.utexas.edu

#### MS197

##### **Double Well Potentials and Nonlocal Brittle Fracture Modeling**

We formulate a nonlocal cohesive model of peridynamic type for calculating dynamic fracture. The force interaction is derived from a double well strain energy density function, resulting in a nonmonotonic material model. The material properties change in response to evolving internal forces eliminating the need for a separate phase field to model the fracture set. The model has the capacity to simulate nucleation and growth of multiple, mutually interacting dynamic fractures. In the limit of zero region of integration, the model recovers a sharp interface evolution characterized by the classic Griffith free energy of brittle fracture with elastic deformation satisfying the linear elastic wave equation off the crack set. These results are reported in Lipton, R. J *Elast* (2016) 124: 143, Lipton, R. J *Elast* (2014) 117: 21.

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

#### MS197

##### **Optimization-Based Coupling for Local and Nonlocal Models**

Peridynamics, a nonlocal extension of continuum mechanics, is well suited to capture material failure. The ability to include peridynamics in engineering analyses, however, remains limited, due largely to the need for an effective method to couple local and nonlocal models. In this work, we formulate an optimization-based coupling strategy for local and nonlocal elasticity. Our approach, as previously demonstrated on local and nonlocal diffusion equations, is to formulate the coupling as a control problem in which the objective is to minimize model mismatch in an overlap region [D'Elia, et al., 2016]. The states in the control problem are the solutions of the local and nonlocal equations, and the controls are boundary conditions for the local model

and volume constraints for the nonlocal model. We provide mathematical analysis of the coupling strategy and present computational results for three-dimensional static elasticity problems in which the local model is discretized with the standard finite element method and the peridynamic model is discretized with a meshfree approach. The optimization-based coupling approach is shown to be a viable means for combining local and nonlocal models within a single computational simulation, and for a number of configurations eliminates the need to apply (non-control) volume constraints to the nonlocal model, which is often problematic in practice.

David Littlewood, Marta D'Elia  
Sandia National Laboratories  
djlittle@sandia.gov, mdelia@sandia.gov

Mauro Perego  
CSRI Sandia National Laboratories  
mperego@sandia.gov

Pavel Bochev  
Sandia National Laboratories  
Computational Math and Algorithms  
pboche@sandia.gov

#### MS197

##### Material Stability and Numerical Stability in Peridynamics

The peridynamic mechanical theory is a nonlocal generalization of classical continuum mechanics that applies to systems of discrete particles as well as continuous media with or without defects. The most commonly used numerical discretization of the peridynamic equations, when applied to continuous bodies, uses midpoint quadrature to sum up the force interactions between the material represented by the nodes. The discretized equations for a continuum turn out to be identical to the original peridynamic equations applied to a system of discrete particles. In this sense, the discretized model is exact. The practical implications of this observation include the manifestation of material instability as numerical instability, particularly with regard to zero-energy modes. By exploring material stability conditions, we arrive at a stabilization method for the particle discretization.

Stewart Silling  
Sandia National Laboratories  
sasilli@sandia.gov

#### MS198

##### Approximate Marginal MCMC for Goal-Oriented Inference with Expensive Models

We develop methods to characterize marginals of high-dimensional probability distributions in a setting where evaluations of the unnormalized joint density are computationally demanding. Our primary motivation is Bayesian computation: in many inference problems, only the marginal distribution of a low-dimensional quantity is of interest. While exact evaluation of the marginal density is often impossible, sample estimates of this density may be obtained via importance sampling or other means. This is the setting for pseudo-marginal algorithms, but our approach is different: we use consistent local polynomial approximations, built from noisy estimates, to approximate the underlying marginal density. The approximation is continually refined in conjunction with an MCMC

scheme that ensures asymptotically exact sampling from the desired marginal. Overall our approach exploits regularity in the marginal density in order to significantly reduce computational expense relative to both regular and pseudo-marginal MCMC. We also determine an “optimal” rate for refining the local polynomial approximation of the marginal density. While asymptotic results provide a useful exactness guarantee, in the non-asymptotic regime excessive refinement can waste computation while infrequent refinement can introduce a large bias. We develop a refinement scheme that ensures that the bias due to the approximation decays at roughly the same rate as the error due to finite MCMC sampling.

Andrew D. Davis  
MIT  
davisad@mit.edu

Natesh Pillai  
Statistics  
Harvard  
pillai@fas.harvard.edu

Aaron Smith  
University of Ottawa  
asmith3@math.stanford.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

#### MS198

##### Improving the Efficiency of Implicit Sampling Using Surrogate Models

Bayesian inverse modeling techniques are computationally expensive because many forward simulations are needed when sampling the posterior distribution of the parameters. We combine the implicit sampling method and generalized polynomial chaos expansion (gPCE) to significantly reduce the computational cost of performing Bayesian inverse modeling. There are three steps in this approach: (1) find the maximizer of the likelihood function using deterministic approaches; (2) construct a gPCE-based surrogate model using the results from a limited number of forward simulations; and (3) efficiently sample the posterior distribution of the parameters using implicit sampling method. The cost of constructing the gPCE-based surrogate model is further decreased by using sparse Bayesian learning to reduce the number of gPCE coefficients that have to be determined. We demonstrate the approach for a synthetic ponded infiltration experiment simulated with a subsurface flow model. The surrogate model is highly accurate with mean relative error that is  $< 0.035\%$  in predicting saturation and  $< 0.25\%$  in predicting the likelihood function. The posterior distribution of the parameters obtained using our proposed technique is nearly indistinguishable from the results obtained from either an implicit sampling method or a Markov chain Monte Carlo method utilizing the full model.

George Shu Heng Pau  
LBNL  
gpau@lbl.gov

Yaning Liu  
Lawrence Berkeley National Laboratory



yaningliu@lbl.gov

### MS198

#### Sequential Implicit Sampling Methods for Bayesian Inverse Problems

The solution to the inverse problems, under the Bayesian framework, is given by a posterior probability density. For large scale problems, sampling the posterior can be an extremely challenging task. Markov Chain Monte Carlo (MCMC) provides a general way for sampling but it can be computationally expensive. Gaussian type methods, such as the Ensemble Kalman Filter (EnKF), make Gaussian assumptions even for the possible non-Gaussian posterior, which may lead to inaccuracy. In this talk, the implicit sampling method and the newly proposed sequential implicit sampling method are investigated for the inverse problem involving time-dependent partial differential equations (PDEs). The sequential implicit sampling method combines the idea of the EnKF and implicit sampling and it is particularly suitable for time-dependent problems. Moreover, the new method is capable of reducing the computational cost in the optimization, which is a necessary and the most expensive step in the implicit sampling method. The sequential implicit sampling method has been tested on a seismic wave inversion. The numerical experiments show its efficiency by comparing it with the MCMC and some Gaussian approximation methods.

Xuemin Tu

University of Kansas  
xuemin@ku.edu

### MS198

#### Aspects of Particle Filtering in Very High Dimensional Systems

Sequential Bayesian Inference in the geosciences is called data assimilation. It is characterised by high-dimensional nonlinear models, with state-vector dimensions up to  $10^9$  and observation vectors of dimension  $10^7$  at specific time intervals. Often the map from model to observation space is also nonlinear. Present-day data-assimilation methods are based on linearisations, but with every increasing complexity of the models and observations fully nonlinear data assimilation methods are needed. Here we explore extremely efficient particle filters that beat the curse of dimensionality. There is a growing interest in so-called localised particle filters that effectively split the full problem up in many low-dimensional problems. However, there are both practical and theoretical issues with this localisation. We will instead explore the proposal density freedom in particle filters. Specifically we will look at particle filters that operate in two steps. They first target the area of high-probability mass, the so-called typical set. Then they explore Hamiltonian Monte Carlo methods to move the particles over the typical set such that they all obtain the same weight. Both steps are integrated in the proposal density. The new particle filter will be explored in high-dimensional applications, including large numbers of independent Lorenz1963 models and climate models, and methods to estimate the essential covariance of errors in the model equations are discussed.

Peter Jan van Leeuwen  
University of Reading

p.j.vanleuween@reading.ac.uk

### MS199

#### Shape Constrained Tensor Factorizations

We consider  $N$ -way data arrays and low-rank tensor factorizations where the time mode is coded as a sparse linear combination of temporal elements from an over-complete library. Our method, Shape Constrained Tensor Decomposition (SCTD) is based upon the CANDECOMP/PARAFAC (CP) decomposition which produces  $r$ -rank approximations of data tensors via outer products of vectors in each dimension of the data. By constraining the vector in the temporal dimension to known analytic forms which are selected from a large set of candidate functions, more readily interpretable decompositions are achieved and analytic time dependencies discovered. The SCTD method circumvents traditional *flattening* techniques where an  $N$ -way array is reshaped into a matrix in order to perform a singular value decomposition. A clear advantage of the SCTD algorithm is its ability to extract transient and intermittent phenomena which is often difficult for SVD-based methods. We motivate the SCTD method using several intuitively appealing results before applying it on a number of high-dimensional, real-world data sets in order to illustrate the efficiency of the algorithm in extracting interpretable spatio-temporal modes. With the rise of data-driven discovery methods, the decomposition proposed provides a viable technique for analyzing multitudes of data in a more comprehensible fashion.

Eric Chi

Department of Statistics  
North Carolina State University  
eric\_chi@ncsu.edu

Bethany Lunch, Nathan Kutz

University of Washington  
Dept of Applied Mathematics  
herwaldt@uw.edu, kutz@uw.edu

### MS199

#### Distributed Nonnegative Tensor Low Rank Approximation for Large-Scale Clustering

In Nonnegative Tensor Factorization (NTF) for a nonnegative order  $d$  tensor, we seek for  $k$  rank-1 tensors whose sum approximates the tensor, where each of the  $d$  factor matrices with  $k$  columns is nonnegative. We explored a general framework of block coordinate descent (BCD) method for NTF in our earlier work. As the data grows larger and larger, there is a need for distributed NTF algorithms in both more classical scientific applications as well as data analysis problems. In this talk, we present some of our recent results for distributed NTF and applications. The computationally expensive and communication intensive step in NTF is Matricized Tensor Times Khatri Rao Product (MTTKRP). We designed an efficient technique for efficiently setting up a message passing interface (MPI) communications. To tame the communication cost of MTTKRP, similar to 2D communicators across rows and columns for matrices, we propose to have communicators across processor fibers and processor hyperslices. We will present the initial benchmarking of our distributed NTF on the Oak Ridge Leadership Computing Facility (OLCF) cluster for both sparse and dense tensors and illustrate the results for clustering applications.

Ramakrishnan Kannan

Oak Ridge National Laboratories  
kannanr@ornl.gov

Grey Ballard  
Sandia National Laboratories  
ballard@wfu.edu

Barry Drake  
Georgia Tech Research Institute  
barry.drake@gtri.gatech.edu

Haesun Park  
Georgia Institute of Technology  
hpark@cc.gatech.edu

### MS199

#### The Tensor Approximation Problem from a Dynamical Systems Perspective

The computation of low-rank approximations to tensors is a challenging optimization problem, for which current algorithms are routinely slow and unreliable. To enable the creation of better algorithms, the optimization problem itself needs to be better understood. To develop this understanding, we have been analyzing the simplest non-trivial cases of the tensor approximation problem. Using methods from dynamical systems, we identify features in the optimization landscape that are difficult for algorithms to handle.

Martin J. Mohlenkamp  
Ohio University  
mohlenka@ohio.edu

### MS199

#### Fast Algorithms for Higher-Order Singular Value Decomposition with Missing Data

Higher-order singular value decomposition (HOSVD) is an efficient way for data reduction and also eliciting intrinsic structure of multi-dimensional array data. It has been used in many applications, and some of them involve incomplete data. We formulate an incomplete HOSVD problem, which simultaneously achieves imputation of missing values and also tensor decomposition. We also present one algorithm for solving the problem based on block coordinate update (BCU). Global convergence of the algorithm is shown under mild assumptions and implies that of the popular higher-order orthogonality iteration (HOOI) method, and thus we, for the first time, give global convergence of HOOI. In addition, we compare the proposed method to state-of-the-art ones for solving incomplete HOSVD and also low-rank tensor completion problems and demonstrate the superior performance of our method over other compared ones. Furthermore, we apply it to face recognition and MRI image reconstruction to show its practical performance.

Yangyang Xu  
Institute for Mathematics and its Application  
yxu76@ua.edu

### MS200

#### Time-Evolving Graph Processing on Commodity Clusters

Real-world graphs are seldom static. Applications that generate graph-structured data today do so continuously, giving rise to an underlying graph whose structure evolves

over time. Mining these time-evolving graphs can be insightful, both from research and business perspectives. While several works have focused on some individual aspects of time-evolving graph processing such as efficient incremental computations (e.g., Naiad and its GraphLINQ library), consistent snapshot generation (e.g., Kineograph, LLAMA, DeltaGraph) or storage support for highly dynamic graphs (e.g., Weaver), there exists no general purpose distributed time-evolving graph processing engine. In this paper, we present Tegra, a time-evolving graph processing system built on a data flow framework. Tegra enables three broad classes of operations on evolving graphs: first, it enables storage, retrieval and bulk transformation of multiple graph snapshots efficiently using a structure-sharing persistent data structure based index. Second, it supports temporal graph analysis tasks such as evolutionary queries using a novel timelapse abstraction that lets it process multiple snapshots simultaneously with low overhead. Finally, Tegra enables a lightweight dynamic computation model, based on the Gather-Apply-Scatter (GAS) decomposition model, that lets it do sliding window analytics on streaming graphs. Our evaluation of Tegra on real-world datasets show promising results.

Anand Iyer  
UC Berkeley  
api@cs.berkeley.edu

Ion Stoica  
University of California, Berkeley, USA  
istoica@cs.berkeley.edu

### MS200

#### High-Performance Analysis of Streaming Graphs

Graph-structured data in social networks, finance, network security, and others not only are massive but also under continual change. These changes often are scattered across the graph. Stopping the world to run a single, static query is infeasible. Repeating complex global analyses on massive snapshots to capture only what has changed is inefficient. We discuss requirements for single-shot queries on changing graphs as well as recent high-performance algorithms that update rather than recompute results. These algorithms are incorporated into our software framework for streaming graph analysis, STING (Spatio-Temporal Interaction Networks and Graphs).

Jason Riedy  
Georgia Institute of Technology  
School of Computational Science and Engineering  
jason.riedy@cc.gatech.edu

### MS200

#### Dense Subgraphs in Temporal Networks: Algorithms and Analysis

The underlying algorithmic challenge with finding dense graphs is that most standard formulations of this problem (like clique, quasi-clique, k-densest subgraph) are NP-hard. Furthermore, current dense subgraph finding algorithms usually optimize some objective in stationary settings, and only find a few such subgraphs without providing any structural relations among them. However, the goal is rarely to find the "true optimum," but to identify many (if not all) dense substructures. On the other hand, most real-world networks are highly dynamic, but existing practical dense subgraph finding algorithms are only able to handle stationary graphs. We present algorithms to find dense

subgraphs in temporal networks and build the hierarchy structure among them. We also use our algorithms to analyze some real-world networks and present new insights.

A. Erdem Sariyuca  
The Ohio State University  
asariyu@sandia.gov

Ali Pinar  
Sandia National Labs  
apinar@sandia.gov

## MS200

### Parallel and Streaming Methods for Real-Time Analysis of Dense Structures from Graphs

Abstract not available

Srikanta Tirthapura  
Department of Electrical and Computer Engineering  
Iowa State University  
snt@iastate.edu

Apurba Das, A. Pavan, Michael Svendsen  
Iowa State University  
adas@iastate.edu, pavan@cs.iastate.edu,  
michael.sven5@gmail.com

Kanat Tangwongsan  
Mahidol University International College  
kanat.tan@mahidol.edu

Kung-Lun Wu  
IBM T.J. Watson Research Center  
klwu@us.ibm.com

## MS201

### Scalable Methods for Optimization of Systems Governed by PDEs with Random Parameter Fields

We focus on risk-averse optimal control of systems governed by PDEs with uncertain coefficient functions. In particular, we seek controls that minimize an objective function incorporating the mean and variance of the control objective. To make the optimization problem computationally tractable, we use a second order Taylor series approximation of the control objective with respect to the uncertain parameter field. This enables deriving closed-form expressions for the mean and variance of the control objective in terms of its gradient and Hessian with respect to the uncertain parameter. For illustration, we apply our method to optimal control of a porous medium flow model with a random permeability field.

Alen Alexanderian  
NC State University  
alexanderian@ncsu.edu

Noemi Petra  
University of California, Merced  
npetra@ucmerced.edu

Georg Stadler  
Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

## MS201

### Multifidelity Computation of Failure Probabilities for Systems with Uncertain Parameters

Failure probability estimation is concerned with computing the probability that a quantity of interest (QoI), modeled as a random variable, exceeds a predefined threshold. Generally, such failure probabilities are small, and standard Monte Carlo estimation of the failure integral requires many samples. This gets complicated by the fact that we consider cases where evaluating the QoI requires solving a computationally expensive model (e.g., a PDE). We thus propose to use importance sampling to reduce the number of expensive model evaluations. Using a suite of surrogate models, we evaluate the QoI, and design a biasing distribution, that is biased towards failure events. The evaluation of the surrogates is computationally cheap. With this biasing distribution, we can evaluate the original failure probability with much fewer samples - and hence much fewer evaluations of the expensive model - while achieving high accuracy. Numerical examples on a combustion-type test problem illustrate our approach.

Boris Kramer, Karen E. Willcox  
Massachusetts Institute of Technology  
bokramer@mit.edu, kwillcox@MIT.EDU

Benjamin Peherstorfer  
ACDL, Department of Aeronautics & Astronautics  
Massachusetts Institute of Technology  
pehersto@mit.edu

## MS201

### Low Rank Solvers for PDE-Constrained Optimization with Uncertain Parameters

We consider the simulation of an optimal control problem constrained by an unsteady partial differential equation involving random data. We consider a generalized polynomial chaos approximation of these random functions in the stochastic Galerkin finite element method. The discrete problem yields a prohibitively high dimensional saddle point system with Kronecker product structure. We discuss how such systems can be solved using tensor-based methods. These methods allow the solution of previously untractable problem sizes and typically only require a small fraction of the storage when compared to traditional techniques. The performance of our approach is illustrated with extensive numerical experiments on both the heat equation and the Stokes-Brinkman equations. We further show that this approach can be used for a nonlinear problem involving the Navier-Stokes equations as constraints.

Martin Stoll  
Max Planck Institute, Magdeburg  
stollm@mpi-magdeburg.mpg.de

Peter Benner  
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany  
benner@mpi-magdeburg.mpg.de

Sergey Dolgov  
University of Bath

Department of Mathematical Sciences  
s.dolgov@bath.ac.uk

Akwum Onwunta  
Max Planck Institute, Magdeburg, Germany  
onwunta@mpi-magdeburg.mpg.de

### MS201

#### Multilevel Monte Carlo Analysis for Optimal Control of Elliptic PDEs with Random Coefficients

This work is motivated by the need to study the impact of data uncertainties and material imperfections on the solution to optimal control problems constrained by partial differential equations. We consider a pathwise optimal control problem constrained by a diffusion equation with random coefficient together with box constraints for the control. For each realization of the diffusion coefficient we solve an optimal control problem using the variational discretization [M. Hinze, *Comput. Optim. Appl.*, 30 (2005), pp. 45-61]. Our framework allows for lognormal coefficients whose realizations are not uniformly bounded away from zero and infinity. We establish finite element error bounds for the pathwise optimal controls. This analysis is nontrivial due to the limited spatial regularity and the lack of uniform ellipticity and boundedness of the diffusion operator. We apply the error bounds to prove convergence of a multilevel Monte Carlo estimator for the expected value of the pathwise optimal controls. In addition we analyze the computational complexity of the multilevel estimator. We perform numerical experiments in 2D space to confirm the convergence result and the complexity bound.

Ahmad Ahmad Ali, Michael Hinze  
Universität Hamburg  
Department Mathematik  
ahmad.ali@uni-hamburg.de,  
michael.hinze@uni-hamburg.de

Elisabeth Ullmann  
TU München  
elisabeth.ullmann@ma.tum.de

### MS202

#### Expansion Mechanisms for Volume and Layer Potentials in Quadrature By Expansion

The rapid evaluation of layer and volume potentials forms the cornerstone of the numerical solution of PDEs with the help of integral equation methods. Quadrature by Expansion is an effective method for evaluating these potentials with high order accuracy up to controlled precision. To do so, the method relies on series expansions of the potential. It offers tremendous flexibility in the choice of series expansions, and I will discuss a variety of novel choices in this talk, along with estimates of truncation error and fast algorithms allowing their practical use.

Andreas Kloeckner  
University of Illinois  
andreask@illinois.edu

### MS202

#### Accurate Derivative Evaluation for Grad-Shafranov Solvers Using Quadrature by Expansion (QBX)

For several applications in magnetic confinement fusion,

such as the design of tokamaks, the stream function associated with the magnetic flux satisfies a semi-linear Poisson equation, the Grad-Shafranov equation. In these devices, the physical quantities of interest are the magnetic field, the parallel current density, and the magnetic curvature, which depend on the first and second derivatives of the stream function. This poses a major challenge for traditional finite element based solvers, as they tend to lose up to two orders of convergence in computing these physical quantities. We present a numerical method for the computation of first and second derivatives of the stream function associated with the magnetic flux; the method has the same order of convergence for the first and second derivatives of the stream function as the stream function itself. The key idea of this method is to construct ‘Dirichlet-to-Neumann’ maps for the solution operator. In this process, high order computation of the derivatives of a volume layer potential on the boundary of the domain is required. We present a variant of Quadrature by Expansion (QBX) which computes these derivatives to high order on the boundary. We illustrate the performance of our method through several numerical examples.

Manas N. Rachh  
Applied Mathematics Department  
Yale University  
manas.rachh@yale.edu

Lee F. Ricketson  
NYU  
Courant Institute  
lfr224@cims.nyu.edu

Antoine Cerfon  
NYU  
cerfon@cims.nyu.edu

Jeffrey Freidberg  
Massachusetts Institute of Technology  
jpfreid@mit.edu

### MS202

#### A Local Target-Specific QBX Method for Laplace’s Equation in 3D Multiply-Connected Domains

We present a new version of the QBX method for accurate evaluation of layer potentials at target points on or near the surface of integration. The QBX method makes use of the fact that the function described by a layer potential is smooth in the interior or exterior region and can be expanded in a Taylor’s (or other) series about an interior/exterior point or ‘center’. Such expansions about multiple centers are then used to accurately evaluate the boundary integral on or near the boundary surface. Our version of the method uses several key ideas to improve its efficiency, making it useful for some large scale problems. For one, we utilize a special ‘target-specific’ Taylor’s expansion of the Green’s function  $G(x,y)$  in which the 3D expansion of order  $p$  about a center  $c$  is computed in only  $O(p)$  operations. The method is applied to some challenging boundary value problems for Laplace’s equation in 3D multiply-connected domains. This is joint work with Anna-Karin Tornberg.

Michael Siegel  
New Jersey Institute of Technology

misieg@njit.edu

### MS202

#### Deriving Estimates for the Quadrature Error in Quadrature by Expansion (QBX)

Quadrature by expansion (QBX) is a method for evaluating layer potentials on or close to the boundary, where the underlying integral can be singular or nearly singular. The method is based on forming local expansions of the layer potential at a distance from the boundary, where the potential is smooth. However, the expansion coefficients are evaluated using integrals which are themselves nearly singular, and sufficient upsampling of the boundary data is necessary to suppress the error due to this. We will discuss how nearly singular quadrature errors of standard quadrature rules can be accurately estimated using a method based on contour integration and a calculus of residues. These error estimates can in practical applications be used for (i) determining the distance from the boundary inside which special quadrature is required, and (ii) determining the upsampling rate required for computing the QBX expansion coefficients to desired accuracy.

Ludvig af Klinteberg

Royal Institute of Technology (KTH)  
ludvigak@kth.se

### MS203

#### A Comprehensive Study of Differences Between Sequential Task Flow and a Real Data Flow Representation

Abstract not available

George Bosilca

University of Tennessee - Knoxville  
bosilca@icl.utk.edu

Thomas Herault

Univ. of Tennessee - Knoxville  
herault@eecs.utk.edu

### MS203

#### Using Kokkos for Performance Portability of the Tpetra Sparse Linear Algebra Library on Intel KNL and Nvidia GPUs

The Kokkos shared-memory programming model has made it easier for Trilinos to make its sparse linear algebra and linear solvers support thread parallelism on many different computer architectures. This talk will summarize the porting effort, focusing on software challenges that other scientific libraries and applications may face.

Mark Hoemmen

Sandia National Laboratories  
mhoemme@sandia.gov

### MS203

#### Overview of a Sequential Task Flow Sparse QR Multifrontal Solver

Abstract not available

Emmanuel Agullo

INRIA  
emmanuel.agullo@inria.fr

Alfredo Buttari

CNRS-IRIT-Université de Toulouse, France  
alfredo.buttari@enseeiht.fr

Abdou Guermouche

LaBRI-INRIA futurs  
abdou.guermouche@labri.fr

Ian Masliah

University Paris-Sud  
ian.masliah@lri.fr

### MS203

#### Programming Model, Performance Analysis and Optimization Techniques for the Intel Knights Landing Xeon Phi

A wide variety of heterogeneous compute resources, ranging from multicore CPUs to GPUs and coprocessors, are available to modern computers, making it challenging to design unified numerical libraries that efficiently use all these varied resources. For example, in order to efficiently and productively use Intel's Knights Landing (KNL), one must design and schedule an application in multiple degrees of parallelism and task grain sizes in order to obtain efficient performance. We propose a productive and portable programming model that allows us to write a serial-looking code, which achieves parallelism and scalability by using a lightweight runtime environment to manage the resource-specific workload, and to control the dataflow and the parallel execution. This is done through multiple techniques ranging from multi-level data partitioning to adaptive task grain sizes, and dynamic task scheduling. We outline the strengths and the effectiveness of this approach on heterogeneous systems.

Azzam Haidar

Department of Electrical Engineering and Computer Science  
University of Tennessee, Knoxville  
haidar@icl.utk.edu

Stanimire Tomov

Computer Science Department  
University of Tennessee  
tomov@icl.utk.edu

Konstantin Arturov, Murat E. Guney, Shane Story

Intel Corporation  
konstantin.i.arturov@intel, murat.e.guney@intel.com, shane.story@intel.com

Jack J. Dongarra

University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

Tyler McDaniel

University of Tennessee  
bmcDan16@utk.edu

### MS204

#### Eigenvalues of Sturm-Liouville Problems with Real Coupled Boundary Conditions

It has been shown in [Bailey and Zettl, Sturm-Liouville Eigenvalue Characterizations, Electronic Journal of Differential Equations 2012 (2012), No. 123, pp. 1-13] that all

the eigenvalues of Sturm-Liouville problems with real coupled boundary conditions are extrema of a one-parameter family of problems with separated boundary conditions. Here we show how this characterization can be used to compute the eigenvalues of the coupled problem using any computer code which works for problems with separated boundary conditions.

Bailey (Emeritus), Sandia National Laboratories, Albuquerque, New Mexico

Anton Zettl (Emeritus), Northern Illinois University, DeKalb, Illinois

Paul B. Bailey

Northern Illinois University, Assoc. Professor  
paulbailey10950@comcast.net

Anton Zettl

Northern Illinois University, Emeritus Prof  
zettl@msn.com

#### MS204

##### Spectral Pollution Arising in Eigenvalue Computations for Schroedinger Operators

We study numerical computations of spectra of Schrödinger operators  $T$  in  $L^2(\{R\}^d)$ . In particular if the spectrum is not purely discrete and/or the potential is not real-valued, numerical values may suggest a completely wrong location of the spectrum. First we consider the approximation of  $T$  by compressing it to an  $n$ -dimensional subspace of the domain of  $T$  (called Galerkin/finite section/projection method). It is well known that, in the limit  $n \rightarrow \infty$ , the eigenvalues of the  $n \times n$  matrix  $T_n$  may accumulate at a point that does not belong to the spectrum of  $T$ . The occurrence of such a spurious eigenvalue is commonly known as spectral pollution. We present methods to identify regions in the complex plane that enclose the set of spectral pollution (as tightly as possible). In the second part of the talk we present spectral convergence results for approximations of  $T$  by truncating it to bounded but expanding domains in  $\{R\}^d$ , subject to various boundary conditions at the artificial boundary.

Sabine Bögli, Universität München, München, Germany  
Marco Marletta, Cardiff University, Cardiff, United Kingdom

Petr Siegl, University of Bern, Bern, Switzerland

Christiane Tretter, University of Bern, Bern, Switzerland

Sabine Boegli

Universitaet Muenchen, Munich, Germany  
boegli@math.lmu.de

Marco Marletta

Wales Institute of Mathematical and Computational Sciences  
Cardiff University  
marco.marletta@cs.cardiff.ac.uk

Christiane Tretter, Petr Siegl

University of Bern, Switzerland

tretter@math.unibe.ch, petr.siegl@math.unibe.ch

#### MS204

##### Spectral Density Functions Associated with Fokker-Planck Collisions

We study a parabolic equation on the half-line that arises in continuum kinetic calculations of energy diffusion in veloc-

ity space due to Fokker-Planck collisions in plasma physics. The exact solution is expressed as a continuous superposition of (non-normalizable) eigenfunctions via a generalized Fourier transform. We show how to efficiently compute these eigenfunctions, as well as the spectral transform of the initial data and the spectral density function of the inverse transform, using high-order timestepping schemes (in the space-like coordinate). We compare two methods of computing the spectral density function, one based on extrapolating the value of the Titchmarsh-Weyl  $m$ -function from the complex upper half-plane to the real axis using Chebyshev interpolation, due to Wilkening and Cerfon, and the other based on a recursive method of Fulton, Pearson and Pruess.

Rockford Foster, University of California, Berkeley  
Jon Wilkening, University of California, Berkeley

Rockford D. Foster

University of California, Berkeley  
rocky@math.berkeley.edu

Jon Wilkening

UC Berkeley Mathematics  
wilken@math.berkeley.edu

#### MS204

##### Spectral Density Functions for Periodic Schrödinger Potentials

We consider the Schrödinger equation

$$-y'' + q(x)y = \lambda y,$$

with potential  $q(x)$ , periodic of period  $\ell$ , on the half line and on the whole line (see Fulton, Pearson, Pruess, arxiv:1303.5878). It is well known that there is a band-gap structure: Namely, if the eigenvalues of the periodic problem ( $y(0) = y(\ell), y'(0) = y'(\ell)$ ) are ordered by  $\lambda_0 \leq \lambda_1 \leq \dots$ , and the eigenvalues of the semi-periodic problem ( $y(0) = -y(\ell), y'(0) = -y'(\ell)$ ) are ordered by  $\mu_0 \leq \mu_1 \leq \dots$ , then  $-\infty < \lambda_0 < \mu_0 \leq \mu_1 < \lambda_1 \leq \lambda_2 \dots$ , and the a.c. spectrum on the half line consists of the closed intervals  $[\lambda_{2i}, \mu_{2i}]$  and  $[\mu_{2i+1}, \lambda_{2i+1}]$ ,  $i = 0, 1, \dots$ . The eigenvalues  $\mu_j$  and  $\lambda_j$  may be computed using SLEIGN2 by the method of Bailey and Zettl. In this talk, we present an algorithm, and numerical results, for computation of the a.c. spectral density function on the bands. This has the feature that a numerical integration is required only over one period. For this we implement the Pruess method (that is, MG2), and the Magnus Method MG4 of Iserles, et. al.

Charles T. Fulton

Florida Institute of Technology, USA  
cfulton@fit.edu

Steven Pruess

Colorado School of Mines, Emeritus Prof  
spruess6@cox.net

David Pearson

University of Hull, U.K. Emeritus Prof  
d.b.pearson@hull.ac.uk

#### MS205

##### Simulation of Flow in Artificial Blood Pumps: Turbulence Modeling and Consideration of Blood Microstructure

Modeling and computational analysis play an increasingly-

important role in bioengineering, particularly in the design of implantable ventricular assist devices (VAD). Numerical simulation of blood flow and associated physiological phenomena has the potential to shorten the design cycle and give the designers important insights into causes of blood damage and suboptimal performance. A set of modeling techniques is presented which are based on stabilized space-time finite element formulation of the Navier-Stokes equations. Alternate methods that represent the rotating components in an averaged sense using a rotating frame of reference, as well as influence of turbulence modeling, will be discussed. In order to obtain quantitative hemolysis prediction, cumulative tensor-based measures of strain experienced by individual blood cells are being developed; red blood cells under shear can be modeled as deforming droplets, and their deformation tracked along pathlines of the computed flow field. An efficient continuum-based approach with rescaling of the variables using logarithmic transformation will be examined. The methods are applied to simplified rotary blood pumps, one of which is currently a subject of an FDA round-robin study.

Marek Behr

RWTH Aachen University  
Chair for Computational Analysis of Technical Systems  
behr@cats.rwth-aachen.de

#### MS205

##### **On the Asymptotic Behavior of Subgrid Scale Models for Large-Eddy Simulations in Complex Geometries**

The capability of several eddy-viscosity based subgrid scale models to mimic the expected turbulence damping near solid boundaries is discussed. Among the variety of static models available in the literature, only a few reproduce the proper cubic behavior of the eddy-viscosity in the case of an incompressible fluid. Only one (the  $\sigma$ -model of Nicoud et al., Phys. of Fluid, 2011) is found to also achieve the theoretical quadratic damping in the case where the divergence-free condition is not met at the wall (e.g.: because of density fluctuations induced by either temperature or mixing variations). When the dynamic procedure is used to adapt the model constant in complex geometries (as in the dynamic Smagorinsky model, Germano et al., Phys. Fluids, 1991), the traditional stabilization techniques (local averaging, clipping) may lead to a mispredicted near-wall behavior. This is illustrated by using a general purpose Large-Eddy Simulation solver to compute the flow in a well-controlled experiment where a pulsatile hot-jet impinges a flat-plate in the presence of a turbulent cross-flow (Baya Toda et al, Phys. Fluids, 2014). The lack of damping produced by the dynamic Smagorinsky model introduces significant errors during the vortex ring/wall interaction while the agreement with the experimental data is much better when the  $\sigma$ -model is used.

Franck Nicoud

CNRS - UMR5149  
University of Montpellier  
franck.nicoud@umontpellier.fr

#### MS205

##### **Large Eddy Simulation of Aortic Flow Using Deconvolution-Based Nonlinear Filter – Theory and Application**

Type B Aortic Dissection (TBAD), one of the highly lethal aortic syndrome, is normally managed conservatively with

medical treatment unless the patients demonstrate complications. However, long-term outcomes indicate significant late mortality in patients with uncomplicated TBAD (uTBAD). Clinical outcomes have shown that patients receiving early thoracic endovascular aortic repair (TEVAR) are associated with less adverse events compared to those who did not. Unfortunately, well-defined clinical indication factors do not currently exist for early TEVAR due to large individual variations. To identify well-defined prognostic factors and the possible risk of late complications, we simulate the hemodynamics of uTBAD on a patient-specific basis, using a Large Eddy Simulation (LES) model [bertagna, Deconvolution-based nonlinear filtering for incompressible flows at moderately large Reynolds numbers, 2015]. We find that the LES solver is able to resolve the high flow rates and associated numerical instabilities prone to occur near the branch of the false lumen and narrowed true lumen using only a fraction of the computational resources required by the DNS solver. These results indicate the LES approach is extremely promising for clinical applications where large volumes of patient-specific CFD simulations need to be executed in a cost-efficient, stable and accurate manner.

Huijuan Xu

GA Tech and Emory University  
huijuanxu.xu@gmail.com

Marina Piccinelli

Department of Radiology, Emory University  
Atlanta, USA  
mpiccin@emory.edu

Bradley Leshnower

School of Medicine  
Emory University  
bleshno@emory.edu

Luca Bertagna

Florida State University  
lbertagna@fsu.edu

Annalisa Quaini

Department of Mathematics, University of Houston  
quaini@math.uh.edu

Robert Nerem

Department of Mechanical Engineering  
Georgia Institute of Technology  
robert.nerem@me.gatech.edu

Alessandro Veneziani

Department of Mathematics and Computer Science  
Emory University  
avenez2@emory.edu

#### MS205

##### **The Reduced Nsalpha Model for Turbulent Channel Flow**

Abstract not available

Camille Zervas

Clemson University  
czervas@clemson.edu

#### MS206

##### **Design Sensitivities in Topology Optimization via**

## Automatic Differentiation

Topology optimization is an iterative design process which distributes a given amount of material in a design domain by minimizing an objective function and by fulfilling a set of constraints. The gradients of the objective and the constraints are often obtained by the adjoint method after discretization of the state equations. The approach is well established for problems with elliptic state equations, where the discretization is obtained by finite element techniques. Discretization of alternative physical processes, like flow problems, conjugate heat transfer, or complex fully coupled multi-physics problems, are often based on alternative schemes (finite volume, DG-FEM). The solutions are obtained either explicitly or iteratively with multiple corrections accounting for the discretization, current solution states or the coupling between the different physics. For such problems the direct application of the adjoint approach is practically impossible due to hidden amendments of the computational process and the high computational complexity. Therefore, the focus of this work is to present, compare and conclude on the applicability of automatic differentiation (AD) techniques for handling the sensitivity analysis in topology optimization. The applicability of AD in topology optimization is demonstrated for a simple wave propagation problem, a turbulent NavierStokes optimization problem and an optimization of unsteady flow based on the lattice Boltzmann method.

Boyan S. Lazarov  
Department of Mechanical Engineering  
Technical University of Denmark  
bsl@mek.dtu.dk

Sebastian Noergaard  
Technical University of Denmark  
sebnorg@mek.dtu.dk

## MS206

### Using a Fully Non-Conformal Geometry Description to Enable Optimization-Based Design

In a typical design/analysis workflow, solid models are constructed in a surface-based representation then converted into a discrete volume-based representation for the simulation phase, i.e., the model is meshed. Decades of research have gone into developing viable tools for this conversion, but meshing of complicated models continues to pose a challenge. Further, during optimization-based design, i.e., shape, topology, material optimization (STMO), parts of the design are computed or modified as a field on the volume-based representation. If the designer is to modify the computed design, the model must be converted back to a form that traditional CAD packages can manipulate. Considerable effort has been invested in developing this capability, but currently no robust, general-use solution exists. Work will be presented that attempts to unify modeling operations and simulation/STMO operations by using a common volume-based geometric description. This solution consists of a non-conformal background discretization and a collection of level set functions that accurately defines the system shape and topology on the background mesh. To this end, the Generalized Finite Element Method has been implemented and integrated with an existing topology optimization capability. This talk will present the details of this approach in the context of combined shape and topology optimization, describe the implementation, and present example applications to demon-

strate the approach.

Joshua Robbins  
Sandia National Laboratories  
jrobbin@sandia.gov

## MS206

### Stochastic Sampling for Structural Topology Optimization with Many Load Cases: Density-Based and Ground Structure Approaches

We propose an efficient probabilistic method to solve a deterministic problem – we present a randomized optimization approach that drastically reduces the enormous computational cost of optimizing designs under many load cases for both continuum and truss topology optimization. Practical structural designs by topology optimization typically involve many load cases, possibly hundreds or more. The optimal design minimizes a, possibly weighted, average of the compliance under each load case (or some other objective). This means that in each optimization step a large finite element problem must be solved for each load case, leading to an enormous computational effort. On the contrary, the proposed randomized optimization method with stochastic sampling requires the solution of only a few (e.g., 5 or 6) finite element problems (large linear systems) per optimization step. Based on simulated annealing, we introduce a damping scheme for the randomized approach. Through numerical examples in two and three dimensions, we demonstrate that the stochastic algorithm drastically reduces computational cost to obtain similar final topologies and results (e.g., compliance) compared with the standard algorithms. The results indicate that the damping scheme is effective and leads to rapid convergence of the proposed algorithm.

Xiaoja Zhang  
Georgia Tech  
xzhang645@gatech.edu

Eric De Sturler  
Virginia Tech  
sturler@vt.edu

Glaucio Paulino  
Georgia Tech  
glaucio.paulino@ce.gatech.edu

## MS206

### Algorithms for the Topology Optimization of Electrical Conductors

Recent research has shown that for large scale topology optimization problems, primal-dual interior point (PDIP) methods can offer favorable convergence when compared with mesh-dependent approaches such as the Method of Moving Asymptotes and its variants [Susana Rojas-Labanda and Mathias Stolpe, Benchmarking optimization solvers for structural topology optimization, Struct. Multidisc. Optim. (2015) 52:527547]. As PDIP methods involve solving a large sparse KKT system with iterative solvers, matrix conditioning can become a serious impediment to convergence of Krylov methods. Multigrid preconditioners have been recently used to accelerate PDIP convergence [Michal Kocvara and Sudaba Mohammed, Primal-Dual Interior Point Multigrid Method for topology optimization, SIAM J. Sci. Comput., Vol. 38, No. 5, pp. B685-B709]. In the current work, we consider preconditioners based on Schur complements which use forward and



adjoint PDE solves as well as use of merit functions and filters. Comparisons are made with reduced space formulations based on trust regions with bound constraints and we identify difficulties associated with primal interior-point methods using barrier functions. PDIP and reduced space methods are applied to the problem of designing electrically conducting structures.

Gregory J. von Winckel  
Sandia National Laboratories  
gvonwin@sandia.gov

#### MS207

##### Advances in Smoothed Particle Hydrodynamics

Smooth particle hydrodynamic methods approximate fluid flow density and velocity as a moving collection of interacting particles. Originally developed for astrophysical simulations, these methods have proven to be robust, naturally adaptive mesh-free methods for challenging fluid flow problems. Unfortunately, many implementations are limited to low orders of accuracy, and a systematic means of analyzing and addressing these issues remains elusive. One promising procedure is to modify the kernel function using nearby particle positions to correct for kernel deficiencies. While this technique boosts the order of accuracy, it requires solving a linear system of equations for each particle which increases the cost of the algorithm considerably. This paper will provide an overview of the state of the art in boosting the accuracy of these methods and provide simple examples of improved rates of convergence.

Louis F. Rossi  
University of Delaware  
rossi@math.udel.edu

Zhenyu He  
W L Gore and Associates  
zhhe@wlgore.com

#### MS207

##### Localized Kernel Methods for Nonlocal Diffusion

We present recently developed meshless methods leveraging localized bases of radial basis functions. We explore problems on both  $\mathbb{R}^n$  and on spheres  $\mathbb{S}^n$ . A localized basis is generated by combinations of polyharmonic splines, which produce a stable, localized, rapidly decaying basis. We apply recently developed quadrature formulas unique to the localized bases to easily evaluate stiffness matrices for partial differential equations and integral equations. We discuss a meshless Galerkin method for solving a volume constrained, nonlocal diffusion problem in  $\mathbb{R}^n$ . The method enables the use of scattered data points to generate a well-conditioned stiffness matrix. We also present a meshless Galerkin method for solving partial differential equations on the sphere.

Stephen Rowe  
Sandia National Laboratory  
srowe@sandia.gov

Richard B. Lehoucq  
Sandia National Laboratories  
rblehou@sandia.gov

Joseph Ward  
Texas A&M University  
Department of Mathematics

jward@math.tamu.edu

Francis J. Narcowich  
Department of Mathematics  
Texas A&M University  
fnarc@math.tamu.edu

#### MS207

##### Predicting Cavitation in Fuel Injectors using Smoothed Particle Hydrodynamics

Abstract not available

Emily Ryan  
Boston University  
ryanem@bu.edu

#### MS207

##### High-Order Staggered Moving Least Squares Schemes for the Steady Stokes with Applications to Suspension Flows

Abstract not available

Nathaniel Trask  
Brown University  
Division of Applied Mathematics  
nat.trask@gmail.com

#### MS208

##### Targeting Next-Generation Software Stacks with the Nabla DSL

The hardware complexity and heterogeneity of incoming architectures involves a redesign of our software stacks to address the flexibility and scalability challenges. There is a growing demand for new programming environments in order to improve scientific productivity, to facilitate better design and implementation, and to optimize next-generation production codes. Taking into account the feedback of scientific application developers and end-users, I will present an update of the Nabla domain specific language (DSL). This programming model comes as an open-source mesh-based numerical toolchain whose purpose is to perform source-to-source transformations in order to generate optimized code for different runtimes and architectures: RAJA, KOKKOS, OKINA (OpenMP w/vectorization) and XeonPHIs, GPUs, FPGAs.

Jean-Sylvain Camier  
CEA, DAM, DIF  
Jean-Sylvain.Camier@cea.fr

#### MS208

##### Experiences with Kokkos in the Production Codebases

Abstract not available

Simon D. Hammond  
Scalable Computer Architectures  
Sandia National Laboratories  
sdhammo@sandia.gov

#### MS208

##### Adapting Proxy Lessons Learned to Integrate C++

**Kokkos into a Production Fortran Code**

The goal of this talk is to present research on the implementation, performance, and optimization of a complex application kernel, dim3 sweep of SNAP, a neutral particle transport proxy, in CUDA through the use of the Kokkos programming model. Examples will be given of kernel performance measurements and optimization techniques enabled through the use of Kokkos. In addition, we will discuss efforts to couple the coarse-grained parallelism of SNAP, as implemented in Legion, a task-based programming model, and the fine-grained aspects, as implemented in Kokkos and CUDA, and how that coupling compares and contrasts to the native MPI+OpenMP of SNAP.

Joshua Payne, Geoff Womeldorff, Ben Bergen  
Los Alamos National Laboratory  
payne@lanl.gov, womeld@lanl.gov, bergen@lanl.gov

**MS208****Performance Portable Production Hydrodynamics Using Raja and Chai**

Abstract not available

Peter Robinson  
Lawrence Livermore National Laboratory  
robinson96@llnl.gov

**MS209****Coupling a Nano-Particle with Fluctuating Hydrodynamics**

We derive a coarse-grained description of the dynamics of a nanoparticle immersed in an isothermal simple fluid by performing a systematic coarse graining of the underlying microscopic dynamics. As coarse-grained or relevant variables we select the position of the nanoparticle and the total mass and momentum density field of the fluid, which are locally conserved slow variables because they are defined to include the contribution of the nanoparticle. The theory of coarse graining based on the Zwanzig projection operator leads us to a system of stochastic ordinary differential equations (SODEs) that are closed in the relevant variables. We demonstrate that our discrete coarse-grained equations are consistent with a Petrov-Galerkin finite-element discretization of a system of formal stochastic partial differential equations (SPDEs) which resemble previously-used phenomenological models based on fluctuating hydrodynamics. Under suitable approximations we obtain closed approximations of the coarse-grained dynamics in a manner which gives them a clear physical interpretation, and provides explicit microscopic expressions for all of the coefficients appearing in the closure. Our work leads to a model for dilute nanocolloidal suspensions that can be simulated effectively using feasibly short molecular dynamics simulations as input to a FEM fluctuating hydrodynamic solver.

Aleksandar Donev  
Courant Institute of Mathematical Sciences  
New York University  
donev@courant.nyu.edu

Pep Español  
Universidad Nacional de Educación a Distancia (UNED)

pep@fisfun.uned.es

**MS209****Non-Isothermal Coarse-Graining of Complex Molecules**

Coarse-graining (CG) of complex molecules is a way to reach time scales that would be impossible to access through brute force molecular simulations. We formulate a coarse-grained model for complex molecules from first principles that ensures energy conservation. Each molecule is described in a coarse way by a thermal blob characterized by the position and momentum of the center of mass of the molecule, together with its internal energy as an additional degree of freedom. This level of description gives rise to an entropy-based framework instead of the usual one based on the configurational free energy (i.e. potential of mean force). The resulting dynamic equations, which account for a proper description of heat transfer at the coarse grained level, have the structure of the Dissipative Particle Dynamics with Energy conservation (DPDE) model but with a clear microscopic underpinning. We provide explicit microscopic expressions for each component (entropy, mean force, friction and conductivity coefficients) appearing in the coarse-grained model. These quantities can be computed directly from MD simulations. The proposed non-isothermal coarse-grained model is thermodynamically consistent and opens up a first principles CG strategy to the study of energy transport issues that are not accessible with current isothermal models. Pep Español

Pep Español  
Universidad Nacional de Educación a Distancia (UNED)  
pep@fisfun.uned.es

**MS209****Investigation of the Molecular Aspects of Fluctuating Hydrodynamics Through the Memory Function Approach**

Fluctuating hydrodynamics (FHD) provides not only a theoretical background for the understanding of fluctuation phenomena in fluid at the mesoscopic level but also a systematic framework for the development of numerical methods. The FHD description is formally given as stochastic partial differential equations (SPDEs) where stochastic fluxes are expressed as multiplicative Gaussian white noise (GWN). While the FHD SPDEs are ill-defined in the sense that mathematically precise interpretation of multiplicative GWN should be provided, the use of GWN reveals that the molecular time scale is ignored in the FHD description. In this talk, we discuss the molecular aspects of the FHD description by investigating the density fluctuations in molecular dynamics systems and the FHD description of diffusion. The memory effects in the density fluctuations are obtained from the memory function approach.

Changho Kim  
Lawrence Berkeley National Laboratory  
changhokim@lbl.gov

**MS209****On Memory Estimation in the Mori-Zwanzig Equation**

In this paper we develop new error estimates and convergent approximations for the memory integral appearing in the Mori-Zwanzig formulation. The new theory is build

upon rigorous mathematical foundations, and it is developed both in phase space and in probability density function space. In particular, we derive error bounds for short-memory approximations and  $t$ -models. We also propose a new hierarchical finite-memory approximation which we prove convergent under appropriate conditions. An application to a high-dimensional dynamical system evolving from a random initial state is presented and discussed.

Yuanran Zhu

University of California, Santa Cruz  
yzhu22@ucsc.edu

Jason Dominy

University of California Santa Cruz  
jdominy@ucsc.edu

Daniele Venturi

Department of Applied Mathematics and Statistics  
University of California Santa Cruz  
venturi@ucsc.edu

## MS210

### Discovering Fast Matrix Multiplication Algorithms Via Tensor Decomposition

The computational complexity of matrix multiplication is an open question. Fast algorithms for matrix multiplication are those that require  $o(n^3)$  flops (that is, the exponent is less than 3) for  $n \times n$  matrices. Nearly all fast algorithms are based on a clever method for multiplying very small matrices that can be applied recursively. Discovering such a method corresponds to finding an exact, low-rank decomposition of a particular tensor. In this talk, I'll discuss numerical techniques for computing these decompositions and the practicality of the algorithms that have been discovered.

Grey Ballard

Sandia National Laboratories  
ballard@wfu.edu

## MS210

### A Practical Randomized CP Tensor Decomposition

The CANDECOMP/PARAFAC (CP) decomposition is a leading method for the analysis of multi-way data. The standard alternating least squares algorithm for the CP decomposition (CP-ALS) involves a series of highly overdetermined least squares problems. We show that, by extending randomized least squares methods to tensors, the workload of CP-ALS can be drastically reduced without a sacrifice in quality. We introduce techniques for efficiently preprocessing, sampling, and computing randomized least squares on a dense tensor of arbitrary order, as well as an efficient sampling-based technique for checking the stopping condition. We also show more generally that the Khatri-Rao product (used within the CP iteration) produces conditions favorable for direct sampling without any preprocessing. In numerical results, we see improvements in speed, storage, and robustness.

Casey Battaglini

Georgia Institute of Technology  
cbattaglini3@gatech.edu

## MS210

### Fast Recompression of Hadamard Products of Ten-

sors

The Hadamard product features prominently in tensor-based algorithms in scientific computing and data analysis. Due to its tendency to significantly increase ranks, the Hadamard product can represent a major computational obstacle in algorithms based on low-rank tensor representations. It is therefore of interest to develop recompression techniques that mitigate the effects of this rank increase. In this work, we investigate such techniques for the case of the Tucker format, which is well suited for tensors of low order and small to moderate multilinear ranks. Fast algorithms are attained by combining iterative methods, such as the Lanczos method and randomized algorithms, with fast matrix-vector products that exploit the structure of Hadamard products. The resulting complexity reduction is particularly interesting for tensors featuring large mode sizes and small to moderate multilinear ranks. To implement our algorithms, we have created a new Julia library for tensors in Tucker format.

Daniel Kressner

EPFL Lausanne  
Mathicse  
daniel.kressner@epfl.ch

Lana Perisa

University of Split, Croatia  
lana.perisa@gmail.com

## MS210

### Tensor Based Approaches in Magnetic Resonance Spectroscopic Signal Analysis

In recent years, magnetic resonance spectroscopic imaging (MRSI) is being used for brain tumor diagnosis. The processing of the MRSI signals are mainly performed using matrix based approaches. Tensorizing the matrix and using suitable tensor decompositions provides certain advantages. We focus on tensor based approaches to perform residual water suppression in the MRSI signal and to differentiate various tissue types in glioma patients from MRSI signals. To suppress the residual water, a Loewner tensor is constructed from the MRSI signal. Canonical polyadic decomposition (CPD) is applied on the tensor to extract the water component, and to subsequently remove it from the original MRSI signal. Tensor based tumor tissue type differentiation consists of building a covariance structured 3-D tensor from the MRSI spectra and then applying a non-negative CPD with common factor in mode-1 and mode-2 and  $l_1$  regularization on mode-3 to extract tissue-specific spectra and its corresponding distribution in the MRSI grid. An in-vivo study shows that our tensor based approach has better performance in identifying tumor and necrotic tissue types in glioma patients compared to matrix based approaches. This tensor based tissue characterization approach can be further extended to multi-parametric magnetic resonance imaging consisting of conventional magnetic resonance imaging, perfusion-weighted imaging, diffusion-weighted imaging and MRSI modalities

Bharath Halandur Nagaraja

ESAT/STADIUS, KU Leuven  
bharath.halandurnagaraja@esat.kuleuven.be

Diana Sima

KU Leuven  
STADIUS  
diana.sima@esat.kuleuven.be

Sabine Van Huffel  
 ESAT-SCD(SISTA) Department of Electrical Engineering  
 Katholieke Universiteit Leuven  
 sabine.vanhuffel@esat.kuleuven.ac.be

**MS211****Speed vs. Accuracy: The Next Frontier for Auto-Tuning?**

Abstract not available

Jeffery Hollingsworth  
 Computer Science Department  
 University of Maryland  
 hollings@cs.umd.edu

**MS211****A Directive-Based Data Layout Auto-Tuning for OpenACC Applications**

OpenACC applications can be executed on various computational environments with accelerators. However, an OpenACC application optimized for a specific device may not achieve enough performance for different devices. We propose a set of directive extensions to allow compilers to adapt data layout, which has a great effect on application performance, for the target device. OpenACC applications can be executed on various computational environments with accelerators. However, an OpenACC application optimized for a specific device may not achieve enough performance for different devices. We propose a set of directive extensions to allow compilers to adapt data layout, which has a great effect on application performance, for the target device.

Tetsuya Hoshino  
 The University of Tokyo  
 hoshino@cc.u-tokyo.ac.jp

Naoya Maruyama  
 RIKEN Advanced Institute for Computational Science  
 nmaruyama@riken.jp

Satoshi Matsuoka  
 Tokyo Institute of Technology  
 matsu@is.titech.ac.jp

**MS211****Algorithmic Revolution and Auto-Tuning for Matrix Computations in Post Moores Era**

It is said that Moores law will be broken within next two decades due to a physical performance limit of semiconductors. On the other hand, some scientists mention that ability of data access (BYTES) still increases relatively with comparison to ability of computations (FLOPS) by adapting upcoming technologies, such as 3D stacking technologies for memory. etc. In this presentation, we focus on this statement, and we discuss a change of numerical algorithms for matrix computations. We also mention that auto-tuning (AT) will be still one of key technologies to establish high performance in era of post Moore. State-of-the-art technologies of AT are shown to enumerate issues of current AT toward to the era of Post Moore. The authors think that computer language to adapt AT technology easier is crucial technologies. To establish this requirement, we are studying a computer language for AT, named ppOpen-AT. In this presentation, an example of adaptation

of ppOpen-AT will be presented with respect to actual simulation code by Finite Difference Method (FDM). Some evidences to show the statement of requirement in Post Moores era will be shown with the Fujitsu PRIMEHPC FX100, which is a supercomputer with a 3D stacking memory.

Takahiro Katagiri, Takahiro Katagiri  
 Information Technology Center  
 Nagoya University  
 katagiri@cc.nagoya-u.ac.jp, katagiri@cc.nagoya-u.ac.jp

Satoshi Ohshima  
 The University of Tokyo  
 Information Technology Center  
 ohshima@cc.u-tokyo.ac.jp

Masaharu Matsumoto  
 Information Technology Center  
 The University of Tokyo  
 matsumoto@cc.u-tokyo.ac.jp

**MS211****User-Defined Directive Translation Using the Xevolver Framework**

As the diversity of high-performance computing (HPC) systems increases, the programming considering the strong features of each HPC system becomes more important. Even in the case of directive-based programming, an appropriate directive should be used for each system. This presentation proposes that a special placeholder triggers insertions of appropriate directives for each system. A user describes the special placeholder that is used to specify the code line where one or more directives are potentially inserted. The special placeholder is translated into an appropriate directive by using user-defined code translation on Xevolver. Therefore, the appropriate directive can be inserted by keeping the maintainability and readability of the original code.

Kazuhiko Komatsu, Ryusuke Egawa, Hiroyuki Takizawa, Hiroaki Kobayashi  
 Tohoku University  
 komatsu@tohoku.ac.jp, egawa@cc.tohoku.ac.jp, takizawa@tohoku.ac.jp, koba@tohoku.ac.jp

**MS212****A Cache-Efficient Rank-Structured Elliptic PDE Solver**

We describe recent updates to a “superfast” Cholesky code based on CHOLMOD, organized around level 3 BLAS for speed. Our approach combine sparsity and low-rank structure and directly factor low-rank blocks using randomized algorithms, while retaining a “black-box” user interface compatible with standard sparse solvers. For a nearly-incompressible elasticity problem, CG with our rank-structured Cholesky converges faster than with incomplete Cholesky and the ML multigrid preconditioner, both in iteration counts and in run time. At  $10^6$  degrees of freedom, we use 3 GB of memory; exact factorization takes 30 GB.

David Bindel  
 Cornell University

bindel@cs.cornell.edu

### MS212

#### Developing Preconditioners with Guaranteed Convergence Rate Using Hierarchical Matrices

Hierarchical matrices have recently emerged as a tool of choice to build robust and general purpose preconditioners. They extend other methods such as multigrid, incomplete LU, and have been shown to remain efficient in cases where techniques like multigrid fail. In this talk, we will review recent progress on these solvers, discussing in particular the convergence of iterative methods with hierarchical preconditioners, as a function of the PDE type, system size, and condition number.

Eric F. Darve

Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

Chao Chen

Stanford University  
cchen10@stanford.edu

Hadi Pouransari

Stanford  
hadip@stanford.edu

Kai Yang

Stanford University  
kyang15@stanford.edu

### MS212

#### Divide-and-Conquer Approximate Inverse Type and Schur Complement Based Preconditioners with Low-Rank Corrections

This talk will present two preconditioning methods based on low-rank approximations in multilevel recursive frameworks for sparse symmetric matrices. The first method is an approximate inverse of the original matrix that exploits the Sherman-Morrison-Woodbury formula. The second method is based on the Schur complement technique and it is based on approximating the inverse of the Schur complement. Low-rank approximations are computed by the Lanczos procedure. Numerical experiments with Krylov subspace accelerators will be presented for model problems and general linear systems. These methods were found to be more efficient and robust than ILU-type preconditioners.

Ruipeng Li

Lawrence Livermore National Laboratory  
li50@llnl.gov

### MS212

#### Distributed-Memory Hierarchical Interpolative Factorization

The hierarchical interpolative factorization (HIF) offers an efficient way for solving or preconditioning elliptic partial differential equations. By exploiting locality and low-rank properties of the operators, the HIF achieves quasi-linear complexity for factorizing the discrete elliptic operator and linear complexity for solving the associated linear system. In this talk, the distributed-memory HIF (DHIF) is introduced as a parallel and distributed-memory implementa-

tion of the HIF. The DHIF organizes the processes in a hierarchical structure and keep the communication as local as possible. The computation complexity is  $O\left(\frac{N \log N}{P}\right)$  and  $O\left(\frac{N}{P}\right)$  for constructing and applying the DHIF, respectively, where  $N$  is the size of the problem and  $P$  is the number of processes. The communication complexity is  $O\left(\sqrt{P} \log^3 P\right) \alpha + O\left(\frac{N^{2/3}}{\sqrt{P}}\right) \beta$  where  $\alpha$  is the latency and  $\beta$  is the inverse bandwidth. Extensive numerical examples are performed on the TACC Stampede system with up to 4096 processes. The numerical results agree with the complexity analysis and demonstrate the efficiency and scalability of the DHIF.

Yingzhou Li

Stanford  
ryanli@stanford.edu

Lexing Ying

Stanford University  
Department of Mathematics  
lexing@stanford.edu

### MS213

#### A Hybrid Mixed-Integer Approach to Design Basin Networks for Water Resources Management

We discuss simulation-based optimization to guide agricultural management decisions in the face of limited availability of water. Specifically, we consider the design of catchment basins to infiltrate storm water and recharge depleted aquifers. The problem is formulated as a MINLP which we approach with a hybridized GA with implicit filtering. The One Water Hydraulic Model extension to the MODFLOW groundwater software is the simulation tool. We compare our approach to several different optimization strategies and outline future work in terms of realistic objective function formulations and inherent challenges in these types of applications.

Kathleen Kavanagh

Clarkson University  
kkavanag@clarkson.edu

### MS213

#### Optimizing the Design of Concentrated Solar Power Plants

Concentrated solar power plants utilize a field of mirrors to redirect sunlight toward a central tower where molten salt is heated. The thermal energy within the salt can then be used to directly produce electricity, or it can be stored for later use. Modeling the financial aspects of such plants involves deterministic simulations and stochastic simulations, along with explicit functional forms for different known costs. Some of the simulations can take tens of minutes to evaluate and all parts of the objective are defined by both discrete and continuous variables. In this talk, we present an approach for optimizing such a multifaceted objective function that attempts to evaluate the computationally expensive simulations as infrequently as possible. We apply our approach to the System Advisor Model developed at the National Renewable Energy Laboratory in order to maximize the expected profit over the lifetime of a concentrated solar power plant.

Jeffrey Larson

Argonne National Laboratory  
jmlarson@anl.gov

Sven Leyffer  
Argonne National Laboratory  
leyffer@mcs.anl.gov

Michael Wagner  
National Renewable Energy Laboratory  
michael.wagner@nrel.gov

### MS213

#### Three Dimensional Variational Data Assimilation Based on Derivative-Free Optimization

This research explores the use of Gaussian radial basis functions in order to build numerical models for Trust Region based methods in the context of data assimilation. In our context, cost functions are of the 3D-Var form wherein observational operators are potentially non-linear. The proposed trust region method proceed as follows: sample points are taken about the background state and for each sampled point, a local model is built based on a first (second) order Taylor expansion. Within the trust region, the local models are merged based on the theory of Gaussian radial basis functions. The trust region validation scheme is then used in order to assess the quality of the numerical model. Experimental tests are performed making use of a quasi-geostrophic model and three different observational operators from the specialized literature. The results reveal that, the impact of spurious innovations owing to linearization can be mitigated when the proposed method is utilized and even more, accurate analysis states are obtained in terms of Root-Mean-Square-Error when results are compared to those obtained via the local ensemble Transform Kalman filter wherein, observational operators are linearized.

Elias Niño-Ruiz  
Universidad del Norte  
enino@vt.edu

Adrian Sandu  
Virginia Polytechnic Institute  
and State University  
asandu7@vt.edu

### MS213

#### Bayesian Optimization under Mixed Constraints with a Slack-Variable Augmented Lagrangian

An augmented Lagrangian (AL) can convert a constrained optimization problem into a sequence of simpler (e.g., unconstrained) problems, which are then usually solved with local solvers. Recently, surrogate-based Bayesian optimization (BO) sub-solvers have been successfully deployed in the AL framework for a more global search in the presence of inequality constraints; however, a drawback was that expected improvement (EI) evaluations relied on Monte Carlo. Here we introduce an alternative slack variable AL, and show that in this formulation the EI may be evaluated with library routines. The slack variables furthermore facilitate equality as well as inequality constraints, and mixtures thereof. We show how our new slack "ALBO" compares favorably to the original. Its superiority over conventional alternatives is reinforced on several mixed constraint examples.

Victor Picheny  
INRA  
victor.picheny@toulouse.inra.fr

Robert Gramacy  
Virginia Tech  
rbg@vt.edu

Stefan Wild  
Argonne National Laboratory  
wild@mcs.anl.gov

### MS214

#### Accelerated Cyclic Reduction (ACR): A Distributed Memory Fast Direct Solver for 3D Structured Linear Systems

We present ACR: Accelerated Cyclic Reduction, a distributed-memory fast direct solver for rank-compressible block tridiagonal linear systems arising from the discretization of 3D elliptic operators. Algorithmic synergies between Cyclic Reduction and Hierarchical matrix arithmetic operations result in a solver with  $O(N \log^2 N)$  arithmetic complexity and  $O(N \log N)$  memory footprint. We provide a baseline for performance and applicability by comparing with the multifrontal method accelerated with hierarchical semi-separable matrices and algebraic multigrid. Numerical experiments consider a range of elliptic problems including the constant-coefficient Poisson equation and the indefinite Helmholtz equation at large processor counts. Results reveal that this method is competitive with state-of-the-art distributed-memory fast direct solvers and that it can tackle problems where algebraic multigrid fails to converge.

Gustavo Chavez  
King Abdullah University of Science and Technology  
gustavo.chavezchavez@kaust.edu.sa

George M Turkiyyah  
King Abdullah University of Science and Technology  
American University of Beirut  
gt02@aub.edu.lb

Hatem Ltaief  
King Abdullah University of Science & Technology  
(KAUST)  
hatem.Ltaief@kaust.edu.sa

David E. Keyes  
KAUST  
david.keyes@kaust.edu.sa

### MS214

#### Parallel Hierarchical Solvers, Convergence for Large Mesh Sizes

Solving sparse linear systems from the discretization of elliptic partial differential equations is an important building block in many engineering applications. Sparse direct solvers can solve general sparse matrices, but are usually slower and use much more memory than effective iterative solvers. To overcome these two disadvantages, a hierarchical solver based on  $\mathcal{H}^2$ -matrices was introduced in [Pouransari, Hadi, Pieter Coulier, and Eric Darve. "Fast hierarchical solvers for sparse matrices." arXiv preprint arXiv:1510.07363 (2015).], which is named LoRaSp. Here, we have developed a parallel algorithm based on LoRaSp to solve large sparse matrices on distributed memory machines. On a single processor, the factorization time of our parallel solver scales almost linearly for three-dimensional problems, as opposed to the quadratic time complexity of

many existing sparse direct solvers. On more than one processor, our algorithm exhibits significant speedups compared to sequential runs. Moreover, our parallel solver leads to almost constant number of iterations for solving problems on a wide range of mesh sizes, when used as a preconditioner. With this parallel solver, we are able to solve large problems much faster than many existing packages as demonstrated by the numerical experiments.

Chao Chen  
Stanford University  
cchen10@stanford.edu

Siva Rajamanickam  
Sandia National Laboratories  
srajama@sandia.gov

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

Eric F. Darve  
Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

#### MS214

##### Preconditioning using Hierarchically Semi-Separable Matrices and Randomized Projection

We present a preconditioner for general linear systems, with a focus on problems arising from PDE discretizations. The preconditioner is constructed from an approximate factorization, based on a multifrontal version of classical Gaussian elimination. The fill-in introduced during the factorization is compressed using Hierarchically Semi-Separable (HSS) matrices, a matrix format with low-rank off-diagonal blocks. For construction of the HSS representation we use a random projection technique and column-pivoted QR to reveal the rank. Factorization of an HSS matrix is done with a ULV-like decomposition. The approximate LU factorization is used as a preconditioner for GMRES and we compare with a number of other common preconditioners such as ILU and AMG. Our code is released as the package STRUMPACK and exploits hybrid parallelism through OpenMP and MPI.

Pieter Ghysels, X. Sherry Li  
Lawrence Berkeley National Laboratory  
pghysels@lbl.gov, xsli@lbl.gov

Francois-Henry Rouet  
Livermore Software Technology Corp.  
frouet@lstc.com

Chris Gorman  
UC Santa Barbara  
gorman@math.ucsb.edu

#### MS214

##### Distributed-Memory Hierarchical Matrix Algebra

We introduce a distributed-memory algorithm for  $\mathcal{H}$ -matrix application, composition and other related operations which avoids the need for the  $\Omega(p^2)$  scheduling procedure used in previous work, where  $p$  is the number of processes. Furthermore, we demonstrate that our algo-

rithm is able to effectively use  $O(N)$  processes for  $N \times N$   $\mathcal{H}$ -matrix and achieve strong scalability to more than an order of magnitude processes than previous results. We also show some recent advance in the large scale iterative solver for  $\mathcal{H}$ -matrix.

Yingzhou Li  
Stanford University  
yingzhouli0417@gmail.com

Jack L. Poulson  
Department of Mathematics and ICME  
Stanford University  
jack.poulson@gmail.com

Lexing Ying  
Stanford University  
Department of Mathematics  
lexing@stanford.edu

#### MS215

##### Multiscale Modeling in Coastal Ocean Hydrodynamics

Accurate modeling of coastal hydrodynamics requires modeling flow and transport from ocean basin scales to the continental shelf, to inland bays, estuaries and coastal floodplains. High resolution meshes and/or higher order solutions are required to capture these features. Complicated coastal topography suggests the use of unstructured finite element meshes. Capturing complex flow features requires the use of robust and accurate finite element techniques; e.g. discontinuous Galerkin methods. In this talk, we will discuss recent developments in finite element based models for simulating 2d and 3d flows in complex coastal regions, with applications to hurricane storm surges and transport of chemical constituents.

Clint Dawson  
Institute for Computational Engineering and Sciences  
University of Texas at Austin  
clint@ices.utexas.edu

#### MS215

##### A Mathematical and Geological Approach for Fractured Geothermal System

Fractures, in the underground, are planar or sub-planar discontinuities along which a rock has been broken, and represent important conduits or barriers for fluid flow. This is of particular importance for the exploitation of deep geothermal system, as matrix porosity and permeability is generally on a negligible scale. Furthermore, fractures increase the contact surface between rocks and fluids, facilitating a crucial contribution to heat transport by conductive heat exchange between fluids and the surrounding rock mass and increasing the economical efficiency of the system. Our work is focused on coupling mathematical and geological aspects. Simulations with real fracture geometries, as well as synthetic, are performed in a three-dimensional discrete fracture network setting, which can be viewed as a screening tool in data processing. The numerical method consider the fractures as objects of codimension one and include the aperture in their mathematical description moreover is able to handle non-conforming fractures intersections giving a high freedom in the simulations. Finally, we employ a mixed-virtual element scheme to compute scalar and vector unknowns on a mesh composed by cells of, almost, arbitrary shape. The long term

vision of our work is to develop an integrated geological and mathematical framework for collection and mutual transfer of data between geological and simulation models for enhanced geothermal system.

Alessio Fumagalli, Eirik Keilegavlen  
University of Bergen  
alessio.fumagalli@uib.no, Eirik.Keilegavlen@uib.no

### MS215

#### Multiscale Models of Sea Ice

The precipitous declines of Arctic sea ice have outpaced the predictions of most global climate models. I will discuss how we are using methods from theories of composite materials and statistical mechanics to address multiscale problems in sea ice physics. Our work is helping to advance how sea ice is represented in climate models and improve climate projections.

Kenneth M. Golden  
University of Utah  
Department of Mathematics  
golden@math.utah.edu

### MS215

#### Multiphysics Modeling of Microseismic Events for Elastic Wavefield Synthesis

Understanding the flow and deformation that occurs in a reservoir during fluid injection for enhanced oil and gas recovery is critical for maximizing the efficiency of production from unconventional reservoirs. However, the time scales of interest for flow and deformation are extremely long as pressure builds up in the reservoir leading to generation of a microseismic release of energy. In contrast, the duration of the resulting propagating wave generated by this source can be modeled with a very short time window (a few seconds). Further, most models assume the simplified model of a Gaussian earthquake source. A one-way coupled model for flow and deformation simulation is used for microseismic nucleation and subsequent wavefield modeling of the elastic emissions from the microseismic events. We examine this coupled phenomenon in an effort to more realistically capture the time evolution of microseismic sources.

Susan Minkoff  
Department of Mathematical Sciences  
University of Texas at Dallas  
sminkoff@utdallas.edu

Matthew McChesney, George McMechan  
Department of Geosciences  
University of Texas at Dallas  
mdm075000@utdallas.edu, mcmec@utdallas.edu

### MS216

#### Curvature Driven Erosion and Channelization of Sedimentary Beds by Fluid Flow

We discuss a closely coupled experimental and numerical study of how channels develop in a porous granular bed as a result of fluid flow. A model of a granular bed was developed by filling monodisperse glass beads homogeneously between parallel glass plates representing a thin horizontal fracture [Arshad Kudrolli and Xavier Clotet, Evolution of Porosity and Channelization of an Erosive Medium Driven

by Fluid Flow, Phys. Rev. Lett. 117, 028001 (2016)]. A fluid was then injected into the fracture with a prescribed rate to examine the evolution of the porous medium. We demonstrate that the porous medium begins to evolve at the interface between regions with high and low porosity corresponding to a mean flux which is lower than the value associated with erosion of an isolated grain. We show that the mean fluid flow inside the system can be modeled with Darcy law, and then demonstrate that the erosion patterns observed in the experiments can be modeled with erosion and deposition laws that dependent only on the local fluid flux. In particular, we will discuss the effect of interface curvature on the evolution of the channels, and the further effect of sediment advection on channel splitting.

Arshad Kudrolli  
Clark University  
Department of Physics  
akudrolli@clarku.edu

### MS216

#### How Bodies Erode and Dissolve in Fluid Flows

A variety of landscapes are formed by the action of flowing fluids, be it air or water flows. In these settings, the development of morphology is a reciprocal process: as a structure is carved by a fluid, its changing shape can alter the local flow. In this talk, I will discuss an idealization of such natural processes. By immersing an erodible or soluble body in fast flowing water, we examine the coevolution of shape and flow in a controlled setting. A fluid-flow model, based on Prandtl boundary-layer theory, yields a new class of scaling laws for the body's vanishing rate. The scaling laws apply to both erosion and dissolution, but shape evolution differs in the two cases. Erodible bodies develop sharp, angular features, while soluble bodies maintain a rounded front throughout. In both cases, the system tends to a state in which the material-removal rate is uniform along the body, offering a principle that may aid our understanding of erosion and dissolution in nature. Combined with the flow model, this principle also allows us to understand the shapes that arise in our experiments. For dissolution, we are led to a Riemann-Hilbert problem, whose solution gives the terminal body shape that is seen in the experiments. Subsequent analysis reveals that erosion, dissolution, and melting can all be linked under the Riemann-Hilbert framework.

Nick Moore  
Florida State University  
nm Moore2@fsu.edu

### MS216

#### Geometry of Channel Networks Incised by Subsurface Flow

When groundwater returns to the surface at springs, channels may be incised and ramified networks can form. Here we focus on two aspects of the geometry of growing groundwater-fed networks: the direction in which channels advance and the angle at which they branch. First, by relating the groundwater flow to a 2D Poisson field, we remark that the direction taken by moving channel tips may be equivalently understood from the maintenance of local symmetry in the groundwater field, the maximization of flux to tips, or motion along the groundwater streamline into tips. Next, we use these ideas to show that a bifurcated tip results in a branching angle of  $2\pi/5 = 72$  degrees. Our studies of a groundwater-fed stream network



on the Florida Panhandle accord well with this prediction. We also present the results of an empirical study of branching angles throughout the contiguous United States. Our results show that in humid climates, branching angles appear to asymptotically approach 72 degrees as the so-called aridity index—the ratio of rainfall rate to potential evapotranspiration—increases. The tendency toward this angle for aridity indices greater than unity is evident over approximately one-half of the country, suggesting that groundwater may play a significant role in channelization processes wherever climates are sufficiently humid.

Daniel Rothman, Yossi Cohen  
MIT  
dhr@mit.edu, ycohen@mit.edu

Olivier Devauchelle  
Institut de Physique du Globe, Paris, France  
devauchelle@ipggp.fr

Hansjoerg Seybold  
ETH, Zurich  
hseybold@ethz.ch

Robert Yi  
MIT  
ryi@mit.edu

#### MS216

##### Viscous Erosion of Immersed Bodies at Low and Intermediate Reynolds Numbers

A body which erodes due to viscous shear stress in a high Reynolds number flow is known to change shape with fore-aft asymmetry and vanish in a nearly self-similar fashion. Eventually, however, such a body will be small enough to create an intermediate or even very low Reynolds number disturbance flow. We will discuss the effect of Reynolds number on shape selection in viscous erosion, starting with theory and a new traction integral equation for studying erosion at zero Reynolds number. The flows and resulting body shapes are qualitatively different in character from their high Reynolds number analogues. A spherical body in a uniform flow first reduces to nearly (but not exactly) the drag minimizing profile of a body in a Stokes flow and then vanishes in finite time. Erosion in a shear flow, erosion near a wall, and erosion of multiple bodies will also be discussed. Moving to intermediate Reynolds numbers, we find striking transitions in shape formation and erosion rate associated with critical transitions in flow structure, and intricate dynamics for freely moving erodible bodies.

Saverio E. Spagnolie  
University of Wisconsin  
Department of Mathematics  
Spagnolie@math.wisc.edu

#### MS217

##### Coupling Between Ice Dynamics and Subglacial Hydrology

When modeling the sliding of large masses of ice over solid bedrocks, a common choice for the ice-bedrock boundary condition is to use a regularized Coulomb friction condition. Such sliding law includes scalar as well as distributed parameters (which need to be properly identified). Among the latter, the effective pressure can be computed by solving an additional set of PDEs on the two-dimensional ice-bedrock interface, modeling the subglacial hydrology

system. In this presentation we consider a coupled ice-hydrology model. We first discuss some theoretical results for the existence of a solution, and then we investigate the impact of the hydrology model on the simulation of the ice dynamics.

Luca Bertagna  
Florida State University  
lbertagna@fsu.edu

Mauro Perego  
CSRI Sandia National Laboratories  
mperego@sandia.gov

Max Gunzburger  
Florida State University  
Department of Scientific Computing  
mgunzburger@fsu.edu

Konstantin Pieper  
Florida State University  
Dept of Scientific Computing  
kpieper@fsu.edu

#### MS217

##### Implicit Time-Stepping for the Cryosphere

Coupled mass and momentum conservation problems for glaciers, subglacial liquid water, and sea ice are all characterized by multiple time scales and free boundaries. In contrast to the fully-hyperbolic coupled evolution problems, which are typical of faster fluids, numerical momentum solutions in these slow-motion cases typically already require solving systems of equations, e.g. those representing Stokes flow. However, current time-stepping schemes update geometry explicitly, using the (velocity) results of the momentum solve in the mass conservation equation. Short time steps are then difficult to avoid, in part because of the free-boundary nature of the geometry updates. This talk will explain this set of problems in thin-layer cases, and then address a variational inequality theory supporting such implicit time-stepping. Numerical examples, including solver aspects and preconditioning, will be shown.

Ed Bueler  
Dept. of Mathematics and Statistics  
University of Alaska Fairbanks  
elbueler@alaska.edu

#### MS217

##### Ice-Ocean Coupled Modeling with POPSICLES

Interactions between ice sheets and the ocean due to warmwater incursions into subshelf cavities are thought to be an important driver of marine ice sheet instability. Understanding this interaction requires coupled ice-ocean models. The POPSICLES model couples the POP2x ocean model, a modified version of the Parallel Ocean Program, and the BISICLES ice-sheet model. POP2x includes sub-ice-shelf circulation using partial top cells and the commonly used three-equation boundary layer physics. Standalone POP2x output compares well with standard ice-ocean test cases (e.g., ISOMIP) and other continental-scale simulations and melt-rate observations. BISICLES makes use of adaptive mesh refinement and a 1st-order accurate momentum balance similar to the L1L2 model of Schoof and Hindmarsh to accurately model regions of dynamic complexity, such as ice streams, outlet glaciers, and

grounding lines. Results of BISICLES simulations have compared favorably to similar simulations with a Stokes momentum balance in both idealized tests (MISMIP-3d) and realistic configurations. We also present results which demonstrate the effectiveness and usefulness of our approach.

Daniel Martin

Lawrence Berkeley National Laboratory  
dfmartin@lbl.gov

Xylar Asay-Davis  
Potsdam Climate Institute  
xylarstorm@gmail.com

## MS217

### Simulating Sea Ice Interactions with Ice Sheets Through Granular Iceberg Mélange in a Discrete Element Model

The recent collapse of permanent ice shelves has been accompanied by rapid breakup of iceberg mélange, a dense aggregation of icebergs and sea ice floating at glacier termini, which may suppress calving through the direct application of back force on the glacier. Current ice flow models are not capable of simulating mélange to test this speculation. I simulate mélange as a granular material with cohesive elastic bond formation and breaking using a specially-adapted version of the Discrete-Element bonded-particle Sea Ice model (DESIGN), a toolbox of LAMMPS-LIGGGHTS, an open-source discrete element method simulator. Icebergs are simulated as semi-rigid cylinders floating in seawater and sea ice is simulated as compressible elastic plates that form between proximate icebergs and break upon exceedance of material strength or length thresholds. Simulations show that mélange laden with thick sea ice can exert sufficient back force on the glacier terminus to shut down calving and thinning of sea ice reduces this back force, increasing the likelihood of calving. Calving events initiate jamming waves within mélange and cause the fracture of the sea-ice matrix that bonds mélange, increasing the likelihood of subsequent calving events. I discuss the comparison between simulations and observations of mélange, the benefits and shortcomings of using a discrete element approach and the potential for mélange coupling to ice sheet flow models.

Alexander Robel

California Institute of Technology  
robel@caltech.edu

## MS218

### Dense-Norm Multi-Dimensional Summation-by-Parts Operators

The solution of partial-differential equations on high-performance computing architectures necessitates flexible and robust numerical methods. The summation-by-parts (SBP) framework is advantageous in the construction of such schemes as it leads to conservative methods of arbitrary high-order that are provably linearly or non-linearly stable. Since these methods are not explicitly tied to basis functions, it is possible to tailor difference operators for different needs by optimizing the operators and/or the nodal distributions thereby providing flexibility not easily afforded by alternative discretization methodologies. The vast majority of work on SBP methods has focused on classical finite-difference operators that are applied to multi-dimensional problems using tensor-products and struc-

tured multi-block meshes. However, we have previously shown that it is possible to extend the SBP framework to unstructured meshes tessellated with arbitrary elements. In those works, our concentration was on diagonal-norm (mass matrix) SBP operators. In this presentation, we will discuss our recent work to extend these ideas to dense-norm SBP operators.

David C. Del Rey Fernandez

University of Toronto Institute For Aerospace Studies  
dcdelrey@gmail.com

Jason E. Hicken  
Rensselaer Polytechnic Institute  
Assistant Professor  
hickej2@rpi.edu

David W. Zingg

University of Toronto Institute for Aerospace Studies  
dwz@oddjob.utias.utoronto.edu

## MS218

### Super Convergence of Summation-By Parts Methods

A stable finite difference method applied to a problem with a smooth solution allows for a straightforward estimate of convergence rate based on the convergence rate of the local truncation error. In many high order cases the local truncation error at a few points near boundaries is significantly larger than at interior points. For standard SBP-SAT based methods this is the case. The straight forward estimate will predict a rate determined by the slowest converging local truncation error. Convergence in numerical computations is often faster than this prediction. We explore ways to improve our understanding of such super-convergence. We will in particular consider the second order wave equation in a two dimensional domain.

Gunilla Kreiss

Division of Scientific Computing  
Uppsala University  
gunilla.kreiss@it.uu.se

## MS218

### Improved Numerical Performance Using the SBP-SAT Technique As the Main Building Block

The stability properties of the various form of SBP-SAT schemes are well known. The question is: "how can we use that stability? In this presentation We will discuss methods for deforming geometries, non reflecting boundary conditions, multi-grid formulations, interface procedures, domain decomposition techniques, dispersion reducing schemes, variance reducing formulations in UQ, divergence free solutions of the incompressible NS equations, data assimilation, etc.

Jan Nordström

Department of Mathematics  
Linköping University  
jan.nordstrom@liu.se

## MS218

### Stable And Accurate Grid Interfaces For The Dynamic Beam Equation Using Summation By Parts

## Finite Differences

The dynamic beam equation is a standard model of flexible body dynamics and is thus of interest in many engineering applications. In [K. Mattsson, High-fidelity numerical simulation of the dynamic beam equation, JCP, 2015] a higher order finite difference method was derived for this equation. In this work the method is extended to accurately grid interfaces. The dynamic beam equation in one dimension,

$$\frac{\partial^2 u}{\partial t^2} = -\frac{\partial^4 u}{\partial x^4},$$

can be viewed as a simplified form of the equations governing Kirchhoff plates. These equations model elastic beams and plates which can describe many physical phenomena, for example large sheets of sea ice. Efficient and robust numerical solution of them can therefore be an invaluable tool in understanding the behavior of such systems. The appearance of the 4th derivative poses unique challenges for implementation of boundary conditions and interface couplings. In [K. Mattsson, Diagonal-norm summation by parts operators for finite difference approximations of third and fourth derivatives, JCP, 2014] summation by parts (SBP) operators for the 4th derivative were constructed and in K.Mattsson (2015) these were used to discretize (1). In this work we employ the SBP framework together with weakly imposed interface conditions to derive a provably stable and accurate grid interface treatment. The result is a robust finite difference method capable of handling discontinuous parameters and grid refinement.

Jonatan Werpers  
Department of Information Technology  
Uppsala University  
jonatan.werpers@it.uu.se

Ken Mattsson  
Uppsala University, Sweden  
ken.mattsson@it.uu.se

## MS219 Computation of Free Energy of Defects in Atomistic Systems

The macroscopic behavior of a crystalline material is strongly dependent on the type and distribution of defects present. This talk describes the analysis of the free energy of defect formation for the model problem of a constrained 1D system, and its convergence properties in the thermodynamic limit. We also create of a coarse-grained model with greatly reduced computational cost.

Matthew Dobson  
University of Massachusetts Amherst  
dobson@math.umass.edu

Hong Duong, Christoph Ortner  
University of Warwick  
m.h.duong@warwick.ac.uk, c.ortner@warwick.ac.uk

## MS219 Force-Based Atomistic-to-Continuum Coupling Methods for Multilattices Materials

Crystal defects play an important role in determining the mechanical and electrical properties of crystalline materials. In this talk, we formulate a model for a point defect in a multilattice crystal and introduce the blended force-based quasicontinuum (BQCF) method to approximate the

mechanics of the defect. We present numerical results establishing the convergence of the BQCF method in the context of a Stone-Wales defect in graphene and discuss possible extensions to bilayer materials.

Derek Olson  
School of Mathematics  
University of Minnesota  
olsond2@rpi.edu

Xingjie Li  
university of north carolina charlotte  
xingjieli5237@gmail.com

Christoph Ortner  
University of Warwick  
c.ortner@warwick.ac.uk

## MS219 Blending Methods for Effective Local/Nonlocal Coupling in Materials Modeling

Materials modeling often requires the combination of non-local models (such as atomistic or generalized continuum ones) with classical (local) continuum models, based on partial differential equations (PDEs), for accurate yet efficient materials simulation. Particular instances of this are problems in fracture mechanics. In such problems, direct application of classical PDEs is challenging, due to the differentiability assumptions on displacement fields not valid along crack surfaces. Local/nonlocal coupling, on the other hand, overcomes this limitation by employing nonlocal models (well defined on crack surfaces) to describe the evolution of cracks, while utilizing classical PDE-based models (which are more computationally efficient) in regions characterized by smooth deformations. In this presentation, we will present blending methods for local/nonlocal coupling in computational mechanics, and we will demonstrate their performance analytically and numerically.

Pablo Seleson  
Oak Ridge National Laboratory  
selesonpd@ornl.gov

## MS219 Seamless Coupling of Nonlocal and Local Models

Nonlocal models are gaining popularity in multiscale modeling. To reduce the added computational cost due to nonlocal interactions, there are advantages to employ local models wherever they provide valid descriptions of the physical processes. To implement such concurrently coupled local and nonlocal models, it is important to construct suitable interfaces between nonlocal models (given by integral equations) and local ones (represented by PDEs). In this talk, we first present some analytical results that help us formulating some well-defined coupled models. We then demonstrate how the recently developed framework of asymptotically compatible schemes can offer robust simulations of such models.

Yunzhe Tao  
Department of Applied Physics and Applied Mathematics  
Columbia University  
y.tao@columbia.edu

Qiang Du, Xiaochuan Tian  
Columbia University  
Department of Applied Physics and Applied Mathematics

qd2125@columbia.edu, xt2156@columbia.edu

lcgraham@fsu.edu, mgunzburger@fsu.edu

## MS220

### Adaptive Construction of Spatially Varying Polynomial Chaos for Uncertainty Quantification in Transport Problems

For transport problems of interest, the solution is a complex function of uncertain input parameters (e.g., in case of a distributed array of flow streams or jets). Further, the dependence of the solution is often strong on a subset of random variables and this dependence typically varies over the spatial domain. For such problems, a spatially varying setting for polynomial chaos (i.e., a different order of approximation at different spatial locations) will be more efficient. In this talk, we will discuss adaptive construction of spatially varying polynomial chaos approximations. Additionally, the details of our software implementation of this strategy will be shared. In particular, we will discuss the use of an interlaced setup for degrees-of-freedom in the physical and stochastic domains. Results will be demonstrated on some prototypical transport problems.

Eric C. Cyr

Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

Jason Li

Department of Mechanical, Aerospace and Nuclear  
Engineering  
RPI  
lij14@rpi.edu

Assad Oberai

Department of Mechanical, Aerospace and Nuclear  
Engineering  
Rensselaer Polytechnic Institute  
oberaa@rpi.edu

Eric Phipps

Sandia National Laboratories  
Optimization and Uncertainty Quantification Department  
etphipp@sandia.gov

Onkar Sahni

Rensselaer Polytechnic Institute  
sahni@rpi.edu

## MS220

### Adaptive Measure-Theoretic Augmentation of Multifidelity Monte Carlo Estimation

Estimating statistical quantities of complex computational models via Monte Carlo integration can be very expensive due to the large number of samples required. Often we have access to multifidelity surrogate models of various types. We demonstrate a method which uses inexpensive low fidelity models to aid the the optimal placement of higher fidelity model evaluations through the use of centroidal Voronoi tessellations. We employ this method within a measure-theoretic implementation of sample-based integration.

Lindley C. Graham, Max Gunzburger  
Florida State University  
Department of Scientific Computing

## MS220

### Error Estimation and Adaptive Error Control in Measure-Theoretic Stochastic Inversion

In this work, we consider a recently developed measure-theoretic approach for solving stochastic inverse problems. We prove that a sample based, non-intrusive, computational algorithm produces exact solutions to the stochastic inverse problem using a certain class of surrogate response surfaces. We use adjoint based techniques to estimate and correct for numerical error in the surrogate while simultaneously increasing the local order of the surrogate response surface. The use of the resulting enhanced surrogates are two-fold where we observe an increase in accuracy and decrease in computational complexity in computation of probabilities of specified events. The methodology may also be utilized in adaptive error control.

Steven Mattis

University of Texas at Austin  
steve.a.mattis@gmail.com

Troy Butler

University of Colorado Denver  
troy.butler@ucdenver.edu

## MS220

### Localizing Uncertainty with Gaussian Markov Random Field Models

The high computational cost of stochastic simulations involving partial differential equations (PDEs) with uncertain input parameters is often attributable to a combination of two bottlenecks: i) the steep cost of evaluating sample paths and ii) the complexity of the underlying parameter space. In this talk we relate both of these problems to the computational mesh, by using Gaussian Markov random fields to model the spatially varying input parameters for a simple PDE. This allows us to exploit readily available local dependency information of the parameter field in conjunction with standard finite element error estimates to identify spatial regions that contribute statistically to the error in the computed quantity of interest. The estimates can then be used to inform local mesh refinement.

Hans-Werner Van Wyk

Department of Mathematics and Statistics  
Auburn University  
hzw0008@auburn.edu

## MS221

### A Finite Element Formulation of Fluctuating Hydrodynamics for Fluids Filled with Rigid Particles Using Boundary Fitted Meshes

When dealing with the motion of submicron particles suspended in fluids, the Brownian motion arising from the thermal fluctuations becomes important. Diffusion is, indeed, relevant in colloidal suspensions, biological processes both at the cell and the sub-cell level, and in several microfluidic applications, in microrheology to determine local properties of complex fluids through the particle motion. In this contribution, we present a finite element implementation of fluctuating hydrodynamics with a moving boundary fitted mesh for treating the suspended particles. The thermal fluctuations are incorporated into the continuum

equations using the Landau and Lifshitz approach. The proposed implementation fulfills the fluctuation-dissipation theorem exactly at the discrete level. Since we restrict the equations to the creeping flow case, this takes the form of a relation between the diffusion coefficient matrix and friction matrix both at the particle and nodal level of the finite elements. Brownian motion of arbitrarily shaped particles in complex confinements can be considered within the present formulation. A multi-step time integration scheme is developed to correctly capture the drift term required in the stochastic differential equation describing the evolution of the positions of the particles.

#### Marco De Corato

Department of Chemical, Materials and Production Engineering  
University of Naples Federico II  
m.de-corato@imperial.ac.uk

#### Johan Slot

Department of Mathematics and Computer Science  
Eindhoven University of Technology  
j.j.m.slot@tue.nl

#### Markus Hutter

Department of Mechanical Engineering,  
Eindhoven University of Technology  
m.huetter@tue.nl

#### Gaetano D'Avino, Pier Luca Maffettone

Department of Chemical, Materials and Production Engineering  
University of Naples Federico II  
gadavino@unina.it, p.maffettone@unina.it

#### Martien Hulsen

Department of Mechanical Engineering,  
Eindhoven University of Technology  
m.a.hulsen@tue.nl

### MS221

#### Fluctuating Hydrodynamics of Reaction-Diffusion Systems

While it is widely appreciated that spatiotemporal fluctuations in the concentration of chemical species play an essential role in reaction-diffusion phenomena, stochastic simulation of a reaction-diffusion system is still computationally expensive. In this talk, we present a fluctuating hydrodynamics (FHD) formulation for stochastic reaction-diffusion systems and corresponding numerical schemes. While the FHD formulation is formally described by stochastic partial differential equations (SPDEs), it becomes similar to the reaction-diffusion master equation (RDME) description when those SPDEs are spatially discretized. In our FHD formulation, reactions are included as a source term having Poisson fluctuations. Higher-order numerical schemes are constructed in a systematic manner through the structure factor analysis. By treating diffusion implicitly, severe restrictions on time step size due to fast diffusion, which is the major computational issue of the RDME, are resolved. Also, depending on how accurately reactions are treated, reaction contributions can be sampled by either the stochastic simulation algorithm of Gillespie or the tau-leaping method. In addition, our numerical schemes work very well even in the case of a small number of molecules. We demonstrate the advantages of our numerical schemes by using several examples such as Turing pattern formation

and the front propagation.

Changho Kim, Andrew Nonaka  
Lawrence Berkeley National Laboratory  
changhokim@lbl.gov, ajnonaka@lbl.gov

Alejandro Garcia  
San Jose State University  
alejandrogarcia@sjsu.edu

John B. Bell  
CCSE  
Lawrence Berkeley Laboratory  
jbbell@lbl.gov

Aleksandar Donev  
Courant Institute of Mathematical Sciences  
New York University  
donev@courant.nyu.edu

### MS221

#### Low Mach Number Fluctuating Hydrodynamics for Electrolytes

We use the fluctuating hydrodynamics framework to simulate multispecies electrolyte solutions. The usual Landau-Lifshitz equations are completed by taking into account the effect of the electrical potential, calculated by solving the Poisson equation. The base of the algorithm is a finite volume scheme that was developed for low Mach number fluctuating multispecies transport [A. Donev & al., Low Mach Number Fluctuating Hydrodynamics of Multispecies Liquid Mixtures, *Phys. Fluids*, 27, 3, 2015]. We show that although the additional forcing terms resulting from the electric potential can be treated explicitly in time, this introduces a restrictive stability limit and we therefore also develop an implicit approach which overcomes this limitation. The structure factors of the density and velocity fluctuations are derived and compared to the numerical results. We compare the results obtained with our approach to results derived from the electroneutral limit. Finally, we extend our method to the situation where the electric permittivity of the solution is non-constant, and apply the methodology to study a type of electrokinetic instability.

#### Jean-Philippe M. Peraud

CCSE  
Lawrence Berkeley National Labs  
jperaud@lbl.gov

Andrew Nonaka  
Lawrence Berkeley National Laboratory  
ajnonaka@lbl.gov

Anuj Chaudhri, John B. Bell  
CCSE  
Lawrence Berkeley Laboratory  
achaudhri@lbl.gov, jbbell@lbl.gov

Aleksandar Donev  
Courant Institute of Mathematical Sciences  
New York University  
donev@courant.nyu.edu

Alejandro Garcia  
San Jose State University

algarcia@algarcia.org

imbortgerard@cims.nyu.edu

**MS221****DSMC simulations of Brownian Motion of a Small Particle in Rarefied Gas**

We present simulations of small (micro/nano) scale rigid particles suspended in a gas in small scale geometries. At this scale the Knudsen number, the ratio of mean free path and the characteristic length, is of order one. Therefore, continuum equations like the compressible Navier-Stokes equations cannot predict the gas flow correctly. So, one has to solve a kinetic equation, like the Boltzmann equation. We solve the Boltzmann equation using a Direct Simulation Monte Carlo (DSMC) method. The motion of the particle is governed by the Newton-Euler equations with the force and torque on the rigid body modeled from the momentum transfer of the gas molecules colliding with the body. We validate the numerical scheme by considering a moving piston in 1D and the Einstein relation for Brownian motion of a suspended particle in 3D. In the case of Brownian motion the translational as well as the rotational degrees of freedom are taken into account. Moreover, we present the motion of a Janus particle and of an object of complex shape under the influence of a thermal gradient in the gas.

Sudarshan Tiwari

FB Mathematik

TU Kaiserslautern

tiwari@mathematik.uni-kl.de

Axel Klar

TU Kaiserslautern

Department of Mathematics

klar@mathematik.uni-kl.de

Steffen Hardt

Nano-Microfluidics, Mechanical Engineering, TU

Darmstadt

hardt@nmf.tu-darmstadt.de

Samir Shrestha

Mathematics, Kathmandu University

samirstha@ku.edu.np

Tobias Baier

Nano-Microfluidics, Mechanical Engineering, TU

Darmstadt

baier@nmf.tu-darmstadt.de

**MS222****Volume Integral Methods for Waves in Plasmas**

Models for wave propagation in inhomogeneous plasma couple electromagnetic fields to the movement of charged particles. In such models the coefficients of the PDEs depend on the spatially variable density of the plasma particles. We will discuss integral-based methods for such problems, including the derivation of an integral formulation as well as the corresponding numerical method. In this framework the volume integral terms come from the variable coefficients, and their discretization is the key to obtain a high order accurate numerical method.

Lise-Marie Imbert-Gerard

CIMS, New York University

**MS222****DG Schemes for Collisional Plasma Models with Insulating Conditions on Rough Boundaries**

We consider in this paper the mathematical and numerical modeling of reflective boundary conditions, including diffusive reflection in addition to specularity, in the context of electron transport in plasmas, and their implementation in Discontinuous Galerkin (DG) schemes that solve the related Boltzmann kinetic model. We study the specular, diffusive and mixed (specular plus diffusive) reflection BC on physical boundaries of the problem. We develop a numerical approximation to model an insulating boundary condition, or equivalently, a zero flux mathematical condition for the electron transport equation. This condition balances the incident and reflective momentum flux at the microscopic level pointwise at the boundary, for the case of a more general mixed reflection with momentum dependant specularity probability  $p(\vec{k})$ . We compare the computational prediction of physical observables given by the numerical implementation of these different reflection conditions in our DG scheme for the Boltzmann model, and observe the influence of the diffusive condition in the kinetic moments over the position domain.

Jose A. Morales Escalante

ICES, The University of Texas at Austin

jose.a.morales.e@utexas.edu

Irene M. Gamba

Department of Mathematics and ICES

University of Texas

gamba@math.utexas.edu

**MS222****A Hybrid DG/Spectral Method for Micro-Macro Partitioned Kinetic Models**

The dynamics of gases can be simulated using kinetic or fluid models. Kinetic models are valid over most of the spatial and temporal scales that are of physical relevance in many application problems; however, they are computationally expensive due to the high dimensionality of phase space. Fluid models have a more limited range of validity, but are generally computationally more tractable than kinetic models. One critical aspect of fluid models is the question of what assumptions to make in order to close the fluid model. In this work we develop a high-order discontinuous Galerkin finite element method (DG-FEM) for a so-called micro-macro decomposition approximation of the kinetic equations. The micro-macro decomposition approach allows us to obtain accurate solutions of the fluid model, but instead of forcing a particular moment-closure approximation, which would typically only have a limited range of validity, this approach directly solves a version of the kinetic equations and uses this solution to provide a closure for the fluid equations. The proposed approach in this work makes use of Hermite spectral method for solving the kinetic portion of the update. The resulting numerical method is validated on several standard test cases.

James A. Rossmannith

Iowa State University

Department of Mathematics

rossmani@iastate.edu

## MS222

### Towards Scalable and Efficient Solution of IMEX Full-Maxwell / Multifluid Plasma Models\*

The mathematical basis for the continuum modeling of multifluid plasma physics systems is the solution of the governing partial differential equations (PDEs) describing conservation of mass, momentum, and total energy for each fluid species, along with Maxwells equations for the electromagnetic fields. To enable accurate and stable approximation of these systems a range of spatial and temporal discretization methods are commonly employed. In the context of finite element spatial discretization these include nodal and discontinuous Galerkin methods of the fluid sub-systems, and structure-preserving (physics-compatible) approaches for the electromagnetics system. For effective time integration of the longer time-scale response of these systems some form of implicitness is often required. Two well-structured approaches, of recent interest, are fully-implicit and implicit-explicit (IMEX) type methods. The requirement to accommodate disparate spatial discretizations, and allow the flexible assignment of mechanisms as explicit or implicit operators, implies a wide variation in unknown coupling, ordering, and the conditioning of the implicit sub-system. These characteristics make the scalable and efficient iterative solution of these systems extremely challenging. In this talk an overview of our approach to the development of scalable block preconditioners for these systems and initial weak parallel scaling results will be presented.

#### John Shadi

Sandia National Laboratories  
Albuquerque, NM  
jnshadi@sandia.gov

Edward G. Phillips  
Sandia National Laboratories  
egphill@sandia.gov

Eric C. Cyr  
Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

Roger Pawlowski  
Multiphysics Simulation Technologies Dept.  
Sandia National Laboratories  
rppawlo@sandia.gov

Matthew Bettencourt, Paul Lin, Sean Miller  
Sandia National Laboratories  
mbetten@sandia.gov, ptlin@sandia.gov,  
seamill@sandia.gov

## MS223

### Predictability of Oscillatory Cracks in Glass: A Peridynamic Study

Oscillatory and stable crack growth in glass has been demonstrated experimentally by slowly immersing a hot thin glass plate in a bath of cold water. The wave-length of the oscillatory crack depends on the plate width and immersion speed. These parameters also influence the transition between different regimes of crack growth: a straight propagating crack, an oscillatory crack, or an un-

stable/branching crack. We present a peridynamic model for thermomechanical failure and use it to test its ability of reproducing the fracture phase-diagram of the above-mentioned experiments. We show that the conventional failure criterion used in peridynamics requires a certain enhancement for this type of problems. With the new criterion, the peridynamic model matched extremely well the experimentally observed behavior. We also explain why phase-field or XFEM models of this problem have not been able to fully predict this crack growth behavior.

#### Florin Bobaru

University of Nebraska-Lincoln  
fbobaru2@unl.edu

Zhanping Xu, Guanfeng Zhang  
University of Nebraska  
zhanpingxu@gmail.com, gzhang@unl.edu

## MS223

### Bridging Scales Through Nonlocal Modeling

Nonlocal integral-differential equations and nonlocal balance laws have been proposed as effective continuum models in place of PDEs for a number of anomalous and singular processes. They may also be used to bridge multiscale models, since nonlocality is often a generic feature of model reduction. An example is the theory of peridynamics that has motivated our work. We discuss a few relevant modeling, computational and analysis issues, including robust simulation codes for validation and verification and seamless coupling of local (PDEs) and nonlocal models.

#### Qiang Du

Columbia University  
Department of Applied Physics and Applied Mathematics  
qd2125@columbia.edu

## MS223

### A Multi-Time-Step Method for Partitioned Time Integration of Peridynamics

Peridynamics is a nonlocal reformulation of continuum mechanics that is suitable for representing deformations with discontinuities. We extend the peridynamic formulation to allow the use of multiple time steps within a single problem domain by decomposing that domain into a number of smaller subdomains, where the critical regions of interest are solved using a small time step and the rest of the problem domain is solved using a larger time step. We explore the numerical properties and computational cost of the proposed approach, and demonstrate through numerical examples that a multi-time-step discretization of peridynamics can be solved much faster than a uniform time step discretization, and without adversely affecting the accuracy of the computed solution.

Payton Lindsay  
Purdue University  
plindsay@purdue.edu

#### Michael L. Parks

Sandia National Laboratories  
mlparks@sandia.gov

Arun Prakash  
Purdue University

aprakas@purdue.edu

tanbui@ices.utexas.edu

**MS223****A Fast Numerical Method for a Linear Peridynamic Model**

We develop a fast numerical method for a two-dimensional bond-based linear peridynamic model, which provides an appropriate description of the planar deformation of a continuous elastic body involving discontinuities or other singularities. The method reduces the computational cost of evaluating and assembling the stiffness matrix from  $O(N^2)$  to  $O(N)$ , where  $N$  is the number of unknowns in the discrete system. The method also reduces the computational work from  $O(N^2)$  to  $O(N \log N)$  and the memory requirement from  $O(N^2)$  to  $O(N)$ . All of this is achieved by carefully exploring the structure of the stiffness matrix of the collocation scheme, without any lossy compression involved.

Hong Wang

University of South Carolina  
Department of Mathematics  
hwang@math.sc.edu

**MS224****Active Subspaces for Low-Dimensional Response Surfaces on High-Dimensional Parameter Spaces**

The cost for constructing a response surface increases exponentially as the number of inputs increases. I will discuss active subspace-based strategies for reducing the dimension to enable otherwise infeasible response surface construction. The strategies often yield insight into the input/output relationship in the model.

Paul Constantine

Colorado School of Mines  
Applied Mathematics and Statistics  
pconstan@mines.edu

**MS224****A Unified Framework for Randomized Methods in Large-Scale Inverse Problems**

We propose a unified approach for constructing scalable randomized methods for large-scale inverse problems. From this broader general framework, four different existing randomized methods for solving inverse problems (EnKF, RML, PCGA, and RMA) as well as new methods can be derived. This new unified theoretical understanding will help further the development and analysis of randomized methods for solving large-scale inverse problems. In particular, from our numerical comparisons and past work, we will demonstrate that misfit randomization methods lead to improved solutions compared to methods that randomize the prior part, and offer an intuition to support this idea.

Ellen B. Le

ICES, UT Austin  
ellenble@gmail.com

Tan Bui-Thanh

The University of Texas at Austin

**MS224****Bayesian Inference for Preconditioned Inverse Ice Sheet Problems**

We address the problem of quantifying uncertainty in the solution of inverse problems governed by Stokes models of ice sheet flows within the framework of Bayesian inference. Computing the general solution of the inverse problem—i.e., the posterior probability density—is intractable with current methods on today's computers, due to the expense of solving the forward model (3D full Stokes flow with nonlinear rheology) and the high dimensionality of the uncertain parameters (which are discretizations of the basal sliding coefficient field). In this talk, we exploit the local sensitivity of data to parameters and build a hierarchically off-diagonal low-rank approximation for the Hessian (of the log posterior). This approximation will be applied as a preconditioner for a Newton-CG type method and for Hessian-based sampling in the inference of basal boundary conditions for ice sheet models.

Ruanui Nicholson

Applied Mathematics  
University of California, Merced  
rnic052@aucklanduni.ac.nz

Georg Stadler

Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

Omar Ghattas

The University of Texas at Austin  
omar@ices.utexas.edu

Noemi Petra

University of California, Merced  
npetra@ucmerced.edu

**MS224****Fast Hessian Approximation in Bayesian Inverse Wave Propagation**

A method for the fast approximation of Hessians in full-waveform inversion is proposed. The method exploits locality properties of the parameter-to-observable map and has the potential to outperform Krylov methods or the randomized SVD. Numerical results for frequency-domain wave inversion are presented.

Georg Stadler

Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

Hejun Zhu

Department of Geosciences  
Princeton University  
hejunzhu@princeton.edu

Omar Ghattas

The University of Texas at Austin



omar@ices.utexas.edu

### MS225

#### Decomposition and Rank Approximation for Symmetric Tensors

In this talk, we will discuss a computable strategy on calculating the rank of a given tensor. This is based on a formulation of a sparse optimization problem via an  $l_1$ -regularization for finding a low-rank approximation of tensors. In addition, we will also describe an extension to the symmetric case.

Carmeliza Navasca

Department of Mathematics  
University of Alabama at Birmingham  
cnavasca@gmail.com

### MS225

#### Generating Polynomials and Symmetric Tensor Decompositions

Symmetric tensors are multi-indexed arrays whose entries are invariant with respect to permutations of multi-indices. Generating polynomials are linear relations of recursive patterns about tensor entries. A set of generating polynomials can be represented by a matrix, which is called a generating matrix. Generally, a symmetric tensor decomposition can be uniquely determined by a generating matrix. We characterize the sets of such generating matrices and investigate their properties. Using these properties, we propose computational methods for symmetric tensor decompositions.

Jiawang Nie

University of California, San Diego  
njw@math.ucsd.edu

### MS225

#### Semialgebraic Geometry of Nonnegative Tensor Rank

In this talk we study the semialgebraic structure of the set of nonnegative tensors of nonnegative rank not more than  $r$ , denoted by  $D_r$ . We determine all nonnegative typical ranks for cubical nonnegative tensors. Under some mild condition (non-defectivity), we show that nonnegative, real, and complex ranks are all equal for a general nonnegative tensor of nonnegative rank strictly less than the complex generic rank. In addition, such nonnegative tensors always have unique nonnegative rank- $r$  decompositions if the real tensor space is  $r$ -identifiable. We determine conditions under which a best nonnegative rank- $r$  approximation has a unique nonnegative rank- $r$  decomposition: for  $r \leq 3$ , this is always the case; for general  $r$ , this is the case when the best nonnegative rank- $r$  approximation does not lie on the boundary of  $D_r$ .

Yang Qi

Gipsa Lab  
yang.qi@gipsa-lab.grenoble-inp.fr

### MS225

#### Leveraging Linear Constraints when Decomposing Large-scale, Incomplete Tensors

When decomposing a tensor into a polyadic decomposition (a sum of rank-1 terms), constraints are often added to

facilitate the computation or to improve the interpretability of the model. For example, it is often reasonable to assume that a factor matrix is defined by coefficients in a known basis. Each factor vector can, for example, be a polynomial with a known maximal degree evaluated in some points. The choice of these basis matrices often leads to large-scale and high-order tensors. The curse of dimensionality quickly prohibits measuring or computing the full tensor. We deal with this computational burden by using incomplete tensors and show that the decomposition can be recovered using only very few known entries. In this talk we discuss how the linear constraints can be exploited in the computation of a polyadic decomposition (PD). More specifically, we present both a data dependent and a data independent method to compute the PD of large and incomplete tensors. The data dependent variant uses all known entries in all iterations of the optimization algorithm. The data independent variant, on the other hand, leverages the linear constraints a priori, removing the need to use the tensor in the optimization algorithm. We finally show how the algorithms can be used to speed up simulations in computational materials sciences.

Nico Vervliet, Otto Debals

Department of Electrical Engineering (ESAT)  
KU Leuven  
nico.vervliet@esat.kuleuven.be,  
otto.debals@esat.kuleuven.be

Lieven De Lathauwer

KU Leuven - KULAK  
lieven.delathauwer@kuleuven.be

### MS226

#### On Betweenness Centrality Problems in Dynamic Graphs

Abstract not available

Elisabetta Bergamini

Institute of Theoretical Informatics  
Karlsruhe Institute of Technology (KIT)  
elisabetta.bergamini@kit.edu

Henning Meyerhenk

Karlsruhe Institute of Technology (KIT)  
meyerhenke@kit.edu

### MS226

#### Large-Scale Dynamic Graph Processing on HPC Systems

In many graph applications, the structure of the graph changes dynamically over time and may require real time analysis. However, most prior work for processing large graphs on HPC, e.g., Graph500, has not focused on a dynamic graphs, rather static only. To address this issue, we have developed algorithms, data structures, and infrastructure management necessary to support dynamic graph analysis at large scale on distributed HPC platforms, including next generation supercomputers which have locally attached NVRAM. We have explored support for online dynamic graph processing using an event-centric infrastructure in HavoqGT. When the graph structure or vertex/edge attributes change, HavoqGT triggers an algorithmic event that allows user-defined callbacks to perform the necessary application updates. In this talk we discuss our DegAwareRHH data-store designed for storing and indexing dynamic graphs persistently on node-local NVRAM

storage. We demonstrate its performance scaling out to store large, scale-free graphs by leveraging compact hash tables with high data locality. We extend DegAwareRHH for distributed memory using the event-centric infrastructure in HavoqGT. Using DegAwareRHH, we demonstrate performance of a dynamic graph-colouring algorithm using a large scale real-world dynamic graph from Wikipedia's edit history.

Keita Iwabuchi

Tokyo Institute of Technology  
iwabuchi.k.ab@m.titech.ac.jp

Roger Pearce  
Texas A&M and  
Lawrence Livermore Lab  
rpearce@llnl.gov

Maya Gokhale  
Lawrence Livermore National Laboratory  
maya@llnl.gov

Satoshi Matsuoka  
Tokyo Institute of Technology  
matsu@is.titech.ac.jp

#### MS226

##### **Predicting Movement of Vertices Across Communities in Dynamic Networks**

As dynamic networks evolve, a certain percentage of vertices can migrate to a different community or form a new community(ies). In this presentation, we present an algorithm on how to identify such vertices. We use a metric known as permanence, which measures by how much a vertex belongs to a community. Identifying the migrating vertices can be used to predict the structure of the dynamic communities as well as the effect of the perturbation in the network.

Sriram Srinivasan

Department of Computer Science  
University of Nebraska, Omaha  
sriram882004@gmail.com

Sanjukta Bhowmick  
Department of Computer Science  
University of Nebraska, Omaha  
sbhowmick@unomaha.edu

#### MS226

##### **Creating Dynamic Graphs from Temporal Data**

In recent years, there has been a growing interest in the analysis of dynamic graphs, which represent relationships and interactions changing over time. Often, analysis is performed under the assumption that a dynamic graph already exists. However, dynamic graphs usually must be created from an underlying temporal stream of edges and it is necessary to choose a method of doing so. Specifically, new data from the stream must be added and old or less relevant data must be aged off. The approach used affects the structure of the resulting dynamic graph and will therefore affect both data analysis results and performance. This talk analyzes methods of aging off old data to create dynamic graphs.

Anita Zakrzewska

School of Computational Science and Engineering

Georgia Institute of Technology  
azakrzewska3@gatech.edu

#### MS227

##### **Adaptive Methods for PDE Constrained Optimization with Uncertain Data**

I present an approach to solve risk averse PDE constrained optimization problems with uncertain data that uses different PDE model fidelities and adaptive sampling to substantially reduce the overall number of costly, high fidelity PDE. The approximation qualities of the optimization sub-problems due to sampling and lower fidelity PDE models is adjusted to the progress of the overall optimization algorithm.

Matthias Heinkenschloss

Department of Computational and Applied Mathematics  
Rice University  
heinken@rice.edu

#### MS227

##### **DDM for Stochastic Problems**

Abstract not available

Jangwoon Lee

University of Mary Washington  
llee3@umw.edu

#### MS227

##### **On the Efficient Treatment of Uncertainties Within Optimization Problems**

The treatment and minimization of uncertainties in design and control tasks is indispensable to address, characterize and minimize the impact of limited knowledge or uncertainty in parameters in complex physical processes. In this talk, we will discuss an efficient method which quantifies the uncertainty by adaptive sparse grids. The proposed approach allows to exploit the underlying structure of the forward problem and has the potential to significantly reduce the overall computational costs.

Claudia Schillings

University of Warkick, UK  
claudia.schillings@hu-berlin.de

#### MS227

##### **Uncertainty Regions in Shape Optimization**

Shape optimization tries to find shapes, which exhibit optimal properties mostly in the framework of partial differential equations. Novel approaches in shape optimization aim at optimization on shape manifolds. Unavoidable problem uncertainties render optimal shapes also as uncertain. In this context, novel strategies for the evaluation and treatment of these uncertainties are discussed.

Volker H. Schulz

University of Trier  
Department of Mathematics  
Volker.Schulz@uni-trier.de

Martin Siebenborn  
Department of Mathematics  
University of Trier

siebenborn@uni-trier.de

### MS228

#### Computing Singular or Nearly Singular Integrals on Smooth Closed Surfaces

We will describe a simple method for computing a singular or nearly singular integral, such as a harmonic function given by a single or double layer potential on a smooth closed curve surface, evaluated on or near the surface. The kernel is regularized, and the integral is replaced by a sum resulting from a quadrature rule which uses surface points that project onto grid points in coordinate planes. Corrections, derived analytically, are added for the errors due to regularization and discretization. The method does not require coordinate charts on the surface or special treatment of the singularity other than the corrections. The accuracy is about third order, uniformly for points of evaluation near the surface. Improved accuracy is obtained for points on the surface and can be used for solving integral equations. Examples have been computed with a variety of surfaces. The method should be useful for moving surfaces since only limited information about the surface is needed.

J. Thomas Beale

Duke University  
beale@math.duke.edu

Wenjun Ying

Department of Mathematics and Institute of Natural Sciences  
Shanghai Jiao Tong University  
wying@sjtu.edu.cn

Jason R. Wilson

Virginia Tech  
jasonwil@math.vt.edu

### MS228

#### High-Accuracy Discretization of Integral Operators on Smooth Surfaces

I will describe an approach to evaluating the singular integrals which arise in the discretization of certain classes of integral operators given on surfaces. This method is designed to achieve high accuracy with reasonable efficiency under adverse conditions; e.g., when the surface is non-smooth or given by a badly behaved parameterization. This is joint work with Zydrunas Gimbutas.

James Bremer

UC Davis  
bremer@math.ucdavis.edu

### MS228

#### Generalized Gaussian Quadratures for Singular and Hypersingular Kernels

In this talk, we present a procedure for the design of high-order quadrature rules that efficiently integrate a mixture of hypersingular and Hilbert kernels, and integrals involving both logarithmic singularities and smooth functions. The resulting generalized Gaussian quadratures only use a few extra nodes to handle the hypersingular and Hilbert parts.

Zydrunas Gimbutas

Courant Institute  
New York University  
zydrunas.gimbutas@gmail.com

### MS228

#### On the Solution of Elliptic Partial Differential Equations on Regions with Corners

The solution of elliptic partial differential equations on regions with non-smooth boundaries (edges, corners, etc.) is a notoriously refractory problem. In this talk, I observe that when the problems are formulated as boundary integral equations of classical potential theory, the solutions (of the integral equations) in the vicinity of corners are representable by series of elementary functions. In addition to being analytically perspicuous, the resulting expressions lend themselves to the construction of accurate and efficient Nyström discretizations of the associated boundary integral equations. The results are illustrated by a number of numerical examples.

Kirill Serkh

Yale University  
Applied Mathematics Program  
kirill.serkh@yale.edu

### MS229

#### A Variational Shifted Boundary Method for CFD and Fluid/Structure Interaction Simulations

Embedded boundary methods obviate the need for continual re-meshing in many applications involving rapid prototyping and design. Unfortunately, many finite element embedded boundary methods for incompressible flow are also difficult to implement due to the need to perform complex cell cutting operations at boundaries. We present a new, stable, and simple embedded boundary method, which we call the shifted boundary method, which eliminates the need to perform cell cutting, and demonstrate it on large-scale incompressible flow problems, using a variational approach based on the Nitsche method. We also show preliminary results on fluid/structure interaction problems.

Alex Main, Ting Song, Nabil Atallah, Guglielmo Scovazzi  
Duke University

geoffrey.main@duke.edu, ting.song@duke.edu,  
nabil.atallah@duke.edu, guglielmo.scovazzi@duke.edu

### MS229

#### The Shifted Boundary Method for Embedded Domain Mechanics: A Variational Implementation Using the Nitsche Approach

Embedded boundary methods obviate the need for continual re-meshing in many applications involving rapid prototyping and design. Unfortunately, many finite element embedded boundary methods for incompressible flow are also difficult to implement due to the need to perform complex cell cutting operations at boundaries. We present a new, stable, and simple embedded boundary method, which we call the shifted Nitsche method, which eliminates the need to perform cell cutting, and demonstrate it in applications in heat conduction, solid mechanics, wave propagation, fluid mechanics and fluid/structure interaction.

Guglielmo Scovazzi, Alex Main, Ting Song, Nabil Atallah  
Duke University

guglielmo.scovazzi@duke.edu, geoffrey.main@duke.edu,

ting.song@duke.edu, nabil.atallah@duke.edu

kangan.li@duke.edu, guglielmo.scovazzi@duke.edu

### MS229

#### A Shifted Boundary Method for Shallow Water Flows

The paradigm of embedded/immersed boundary methods is very powerful, since the limitations of grid generation can be avoided or greatly reduced. However, a direct implementation of an embedded method will suffer from issues related to small cut cells in the discretization, with adverse effects on the conditioning of the overall system of equations or the size of the time step. We propose a new approach, named shifted boundary method, in which, rather than imposing boundary conditions on the cut surface, we modify them and enforce them on a surrogate boundary that has no cut cells. We implemented the method in the context of unstructured finite elements and weak imposition of boundary conditions of Nitsche type. The resulting method is stable, provides second-order accurate solutions and is easy to implement owing to the fact that there is no need to perform cumbersome mesh subdivision near the boundaries. The robustness and accuracy of the method is demonstrated by a series of tests for hyperbolic systems, in particular waves and shallow water flows.

Ting Song, Alex Main, Guglielmo Scovazzi  
Duke University

ting.song@duke.edu, geoffrey.main@duke.edu,  
guglielmo.scovazzi@duke.edu

### MS229

#### A Ghost-Fluid Method for Shock Hydrodynamics of Multi-Phase Flows

We present an embedded boundary method for shock hydrodynamics of multi-phase flows. In present work we consider inviscid compressible flows and stiffened gas equations of state; but there is no severe obstacle to extend the method for more general cases. This method combines a variational multiscale (VMS) stabilized finite element method for single-material shock hydrodynamics and the ghost-fluid technique to capture the transmission conditions across the material interface. The material interface is captured implicitly by a level set method, so that it does not have to coincide with grid lines. For this reason, the proposed method is particularly suitable for problems in which the material interface may undergo large motions and topological changes, such as in shock-bubble interactions. A particular difficulty with the ghost-fluid method is in appearance of very strong shocks. In these scenarios, the conventional extrapolation based ghost value population technique tends to lead to unphysical solutions near the material interface. To resolve this issue, we define and solve a two-material Riemann problem between two fluids, and use its solutions to construct ghost values. This technique is shown to be able to handle strong discontinuities, such as very large density ratios very well.

Xianyi Zeng  
University of Texas El Paso  
Mathematical Sciences Dept.  
xzeng@utep.edu

Kangan Li, Guglielmo Scovazzi  
Duke University

### MS230

#### Recent Advances in Embedded Finite Element Methods

An emerging class of embedded finite element methods for evolving boundary value problems in mechanics will be presented. These methods have been designed to circumvent long-standing difficulties with finite elements for Lagrangian simulations of deformable media with complex geometry. Particularly for problems with significant changes in topology, continuous remeshing strategies have simply not proven sufficiently viable or robust. The embedded methods provide a means for the geometry of features of interest, such as sharp phase interfaces or fracture surfaces, to be represented independently of the mesh. This relaxation between mesh and geometry obviates the need for remeshing strategies in many cases and greatly facilitates adaptivity in others. The approach is very similar to the Eulerian methodologies developed by the finite difference and level-set communities, but within a variational setting that facilitates error and stability analysis. This talk will describe the theory behind the embedded method and methodological advances, with an emphasis on recent developments. In particular, we will describe the use of ghost-penalty stabilization for higher-order basis functions such as B-splines.

John Dolbow  
Duke University  
john.dolbow@duke.edu

### MS230

#### Solving an Inverse Interface Problem with a Fixed Mesh

We consider an inverse problem for the typical boundary problem of a second order elliptic equation whose coefficient is discontinuous across an interface. The inverse problem for this boundary value problem is to use a measurement of the solution of the boundary value problem to identify the coefficient in the differential equation including the interface for the discontinuity. We seek a solution to the inverse problem through the standard least squares formulation. The objective functional is discretized with the finite element method that can use a mesh independent of the interface. The gradient of the discrete objective functional is derived which can be computed within the chosen finite element framework such that the inverse problem can efficiently be solved by a standard minimization method such as the quasi-Newton method. Numerical examples will be presented to demonstrate features of this finite element method for solving the inverse problem.

Tao Lin  
Department of Mathematics, Virginia Tech  
tlin@math.vt.edu

Min Lin  
College of Science, Southwest Petroleum University  
Chengdu, China, 610500  
405090450@qq.com

### MS230

#### On the Error Estimates of An Unfitted Nitsche

### Method Applied to Interface Problem

We prove optimal error estimates for the flux variable for a stabilized Nitsches method applied to an elliptic interface problem with discontinuous constant coefficients. These error estimates are totally in- dependent of the contrast between diffusion coefficients.

Manuel Sanchez-Uribe

University of Minnesota

manuel\_sanchez\_uribe@alumni.brown.edu

Erik Burman

University College London

e.burman@ucl.ac.uk

Johnny Guzman

Brown University

johnny\_guzman@brown.edu

Marcus Sarkis

Worcester Polytechnic Institute

Instituto de Matematica Pura e Aplicada (Brazil)

msarkis@wpi.edu

### MS230

#### Gradient Recovery for Elliptic Interface Problem

In this talk, we present two types of novel gradient recovery methods for elliptic interface problem: 1. Finite element methods based on body-fitted mesh; 2. Immersed finite element methods. Due to the lack of regularity of solution at interface, standard gradient recovery methods fail to give superconvergent results, and thus will lead to over-refinement when served as *a posteriori* error estimator. This drawback is overcome by designing an immersed gradient recovery operator in our methods. We analyze the superconvergence of these methods, and provide several numerical examples to verify the superconvergence and its robustness as *a posteriori* error estimator.

Hailong Guo, Xu Yang

Department of Mathematics

University of California, Santa Barbara

hlguo@math.ucsb.edu, xuyang@math.ucsb.edu

### MS231

#### Domain Decomposition of Parabolic Equations - A Splitting Approach

We will analyze temporal approximation schemes based on overlapping domain decompositions. As such schemes enable computations on parallel and distributed hardware, they are well suited for integrating large-scale parabolic systems. Our analysis is conducted by first casting the domain decomposition procedure into a variational framework based on weighted Sobolev spaces. The time integration of a parabolic system can then be interpreted as an operator splitting scheme applied to an abstract evolution equation governed by a maximal dissipative vector field. By utilizing this abstract setting, we derive an temporal error analysis for the most common choices of domain decomposition based integrators. This is joint work with Erik Henningson.

Eskil Hansen

Lund University

Box 118, 221 00 Lund, Sweden

eskil@maths.lth.se

### MS231

#### Faster and More Accurate Pseudospectral DNS Through Adaptive High-Order Time Integration

Time integration of Fourier pseudo-spectral DNS is usually performed using the classical fourth-order accurate Runge–Kutta method, or other methods of second or third order. We investigate alternative time integrators in the large-scale parallel code SpectralDNS. We find that the fifth-order accurate Runge–Kutta pair of Bogacki & Shampine gives much greater accuracy at a significantly reduced computational cost. Furthermore, the use of adaptive time stepping using an embedded pair yields enormous cost savings in simulations involving the development of turbulence from a laminar flow. Extensive numerical tests of incompressible turbulent flows confirm the reliability and efficiency of the method.

David I. Ketcheson

CEMSE Division

King Abdullah University of Science & Technology

david.ketcheson@kaust.edu.sa

Matteo Parsani

KAUST

matteo.parsani@kaust.edu.sa

Nathanael Schilling

TUM

nathanael.schilling@in.tum.de

### MS231

#### Flexible and Accurate Multiphysics Time Integration with Arkode

Increasingly, modern computational science requires large-scale simulations that consistently and accurately couple distinct physical processes. While the mathematical models for each individual process often have well-known type (hyperbolic, parabolic, etc.), may be either linear or nonlinear, and are suitable for classical numerical integrators, the same cannot be said for the coupled models. These multiphysics models are often of mixed type, involve both linear and nonlinear sub-components, may have limited differentiability, and involve processes that evolve at dissimilar rates. As such, many multiphysics simulations require newer and more flexible time integrators that may be tuned for these complex problems. In this talk we discuss recent work in ARKode, a library providing flexible and accurate solvers for stiff, nonstiff, mixed implicit-explicit, and multirate systems of differential equations. Based on time-adaptive additive Runge Kutta methods, ARKode supports a variety of flexibility enhancements, including application-specific data structures, RK coefficients, and adaptivity functions, among others. We will discuss these flexibility enhancements in the context of two large scale parallel case-studies: dislocation dynamics simulations in materials science, and non-hydrostatic 'dycore' simulations in global climate modeling. We will conclude by discussing current work on adaptive multi-rate time integration for partitioned multiphysics systems.

Daniel R. Reynolds

Southern Methodist University

Mathematics

reynolds@smu.edu

David J. Gardner  
Lawrence Livermore National Laboratory  
gardner48@llnl.gov

Carol S. Woodward  
Lawrence Livermore Nat'l Lab  
woodward6@llnl.gov

Jean Sexton  
Southern Methodist University  
jmsexton@smu.edu

### MS231

#### **SUNDIALS Time Integrators and their Application**

The SUNDIALS Suite of Nonlinear Differential-Algebraic Solvers and Integrators includes highly robust and adaptive time integration methods and software for ODEs and DAEs. In this talk, we will overview SUNDIALS and present applications of the packages for simulations in materials science, power grid, and climate. Lastly, we will discuss future directions for SUNDIALS. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC. LLNL-ABS-702576

Carol S. Woodward  
Lawrence Livermore Nat'l Lab  
woodward6@llnl.gov

Daniel R. Reynolds  
Southern Methodist University  
Mathematics  
reynolds@smu.edu

David J. Gardner, Slaven Peles  
Lawrence Livermore National Laboratory  
gardner48@llnl.gov, peles2@llnl.gov

Alan Hindmarsh  
Lawrence Livermore National Laboratory  
Center for Applied Scientific Computing  
alanh@llnl.gov

L. Edward Banks  
Lawrence Livermore National Laboratory  
banks12@llnl.gov

### MS232

#### **Topology Optimization of Compact Wideband Coaxial-to-Waveguide Transitions with Minimum-Size Control**

We present a density-based topology optimization approach to design a compact, multilayer element that interfaces a coaxial cable with a rectangular waveguide. The conductivity distribution in each layer is designed using a gradient-based optimization method that relies on finite-difference time-domain solutions to the 3D Maxwell equations. The optimization problem is strongly self-penalized towards designs solely consisting of a good conductor or a good dielectric, which unfortunately leads to convergence to a local optimum with unsatisfactory performance if no precautions are taken. To address the self-penalization issue and to avoid small scattered features that can result in ohmic losses, we have developed a nonlinear filtering ap-

proach that relies on a sequence of harmonic mean filters operating in two phases. The first phase involves a continuation over the filter radius, which is decreased until it reaches a specified strictly positive minimum value. This minimum value determines the size of the smallest features that might appear in the final design. In the second phase, we fix the filter radius to the final value used in phase one and use a second continuation approach to gradually increase the nonlinearity of the filter in order to obtain binary designs. The numerical experiments demonstrate that the two-stage approach makes it possible to impose minimum size control while still obtaining a transitional device with good performance.

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

Emadeldeen Hassan, Linus Hagg, Eddie Wadbro  
Department of Computing Science  
Umea University  
emad@cs.umu.se, linush@cs.umu.se, eddie@cs.umu.se

### MS232

#### **Topology Optimization for Design of Coaxial Cables**

Coaxial cables, and transmissions lines in general, are critical elements in transmission of high frequency alternating currents over relatively long distances. They appear in modern communications equipment, but also in devices that require transmission of a large amount of current applied over a short period of time. This talk will explore topology optimization formulations for design of coaxial cables with an eye to the requirements of complex high energy devices. These problems can be reduced to the problem of placing conductive material in a void to carry current through the domain between specified in and out ports. Two simplified models of the physics will be considered. The first a electrostatic model is intended to satisfy the basic power flow requirements. The second application model, uses a quasi-static Maxwell formulation to achieve realistic short time current flows.

Eric Cyr, Gregory J. von Winckel, Thomas Gardiner  
Sandia National Laboratories  
eric.c.cyr@gmail.com, gvonwin@sandia.gov,  
tagardi@sandia.gov

Drew P. Kouri  
Optimization and Uncertainty Quantification  
Sandia National Laboratories  
dpkouri@sandia.gov

John Shadid  
Sandia National Laboratories  
Albuquerque, NM  
jnshadi@sandia.gov

### MS232

#### **Adaptive Finite Element Method for Topology Optimization with Stress-Based Constraints**

We present a topology structural optimization framework with adaptive mesh refinement and stress-constraints. Finite element approximation and geometry representation benefit from such refinement by enabling more accurate stress field predictions and greater resolution of the optimal structural boundaries. We combine a density filter to

impose a minimum design feature size, the solid isotropic material with penalization (SIMP) to generate black-and-white designs and a SIMP-like stress definition to resolve the stress singularity problem. Regions with stress concentrations dominate the optimized design. As such rigorous simulations are required to accurately approximate the stress field. To achieve this goal, we invoke mesh refinement during the optimization. We do so in an optimal fashion, by applying adaptive mesh refinement techniques that use error indicators to refine and un-refine the mesh as needed. These error estimators are based on the Dual-Weighted Residual method to provide an optimal mesh for the optimality conditions that drive the design. In this way, we obtain more accurate simulations and greater resolution of the design domain. We present results and compare them to those obtained with a uniformly refined mesh for verification purposes and to demonstrate the efficiency of our method.

Miguel Salazar de Troya  
Lawrence Livermore National Laboratory  
University of Illinois at Urbana-Champaign  
salazardetro1@llnl.gov

Daniel Tortorelli  
Mechanical & Industrial Engineering Department  
University of Illinois, Urbana-Champaign  
dtortore@uiuc.edu

### MS233

#### **Meshless Approximation Using RBF-Partition of Unity Method with Applications to the Reconstruction of Basins of Attraction in Dynamical Systems**

The problem of approximating large scattered data sets is rather common in applications. Here we focus the attention on the partition of unity method, performed by blending Radial Basis Functions (RBFs) as local approximants and locally supported weight functions. In particular, we present space-partitioning data structures based on a partition of the underlying generic domain. This approach allows us to optimize the search process of the nearest neighbour points. Moreover, we consider as application a problem in population dynamics models. In dynamical systems saddle points partition the domain into basins of attractions of the remaining locally stable equilibria. More precisely, a particular solution of a dynamical system is completely determined by its initial condition and by the parameters involved in the model. Therefore, accurately reconstructing the attraction basins is a crucial point. In this talk we focus on dynamical systems of ordinary differential equations presenting two and three stable equilibria. We use a bisection technique for the detection of the points located on the separatrix manifolds determining the basins of attraction. Then, we propose algorithms for the reconstruction of such manifolds. The separatrices are reconstructed by means of the RBF-partition of unity method. Also the case of positivity preservation using constrained RBF interpolation is considered. Finally, numerical experiments are presented.

Alessandra De Rossi, Roberto Cavoretto, Emma Perracchione  
Department of Mathematics  
University of Torino  
alessandra.derossi@unito.it, roberto.cavoretto@unito.it,

emma.perracchione@unito.it

### MS233

#### **Numerical Study of Space-Time RBF Method for PDEs**

We are experimenting with space-time radial basis function methods for the numerical solutions of time-dependent PDEs. Unlike common numerical schemes where space and time are treated independently (known as the method of lines), the time variable in space-time formulation is treated as another space variable. The discretized version of PDEs are then solved as boundary value problems in space-time domain. Several numerical examples will be presented.

Alfa Heryudono  
University of Massachusetts Dartmouth  
Department of Mathematics  
aheryudono@umassd.edu

### MS233

#### **High-Order Radial Basis Function Finite Difference Methods for the Meshfree Solution of PDEs on Time-Varying Irregular Domains**

We present numerical methods based on Radial Basis Function (RBF) interpolation for automatic node generation on arbitrary domains in two and three dimensions. We then present the numerical solution of PDEs on these domains using high-order RBF-Finite Difference methods.

Varun Shankar  
Department of Mathematics  
University of Utah  
vshankar@math.utah.edu

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

### MS233

#### **Radial Basis Functions**

Abstract not available

Grady B. Wright  
Department of Mathematics  
Boise State University, Boise ID  
gradywright@boisestate.edu

### MS234

#### **The Flexible Computational Science Infrastructure (flecsi) Project: Interfaces for Multi-Physics Applications Development**

Abstract not available

Ben Bergen, Marc Charest, Irina Demeshko, Nick Moss, Joshua Payne  
Los Alamos National Laboratory  
bergen@lanl.gov, charest@lanl.gov, irina@lanl.gov, nickm@lanl.gov, payne@lanl.gov

### MS234

#### **(Re-)Designing Iterative Solvers for Task-Based**

**Runtime Systems**

Abstract not available

Steven Dalton

NVIDIA

sdalton@nvidia.com

**MS234****Stapl, Spatial Decomposition, and Problem Space Representations for Compact Binary Mergers**

Co-design brings together domain experts from multiple disciplines to solve problems. Our team brought together physicists, applied mathematicians and computer scientists to solve problems in Binary Star Mergers. In order to improve performance for our simulation, we explored different ways to decompose the spatial domain using space-filling curves to improve locality in our computations, and new problem space representations to control and manipulate our data. We used Texas A&M's STAPL library, a Department of Energy-sponsored project, to compare our different alternatives.

Glen Hordemann

Texas A&M University

glen@hordemann.com

Nicolas de Brye

University of Valencia

nicolas.de-brye@uv.es

Daniel George

University of Illinois at Urbana-Champaign

dan7geo@gmail.com

Hyun Lim

Brigham Young University

hyun\_lim@outlook.com

Julien Loiseau

University of Reims

julien.loiseau@univ-reims.fr

Jonah Miller

Perimeter Institute for Theoretical Physics

jonah.maxwell.miller@gmail.com

Jonathan Sharman

Rice University

jonathan.sharman@gmail.com

**MS235****Efficient Exploration of the Conformational Space of Proteins Using the Concurrent Adaptive Sampling Algorithm**

Molecular dynamics (MD) simulations are useful in sampling thermodynamic and kinetic properties of biomolecules. However, our sampling of the bio-molecules properties is severely limited by the timescale barrier, and MD simulations routinely get stuck in metastable free energy minima. While there are several existing methods to overcome these issues, problems remain in regard to being able to sample unknown systems, deal with high-dimensional space, and focus effort on slow timescales. Hence, a new sampling method, called the Concurrent Adaptive Sampling algorithm (CAS), has been developed

to tackle these three main problems and efficiently obtain conformations and pathways. The beauty of CAS is that it adaptively constructs macrostates, requires little a priori knowledge about the system, and considers an arbitrary number of general collective variables. In addition, CAS uses the second eigenvector of the transition matrix to maintain a fine discretization along pathways while using importance sampling in orthogonal directions to control computational cost. In this talk, we introduce the new algorithm and show new results about the triazine polymers [Grate et al, Triazine-Based Sequence-Defined Polymers with Side-Chain Diversity and Backbone-Backbone Interaction Motifs] that were found using CAS.

Surl-Hee Ahn

Stanford University

sahn1@stanford.edu

Jay Grate

Pacific Northwest National Laboratory

jwgrate@pnnl.gov

Eric F. Darve

Stanford University

Mechanical Engineering Department

darve@stanford.edu

**MS235****Nonlocal Models for Nanoscale Heat Conduction**

It has been a typical observation that the heat conduction processes in nano-scale materials exhibit non-Fourier law behavior. The brute force approach of using full atomistic simulations is often limited by the size of the system. In this talk, I will present a coarse-graining approach using the Mori's projection formalism. In particular, we derive nonlocal models for the heat conduction process. We discuss further simplifications, as well as the connection to relaxation type of heat conduction models at the continuum level.

Xiantao Li

Department of Mathematics

Pennsylvania State University

xli@math.psu.edu

**MS235****Brownian Dynamics Approximation to Generalized Langevin Equations**

We present a coarse-grained model from Langevin dynamics which result in the form of generalized Langevin equations. Further, we carry out reduction to a coordinate only stochastic model, which in its exact form, involves a forcing term with memory and a general Gaussian noise. It will be shown that a similar fluctuation-dissipation theorem still holds at this level. We study the approximation by the typical Brownian dynamics as a first approximation. Our numerical test indicates how the intrinsic frequency of the kernel function influences the accuracy of this approximation. In the case when such an approximate is inadequate, further approximations can be derived by embedding the nonlocal model into an extended dynamics without memory. By imposing noises in the auxiliary variables, we show how the second fluctuation-dissipation theorem is still exactly satisfied.

Lina Ma

Pennsylvania State University



linama@psu.edu

Xiantao Li  
Department of Mathematics  
Pennsylvania State University  
xli@math.psu.edu

Chun Liu  
Penn State University  
cxl41@psu.edu

### MS235

#### Dynamic Sub-Grid Scalthe Mori-Zwanzig Formalism Models for Large Eddy Simulations Based on

This work uses the Mori-Zwanzig formalism to recast the Navier-Stokes equations into a coarse-grained non-Markovian system of equations. This low-dimensional system provides a starting point for the development of MZ-based coarse-grained models. The structural form of the MZ-based models are defined by the mathematics of the coarse-graining process. Non-local effects are captured through a finite memory approximation of the MZ memory kernel, the time-scale of which is determined through a dynamic Germano-like procedure. The outcome of this modeling process is a parameter-free, mathematically-derived sub-grid model. Results are presented for Fourier-Galerkin solutions of rotating and wall-bounded turbulent flows.

Eric Parish  
Univ of Michigan  
parish@umich.edu

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

### MS236

#### Performing Successive Tensor-Times-Matrix and -Vector Multiplies using Dimension Trees

Tensors are increasingly employed in many application domains to analyze the data, extract hidden relations within the data, or compress it for efficiency. Tucker and CP tensor decompositions are considered to be the canonical methods to accomplish these goals. They express an  $N$ -dimensional tensor  $\mathcal{X}$  using  $N$  matrices  $\mathbf{M}_1, \dots, \mathbf{M}_N$ , in the form of an outer or a tensor (times matrix) product. The standard algorithm for computing these decompositions is the alternating least squares (ALS) method. Starting with an initial guess on the matrices  $\mathbf{M}_1, \dots, \mathbf{M}_N$ , ALS keeps updating these matrices one-by-one using all other  $N-1$  matrices in an iterative fashion until convergence. For Tucker and CP decompositions, this update involves the multiplication of the tensor  $\mathcal{X}$  with  $N-1$  matrices (TTM), or the column vectors of these matrices (TTV), respectively. Since most matrices do not change from one iteration to the next, some parts of the computation can be stored and reused. To this end, we propose an algorithmic scheme to carry out such successive TTM and TTV operations more efficiently. We do this by storing and reusing partial TTM and TTV results in a tree framework. This approach provides significant computational gains both for Tucker and CP decompositions algorithms that involve successive TTMs and TTVs at their core, and is applicable to both sparse and dense tensors.

Oguz Kaya  
Inria and ENS Lyon

France  
oguz.kaya@ens-lyon.fr

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

### MS236

#### Efficient Parallel Software for Tucker Decompositions of Dense Tensors

As parallel computing tends toward the exascale, scientific data produced by simulations are growing increasingly massive, sometimes resulting in terabytes of data. By viewing this data as a dense tensor, we can compute a Tucker decomposition to find inherent low-dimensional multilinear structure, achieving impressive compression ratios with negligible loss in accuracy. We present recent improvements in our distributed-memory parallel implementation of the Tucker decomposition, whose key computations correspond to parallel linear algebra operations. To demonstrate the compression and accuracy of the method, we apply our software to real-world data sets from combustion simulations. We also provide detailed performance results.

Alicia Klinvex, Grey Ballard, Tamara G. Kolda  
Sandia National Laboratories  
amklinv@sandia.gov, ballard@wfu.edu,  
tgkolda@sandia.gov

### MS236

#### Tensor Decompositions for Bernoulli Data

Tensor decompositions are a powerful tool for multiway factor analysis. We consider the problem of modeling binary data using the canonical polyadic (CP) tensor decomposition. Binary data appears in many scenarios. For instance, a tensor representing a time-evolving graph may have a true in position  $(i, j, k)$  if nodes  $i$  and  $j$  are connected at time  $k$ . We assume the data is Bernoulli distributed and that the parameters in the tensor decomposition represent the logarithm of the odds ratio (i.e., also known as the “logit” function) of a true value. We discuss the formulation of the problem, its connection to standard (Gaussian) tensor decomposition, and give examples of its utility on real-world problems.

Tamara G. Kolda  
Sandia National Laboratories  
tgkolda@sandia.gov

Cliff Anderson-Bergman  
Sandia National Labs  
ciande@sandia.gov

### MS236

#### An Exploration of Optimization Algorithms for High Performance Tensor Completion

Tensor completion is a powerful tool used to estimate or recover missing values in multi-way data. It has seen great success in domains such as product recommendation and healthcare. Tensor completion is most often accomplished via low-rank sparse tensor factorization, a computationally expensive non-convex optimization problem which has only recently been studied in the context of parallel computing. In this work, we study three optimization algorithms that have been successfully applied to tensor com-

pletion: alternating least squares (ALS), stochastic gradient descent (SGD), and coordinate descent (CCD++). We explore opportunities for parallelism on shared- and distributed-memory systems and address challenges such as memory- and operation-efficiency, load balance, cache locality, and communication. Among our advancements are an SGD algorithm which combines stratification with asynchronous communication, an ALS algorithm rich in level-3 BLAS routines, and a communication-efficient CCD++ algorithm. We evaluate our optimizations on a variety of real datasets using a modern supercomputer and demonstrate speedups through 1024 cores. These improvements effectively reduce time-to-solution from hours to seconds on real-world datasets. We show that after our optimizations, ALS is advantageous on parallel systems of small-to-moderate scale, while both ALS and CCD++ will provide the lowest time-to-solution on large-scale distributed systems.

Shaden Smith  
University of Minnesota  
shaden@cs.umn.edu

Jongsoo Park  
Intel  
jongsoo.park@intel.com

George Karypis  
University of Minnesota / AHPARC  
karypis@cs.umn.edu

### MS237

#### Performance of the Quadruple Precision Eigensolver Library QPEigenK on Supercomputer Systems

In many computational science applications, it is frequently required to solve large and dense symmetric eigenvalue problems. The accuracy of the eigenpairs becomes worse as the matrix dimension increases due to rounding errors. In our experience, the accuracy of the computed eigenvalues of a 375,000-dimensional matrix by a double precision eigensolver is only a few digit while it costs an hour or more on a supercomputer system. To overcome the accuracy problem, we have recently developed a quad precision dense standard eigensolver library QPEigenK for distributed parallel computer systems, which supports the double-double precision arithmetic and parallelized in an MPI/OpenMP hybrid fashion. For example, all the eigenpairs of a 5,000-dimensional matrix can be computed in quadruple precision within only 10.8 seconds by using 512 nodes of the K computer. In this talk, we will present the performance evaluation results of the solvers on both the K computer and a Fujitsu FX100 supercomputer, which have different memory architectures. We will also discuss the accuracy and performance bottleneck of the solvers.

Yusuke Hirota  
RIKEN AICS  
yusuke.hirota@riken.jp

Susumu Yamada  
Japan Atomic Energy Agency  
yamada.susumu@jaea.go.jp

Toshiyuki Imamura  
RIKEN Advance Institute for Computational Science  
imamura.toshiyuki@riken.jp

Narimasa Sasa, Yasuhiro Idomura, Takuya Ina, Masahiko Machida  
Japan Atomic Energy Agency  
sasa.narimasa@jaea.go.jp, idomura.yasuhiro@jaea.go.jp,  
ina.takuya@jaea.go.jp, machida.masahiko@jaea.go.jp

### MS237

#### Reproducibility of Linear Algebra Operations

Leveraging the hierarchical and modular structure of linear algebra libraries, we investigate the possibility of building reproducible complex linear algebra operations – such as matrix factorizations – that can be expressed in terms of reproducible versions of low-level kernels like those defined in BLAS (Basic Linear Algebra Subprograms). As a case study, we introduce a reproducible variant of the unblocked LU factorization for graphics processor units (GPUs). For this purpose, we provide Level-1/2 BLAS kernels in the ExBLAS (Exact BLAS) library that deliver correctly-rounded and reproducible results for the dot (inner) product, vector scaling, and the matrix-vector product. In addition, we draw a strategy to enhance the accuracy of the triangular solve via inexpensive iterative refinement. Following a bottom-up approach, we finally assemble a reproducible implementation of the LU factorization for GPUs, which accommodates partial pivoting for stability and can be eventually integrated into a (blocked) high-performance and stable algorithm for the LU factorization.

Roman Iakymchuk  
HPCViz/PDC  
KTH Royal Institute of Technology  
riakymch@kth.se

Stef Graillat  
University Pierre and Marie Curie (UPMC)  
LIP6  
stef.graillat@lip6.fr

David Defour  
DALI/LIRMM  
University of Perpignan  
david.defour@univ-perp.fr

Sylvain Collange  
INRIA – Rennes  
sylvain.collange@inria.fr

Enrique S. Quintana-Orti  
Universidad Jaime I  
quintana@icc.uji.es

Erwin Laure  
KTH Royal Institute of Technology  
CSC, PDC/CST  
erwinl@pdc.kth.se

### MS237

#### Reproducible Parallel Simulations in HPC

Post Moore's era supercomputing will certainly require more hierarchical parallelism and variable precision floating-point arithmetic to satisfy the computing need of exascale numerical simulations. Nevertheless floating-point addition will remain non associative and so parallel computations will still be prone to return results being different from one run to another one. These failures of the numer-

ical reproducibility reduce the simulation reliability and complicate the debugging and the validating steps of large scale software.

We present two case studies to illustrate how to recover this numerical reproducibility without jeopardizing the computing efficiency. Hydrodynamics parallel simulations with the openTelemac code rely on finite element modelization, subdomain decomposition and iterative solvers. Two open-Telemac modules have been modified to provide reproducible results for any number of computing units thanks to targeted compensation techniques. We also describe and analyze generic solutions that are also provided by reproducible and accurately rounded BLAS.

Chemseddine Chohra  
University of Perpignan France  
chemseddine.chohra@univ-perp.fr

Philippe Langlois  
University of Perpignan  
France  
langlois@univ-perp.fr

Rafife Nheili  
University of Perpignan France  
rafife.nheili@univ-perp.fr

David Parelo  
University of Perpignan  
France  
david.parelo@univ-perp.fr

### MS237

#### Faithful Rounding for Matrix Multiplication

This talk concerns accurate matrix multiplication using floating-point arithmetic. Let  $F$  be a set of floating-point numbers. A constant  $u$  is used as a roundoff unit, e.g.,  $u = 2^{-53}$  for binary64 in the IEEE 754 standard. For matrices  $A \in F^{m \times n}$  and  $B \in F^{n \times p}$ , our goal is to obtain a matrix  $C \in F^{m \times p}$  such that

$$|AB - C| \leq \frac{1}{2}u|C|.$$

Here, assume that underflow does not occur in the floating-point arithmetic for the matrix multiplication  $AB$ . The computed result  $C$  is called faithful rounding for  $AB$ , because the computed result is one of the immediate floating-point neighbors of the exact result. For summation and dot product, there is an excellent algorithm for faithful rounding [Rump, Ogita, Oishi, Accurate floating-point summation part I: Faithful rounding, SISC, 2008]. We propose a fast algorithm which produces a faithfully rounded result of  $AB$  based on an error-free transformation of matrix multiplication [Ozaki et al., Error-free transformations of matrix multiplication by using fast routines of matrix multiplication and its applications, Numer. Alg., 2011]. We combine an a priori error estimation for blockwise matrix computations and a posterior validating technique with the error-free transformation for acceleration of computational performance. As a result, faithful rounding is guaranteed by computing only six matrix products in the best case.

Katsuhisa Ozaki  
Shibaura Institute of Technology / JST  
ozaki@sic.shibaura-it.ac.jp

Takeshi Ogita

Tokyo Woman's Christian University  
ogita@lab.twcu.ac.jp

### MS238

#### Fast Construction of Some Hierarchical Rank Structured Matrices with Nested Bases

In this talk, we present an efficient way to construct some hierarchical rank structured matrices with nested bases. Given a kernel function, a tree structure is first constructed based on an adaptive partitioning of the computational domain. In contrast with existing schemes based on either analytic or purely algebraic approximation, the proposed scheme takes advantages of both approaches and greatly improves the efficiency. The algorithm follows a bottom-up traversal of the tree and is able to perform the operations associated with each node on the same level in parallel. The main advantages of this construction scheme are its simplicity of implementation and its flexibility to construct various hierarchical rank structures. A rigorous error analysis is conducted to show that this scheme is both fast and stable. The efficiency and robustness of the proposed scheme are demonstrated through a few test problems arising from integral equations, structured matrices as well as machine learning applications

Difeng Cai  
Purdue University  
cai92@purdue.edu

Yuanzhe Xi  
University of Minnesota  
yxi@cs.umn.edu

Yousef Saad  
Department of Computer Science  
University of Minnesota  
saad@cs.umn.edu

Edmond Chow  
School of Computational Science and Engineering  
Georgia Institute of Technology  
echow@cc.gatech.edu

### MS238

#### Randomized QR with Column Pivoting

Factorizing large matrices using QR with Column Pivoting (QRCP) ordinarily requires substantially more processing time than QR without pivoting. This is due to the communication complexity required to process pivoting decisions. We discuss a modification to QRCP that uses randomized sampling to process blocks of column pivots. Our implementation of this algorithm competes with high-performance library implementations of QR in processing time, but it is as effective as QRCP for most rank-revealing applications as well as low-rank approximations. We extend this to a fast and reliable approximation of the truncated Singular Value Decomposition (SVD).

Jed Duersch  
University of California Berkeley  
jduersch@math.berkeley.edu

### MS238

#### Interconnected Hierarchical Structures for Solving

**Non-Coercive Elliptic PDEs**

We present a direct method for solving discretized elliptic PDEs. Based on a multilevel domain partitioning, we propose a way to form consistent hierarchically semiseparable matrices across different levels of subdomains. The representation reuses the basis from lower-level compression results. This improves the speed and the accuracy by avoiding existing recompression procedures. One major application is the Helmholtz equation with many wavelengths. This is joint work with Maarten V. de Hoop and Jianlin Xia.

Xiao Liu  
Rice University  
xiao.liu@rice.edu

Maarten de Hoop  
Center for Computational & Applied Mathematics  
Purdue University  
mdehoop@math.purdue.edu

Jianlin Xia  
Purdue University  
xij@math.purdue.edu

**MS238****Reliable Randomized Spectrum Revealing Matrix Factorizations**

We develop efficient and effective randomized algorithms for computing pivoted LU factorization and pivoted QR factorization. Our shared memory and distributed memory implementations outperform existing LAPACK and ScaLAPACK routines. The numerical experiments show that these randomized algorithms are as efficient and reliable as state-of-the-art alternatives for computing low-rank approximations. Additionally, we show that these randomized algorithms are as reliable as their deterministic counterparts with very high probability.

Jianwei Xiao  
University of California Berkeley  
jwxiao@math.berkeley.edu

Ming Gu  
University of California, Berkeley  
Mathematics Department  
mgu@math.berkeley.edu

**MS239****Global Optimization with Native Space Semi-Norm Bounds**

We propose a global optimization algorithm for expensive black-box objective functions that is based on radial basis functions (RBFs). Given an upper bound on the semi-norm of the objective function, we prove that our algorithm is globally convergent even though it may not sample densely. We discuss expected convergence rates and illustrate the performance of the method via experiments on a set of test problems.

David Eriksson, David Bindel, Christine A. Shoemaker  
Cornell University  
dme65@cornell.edu, bindel@cs.cornell.edu,

cas12@cornell.edu

**MS239****Multi-Fidelity Surrogate Modeling of Expensive Cosmology Functions**

We consider optimization problems whose objective function evaluations are based on computationally expensive high-fidelity simulation model runs and for which we have access to computationally cheaper low-fidelity models, with, for example, decreased physics complexity or resolution. Throughout the optimization search, we exploit the information from the low-fidelity model in order to inform the search on the high-fidelity model. We use computationally cheap surrogate models to approximate the simulation models of all fidelity levels. The information from the surrogate models is used to select new evaluation points. We sample the low-fidelity model at many more points than the high-fidelity model. The low-fidelity surrogate model suggests promising sample points to the high-fidelity surrogate model. Using the information derived from the high-fidelity surrogate model, we estimate the quality of the suggested parameters and either reject or accept them for evaluation. Vice versa, we use the high-fidelity surrogate model to select promising sample points and we use the low-fidelity simulation to assess the quality of the proposed point. We develop reward functions in order to decide when to sample the low- and the high-fidelity model, respectively. We show the performance of our algorithm on an application that involves computationally expensive cosmology functions.

Juliane Mueller  
Lawrence Berkeley National Lab  
juliane.mueller2901@gmail.com

Zarija Lukic  
Lawrence Berkeley National Laboratory  
zarija@lbl.gov

Dmitriy Morozov  
Lawrence Berkeley National Lab  
dmitriy@mrzv.org

**MS239****Tuning the Global Optimization Solver BARON using Derivative-Free Optimization Algorithms**

Optimization solvers provide users with numerous options that control different algorithmic aspects. These options can have a significant impact on solver performance. Therefore, tuning solver options is often necessary and leads to significant performance improvements. The aim of this work is to identify option settings that result in the best solver performance in terms of execution time and solution quality. Tuning options can be regarded as an optimization problem. This problem is hard to solve for two reasons. First, the relationship between the parameters and solver performance is not explicit. Second, some of the options may take discrete values, so the objective function is complex and non-smooth. Hence, the solver must be treated as a black-box system, whose input is values for the different options and output is a performance metric, such as the execution time. Derivative-free optimization algorithms (DFO) are attractive for this tuning problem since they do not require explicit functional representations of the objective function. We perform a computational study over a set of 126 problems from GLOBALLib and MINLPLib collections in order to identify optimal val-

ues for each one of the problems and also find a single set of options that can improve the performance of BARON across the entire test collection. A total of 27 DFO algorithms are used for this reason. We present extensive computational results.

Nikolaos Ploskas, Jianfeng Liu, Nikolaos Sahinidis  
Carnegie Mellon University  
nploskas@andrew.cmu.edu, liu1172@purdue.edu, sahinidis@cmu.edu

### MS239

#### Efficient Sampling Allocation Using Multi-Fidelity Information in Simulation Optimization

For complex problems intractable to analytical approaches, high-fidelity simulation models are often used to evaluate the performance of alternative system designs and select the best. For large-scale systems, high-fidelity simulations can be time-consuming and using simulation to rank and select the best design often faces a significant computational challenge. In this talk, we present a new framework that integrates information from multi-fidelity models to increase computational efficiency. We propose a Gaussian mixture model to capture important information contained in low-fidelity models that might not accurately approximate the performance of the high-fidelity simulation model. Posterior information obtained by a Gaussian mixture model-based clustering analysis incorporates both cluster-wise information and idiosyncratic information for each design. A new budget allocation method is proposed to efficiently allocate high-fidelity simulation replications, utilizing posterior information. Numerical experiments show that the proposed multi-fidelity framework achieved a significant boost in computational efficiency even when the low-fidelity models have large and heterogeneous model errors.

Yijie Peng  
George Mason University  
Fudan University  
ypeng10@gmu.edu

Jie Xu, Chun-Hung Chen  
George Mason University  
jxu13@gmu.edu, cchen9@gmu.edu

Loo Hay Lee  
National University of Singapore  
iseleelh@nus.edu.sg

### MS240

#### Breaking the $O(N)$ Barrier for Solving the High-Frequency Helmholtz Equation

We present a fast solver for the 3D high-frequency Helmholtz equation with heterogeneous, constant density, acoustic media. The solver is based on the method of polarized-traces, coupled with distributed linear algebra libraries and pipelining to obtain a solver with online runtime  $\mathcal{O}(\max(1, R/n)N)$  where  $N = n^3$  is the total number of degrees of freedom and  $R$  is the number of right-hand sides. In this implementation, we utilize hierarchical semi-separable (HSS) compression, as part of a compressed linear algebra library (STRUMPACK), to preserve linear scaling in the oversampled, high-frequency regime ( $\omega \sim n^{2/3}$ ) while preserving accuracy at any given frequency.

Russell Hewett

Total  
russell.hewett@total.com

Leonardo Zepeda-Nunez  
University of California, Irvine  
lzepeda@math.mit.edu

Adrien Scheuer  
Université catholique de Louvain  
adrien.scheuer@uclouvain.be

Laurent Demanet  
Department of Mathematics, MIT  
laurent@math.mit.edu

### MS240

#### Sparse Supernodal Solver Using Hierarchical Compression over Runtime System

In this talk, we present the PASTIX sparse supernodal solver, using hierarchical compression to reduce the burden on large blocks appearing during the nested dissection process. We compare the numerical stability, and the performance in terms of memory consumption and time to solution of different approaches by choosing when the compression of the factorized matrix should occur. In order to improve the efficiency of the sparse update kernel for both BLR (block low rank) and HODLR (hierarchically off-diagonal low-rank), we investigate the BDLR (boundary distance low-rank) method to preselect rows and columns in the low-rank approximation algorithm.

Gregoire Pichon  
INRIA  
gregoire.pichon@inria.fr

Eric F. Darve  
Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

Mathieu Faverge  
Bordeaux INP - Labri - Inria  
mathieu.faverge@inria.fr

Pierre Ramet  
Bordeaux University - INRIA  
pierre.ramet@inria.fr

Jean Roman  
INRIA  
Jean.Roman@inria.fr

### MS240

#### Hierarchical Matrix Operations on GPUs

We describe a high performance implementation of hierarchical matrix-vector multiplication, orthogonalization and compression using hardware accelerators. Dense matrices arising in scientific computing may often be represented by a hierarchical structure of blocks that can be approximated off the diagonal by low rank matrices, leading to greatly reduced memory footprints and arithmetic complexity. The generators for these blocks form a more complex family of objects than the uniform tiles of traditionally represented dense matrices that possess high performance implementations on GPUs. We describe the underlying tree structures of hierarchical matrices and an effective map-

ping onto GPU architecture. In a hierarchical basic linear algebra library that will ultimately open up new opportunities in matrix factorizations, inversions, and efficiently solving high-dimensional systems, these new GPU kernels represent key first building blocks.

Wajih Halim Boukaram  
KAUST  
wajihhalim.boukaram@kaust.edu.sa

George M Turkiyyah  
King Abdullah University of Science and Technology  
American University of Beirut  
gt02@aub.edu.lb

David E. Keyes  
KAUST  
david.keyes@kaust.edu.sa

Hatem Ltaief  
King Abdullah University of Science & Technology  
(KAUST)  
hatem.Ltaief@kaust.edu.sa

#### MS240

##### **Accelerating Parallel Optimization Methods with Hierarchical Matrix Computations**

We present recent developments in the algorithms, analysis, and applications of parallel optimization algorithms accelerated by hierarchical semiseparable (HSS) matrix computations. The structured eigenvalue solutions and other linear algebra subroutines allow for significantly decreased computational cost and storage over standard algorithms, while exploiting natural structure endemic to many classes of problems. We include theoretical complexity, convexity, and convergence analysis. We support our claims with numerical results, including applications to structural dynamics and other engineering problems.

James Vogel, Jianlin Xia  
Purdue University  
vogel13@purdue.edu, xiaj@math.purdue.edu

#### MS241

##### **Mixed Methods for Two-Phase Darcy-Stokes Mixtures of Partially Melted Materials with Regions of Zero Porosity**

The Earth's mantle, or an ice sheet, involves a deformable solid matrix phase within which a second phase, a fluid, may form due to melting processes. The mechanics of this system is modeled as a dual-continuum, with at each point of space the solid matrix being governed by a Stokes system and the fluid melt, if it exists, being governed by a Darcy law. This system is mathematically degenerate when the porosity (volume fraction of fluid) vanishes. We develop the variational framework needed for accurate approximation of this Darcy-Stokes system, even when there are regions of positive measure where only one phase exists. We then develop an accurate mixed finite element method for solving the system and show some numerical results.

Todd Arbogast  
University of Texas at Austin  
arbogast@ices.utexas.edu

Marc A. Hesse  
University of Texas

Department of Geological Sciences  
mhesse@jsg.utexas.edu

Abraham Taicher  
University of Texas at Austin  
ataicher@gmail.com

#### MS241

##### **Virtual Farming: Combining Knowledge to Enhance Decision-Making in Agricultural Regions**

Recent over pumping of aquifers lying beneath large agricultural regions, combined with the effects of long-term droughts, have stressed these underground water resources to near extinction. Even without extinction, the effects associated with overuse of this resource include salt water intrusion of coastal aquifers, which causes a reduction in the amount of drinking water available, and land subsidence, which compacts the aquifer and leads to a long-term reduction in the availability of water. Resolution of these issues requires community and state based solutions, where a variety of stakeholder objectives needs to be considered. Our team has developed a virtual farming tool to aid decision makers in this region, both from a livelihood and regulatory viewpoint. The tool allows us to consider a multi-component region, where environmental, urban, and agricultural considerations can be evaluated. We are developing the tool by combining existing software frameworks to take advantage of expertise in large-scale farm simulations and optimization algorithms. This has required our team to write wrappers to allow for communication between the tools, and it has required we work closely with a variety of industry partners to ensure the results are representative of the problem we want to resolve. In this talk, we discuss the stages of development of our framework and the cooperative nature of our team that has allowed us to push forward on this problem.

Eleanor Jenkins  
Clemson University  
lea@clemson.edu

Kathleen Kavanagh  
Clarkson University  
kkavanag@clarkson.edu

#### MS241

##### **Multiscale Integration of Blood Flow Mechanics and Neural Signaling to Model Retinal Microcirculation**

Several sight-threatening diseases are associated with impairment of blood flow regulation, which is the mechanism deputed to maintain proper tissue perfusion. Blood flow regulation is particularly important in the retina, where several processes contribute to determine vessel diameter and blood flow, including feedback mechanisms due to metabolic controls and feedforward mechanisms due to neurovascular coupling. Understanding and quantifying the relative importance of these mechanisms in health and disease may help clinical research in devising new strategies to preserve vision in many patients. Here, we address this issue from a theoretical viewpoint and we propose a multiscale model of retinal microcirculation that integrates: (i) blood flow; (ii) chemical reaction and transport across the arterial wall; and (iii) neural signaling from glial cells surrounding the blood vessel. The model aims at predicting alterations in arterial diameters as a function of space, time, biomechanical parameters and biochemical signals.

Numerical simulations suggest that feedforward glial synthesis plays an important role in retinal functional hyperaemia and that impaired nitric oxide synthesis may significantly affect vascular regulation. In particular, nitric oxide levels lower than baseline do not favor vasoconstriction and increase vasodilation, consistently with experimental data, whereas elevated nitric oxide levels result in a pathologically marked vasoconstriction.

Riccardo Sacco, Aurelio G. Mauri, Alessandra Cardani  
Politecnico di Milano  
riccardo.sacco@polimi.it,  
aureliogiancarlo.mauri@polimi.it,  
alessandra1.cardani@mail.polimi.it

Giovanna Guidoboni  
Indiana University-Purdue University at Indianapolis  
Department of Mathematical Sciences  
gguidobo@math.iupui.edu

Alon Harris  
Indiana University School of Medicine  
alharris@indiana.edu

#### MS241

##### Consolidation of a Sedimentary Basin

We model a sedimentary basin as a saturated granular deforming porous medium of sand or silt consolidating due to its own weight and an overburden pressure. The porosity satisfies a degenerate PDE of pseudo-parabolic type for which we describe the existence and regularity of solutions.

Ralph E. Showalter, Dwight Holland  
Oregon State University  
show@math.oregonstate.edu,  
ladwi@math.oregonstate.edu

hol-

#### MS242

##### Severe Erosion of a Cylinder Within the Subcritical Reynolds Range

The fluid structure interaction of an eroding clay cylinder in cross flow is simulated with a moving mesh method. The cylinder erodes as a function of the local wall shear stress and the flow field is developed for each mesh deformation. This problem involves multiple scales as the oscillating vortices downstream of the cylinder act on significantly separate time and space scales compared to the receding clay boundary. The cylinder profile in our simulations approached a terminal shape before eroding self-similarly until vanishing; the same self-similar form emerged from experiments by Ristroph et al. (2012) and theory by Moore et al. (2013). The self-similar form is a rounded triangular profile pointed towards the direction of flow and this boundary emerges from the initial circular cylinder. A novel method of positioning the mesh nodes on the cylinder boundary is employed to preserve mesh quality throughout the deformations. A second configuration with an array of eroding cylinders is explored by applying periodic boundary conditions. A similar terminal form is found regardless of the initial or intermediate cylinder profiles.

James N. Hewett  
University of Canterbury  
james@hewett.nz

Mathieu Sellier

Department of Mechanical Engineering  
University of Canterbury  
mathieu.sellier@canterbury.ac.nz

#### MS242

##### Sculpting of a Dissolving Body

Dissolution in fluids leads to the occurrence of many natural pattern formations. For example the Karst topography occurs when water dissolves limestone, and travertine terraces form as a balance of dissolution and precipitation. Dissolution, ablation and erosion are related problems involving the close interaction of a fluid domain and the solid morphology. Unlike erosion, dissolution can happen with or without external flow through molecular diffusion. The dissolved solute causes convection and leads to a self-generated flow, which enhances the dissolution and changes the shape evolution. In our study, we consider the shape evolution of a soluble object immersed in water, both with external flow imposed and in quiescent fluid. We find that different conditions lead to one of three possible outcomes during the shape evolution. In one case, the shape maintains a memory of its initial form, which is preserved throughout the entire dissolving process. A second possibility is that the shape converges to a unique final shape, which persists thereafter before vanishing. And the last possibility is a runaway process in which that the shape continues to evolve throughout time without reaching any terminal stage. These different possibilities affect our ability to accurately infer past conditions from observations of the present morphology, an issue of general importance in geomorphology.

Jinzi Mac Huang  
New York University  
jinzi@cims.nyu.edu

Megan Davies Wykes  
Courant Institute  
davieswykes@cims.nyu.edu

Nick Moore  
Florida State University  
mnmooore2@fsu.edu

Leif Ristroph  
Courant Institute, NYU  
lr1090@nyu.edu

#### MS242

##### Do Dissolving Objects Converge to a Universal Shape?

Surprisingly, macroscopic objects such as melting ice cubes and growing stalactites approach nonintuitive geometric ideals. Here we investigate the shape of dissolving cylinders in a large volume of water. The cylinders are oriented vertically and consist of amorphous glucose or poly(ethylene glycol). The dissolution causes density differences in the surrounding fluid, which induce gravity-driven convection downward along the object. The resulting concentration gradient shapes the cylinder according to fast dissolution at the tip and slow dissolution at the base. The contour of the object approaches a power law of the form  $z = aR^2$ , where  $z$  is the vertical distance from the tip and  $R$  is the corresponding radius. We suggest that this paraboloidal shape is the geometric attractor for the dissolution of non-

crystalline objects in the presence of gravity.

Elias Nakouzi

Florida State University  
 enakouzi@chem.fsu.edu ;enakouzi@chem.fsu.edu

#### MS242

##### **Interfacial Dynamics of Dissolving Objects from Discrete and Continuum Perspectives**

A diffusion-limited model of an object being dissolved will be examined from the continuous and discrete perspectives. In the first half, a continuum model of an object being dissolved in a two-dimensional potential flow will be explored. By taking advantage of conformal invariance of the model, a numerical method will be introduced that tracks the evolution of the object boundary in terms of a time-dependent Laurent series. Simulations of several dissolving objects will be shown, all of which show collapse to a single point in finite time. The simulations reveal a surprising connection between the position of the collapse point and the initial Laurent coefficients, which was subsequently derived analytically. In the second half, the discrete analog of the model will be explored, whereby the object is modeled as a cluster of particles that are annihilated by random walkers. Connections between the continuum and discrete perspectives will be explored.

Chris H. Rycroft, Yuexia Lin  
 Harvard University  
 chr@seas.harvard.edu, y\_lin@g.harvard.edu

#### MS243

##### **Demonstration of Unstructured Mesh Adaptation for Ice Sheet Modeling Applications**

Ice sheet dynamics on large continental ice masses, namely Greenland and Antarctica, impact the change in sea-level. The amount of ice sheet mass loss is a function of the ice velocity and thickness, fields which are far from smooth and have significant local variation. In order to estimate accurately these two fields, it is critical that one resolve the length scales that are important to the flow dynamics near local features like river beds. Meanwhile, it is desirable to employ a coarse discretization in areas where the ice flow is effectively constant to keep the problem size as small as possible. These requirements drive a need for time-varying adaptive meshing. This talk will overview the adaptive mesh capabilities implemented in the Albany/FELIX (Finite Elements for Land Ice eXperiments) solver that is under development at Sandia National Laboratories. Ice sheets are often represented as an incompressible non-Newtonian fluid, modeled by the Stokes equations with a nonlinear viscosity. The Albany/FELIX dycore implements a First-Order Stokes approximation to these equations. Albany is directly coupled to the RPI parallel mesh adaptation capabilities. We describe the adaptive meshing capabilities implemented in the Albany/PUMI simulation suite and provide an overview of the analysis capabilities that are available. Adaptive results for a steady-state Greenland ice flow scenario employing an unstructured prism mesh are presented.

Mauro Perego  
 CSRI Sandia National Laboratories  
 mperego@sandia.gov

Cameron Smith  
 Scientific Computation Research Center

Rensselaer Polytechnic Institute  
 smithc11@rpi.edu

Dan A. Ibanez  
 Rensselaer Polytechnic Institute  
 SCOREC  
 ibaned@rpi.edu

Glen Hansen, Irina K. Tezaur  
 Sandia National Laboratories  
 gahanse@sandia.gov, ikalash@sandia.gov

Mark S. Shephard  
 Scientific Computing Research Center  
 Rensselaer Polytechnic Institute  
 shephard@rpi.edu

#### MS243

##### **PISM-FEvoR: A Multi-Scale Ice Flow Model Incorporating Fabric Evolution with Recrystallization**

The deformation of ice within an ice-sheet is anisotropic when there is a preferred orientation to the ice crystallites (grains) due to the strong anisotropy of the crystalline lattice. This preferred orientation (often termed fabric) is widely observed in ice-sheets. In recent years there has been significant effort to connect polycrystalline fabric models with ice-sheet models to better account for the anisotropic response of ice. These resulting ice-sheet models (e.g., CAFFE and GOLF), use external polycrystalline fabric models to tune microstructural parameters that control the fabric development and recrystallization. These ice-sheet models assume the microstructure evolution equations are steady over time. We are developing a combined polycrystalline-flow model, using the Parallel Ice Sheet Model (PISM) and the Fabric Evolution with Recrystallization polycrystalline model (FEvoR). PISM-FEvoR provides the first integrated flow-fabric model that explicitly computes the fabric evolution and includes all three major recrystallization processes; grain growth, rotation recrystallization, and migration recrystallization. Because FEvoR is directly incorporated, we do not require steady material parameters. We use PISM-FEvoR to model the evolution of an idealized glacier, initialized with a variety of fabric profiles. PISM-FEvoR captures the flow enhancement due to fabric and we show that the entire integrated fabric-flow history determines the final simulated flow.

Joseph H. Kennedy  
 Oak Ridge National Laboratories  
 kennedyjh@ornl.gov

Constantine Khroulev  
 University of Alaska Fairbanks  
 c.khroulev@gmail.com

Florian Ziemen  
 Max Planck Institute for Meteorology  
 florian.ziemen@mpimet.mpg.de

Erin Pettit  
 University of Alaska Fairbanks  
 pettit@gi.alaska.edu

#### MS243

##### **A Three-Dimensional, Implicit, Thermo-Mechanical Computational Model for Polythermal**



**Ice**

Simulating ice sheet dynamics is essential for predicting the sea level rise in next decades to centuries. We present an implicit, fully three-dimensional computational model for the simulation of polythermal ice. In particular we focus on the temperature solver and we adopt the enthalpy formulation proposed by Aschwanden et. al. (2012), with the addition of the gravity-driven moisture drainage model proposed by Schoof and Hewitt (2016). The enthalpy model is implicitly coupled with the Blatter-Pattyn ice sheet flow model. The resulting thermo-mechanical model is fully implicit and allows for the solution of the steady state for temperature and velocity without the need of performing a temperature spin-up. We will present results on simplified geometries as well as for large-scale ice sheet problems, and compare them with results from the literature or obtained with other temperature formulations.

Alessandro Barone  
Department of Mathematics and Computer Science  
Emory University, Atlanta  
alessandro.barone@emory.edu

Mauro Perego  
CSRI Sandia National Laboratories  
mperego@sandia.gov

**MS243****Joint Model and Parameter Dimension Reduction for Bayesian Inversion Applied to an Ice Sheet Problem**

Model-based projections of the dynamics of the polar ice sheets play a central role in anticipating future sea level rise. However, a number of mathematical and computational challenges place significant barriers on improving predictability of these models. One such challenge is caused by the unknown model parameters (e.g., in the basal boundary conditions) that must be inferred from heterogeneous observational data, leading to a large-scale ill-posed inverse problem and to the need to quantify uncertainties in its solution. In this talk we discuss the problem of estimating the uncertainty in the solution of (large-scale) ice sheet inverse problems within the framework of Bayesian inference. The focus will be on a data-informed approach that identifies low-dimensional structure in both parameter space and the forward model state space. We show that using only a limited number of forward solves, the resulting subspaces lead to an efficient method to explore the high-dimensional posterior.

Noemi Petra  
University of California, Merced  
npetra@ucmerced.edu

Tiangang Cui  
Monash University  
tiangang.cui@monash.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

Benjamin Peherstorfer

ACDL, Department of Aeronautics & Astronautics  
Massachusetts Institute of Technology  
pehersto@mit.edu

Karen E. Willcox  
Massachusetts Institute of Technology  
kwillcox@MIT.EDU

**MS244****Fourth-Order Entropy Stable Non-Oscillatory Spectral Collocation Schemes**

A fourth-order WENO spectral collocation schemes is developed, that is nonlinearly entropy stable for the one-dimensional Navier-Stokes equations. Individual spectral elements are coupled using penalty type interface conditions. The resulting entropy-stable WENO spectral collocation scheme achieves design order accuracy, maintains the WENO stencil biasing properties across element interfaces, and satisfies the summation-by-parts (SBP) operator convention, thereby ensuring nonlinear entropy stability. Numerical results demonstrating accuracy and non-oscillatory properties are presented.

Mark H. Carpenter  
NASA Langley Research Center  
mark.h.carpenter@nasa.gov

**MS244****Toward Entropy-Stable Gas-Dynamics Simulations Using Tetrahedral SBP Elements**

We present an entropy-stable semi-discretization of the Euler equations. The scheme is based on high-order summation-by-parts (SBP) operators for triangular and tetrahedral elements, although the theory is applicable to multidimensional SBP operators on more general elements. While there are established methods for proving stability of linear equations, such as energy analysis, they are not adequate for nonlinear equations. To address nonlinear stability, we use the matrix properties of the SBP operators combined with entropy-conserving numerical flux functions to prove the semi-discrete scheme conserves entropy. Significantly, the proof does not rely on integral exactness, and, therefore, the discretization has a stronger claim of robustness than a similar finite-element method. This generalizes previous work proving entropy stability for tensor-product elements to more general elements, benefiting problems involving complex geometry. Numerical experiments are conducted to verify accuracy and entropy conservation on an isentropic vortex flow.

Jared Crean  
Rensselaer Polytechnic Institute  
jcrean01@gmail.com

Jason E. Hicken  
Rensselaer Polytechnic Institute  
Assistant Professor  
hickej2@rpi.edu

**MS244****BR1 is not Unstable for the Compressible Navier-Stokes Equations**

The Bassi-Rebay-1 (BR1) scheme for coupling the viscous terms in discontinuous Galerkin approximations of the compressible Navier-Stokes equations is very simple to im-

plement and does not have adverse effects on the time step. It has been associated, however, with unstable computations, which have led some to abandon it for more complicated coupling procedures. We will show how to use a discrete integral calculus that is derived from the summation by parts property of the Legendre-Gauss-Lobatto quadrature to prove that the discontinuous Galerkin spectral element approximation (DGSEM) with the BR1 scheme is stable, provided that the advective terms in the equations are approximated stably.

David A. Kopriva  
Department of Mathematics  
The Florida State University  
kopriva@math.fsu.edu

Gregor Gassner  
Institute for Aerodynamics and Gasdynamics  
Universitaet Stuttgart  
ggassner@math.uni-koeln.de

Andrew R. Winters  
Mathematical Institute  
University of Cologne  
awinters@math.uni-koeln.de

Florian Hindenlang  
Max-Planck Institute for Plasma Physics  
florian.hindenlang@ipp.mpg.de

#### MS244

##### Efficiency Assessment of Split Form Nodal Discontinuous Galerkin Schemes for the Compressible Navier-Stokes Equations

Gassner et al. (*Split form nodal discontinuous Galerkin schemes with summation-by-parts property for the compressible Euler equations, Journal of Computational Physics, 327:39–66, 2016*) recently proposed a split form nodal discontinuous Galerkin (DG) framework. For a non-linear partial differential equation (PDE) different split forms are found by interpreting any non-linearities of the system, e.g. a quadratic term  $\rho u$ , and averaging conservative and non-conservative forms of the PDE. It was shown that the split form DG framework remains conservative and greatly enhanced the robustness of the numerical approximation particularly for under-resolved turbulence computations. This improved robustness comes with the caveat that the computational cost of the split form DG increases. In this presentation we provide implementation details and efficiency comparisons of the split form DG scheme to the classical DG spectral element method.

Andrew R. Winters  
Mathematical Institute  
University of Cologne  
awinters@math.uni-koeln.de

Florian Hindenlang  
Max-Planck Institute for Plasma Physics  
florian.hindenlang@ipp.mpg.de

Gregor Gassner  
Institute for Aerodynamics and Gasdynamics  
Universitaet Stuttgart

ggassner@math.uni-koeln.de

#### MS245

##### Propagating Fluctuations Across Atomistic and Mesoscopic Simulations

In this talk, we focus on continuum-particle hybrid simulations within the framework of domain decomposition method (DDM) for simple and complex fluids. Most of the pioneering hybrid simulations focused on the average quantities of the coarse observables from a perspective of Navier-Stokes (NS) description, whereas thermal fluctuations from discrete molecular dynamics (MD) were considered as unfavorable noise to be filtered out in course of the coupling. However, the fluctuations are indispensable elements for the dynamics of soft matter, for example, a colloid dispersion. We consider an application of a colloid solute surrounded by solvent particles which are described by molecular dynamics while further away from the colloid, it is simulated by a Lagrangian particle solver of the fluctuating hydrodynamics, such as the smoothed dissipative particle dynamics (SDPD). Motivated by applications of such, we devote efforts to propagate fluctuations between MD and SDPD in a hybrid simulation. As a matter of fact, we only need to couple simple fluids between the two subdomains, as the colloid is always in the center of MD and does not interact with the SDPD subdomain directly. Therefore, it is more natural to couple the state variables and flux between the simple liquids described MD and SDPD. The hydrodynamic coupling also enables time-scale separation between MD and SDPD so that both spatial and temporal efficiencies of the SDPD are exploited.

Xin Bian  
Brown University  
xin.bian@tum.de

Weinan E  
Princeton University  
Department of Mathematics  
weinan@math.princeton.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

#### MS245

##### Large-Scale Real-Space Electronic Structure Calculations

In this talk, the development of a real-space formulation for Kohn-Sham density functional theory (DFT) and a finite-element discretization of this formulation, which can handle arbitrary boundary conditions and is amenable to adaptive coarse-graining, will be presented. In particular, the accuracy afforded by using higher-order finite-element discretizations, and the efficiency and scalability of the Chebyshev filtering algorithm in pseudopotential and all-electron Kohn-Sham DFT calculations will be demonstrated. Further, the development of a subquadratic-scaling approach (in the number of electrons) based on a subspace projection and Fermi-operator expansion will be discussed, which will be the basis for the future development of coarse-graining techniques for Kohn-Sham DFT. The developed techniques have enabled, to date, pseudopotential calculations on non-periodic and periodic systems containing  $\sim 10,000$  atoms, as well as all-electron calculations on systems containing  $\sim$

10,000 electrons.

Vikram Gavini

Mechanical Engineering, University of Michigan  
vikramg@umich.edu

#### MS245

##### **Bridging Time and Length-Scales in Dislocation Dynamics**

Dislocations are line defects found in crystals which evolve under loading, and are the carriers of plastic slip in such materials. Moreover, they interact via long-range stress fields, which leads to complex ensemble behaviour. For these reasons, accurately modelling and simulating their evolution is an important problem for Materials Science applications. In this talk, I will discuss my recent work to provide a rigorous foundation for Discrete Dislocation Dynamics, as well as ongoing work to analyse the computational methods used to simulate such models.

Thomas Hudson

University of Warwick  
hudsont@cermics.enpc.fr

#### MS245

##### **Data-Driven Parameterization of Generalized Langevin Equation**

We propose a method based on appropriate parameterization to compute the memory kernel of the generalized Langevin Equation (GLE). The proposed method is data-driven as kernel is constructed by merely using trajectory data. While the first-order approximation is consistent with the Markovian approximation, higher order approximations can be systematically derived. The approximated kernel formulation satisfies the second fluctuation-dissipation conditions naturally and enables us to eliminate the history-dependence by introducing auxiliary variables of Markovian process.

Huan Lei, Nathan Baker

Pacific Northwest National Laboratory  
huan.lei@pnnl.gov, nathan.baker@pnnl.gov

Xiantao Li

Department of Mathematics  
Pennsylvania State University  
xli@math.psu.edu

#### MS246

##### **Combined Uncertainty and A Posteriori Bounds for CFD Calculations**

This presentation discusses the ongoing development of combined uncertainty and error bound estimates for computational fluid dynamics (CFD) calculations subject to imposed random parameters and random fields. An objective of this work is the construction of computable error bound formulas for output uncertainty statistics that guide CFD practitioners in systematically determining how accurately CFD realizations should be approximated and how accurately uncertainty statistics should be approximated for output quantities of interest. Formal error bounds formulas for moment statistics that properly account for the presence of numerical errors in CFD calculations and numerical quadrature errors in the calculation of moment statistics have been previously presented in Barth

(2013). In this past work, hierarchical node-nested dense and sparse tensor product quadratures are used to calculate moment statistics integrals. In the present work, a framework has been developed that exploits the hierarchical structure of these quadratures in order to simplify the calculation of an estimate of the quadrature error needed in error bound formulas. When signed estimates of realization error are available, this signed error may also be used to estimate output quantity of interest probability densities as a means to assess the impact of realization error on these density estimates.

Timothy J. Barth

NASA Ames Research Center  
Timothy.J.Barth@nasa.gov

#### MS246

##### **Goal-Based Adaptive Control for Uncertainties Dominated Compressible Flows**

The purpose of this work is to address the issue of realistic, compressible flows while in presence of uncertainties and errors. More precisely, an adaptive method is proposed to enhance the quality of a functional of interest obtained from the solution of a parametrised system. Adaptive methods has been widely explored in the literature and recent extensions of some approaches to stochastic problems has been proposed using a posteriori estimates. The main issue with these approaches is the difficulty to extract accurate anisotropic information. Our purpose is to fulfill this need by using a priori error estimates based on interpolation errors. The anisotropic information in the deterministic space is extracted using a Riemannian metric based method and a similar approach is extended to stochastic errors.

Our purpose is to control the errors in an automatic, adaptive way. The spatial error is controlled through mesh adaptivity where the "optimal" mesh is generated from an optimal Riemannian metric, a matrix tensor defined for each node as a hessian weighted by an adjoint state. In a similar manner we can control the stochastic error by solving an optimisation problem where the error is controlled through interpolation errors in the stochastic space weighted by the probability density function.

Performance of the proposed approach is illustrated on 2D and 3D CFD problems and we will restrain to low to moderate number of uncertain parameters.

Anca Belme

Sorbonne Univ, UPMC Univ Paris 06  
CNRS, UMR 7190, Inst Jean Le Rond d'Alembert  
belme@dalembert.upmc.fr

#### MS246

##### **Uncertainty Quantification for Approximate P-Quantiles for Physical Models with Stochastic Inputs**

We consider the problem of estimating the p-quantile for a given functional evaluated on solutions of a deterministic model in which model input is subject to stochastic variation. We derive upper and lower bounding estimators of the p-quantile. We perform an a posteriori error analysis for the p-quantile estimators that takes into account the effects of both the stochastic sampling error and the deterministic numerical solution error and yields a computational error bound for the estimators. We also analyze the asymptotic convergence properties of the p-quantile estimator bounds in the limit of large sample size and decreasing numerical

error and describe algorithms for computing an estimator of the p-quantile with a desired accuracy in a computationally efficient fashion.

Don Estep  
Colorado State University  
estep@stat.colostate.edu

Daniel Elfverson  
Umeå University  
Department of Mathematics and Mathematical Statistics  
daniel.elfverson@umu.se

Fredrik Hellman  
Department of Information Technology  
Uppsala University  
fredrik.hellman@it.uu.se

Axel Malqvist  
Chalmers University of Technology  
axel@chalmers.se

#### MS246

##### Output-Based Adaptation for Chaotic Flow Simulations

This talk will present advances in output-based error estimation and adaptation methods for computational fluid dynamics problems exhibiting unsteady phenomena ranging from small-scale unsteadiness to chaos. Of interest is the impact of spatial and temporal mesh resolution on statistical outputs computed from the unsteady flow solution. Multiple adaptive techniques will be discussed and compared, including heuristics based on flow residuals, and adjoint-based methods. The adjoint-based methods will be regularized using steady-state and windowing approximations, and using least-squares shadowing. The methods will be applied to the Lorenz oscillator and to the Navier-Stokes equations.

Krzysztof Fidkowski, Yukiko Shimizu  
University of Michigan  
kfid@umich.edu, ykmizu@umich.edu

#### MS247

##### Brownian Dynamics for a Confined Suspension of Microrollers

Brownian Dynamics of confined rigid bodies is an emerging topic both in numerics and experiments. In this talk we will combine various methods recently developed for Brownian Dynamics to simulate suspensions of active colloidal rollers near a wall. In this setup, Brownian Motion plays an essential role by setting the gravitational height of the particles. Recent theoretical, numerical and experimental studies have shown that the gravitational height controls the dynamics of this system, and more specifically the wavelength of the fingering instability. We will show how to properly include fluctuations efficiently in this system and study the effect of the gravitational height on the fingering instability.

Florencio Balboa Usabiaga  
Courant Institute of Mathematical Sciences  
New York University  
fbalboa@courant.nyu.edu

Blaise Delmotte  
Courant Institute

NYU  
delmotte@cims.nyu.edu

Aleksandar Donev  
Courant Institute of Mathematical Sciences  
New York University  
donev@courant.nyu.edu

#### MS247

##### A Fluctuating Boundary Integral Method for Brownian Suspensions

We present a novel boundary integral method for simulating Brownian Dynamics (BD) of rigid particles immersed in a Stokesian fluid. A key ingredient of performing BD simulation is the accurate and efficient sampling of Brownian displacements with zero mean and covariance proportional to the hydrodynamic mobility matrix, which translates applied forces and torques to particle motion, as dictated by fluctuation-dissipation theorem. We achieve this task by proposing a first-kind boundary integral formulation of the mobility problem with a random slip velocity added to the surface of particles whose covariance is proportional to the periodic Greens function of Stokes flow (Oseen tensor). The efficient sampling of the random velocity makes use of Ewald splitting of Stokeslet into two SPD pieces: a far-field wave-space piece generated by FFT, and a near-field real-space piece generated by a Lanczos iterative method. We develop Krylov iterative method for the resulting linear saddle-point system with a block-diagonal preconditioner which virtually eliminates the inherent ill-conditioning of the first-kind formulation, and the matrix-vector multiplication is accelerated by the Spectral Ewald method. Numerical tests show that our method achieves high-order accuracy and the computational costs scale linearly with the number of particles.

Yuanxun Bill Bao, Aleksandar Donev  
Courant Institute of Mathematical Sciences  
New York University  
billbao@cims.nyu.edu, donev@courant.nyu.edu

Eric Keaveny  
Imperial College  
e.keaveny@imperial.ac.uk

Leslie Greengard  
Simons Foundation and Courant Institute  
New York University  
greengar@cims.nyu.edu

Manas N. Rachh  
Applied Mathematics Department  
Yale University  
manas.rachh@yale.edu

#### MS247

##### Collective Colloid Diffusion Under Soft 2D Confinement

This work presents a numerical and theoretical investigation of the collective dynamics of colloids in an unbounded solution but trapped in a harmonic potential. In the limit of infinitely stiff trap (strict 2D colloidal confinement) the collective diffusion is enhanced and diverges at zero wavenumber (like  $1/k$ ), due to the hydrodynamic propagation of the confining force across the layer. We study several issues on this problem. We first consider

the gradual transition to normal diffusion taking place as the confining potential trap is relaxed and secondly the transition from Stokesian dynamics to inertial dynamics at large wavelengths. Numerical simulations, performed with finite-volume fluctuating hydrodynamics and the immersed boundary method, resolve the fluid inertia and permit to vary the Schmidt number  $Sc$  at will. Decreasing  $Sc$  leads to a transition from Stokesian dynamics to inertial flow dynamics, whereby the colloid density autocorrelation is no longer overdamped and presents anticorrelation at some finite time ruled by feedback momentum transfer of the solvent. Ideal tracers and interacting colloids are considered, and compared with analytical solutions of the collective diffusion coefficient. Comparison with Brownian dynamics without hydrodynamics illustrate the masking by conservative forces, of the hydrodynamic diffusion enhancement.

Rafael Delgado Buscalioni  
 Departamento de Física Teórica de la Materia  
 Condensada  
 Universidad Autónoma de Madrid, Madrid, Spain  
 rafael.delgado@uam.es

#### MS247

##### **A Rigid Multiblob Algorithm for Confined Brownian Suspensions**

We examine the simulation of Brownian motion of rigid bodies with arbitrary shape, suspended in a viscous solvent in both confined and unconfined domains. Existing numerical techniques which capture the correct stochastic drift term require the solution of two saddle point problems per time step and only achieve first order accuracy deterministically. The saddle point systems are required to enforce the rigidity constraint and can be quite expensive to solve for a large number of rigid bodies. We propose a novel method which requires the solution of two saddle point problems per time step but achieves second order accuracy in time.

Brennan Sprinkle  
 Northwestern University  
 brennansprinkle2013@u.northwestern.edu

Florencio Balboa Usabiaga  
 Universidad Autónoma de Madrid  
 Madrid, Spain  
 florencio.balboa@uam.es

Neelesh Patankar  
 Northwestern University  
 neeshpatankar@gmail.com

Aleksandar Donev  
 Courant Institute of Mathematical Sciences  
 New York University  
 donev@courant.nyu.edu

#### MS248

##### **Exact Preservation of Zero Velocity Divergence in a High-Order, Mapped-Grid, Finite-Volume Discretization of a Gyrokinetic System**

Gyrokinetic models of plasmas are developed using asymptotic methods leading to coordinate systems whose time evolution defines the phase space particle velocity. Developed as Hamiltonian systems, the equations of motion are therefore area-preserving, yielding divergence-free velocities. Consequently, gyrokinetic Vlasov models of

phase space advection can be expressed in conservation form. When discretizing these models, preservation of the divergence-free velocity property is important in avoiding the accumulation of truncation error in long time integrations. In this talk, we describe our approach for satisfying the discrete divergence-free property to machine roundoff in a fourth-order, mapped-grid, finite-volume discretization of a gyrokinetic system in 4D (two axisymmetric configuration space coordinates plus two velocity space coordinates) and 5D (three configuration space coordinates plus two velocity space coordinates). Implementation of the approach in our COGENT code applied to tokamak reactor edge plasma simulations will also be described. This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and by Lawrence Berkeley National Laboratory under contract DE-AC02-05CH11231.

Milo Dorr  
 Center for Applied Scientific Computing  
 Lawrence Livermore National Laboratory  
 dorr1@llnl.gov

Phillip Colella  
 Lawrence Berkeley National Laboratory  
 PColella@lbl.gov

Mikhail Dorf  
 Lawrence Livermore National Laboratory  
 dorfl@llnl.gov

Debojyoti Ghosh  
 Lawrence Livermore National Laboratory  
 Center for Applied Scientific Computing  
 ghosh5@llnl.gov

Jeffrey A. Hittinger  
 Center for Applied Scientific Computing  
 Lawrence Livermore National Laboratory  
 hittinger1@llnl.gov

#### MS248

##### **Implementation of Implicit-Explicit Time Integration for the Kinetic Modeling of Tokamak Plasma Edge**

The kinetic simulation of tokamak edge plasma requires the solution of the Vlasov-Poisson equations with an appropriate collision model. Due to sharp density and temperature variations, there is a large variation in the collisional time scale (between hot and cold regions), as well as between the collision and Vlasov time scales. Explicit time-integration is thus inefficient with time steps bounded by the fastest time scales. We describe the implementation of high-order, conservative implicit-explicit time integrators in COGENT, a finite-volume solver for the 4-dimensional Vlasov-Poisson equations on mapped, multi-block grids. We consider the multi-stage additive Runge-Kutta (ARK) methods for our application. In our implementation, the Vlasov (hyperbolic) term is integrated explicitly in time while the collision (parabolic/source) term is integrated implicitly. The implicit stages are solved using the Jacobian-free Newton-Krylov approach; however, a lower-order (hence, sparser) approximation to the Jacobian is assembled and stored for preconditioning the linear solve. We analyze the performance of the semi-implicit time integrators with the Fokker-Planck collision model on problems that are representative of the tokamak edge re-

gion. In particular, we consider cases with both uniform as well as sharply varying collisional time scales. The accuracy and convergence of the ARK methods are verified and the computational expense is compared to that of explicit RK methods.

Debojyoti Ghosh

Lawrence Livermore National Laboratory  
Center for Applied Scientific Computing  
ghosh5@llnl.gov

Milo Dorr

Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory  
dorr1@llnl.gov

Mikhail Dorf

Lawrence Livermore National Laboratory  
dorf1@llnl.gov

Jeffrey A. Hittinger

Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory  
hittinger1@llnl.gov

#### MS248

##### **Astrophysical Plasmas at the Exascale: Challenges and Opportunities**

In this talk, I discuss some of the challenges and opportunities relating to cosmological simulations of galaxy formation in particular, those relating to modeling the diffuse plasma in the interstellar and intergalactic medium. These media are weakly collisional, have a tremendous range of magnetic field properties, and often are dusty and have significant energy contributions from radiation and cosmic rays. I will discuss some of the challenges of modeling this multiphysics, high-dynamic-range environment as we move toward exascale computers.

Brian O'Shea

Department of Physics and Astronomy  
Michigan State University  
oshea@msu.edu

#### MS248

##### **Stabilization of Numerical Interchange in Spectral-Element Magnetohydrodynamics**

Auxiliary numerical projections of the divergence of flow velocity and vorticity parallel to magnetic field are developed and tested for the purpose of suppressing unphysical interchange instability in MHD simulations. Free energy for interchange comes from the alignment of magnetic curvature and the gradient of fluid pressure, and magnetic shear leads to spatial localization and Suydam's threshold condition for physical interchange. The numerical version of the instability arises with equal-order C0 finite- and spectral-element expansions of the flow velocity, magnetic field, and pressure, and it is sensitive to behavior at the limit of resolution. The auxiliary numerical projections are motivated by physical considerations, and coercive responses to the projections are added to the flow-velocity equation. Their incomplete expansions are limited to the highest-order orthogonal polynomial in at least one coordinate of the spectral elements. Cylindrical eigenmode computations show that the projections induce convergence from the stable side with first-order ideal-MHD equations during h-refinement and p-refinement [Sovinec, JCP 319,

61-78 (2016)]. Hyperbolic and parabolic responses are tested and compared. The projections are also shown to be effective in linear and nonlinear time-dependent computations, provided that the responses introduce some numerical dissipation.

Carl Sovinec

University of Wisconsin - Madison  
csovinec@wisc.edu

#### MS249

##### **Experimental Design In Sparse Polynomial Chaos Expansions**

Compressive sampling methods have proven efficient in solving PDEs with random inputs, when the solution-dependent quantities of interest (QoI) admit sparse expansions in orthogonal polynomial bases, e.g. polynomial chaos. The majority of compressive sampling techniques rely on samples of the QoI generated in a Monte Carlo fashion, i.e., random samples. This talk will discuss the use of classical design of optimal experiments in the context of compressive sampling via greedy techniques, and provide empirical results illustrating the advantage of such sampling strategies.

Alireza Doostan

Department of Aerospace Engineering Sciences  
University of Colorado, Boulder  
Alireza.Doostan@Colorado.EDU

Paul M. Diaz

Colorado School of Mines  
pdiaz@mymail.mines.edu

Paul Diaz

University of Colorado, Boulder  
paul.diaz@colorado.edu

#### MS249

##### **Robust Compressive Sensing with Application to Multifidelity Analysis of Complex Turbulent Flows**

Engineering and science models of practical interest often depend on high-dimensional parameter spaces. When combined with expensive simulations that result in a limited number of simulations, direct explorations of the model behavior over this space (often part of its uncertainty quantification analysis) is prohibitive. We thus take an approach to construct *surrogate* models in lower dimensional spaces, by identifying the important dimensions through global sensitivity analysis (GSA). In particular, we enable GSA by first finding sparse linear architectures in orthogonal polynomial basis using sparse regression techniques. One major challenge, when the number of candidate basis functions is much larger than the number of model evaluations, is the phenomenon of overfitting. In this talk, we focus on the development of algorithms that are robust towards overfitting, by directly incorporating cross-validation into the basis selection procedure. This approach is demonstrated on a multifidelity analysis of large eddy simulations of SCRAMJET engine, where sparse polynomial chaos surrogates are built from data across different model fidelities and grid resolutions.

Xun Huan, Cosmin Safta, Khachik Sargsyan

Sandia National Laboratories  
xhuan@sandia.gov, csafta@sandia.gov,  
ksargsy@sandia.gov

Michael S. Eldred  
Sandia National Laboratories  
Optimization and Uncertainty Quantification Dept.  
mseldre@sandia.gov

Zachary Vane, Guilhem Lacaze, Joseph C. Oefelein  
Sandia National Laboratories  
zvane@sandia.gov, gnlacaz@sandia.gov,  
oefelei@sandia.gov

Habib N. Najm  
Sandia National Laboratories  
Livermore, CA, USA  
hnnajm@sandia.gov

#### MS249

##### Polynomial Approximation Using Compressive Sampling and the Spectral Distribution

Given a probability measure, we construct a sampling measure from which to perform compressive sampling reconstruction and approximation. This sampling measure, the Spectral Distribution, has various near-optimality properties for regularization approximations in the collocation framework. We discuss various theoretical properties of this distribution, along with numerical considerations and example. One significant advantage of the spectral distribution is that we obtain near-optimal approximation results that are insensitive to the choice of probability measure. In particular we can sample from the Spectral Distribution with only limited statistical information of the underlying probability measure.

Akil Narayan  
University of Utah  
akil@sci.utah.edu

#### MS249

##### Enhancing Sparsity of Generalized Polynomial Chaos Expansions

Compressive sensing has become a powerful addition to uncertainty quantification in recent years. In this talk we introduce a new method to identify new random variables through linear mappings such that the representation of the quantity of interest is more sparse with the new random variables. This sparsity increases both the efficiency and accuracy of the compressive sensing-based uncertainty quantification method. This method is applicable for generalized polynomial chaos expansions. We demonstrate the effectiveness of the new method with applications in solving stochastic partial differential equations.

Xiu Yang  
Pacific Northwest National Laboratory  
xiu.yang@pnnl.gov

Xiaoliang Wan  
Louisiana State University  
Department of Mathematics  
xlwan@math.lsu.edu

Lin Lin  
University of California, Berkeley  
Lawrence Berkeley National Laboratory  
linlin@math.berkeley.edu

Nathan Baker

Pacific Northwest National Laboratory  
nathan.baker@pnnl.gov

#### MS250

##### A Fractional Laplacian-Based Closure Model for Turbulent Fluid Flows

We study a closure model with the fractional Laplacian of order  $\alpha$ ,  $\alpha \in (0, 1)$  representing the turbulent diffusivity. We investigate the energy spectrum of the model by applying Pao's energy transfer theory. For the case  $\alpha = 1/3$ , the corresponding power law of the energy spectrum in the inertial range has a correction exponent for the standard Kolmogorov  $-5/3$  scaling exponent. For this case, the model corresponds to Richardson's particle pair-distance superdiffusion of fully developed homogeneous turbulent flow and Lévy jumps that lead to the superdiffusion. For other values of  $\alpha$ , the power law of the energy spectrum is consistent with the standard Kolmogorov  $-5/3$  scaling exponent. We also propose and study a modular algorithm for the semi-discrete form of the Navier-Stokes equations. The algorithm is minimally intrusive to a given legacy code for solving the Navier-Stokes equations by decoupling the local and nonlocal parts of the equations. We prove that the algorithm is unconditionally stable and unconditionally, first-order convergent.

Max Gunzburger  
Florida State University  
Department of Scientific Computing  
mgunzburger@fsu.edu

Nan Jiang  
Missouri University of Science and Technology  
jiangn@mst.edu

Feifei Xu  
Florida State University  
winterflyfei@gmail.com

#### MS250

##### Optimization with Respect to Order in a Fractional Diffusion Model: Analysis and Approximation

We consider an identification problem, where the state  $u$  is governed by a fractional elliptic equation and the unknown variable corresponds to the order  $s \in (0, 1)$  of the operator. We study the existence an optimal pair  $(\bar{u}, \bar{s})$  and provide sufficient conditions for its uniqueness. We develop semi-discrete and fully discrete algorithms to approximate the solution and provide an analysis of their convergence properties. We present numerical illustrations that confirm and extend our theory.

Abner J. Salgado  
Department of Mathematics  
University of Tennessee  
asalgad1@utk.edu

Enrique Otarola  
Department of Mathematics  
Universidad Tecnica Federico Santa Maria, Chile  
enrique.otarola@usm.cl

Harbir Antil  
George Mason University  
Fairfax, VA

hantil@gmu.edu

### MS250

#### An Efficient Probabilistic Numerical Method Based on Fourier-Cosine Series for Fractional Laplacian Equations

We develop a probabilistic numerical scheme based on Fourier-cosine series to solve linear and semi-linear fractional Laplacian equations in unbounded domains. Since the fractional Laplacian operator is the infinitesimal generator of the standard symmetric alpha stable process, the temporal discretization leads to an induction time-stepping scheme involving conditional expectations with respect to the alpha stable process. Those expectations are then approximated using the Fourier-cosine series expansions, relying on the availability of the characteristic function of the stochastic process. We provide error estimates of our scheme in the one-dimensional case. The proposed scheme is applied to one- and two-dimensional fractional Laplacian equations in unbounded domains, in order to demonstrate its effectiveness and efficiency.

Guannan Zhang

Oak Ridge National Laboratory  
zhangg@ornl.gov

### MS250

#### Nonlocal Transport in Bounded Domains

The study of nonlocal transport in physically relevant systems requires the formulation of mathematically well-posed and physically meaningful nonlocal models in bounded spatial domains. The main problem faced by fractional diffusion models in this case resides in the treatment of the boundaries. For example, the naive truncation of the Riemann-Liouville fractional derivative in a bounded domain is in general singular at the boundaries and, as a result, the incorporation of generic, physically meaningful boundary conditions is not feasible. In this presentation we discuss alternatives to address the problem of boundaries in fractional diffusion models. Our main goal is to present models that are both mathematically well posed and physically meaningful. Our approach is based on the regularization of the singularities at the boundary. Following the formal construction of the models in 1-dimensional and 2-dimensional bounded domains, we present finite-difference methods to solve the regularized models, and conclude with examples of nonlocal heat transport in nuclear fusion plasmas.

Diego del-Castillo-Negrete

Oak Ridge National Lab.  
delcastillod@ornl.gov

### MS251

#### Data Assimilation with Reduced Models

Due to computationally expensive forward models, ensemble filters often either use the full model with a small ensemble, or use a reduced model with a large ensemble. In either case, the sampling error or the model error poses a challenge to ensemble filter performance. Covariance localization and inflation techniques have been developed to ameliorate the errors. However, tuning localization and inflation is expensive, and they introduce further filter errors. We investigate an alternative approach of using improved reduced prior models. The improved reduced model is a

non-Markovian model obtained by quantifying the model error using a discrete-time stochastic parametrization, so that it captures the statistical and dynamical features of the full model. The improved reduced model achieves good filter performance using a reasonable size of ensemble without any localization or inflation. Tests on the two-layer Lorenz 96 model show that an improved reduced model has better filter performance than the original reduced system with tuned localization and inflation, and it has comparable performance to the full system with a small ensemble and tuned localization and inflation.

Fei Lu

Department of Mathematics  
Lawrence Berkeley National Laboratory  
flu@lbl.gov

Fei Lu

Department of Mathematics, UC Berkeley  
Mathematics group, Lawrence Berkeley National Laboratory  
flu@lbl.gov

### MS251

#### Reduced-Order Modeling of Bayesian Inverse Problems for PDEs

The solution of inverse and uncertainty quantification problems involving systems modeled by partial differential equations (PDEs) is computationally demanding. We show how to take advantage of reduced-order modeling techniques to speed up the numerical approximation of Bayesian inverse problems related with parameter estimation for both stationary and time-dependent PDEs. In the former case, we rely on Markov Chain Monte Carlo (MCMC) methods to characterize the posterior distribution of the parameters; in the latter, we exploit the Ensemble Kalman filter for performing state/parameter estimation sequentially. In both cases, we replace usual high-fidelity techniques (e.g., the finite element method) with inexpensive but accurate reduced-order models (ROMs) to speed up the solution of the forward problem. On the other hand, we develop suitable reduction error models (REMs) - or ROM error surrogates - to quantify in an inexpensive way the error between the high-fidelity and the reduced-order approximation of the forward problem, in order to gauge the effect of this error on the posterior distribution of the identifiable parameters. Numerical results dealing with the estimation of both scalar parameters and parametric fields highlight the combined role played by RB accuracy and REM effectivity.

Andrea Manzoni

EPFL, MATHICSE-CMCS  
Switzerland  
andrea.manzoni@epfl.ch

Stefano Pagani

Politecnico di Milano  
stefano.pagani@polimi.it

Alfio Quarteroni

Ecole Pol. Fed. de Lausanne  
Alfio.Quarteroni@epfl.ch

### MS251

#### On a Formulation for Constructing Optimal Low-Rank Approximations with Respect to Goal Func-



**tionals**

The subject of the talk will be concerned with the development of a mathematical formulation with the goal of constructing reduced-order models tailored for the approximation of quantities of interest. The main idea is to formulate a minimization problem that includes an inequality constraint on the error in the goal functional so that the resulting model be capable of delivering predictions of the quantity of interest within some prescribed tolerance. The formulation will be applied and tested to the so-called proper generalized decomposition (or low-rank approximation) method. Such a paradigm represents a departure from classical goal-oriented approaches in which a reduced model is first derived by minimization of the energy, or of residual functionals, and then adapted via a greedy approach by controlling the error with respect to quantities of interest using dual-based error estimates. Numerical examples will be presented in order to demonstrate the efficiency of the proposed approach.

Serge Prudhomme

École Polytechnique de Montréal  
Montréal, Québec, Canada  
serge.prudhomme@polymtl.ca

Kenan Kergrene

École Polytechnique de Montréal  
kenan.kergrene@polymtl.ca

Ludovic Chamoin

LMT, ENS Cachan, CNRS, Université Paris-Saclay  
chamoin@lmt.ens-cachan.fr

Olivier Le Maître

Laboratoire d'Informatique pour la Mécanique et les  
Science  
olm@limsi.fr

Marc Laforest

École Polytechnique de Montréal  
marc.laforest@polymtl.ca

**MS251****Equation-Driven Probability Density Function Estimators for High-Dimensional Stochastic Dynamical Systems**

We present a new general data-driven method to estimate the PDF of a quantity of interest in high-dimensional stochastic dynamical systems. The key idea is to combine information from sample trajectories with the exact evolution equation governing the PDF of the quantity of interest. Such equation can be derived in rather general cases, e.g., by using the Mori-Zwanzig formulation. The PDF estimate is obtained by maximizing a suitable convex likelihood functional subject to normalization conditions, positivity constraints and the (linear) PDE that governs the PDF we are interested in. To address high-dimensionality, we employ separated series expansions and alternating direction algorithms. The effectiveness of proposed methods will be demonstrated through various examples involving high-dimensional stochastic dynamical systems.

Daniele Venturi

Department of Applied Mathematics and Statistics  
University of California Santa Cruz  
venturi@ucsc.edu

Johannes O. Royset

Operations Research Department  
Naval Postgraduate School  
joroyset@nps.edu

**MS252****Surrogate Construction and Adaptivity in Simulations/Sampling**

In this talk we will focus on the use of adjoint based information in constructing surrogates of large complex simulations. Our work here is motivated by the sensitivity to numerical and model evaluation error of simulations used in surrogate construction. We will look at this using two mechanisms first, we will look at simple gradient enhanced constructions of surrogates. Second, we will also use information from dual weighted residual type computed using adjoints and residual estimates. In both cases we will establish the quality of the surrogates and their fitness for purpose - UQ in our case.

Hossein Aghakhani

University at Buffalo  
haghakha@buffalo.edu

Abani Patra

SUNY Buffalo  
abani@eng.buffalo.edu

Elaine Spiller

Marquette University  
elaine.spiller@marquette.edu

**MS252****Basis and Sample Adaptive Polynomial Chaos Expansions**

Methods for computing accurate approximations of a Quantity of Interest (QoI) via Polynomial Chaos and  $\ell_1$ -regression with low-coherence sampling have developed promising numerical and theoretical results. We present an iterative method which for each iteration identifies a basis set and corresponding coefficients which adapt to the QoI as a function. In addition to investigating this adaptation on its own for several models, we also show a method that adapts previously generated low-coherence samples for computations with subsequent bases by generating samples accounting for coherence adjustments arising from considering iteratively adapting bases, and this additional adaptation is seen to be beneficial for certain problems.

Jerrad Hampton

University of Colorado, Boulder  
jerrad.hampton@colorado.edu

Alireza Doostan

Department of Aerospace Engineering Sciences  
University of Colorado, Boulder  
Alireza.Doostan@Colorado.EDU

**MS252****Variational Multiscale Method Based Error Estimation and Mesh Adaptivity for Intrusive UQ**

In this talk, we develop and apply an adaptive approach based on the variational multiscale (VMS) formulation for stochastic PDEs with uncertain input data. In this approach we employ finite elements and unstructured meshes

in the physical domain and generalized polynomial chaos in the stochastic domain. The VMS method allows in computing an accurate solution in a coarse space while accounting for the missing/fine scales through a model term. This model term is defined by an approximation of the fine-scale stochastic Green's function. Within each element, we consider an algebraic approximation for this model term based on the strong-form residual and a stochastic stabilization parameter, which is used in the variational statement for computing the numerical solution. Similarly, a model term is derived to estimate the error in the numerical solution in a local/element-wise fashion. Further, the model term for the local error is approximated using the same stabilization parameter that is used in computing the numerical solution and thus, making it relatively cheap to compute. We show that the proposed method is effective as a local error estimator and in driving adaptivity. In this talk, we will focus on mesh adaptivity in the physical domain. We demonstrate the effectiveness of our approach on multiple problem cases including the advection-diffusion equation in which both advective and diffusive regimes are spanned and scalar transport in a channel with multiple branches.

Jason Li

Department of Mechanical, Aerospace and Nuclear  
Engineering  
RPI  
lij14@rpi.edu

Assad Oberai

Department of Mechanical, Aerospace and Nuclear  
Engineering  
Rensselaer Polytechnic Institute  
oberaa@rpi.edu

Onkar Sahni

Rensselaer Polytechnic Institute  
sahni@rpi.edu

## MS252

### Intrusive UQ Algorithms for Emerging Computing Platforms

In this work we study the discretization of stochastic boundary value problems in the context of emerging computing platforms. We consider a linear elliptic diffusion equation with uncertain diffusivity modeled via Karhunen-Loeve expansions. The stochastic solution is based on Polynomial Chaos expansions and employs a finite element discretization in the physical space. We will discuss the impact of several sparsity patterns on the construction and solution of the resulting stochastic Galerkin system on manycore architectures.

Cosmin Safta

Sandia National Laboratories  
csafta@sandia.gov

Habib N. Najm

Sandia National Laboratories  
Livermore, CA, USA  
hnnajm@sandia.gov

Eric Phipps

Sandia National Laboratories  
Optimization and Uncertainty Quantification Department

etphipp@sandia.gov

## MS253

### Multigrid Preconditioning in Support of Optimization-Based Domain Decomposition for Elliptic Equations

We present a multigrid preconditioning strategy in support of the optimization-based non-overlapping domain decomposition method for elliptic partial differential equations developed by Gunzburger, Heinkenschloss, and Lee (2000). The proposed method is rooted in our recent work on multigrid preconditioners for boundary control of elliptic equations, and results, as expected, in a computational cost that is decreasing with increasing resolution relative to the cost of solving the elliptic equations on each subdomain.

Andrei Draganescu, Mona Hajghassem

Department of Mathematics and Statistics  
University of Maryland, Baltimore County  
draga@umbc.edu, mona4@umbc.edu

## MS253

### Solver Composition Across the PDE/Linear Algebra Barrier

High-level finite element projects such as FEniCS and now, Firedrake, automate the construction of algebraic systems from a symbolic description of the underlying weak forms. Beneath this lies access to scalable linear algebra libraries such as PETSc. The resulting combination delivers high-performance simulations with low user overhead. However, this chain works best in a "one way" mode – the PDE code produces sparse matrices and residual evaluation and then any (composition of) algebraic solvers present in PETSc can be used. Customized algorithms that make use of PDE-level information, such as preconditioners that discretize auxiliary differential operators, can be difficult to integrate into the system. In this talk, we present recent work on "implicit matrices" in the context of the Firedrake project. In addition to assembled sparse matrices, Firedrake can also construct a Python PETSc matrix that embeds the entire problem description. This option makes several advanced features relatively straightforward to use from Firedrake. For example, matrix-free methods are now possible via a simple option. Block-structured matrices and field-split preconditioners are available via processing a UFL bilinear form rather than assembled monolithic or nested matrices. It is now relatively simple to implement and compose custom preconditioners in Python. We will describe the new technology and present several examples from model fluids problems.

Rob C. Kirby

Baylor University  
robert\_kirby@baylor.edu

Lawrence Mitchell

Department of Computing  
Imperial College London  
lawrence.mitchell@imperial.ac.uk

## MS253

### Accurate Finite Element Solution of the Fully Coupled Thermoacoustic Equations for Modeling of Trace Gas Sensors

Trace gas sensors are utilized in a wide range of applica-

tions, such as monitoring of air quality, breath analysis for medical diagnosis, and industrial process control. In these sensors a sinusoidally modulated laser source excites the trace gas to generate a pressure-temperature wave that is detected by a quartz tuning fork. To date, mathematical modeling of trace gas sensors has focused on either the acoustic or the thermal components of the wave only. A model that fully couples temperature, pressure and fluid velocity, as derived by Morse and Ingard, has only been considered for relatively simple domains. In this paper we present the first 3D finite element model of the thermoacoustic equations in the presence of the tuning fork, coupled with temperature diffusion in the sensor. We employ a perfectly matched layers (PML) method to truncate our computational domain, and benchmark it against the analytic solution in the free-space. As the resulting discretized system is poorly conditioned, we present a parallelizable block preconditioning approach. In order to precondition the highly indefinite Helmholtz block with PML, we test a novel idea based on domain decomposition, where we couple an algebraic multigrid solver in the computational domain to a direct solver in the PML region.

Artur Safin

University of Texas at Dallas  
artur.safin@utdallas.edu

Susan Minkoff

Department of Mathematical Sciences  
University of Texas at Dallas  
sminkoff@utdallas.edu

John Zweck

University of Texas at Dallas  
Department of Mathematical Sciences  
zweck@utdallas.edu

**MS253**

**Preconditioners for Stokes Flow with Highly Heterogeneous Viscosity Structure: Saddle-Point Smoothing Via Local Incomplete Factorization**

The solution of stationary Stokes flow with highly heterogeneous, non grid-aligned viscosity structure is a prevalent computational bottleneck in geodynamics simulation. Scalable, robust multigrid methods are important to allow for the long timescale, 3D, whole-planet simulations used to investigate the formation of structures in the Earth and other planets. To this end, we present results using incomplete local ILDL factorizations as smoothers within a geometric multigrid hierarchy for stationary Stokes flow problems, demonstrating a new scalable and robust technique.

Patrick Sanan

USI - Università della Svizzera italiana  
ICS  
patrick.sanan@gmail.com

Olaf Schenk

USI - Università della Svizzera italiana  
Institute of Computational Science  
olaf.schenk@usi.ch

Matthias Bollhoefer

Institut fuer Mathematik  
Technische Universitaet Berlin  
bolle@math.tu-berlin.de

Karl Rupp

Institute for Microelectronics  
Vienna University of Technology  
me@karlrupp.net

Dave A. May

ETH Zurich  
dave.may@erdw.ethz.ch

**MS254**

**Multipreconditioned GMRES for Shifted Systems**

We propose using the Multipreconditioned Generalized Minimal Residual (MPGMRES) method for solving shifted linear systems  $(A + \sigma_j I)x_j = b$  for  $j = 1, \dots, n_\sigma$  using multiple shift-and-invert preconditioners. The multipreconditioned space is obtained by applying the preconditioners to all search directions and searching for a minimum norm solution over this larger subspace. We show that this space grows linearly and show results for systems arising from an example problem in subsurface imaging.

Tania Bakhos

Institute for Computational and Mathematical Engineering  
Stanford University  
tbakhos@bcamath.org

Peter K. Kitanidis

Dept. of Civil and Environmental Engineering  
Stanford University  
peterk@stanford.edu

Scott Ladenheim

Temple University  
saladenh@temple.edu

Arvind Saibaba

Department of Electrical and Computer Engineering  
Tufts University  
arvind.saibaba@tufts.edu

Daniel B. Szyld

Temple University  
Department of Mathematics  
szyld@temple.edu

**MS254**

**Adaptive Coarse Spaces for FETI-DP Domain Decomposition Methods**

We present a family of adaptively chosen coarse spaces for FETI-DP domain decomposition methods for the equations of linear elasticity in 3D with arbitrarily large jumps in the material coefficients. The coarse spaces are defined by local eigenvectors which are obtained from adaptive eigenvalue computations related to a given tolerance. The new family of algorithms roots in an earlier algorithm by Mandel and Sousedik. For one of our new algorithms a rigorous proof and a condition number estimate will be given. This estimate is independent of the coefficient jumps, which do not have to align with the subdomain boundaries. It also holds for the case of almost incompressible elasticity. This estimate and its proof are new for three dimensions. Additional variants of this algorithms are obtained, using different simplifications and heuristics. For some problems, this leads to computationally more efficient versions. A comparison of the different algorithms

is provided as well.

Axel Klawonn, Kühn Martin  
 Universitaet zu Koeln  
 axel.klawonn@uni-koeln.de, makuehn@math.uni-koeln.de

Oliver Rheinbach  
 Technische Universität Bergakademie Freiberg  
 oliver.rheinbach@math.tu-freiberg.de

#### MS254

##### **Adaptive Multipreconditioning and its Application to Domain Decomposition.**

Domain Decomposition methods are a family of solvers designed for very large linear systems that require parallel computers. They proceed by splitting the computational domain into subdomains and then approximating the inverse of the original problem with local inverses coming from the subdomains. I will present a classical domain decomposition method and show that for realistic simulations (with heterogeneous materials for instance) convergence usually becomes very slow. Then I will explain how this can be fixed by injecting more information into the solver. In particular I will show how using multiple search directions within the conjugate gradient algorithm makes the algorithm more reliable. Efficiency is also taken into account since our solvers are adaptive.

Nicole Spillane  
 Centre de Mathématiques Appliquées  
 Ecole Polytechnique  
 nicole.spillane@cmap.polytechnique.fr

Christophe Bovet  
 LMT, ENS Cachan, CNRS, Université Paris-Saclay,  
 94235 Cachan  
 christophe.bovet@lmt.ens-cachan.fr

Pierre Gosselet  
 LMT, ENS Cachan  
 gosselet@lmt.ens-cachan.fr

#### MS254

##### **Enlarged GMRES**

We propose a variant of GMRES method for the solution of linear systems of equations with one right-hand side and multiple right-hand sides. Our method is based on the idea of the enlarged Krylov subspace introduced by L. Grigori et al. to reduce communication. By the nature of this idea, our methods inherits a block version of GMRES, so we were interested in the detection of inexact breakdowns focusing on the strategy proposed by H. Calandra et al. with a modification on the test of detection. We also propose an eigenvalues deflation technique aiming to have two benefits, on one hand, to avoid the plateau of convergence after the end of a cycle in the restarted version, and on the other hand to have a very fast convergence when solving the same system with different right-hand sides each one given in a time (useful for CPR preconditioner).

Hussam Al Daas  
 INRIA  
 hussam.al-daas@inria.fr

Laura Grigori  
 INRIA  
 France

Laura.Grigori@inria.fr

Pascal Henon  
 TOTAL E&P  
 pascal.henon@total.com

Philippe Ricoux  
 TOTAL SA  
 philippe.ricoux@total.com

Olivier Tissot  
 INRIA  
 olivier.tissot@inria.fr

#### MS255

##### **A Parallel Compensated Horner Scheme**

The Compensated Horner Scheme is an accurate and fast algorithm to evaluate univariate polynomials in floating-point arithmetic. The accuracy of the computed result is similar to the one given by the Horner scheme computed in twice the working precision. This Compensated Horner Scheme runs at least as fast as existing implementations producing the same output accuracy. In this talk, we will present a parallel version of the Compensated Horner Scheme. Some experiments on multicore and Graphics Processor Units (GPU) architectures will be presented to show the efficiency of this algorithm.

Stef Graillat  
 University Pierre and Marie Curie (UPMC)  
 LIP6  
 stef.graillat@lip6.fr

#### MS255

##### **Verified Conic Solutions of Underdetermined Linear Systems**

A wide variety of optimization problems are attacked by using conic linear programming. In particular, this includes linear and semidefinite programming. In cases where the use of floating point arithmetic leads to erroneous results, it is desirable to secure the solving procedure by applying verification methods. However, without sensible a priori bounds on the optimal decision variables or the condition number of the considered programming problem, the computation of verified bounds for the optimal objective value is still challenging since it requires the computation of tight inclusions for (near) optimal feasible points. An algorithm that realizes the computation of such inclusions has been developed by Jansson. Together with Chaykin and Keil, he has shown how rigorous error bounds for the optimal value of semidefinite programming problems can be computed by postprocessing the output of the applied solver. In this work, we present two new postprocessing strategies which can be traced back to the computation of verified solutions of underdetermined linear systems satisfying additional conic inequalities. We are discussing the operation ranges, advantages, and disadvantages of the new procedures, and compare them with the original postprocessing approach due to Jansson.

Marko Lange  
 Waseda University  
 markonoon@gmail.com

#### MS255

##### **Iterative Refinement for Eigenvectors of Symmet-**

### ric Matrices with Clustered Eigenvalues

This study is concerned with numerical methods for solving a standard eigenvalue problem  $Ax = \lambda x$  where  $A$  is a symmetric matrix. It is difficult to obtain accurate eigenvectors associated with clustered eigenvalues. To overcome the problem, we propose efficient algorithms for refinement of eigenvectors. The proposed algorithms are based on matrix multiplication in higher-precision arithmetic. We first present a basic algorithm for improving all the eigenvectors associated with well-separated eigenvalues, which quadratically converges if a modestly accurate initial guess is given. On the basis of the basic algorithm, we propose a refinement algorithm which can also improve the eigenvectors associated with clustered eigenvalues. The proposed algorithm provides accurate eigenvectors of a real symmetric matrix by iteration, up to the limit of computational precision. We present numerical results showing excellent performance of the proposed algorithm in terms of convergence rate and overall computational cost.

Takeshi Ogita  
Tokyo Woman's Christian University  
ogita@lab.twcu.ac.jp

Kensuke Aishima  
Graduate School of Information Science and Technology,  
University of Tokyo  
Kensuke\_Aishima@mist.i.u-tokyo.ac.jp

### MS255

#### Verified Solution of Unconstrained and Constrained Nonlinear Global Optimization Problems

In this talk we will present method for the verified solution of unconstrained and constrained nonlinear global optimization problems. Here verified means that under all circumstances and including all possible sources of errors mathematically correct inclusions global minima will be computed. More precisely, two lists will be generated: The first list comprising of verified global minima, and the second of remaining candidates. The latter applies if, for example, there are several global minima which cannot be distinguished, or if the problem is too ill-conditioned to be solved in double precision floating-point arithmetic. The methods are written in pure Matlab/Octave code as part of INTLAB, the Matlab/Octave toolbox for reliable computing. Preliminary tests show that, despite the interpretation overhead, the methods compare favourably with competitive methods, even when implemented in C. For example, the famous Griewank function for  $n = 50$  unknowns has an estimated number of  $10^{129}$  local minima in the original box of definition  $[-600, 600]^n$ . Sharp bounds for the unique global minimum are computed on a Laptop in about two minutes.

Siegfried M. Rump  
Hamburg University of Technology  
rump@tuhh.de

### MS256

#### An Immersed Structural Potential Method Framework for Incompressible Flexible/Rigid/Multi-Phase Flow Interaction

A general framework for the computational simulation of Fluid-Structure Interaction (FSI) problems involving rigid/flexible solids and multiphase flows is presented. Our methodology builds upon the Immersed Structural Pot-

tential Method (ISPM) developed by the authors for the simulation of single-phase FSI problems. First, the extension for the modelling of rigid bodies is considered for multiphase-phase flows by suitable consideration of constraints on the velocity field which are projected in a weighted least-squares sense. The corresponding constrained system is then solved by use of the augmented Lagrange method in conjunction with a fractional step approach. This approach allows for a simple and unified treatment of a variety of conditions for possible rigid bodies. The extension for the general case with flexible structures is accomplished by adding to the momentum equation the corresponding body force obtained by use of the ISPM. The extension of this methodology for flexible solids of different densities to those of the surrounding fluid phases is also presented within the same framework. The proposed allows for robust and flexible numerical solution of complex multi-phase multi-body systems with arbitrary combinations of rigid and flexible bodies, and the specific representation of rigid bodies does not increase the cost of the projection step in the fractional step method for multi-phase flows.

Antonio Gil  
University of Swansea  
a.j.gil@swansea.ac.uk

L. Yang  
Swansea University  
l.yang@swansea.ac.uk

Javier Bonet  
University of Greenwich, London  
j.bonet@swansea.ac.uk

### MS256

#### Over-Coming Fluid-Structure Instabilities for Incompressible Flows and Light Bodies

The added-mass instability has, for decades, plagued partitioned fluid-structure interaction (FSI) simulations of incompressible flows coupled to light solids and structures. Many current approaches require tens or hundreds of expensive sub-iterations per time-step. In this talk some new stable partitioned algorithms are described for coupling incompressible flows with (1) compressible elastic bulk solids, (2) thin structural beams and (3) rigid bodies. These added-mass partitioned (AMP) schemes require no sub-iterations, can be made fully second- or higher-order accurate, and remain stable even in the presence of strong added-mass effects. These schemes are implemented using moving and deforming overlapping grids with the Overture framework.

William D. Henshaw  
Rensselaer Polytechnic Institute  
henshw@rpi.edu

### MS256

#### Coupled FEM-DEM Methods for Modeling Variably Saturated Granular Media

The Discrete Element Method is commonly used in geotechnical engineering to model the dynamics of granular media. Recently, there has been increased interest in understanding the dynamics of granular media in partially and fully saturated conditions at both low and high particle Reynolds numbers. The larger objective is to develop micro-scale models of processes such as erosion and slope failure. In our related work we have developed two-phase

incompressible flow models for modeling air/water flows interacting with moving vessels and vehicles, porous solids, and flexible vegetation, but the various methods we have applied lack robustness and accuracy for coupling to DEM-based solid dynamics for a variety of reasons, most notably the frequent solid contacts causing topological change of the solid phase. As our two-phase work has relied heavily on the level set methodology, including signed distance function representations of phase interfaces on unstructured meshes, representing the granular phase via signed distances functions is a natural extension of our existing tools. In this work we present results for modeling such immersed solids in two-phase flows using the signed distance representation of the granular material.

Chris Kees

U.S. Army Engineer Research and Development Center  
Coastal and Hydraulics Laboratory  
christopher.e.kees@usace.army.mil

Matthew Farthing

US Army Engineer Research and Development Center  
matthew.w.farthing@usace.army.mil

**MS256**

**An Embedded Boundary Method for Shock-Dominated Fluid-Solid Interaction Problems**

Shock dominated fluid-solid interaction problems arise in a number of modern engineering applications. The abrupt, high amplitude pulses of mechanical energy associated with shock waves can be used to crush a kidney stone, destroy an enemy submarine, clean the biofouling on ship hull, and consolidate powders for 3D printing just to name a few examples. The development of a predictive simulation approach for this type of problems is a formidable challenge. It requires accounting for geometrical nonlinearities and dynamic fracture in the computational model. It also calls for the development of advanced algorithms for the computation and transfer of shock loads on fluid-solid interfaces which often undergo large deformation and topological change. In this talk, a second-order accurate, numerically stable partitioned procedure coupling a 3D, embedded boundary computational fluid dynamics (CFD) solver and a 3D computational solid dynamics (CSD) solver will be presented. Several recently developed numerical methods for tracking complex material interfaces and enforcing the kinematic and dynamic interface conditions will be discussed in detail. The salient features of this computational framework will be demonstrated through applications in shock wave lithotripsy, cavitation erosion, and underwater explosion and implosion.

Kevin Wang

Virginia Tech  
Aerospace & Ocean Engineering  
kevinw3@vt.edu

Yeyue Xiong

Department of Biomedical Engineering and Mechanics  
Virginia Tech  
xiongyy@vt.edu

Shunxiang Cao

Department of Aerospace and Ocean Engineering  
Virginia Tech  
csxtovt@vt.edu

Pei Zhong, Ying Zhang

Department of Mechanical Engineering and Materials  
Science  
Duke University  
pzhong@duke.edu, zhang.ying@duke.edu

**MS257**

**Robust Flux Error Estimation of Nitsches Methods for High Contrast Interface Problems**

We prove optimal error estimates for the flux variable for a stabilized Nitsche's method applied to an elliptic interface problem with discontinuous constant coefficients. These error estimates are totally independent of the contrast between diffusion coefficients.

John Guzman

Brown University  
johnny\_guzman@brown.edu

Erik Burman

University of Sussex  
Department of Mathematics  
e.n.burman@sussex.ac.uk

Manuel Sanchez-Uribe

University of Minnesota  
manuel\_sanchez\_uribe@alumni.brown.edu

Marcus Sarkis

Worcester Polytechnic Institute  
Instituto de Matematica Pura e Aplicada (Brazil)  
msarkis@wpi.edu

**MS257**

**Improved ZZ a Posteriori Error Estimation of Conforming FEM for Diffusion Problems**

In this talk, I will introduce and analyze an improved Zienkiewicz-Zhu (ZZ) error estimator. Due to easy implementation, generality, and ability to produce quite accurate estimation, the ZZ estimator has been widely adapted in engineering practice and has been the subject of mathematical study. Despite popularity of the ZZ estimator, it is also well known that AMR algorithms using the ZZ estimator are not efficient to reduce global error for non-smooth problems because it over-refines regions with relative small errors. Our improved error estimation overcomes the inefficiency of the ZZ estimator for non-smooth diffusion problems, and it is proved theoretically and numerically to be both reliable and efficient. Moreover, the recovery techniques are explicit and, hence, the implementation is simple and highly cost-effective.

Cuiyu He

Purdue University  
he75@purdue.edu

**MS257**

**Composite Finite Elements**

Composite Finite Elements are a new class of finite elements for the discretization of boundary value problems with complicated structures, e.g., in the geometry of the physical object and/or in the coefficients of the differential operator and boundary conditions. In contrast to standard finite elements, the minimal dimension of the approximation space is independent of the geometric details and this is especially advantageous for problems on domains with

complicated micro-structures. In our talk, we will introduce this discretization method for different kinds of applications such as Poisson-type equations, Lam equation, and Stokes equation. We will analyse its convergence in an a-priori and a-posteriori way and illustrate the analysis by numerical experiments.

Stefan Sauter  
University of Zurich  
stas@math.uzh.ch

### MS257

#### Error Estimates for Immersed Finite Element Volume Methods

We will present a class of finite volume element methods (FVEM) for elliptic interface problems. These FVEM are based on the immersed finite element functions, so that the solution meshes do not need to align with interfaces. We will present the a priori error estimates for the immersed FVEM. Moreover, we will compare the immersed FVEM with the immersed finite element methods (FEM) in their computational performance and error analysis. This is a joint work with Qingsong Zou from Sun Yat-Sen University.

Xu Zhang  
Mississippi State University  
xuzhang@math.msstate.edu

Qingsong Zou  
Sun Yat-Sen University  
mcszqs@mail.sysu.edu.cn

### MS258

#### Parallel Exponential Integrators Based on General Linear Methods

Exponential integrators are efficient numerical methods for solving large systems of stiff ordinary differential equations. In this talk, we present a new class of exponential integrators that are parallelizable both across the method and across the steps. These new integrators can be expressed as exponential general linear methods with coefficient matrices constructed to maximize parallelism and minimize communication costs. For sufficiently large systems, these new schemes provide high-order accuracy for similar cost as the Euler method. We will present several performance tests, and discuss advantages of using these methods in situations where exponential matrix functions are computed using Krylov subspace methods.

Tommaso Buvoli  
Department of Applied Mathematics, University of Washington  
Seattle, WA 98195 USA  
buvoli@uw.edu

### MS258

#### Solving the Vlasov-Maxwell System in the Classical Limit

In order to model non-thermalized plasmas in the fully magnetized and relativistic regime, the relativistic Vlasov-Maxwell (RVM) system is most commonly used. This model includes a dimensionless parameter  $c$  that determines the relative strength of relativistic effects. For  $c$  close to unity a number of efficient numerical methods have been

developed (for example, the VALIS algorithm described in [N. J. Sircombe, T.D. Arber, J. Comput. Phys. 228, pp. 4773-4788 (2009)] and the Hamiltonian splitting in [N. Crouseilles, L. Einkemmer, E. Faou, J. Comput. Phys. 283, pp.224-240 (2015)]). In some applications, however, the weakly relativistic regime (i.e. where  $c$  is large) is of interest. In this case numerical simulations show that the methods that have been developed for the RVM system are extremely inefficient. In addition, many numerical methods are not even able to recover the correct limit behavior (i.e. convergence to the Vlasov-Poisson system). In this presentation we describe the asymptotic preserving scheme introduced in [N. Crouseilles, L. Einkemmer, E. Faou, Comput. Phys. Commun., in press (2016)]. This method is robust with respect to the classic limit and imposes no step size restriction. Our approach relies on a time splitting approach for the RVM system and employs an implicit integrator for the linear part of Maxwell's equations. To illustrate the efficiency of the numerical scheme we present the results for a number of numerical simulations.

Lukas Einkemmer  
University of Innsbruck  
lukas.einkemmer@uibk.ac.at

### MS258

#### New Approach to Simulation of Elastodynamics Systems

The aim of this talk is to focus on computational modeling of coupled oscillators which are used to describe elastodynamics systems (such as fibers, textiles, flexible solids, etc). Such highly nonlinear processes usually involve a wide spectrum of spatial and temporal scales (also known as stiffness) and thus challenge numerical solution as a difficult task. A recent paper by D.L. Michels et al. "Exponential Integrators for Stiff Elastodynamic Problems, *ACM Transactions on Graphics*, Vol. 33, No. 1, Article 7, 2014] demonstrated that a Gautschi-type integrator offers significant computational savings in solving these elastodynamic problems. Inspired by this work, we propose a new formulation of the problem as a semilinear system of first-order ordinary differential equations (ODEs), where the linear part is a skew-symmetric matrix. The advantage of such formulation is the known location of the spectrum of the skew-symmetric matrix, which can allow computation of the relevant functions of this matrix in a much more efficient manner. With the help of adaptive Krylov implementations, we then integrate the resulted system of ODEs by using stiffly accurate exponential Rosenbrock schemes. This new approach is shown to be significantly more efficient than current state-of-the-art techniques over a set of numerical experiments and thus enables more complex and realistic models to be explored in computer graphics.

Vu Thai Luan  
School of Natural Sciences, University of California, Merced  
vluan@ucmerced.edu

Dominik Michels  
Department of Computer Science, Stanford University  
mail@dmichels.de

Mayya Tokman  
University of California, Merced  
School of Natural Sciences

mtokman@ucmerced.edu

### MS258

#### EPIRK-K: A Fusion of EPIRK and K-Methods Theory

This presentation is focused on the extension of EPIRK methods into the K-method framework. We provide a brief overview of the benefits of EPIRK-K schemes and the advantages of considering the approximation of matrix exponential vector products and the time integration method as a single computational process. We present the derivation of order conditions of EPIRK-K methods, discuss the implementation of EPIRK-K methods, and include some comparison of new methods with current technology.

Mahesh Narayanamurthi  
Virginia Tech, USA  
maheshnm@vt.edu

Paul Tranquilli  
Virginia Polytechnic Institute and  
State University  
ptranq@vt.edu

Ross Glandon  
Virginia Tech  
rossg42@vt.edu

Mayya Tokman  
University of California, Merced  
School of Natural Sciences  
mtokman@ucmerced.edu

Adrian Sandu  
Virginia Polytechnic Institute  
and State University  
asandu7@vt.edu

### MS259

#### High-Order Adaptive Kernel Compression Time-Stepping Methods for Fractional Differential Equations

The nonlocal nature of the fractional integral and the singularity of its kernel make the numerical treatment of fractional differential equations expensive, and the design of high-order and adaptive methods for such problems, difficult. We present a method for reducing the costs associated with the evaluation of the history term of the fractional integral of a function  $f$  while controlling the accuracy of the scheme. The approach may be interpreted as the approximation of the kernel by a sum of exponentials. Several methods based on this idea have been proposed in the past, e.g., [Li, 2010], [Beylkin and Monzón, 2010] and [López-Fernández, Palencia, and Schädle, 2006]. We adapt the method proposed in [Alpert, Greengard, and Hagstrom, 2000] to the approximation of the history term, and obtain a priori estimates on the relative error of the convolution approximation in the  $L^2$  norm. The approximation requires only local information to advance, which makes it attractive for use within time-stepping schemes. We present numerical results showing the kernel approximation directly, and results obtained with several time-stepping methods – including high-order adaptive methods.

Daniel H. Baffet, Jan S. Hesthaven

EPFL  
daniel.baffet@epfl.ch, jan.hesthaven@epfl.ch

### MS259

#### A FEniCS-HPC Framework for Simulating Diffusion in Heterogeneous Media

The Bloch-Torrey equation [H. Torrey, Prola, 1956] has been played an important role in probing the diffusion characteristics of molecules during a diffusion time in the range of ten thousands of micro seconds and at the spatial resolution of a microscopic to a macroscopic scale. The media heterogeneity is modelled by interface conditions including the membrane permeability and the jump in magnetization. The complex-valued and time-dependent pseudo-periodic conditions [J. Xu et al., Phys. Med. Biol., 2007] can be applied for exterior boundaries to mimic the phenomena that the molecules can enter and exit the computational domain. These issues cause difficulties in solving the equation efficiently and many efforts have been made to develop an efficient numerical method. We propose a framework for the multi-compartment Bloch-Torrey models based on the FEniCS-HPC platform, a branch of the FEniCS project [http://www.fenicsproject.org, 2003], for space discretization and the adaptive Runge-Kutta Chebyshev time stepping [B. P. Sommeijer et al., J. Comput. Appl. Math., 1998]. This work is a continuation of [D. V. Nguyen et al., J. Comput. Phys., 2014] but we propose here a common form for the interface and boundary conditions that allows weakly imposing the pseudo-periodic boundary conditions to avoid periodic meshes and to facilitate for parallelization. Numerical analyses show efficiency and promise in validation against more realistic models.

Dang Van Nguyen  
KTH  
vdnguyen@kth.se

Johan Jansson  
School of Computer Science and Communication  
Royal Institute of Technology KTH  
jjan@kth.se

Johan Hoffman  
Royal Institute of Technology KTH  
jhoffman@kth.se

Jing-Rebecca Li  
INRIA Saclay  
jingrebecca.li@inria.fr

### MS259

#### Fast Galerkin Bem for Parabolic Moving Boundary Problems

Time dependence in boundary integral reformulations of parabolic PDEs is reflected in the fact that the layer potentials involve integrals over time in addition to integrals over the boundary surface. For the numerical solution this implies that a time step involves the summation over space and the complete time history. Thus the naive approach has order  $N^2M^2$  complexity, where  $N$  is the number of unknowns in the spatial discretization and  $M$  is the number of time steps. With a space-time version of the fast multipole method the complexity can be reduced to nearly  $NM$ . The talk will focus on the application of the methodology to problems with moving geometries. Here the major complication is the evaluation of singular integrals of the



discrete layer potentials. We introduce a set of singularity removing transformations that enable the use of standard Gauss quadrature rules.

Johannes Tausch  
Southern Methodist University  
Department of Mathematics  
tausch@mail.smu.edu

### MS259

#### An Efficient Algorithm for Solving Fractional Diffusion Equations

The computational work and storage of numerically solving the time fractional PDEs are generally huge. To overcome this difficulty, we present an efficient algorithm for the evaluation of the Caputo fractional. The algorithm is based on an efficient sum-of-exponentials approximation for the Abel kernel. The resulting algorithm requires only  $O(N_S \log(N_T))$  storage and  $O(N_S N_T \log(N_T))$  work with  $N_S$  and  $N_T$  representing the total number of points in space and time when numerically solving the time fractional PDEs. Furthermore, we also give the stability and error analysis of the new scheme, and present several numerical examples to demonstrate the performance of our scheme.

Jiwei Zhang  
Beijing Computational Science Research Center  
jwzhang@csrc.ac.cn

### MS260

#### Massively Parallel Sph Simulations and Laboratory Experiments of Fracture Flow Dynamics: Toward the Process Understanding of Gravitational Driven Flows in Unsaturated Fractured Media

Abstract Not Available At Time Of Publication

Jannes Kordilla  
Geoscientific Centre  
University of Goettingen  
jkordil@gwdg.de

### MS260

#### SPH Approximations of Stochastic Boundary Value Problems

We will present non-local SPH approximations of stochastic boundary value problems, including stochastic Navier-Stokes and advection-diffusion equations. We will demonstrate that boundary conditions in the form of volumetric constraints are dissipative in nature and give rise to random terms with statistics determined by the Fluctuation-Dissipation Theorem.

Alexander Tartakovsky  
Pacific Northwest National Laboratory  
alexandre.tartakovsky@pnnl.gov

### MS260

#### Coupled Fluid-Solid-Fracture Using SPH

Abstract Not Available At Time Of Publication

John R. Williams  
Massachusetts Institute of Technology  
Cambridge, MA, 02139, USA

jrw@mit.edu

### MS261

#### Code as a Research Product: Open Source for Open Science

Open Science is the idea that scientific knowledge of all kinds should be shared publicly as early as is practical in the discovery process. Reproducibility of experiments is one of the foundations of science. For simple models and small data sets, calculations are reproducible in principle and in practice. As simulations become more complex and data sets become larger, calculations that are reproducible in principle are no longer reproducible in practice without access to the code and data. Reports of numerical experiments should therefore include:

- all source code needed to reproduce the calculation,
- all input data used to perform the calculation, and
- all meta-data required to allow other codes to use the input data

However, it is not (yet) standard for all scientific programs to be released with source code. Many barriers exist, but most of these are social. They include

- software is rarely recognized by citation in academic science,
- attribution metrics don't take software into account,
- institutions often don't recognize software and data curation as intellectual work, and
- academic scientists often don't build software for sustainability.

Open source is essential for reproducible, open science. There are no easy solutions to problems of Recognition, Attribution, and Sustainability, but there are some hopeful signs that these problems are being addressed.

J. Daniel Gezelter  
Notre Dame University  
gezelter@nd.edu

### MS261

#### Continuous Integration For Large-Scale Scientific Software Development

The Multiphysics Object Oriented Simulation Environment (MOOSE) framework sits at the center of a large open source community while supporting very high software quality assurance standards. MOOSE's custom-built CI tool (CIVET) is essential for enabling concurrent development within its hierarchy of applications and the framework while ensuring individual application integrity and functional multi-project integration. This is accomplished through a multi-layer testing model that tests all combinations of development and stable branches of each application to pinpoint defects through simultaneous development efforts. CIVET's distributed client model and version controlled recipes enabled proper configuration management and scalable testing and reporting as the project continues to expand. This model allows mature application developers full control over their testing strategy without inhibiting contributions from less experienced developers. Our development model and custom integration tool will be demonstrated.

Cody J. Permann  
Center for Advanced Modeling and Simulation  
Idaho National Laboratory

cody.permann@inl.gov

### MS261

#### Community Building and its Impact on Sustainable Scientific Software

In this talk, I will explore the aspects of community building most relevant to developing, releasing, sustaining and improving scientific software over its lifetime. This will include topics such as code of conduct, developing systems for communication, bridging barriers between different methods of cooperation and competition, and working to ensure that all participants are both recognized for their contributions and feel valued for their efforts. Finally, I will conclude with a few use cases of how this atmosphere can contribute to overall scientific productivity, rather than detract from the productivity of individual researchers.

Matthew J. Turk  
Columbia University  
matthewturk@gmail.com

### MS261

#### 15 Years of Trilinos: What Has Worked and Not

The Trilinos Project started about 15 years ago. A core purpose of the project has always been improving software quality practices, processes and tools as part of delivering capabilities to customers. 15 years is long enough to see the long-term impact (both positive and negative) of choices we have made over the years. In this talk we discuss observations and lessons learned from our efforts to produce reusable scientific software. We talk about activities in software development, testing, team interactions, and related. We end the presentation with our latest activities and plans for future software quality improvements.

James Willenbring, Michael Heroux  
Sandia National Laboratories  
jmwille@sandia.gov, maherou@sandia.gov

### MS262

#### Duality and Conditional Expectations in Nakajima-Mori-Zwanzig

We an operator algebraic formalism for the Nakajima-Mori-Zwanzig method of projections, in order to clarify and make rigorous the foundations of this important and widely used tool. In doing so, the method is simultaneously developed for both classical and quantum systems, and the natural duality between the two common versions of the formalism – those for “phase space functions” and for probability density functions – is made explicit and connected to the dual Heisenberg and Schrödinger pictures of quantum mechanics. We explore the projections used in the method and find that, under natural assumptions, they must be conditional expectations, in the operator algebraic sense.

Jason Dominy  
University of California Santa Cruz  
jdominy@ucsc.edu

Daniele Venturi  
Department of Applied Mathematics and Statistics  
University of California Santa Cruz

venturi@ucsc.edu

### MS262

#### Data-Driven Stochastic Parameterization of Multi-scale Dynamical Systems

Abstract Not Available At Time Of Publication

John Harlim  
Pennsylvania State University  
jharlim@psu.edu

### MS262

#### Empirical Approaches to the Mori-Zwanzig Formalism

The Mori-Zwanzig (MZ) formalism shows when we project a dynamical system onto a subset of its degrees of freedom, the resulting process generally exhibits memory. How to represent the memory is a major question in model reduction and data-driven modeling. In this talk, I will discuss some theoretical and practical issues in evaluating memory terms from data, and describe a class of simple, effective methods to resolve them. We compare our work to earlier proposals and present sample applications.

Fei Lu  
Department of Mathematics  
Lawrence Berkeley National Laboratory  
flu@lbl.gov

Fei Lu  
Department of Mathematics, UC Berkeley  
Mathematics group, Lawrence Berkeley National Laboratory  
flu@lbl.gov

Kevin K. Lin  
University of Arizona  
klin@math.arizona.edu

Alexandre Chorin  
Department of Mathematics  
University of California at Berkeley  
chorin@math.berkeley.edu

### MS262

#### When Big Computers Are Not Enough: Model Reduction and Mesh Refinement

Despite the remarkable increase in computational power, most real-world systems are still too complex to simulate in full detail. In such cases, the hope is to construct numerical algorithms which can retain the salient features of a system while reducing the complexity. Two different ways of dealing with complexity, namely mesh refinement and model reduction are based on the common concept of transfer of activity from larger scales to smaller ones. This allows the unified development of reduced models and mesh refinement schemes. Examples involving singularity detection and tracking as well as uncertainty quantification for systems exhibiting bifurcations will be presented for illustration purposes.

Panos Stinis  
Pacific Northwest National Laboratory

panagiotis.stinis@pnnl.gov

### MS263

#### Herding Cats and Elks: Big-Plans with Big-Data and HPC Administration

Abstract Not Available At Time Of Publication

Elizabeth Bautista

Lawrence Berkeley National Lab  
ejbautista@lbl.gov

### MS263

#### What the ‘Heka’ is Freeboard? Dynamic Visualizations

The amount of data in our world has been exploding, and analyzing large data sets—so-called big data—will become a key basis of predictive analytics, underpinning new waves of productivity growth, innovation, and machine-learning techniques to identify the likelihood of future outcomes based on historical data. The increasing volume and detail of information captured by NERSC, is pioneering the rise of forward-thinking leaders across sectors, fueling exponential growth in analyzing big data for the foreseeable future. But how is all this big data collected? This is where Heka comes into play. A tool for high performance data gathering, analysis, monitoring, and reporting. Hekas main component is hekad, a lightweight daemon program that can run on nearly any host machine. Until now, the torrent of data flooding our world has been a phenomenon that probably only excited a few data geeks. But we are now at an inflection point where we have a competitive advantage by aggregating and analyzing live data. Lets take a deep dive into all this big data with the tools that will get us there.

Jasmine Farrell

Lawrence Berkeley National Laboratory  
jfarrell@lbl.gov

### MS263

#### Old-School Hipsters Crash Big-Datas Party

Big Data is quite the Buzzword these days. What to keep, why to keep it and how to keep it is a big decision. We are keeping log data from system environmental, performance and weather sensors consisting of some 95 billion items in our database. This will support current and future data analysis for problem solving and improvement seeking. This huge amount ruled out pay services like Splunk, as it would be cost ineffective. We chose Elasticsearch, which is open source, scalable, and is the basis of the ELK Stack, a platform that covers logs to visualization. We include system logs for comprehensive storage as well as the ability to include their data for analysis and visualization.. This phase of the project involved migrating the syslog data from an aging system that was ready to be decommissioned. To do this we utilize Logstash to index and tag the syslog feed along with RabbitMQ to process the data stream. With custom tools using languages such as Python, we could extract and analyze the syslogs. Another advantages was providing the security team with a useful way to analyze the syslogs. As a result, many groups from across the organization now have access to the data in a centralized location that can be archived for future use. Also, data from one system can be cross referenced against data from many others. Ultimately, the data collected can

be used as a basis for machine learning to improve alerts and automate many of the tasks currently performed by staff.

Mark Sing

Lawrence Berkeley National Lab  
msing@lbl.gov

### MS263

#### Laser-Swords into Data-Shares: Game Engines for Interactive Visualization

One of the biggest visualization hurdles is to give an audience digestible, contextual information. This is usually seen as a design problem, but its also a relational problem. Who will the viewer(s) be, what labels and data will make sense to them, and how much context is required to highlight the critical insights? Traditional visualizations require the presenters to know (or guess!) the answers to these questions ahead of time in order to design in relation to the audience. This can require time-consuming customization, or create an unappealing level of dense material in an attempt to be both accommodating and clear. With the increase in free and low-cost game engines, the field is focused on ease of access like never before, and can fit any budget. With interactive access, the audience can explore a wealth of material on an individual level. Now, presentations can pack in as much data as resources allow. The right interface can manage context for a user, empowering them to navigate to the right level of data. This can allow peers and experts to dive into specific details, while still giving other colleagues and novices enough general understanding to ask questions and support further work. We will briefly explore things to consider in selecting a game engine right for you and your project, some basics of interactive design, and a demonstration of a project being developed at NERSC that uses a game engine to display live data from HPC systems.

Eric Zimmerman

Lawrence Berkeley National Lab  
zim@lbl.gov

### MS264

#### Uncertainty Quantification with Dakota and Nasa’s Overflow2 CFD Solver within an in Situ Based Workflow

NASA has called for the integration of VV/UQ tools into engineering workflows as a vital component of its CFD Vision 2030 Study. In that document, the authors assert that as exascale computing becomes available, the capability to perform many highly resolved CFD simulations will become ubiquitous, enabling the ability, and need, to perform UQ and sensitivity analysis. To address this need, this presentation summarizes current work whereby the Dakota code drives an ensemble of OVERFLOW2 jobs to explore a nested epistemic-aleatory parameter space. Concurrently, OVERFLOW2 was instrumented with VisIt/Libsim to export surface extracts (XDB) via in situ processing to reduce the data for post-hoc visualization and analysis. The presentation presents UQ concepts in general and their implementation within the current framework.

Brad Whitlock

Lawrence Livermore National Laboratory  
whitlock2@llnl.gov

Earl P. N. Duque

Intelligent Light  
epd@ilight.com

#### MS264

### An Open-Source Framework for Multi-Physics Simulations and Design Optimization

Open-source technology has become mature enough as to become highly attractive for industrial applications and this is due to its availability, reliability, flexibility, scalability and relative low cost. In this work, we present a simulation framework entirely based on open-source technology and able to run in high performance computing environments. For geometry parametrization we use OpenVSP and SALOME, for meshing and mesh smoothing we use OpenFOAM, SALOME and MESQUITE, as a multi-physics solver we use OpenFOAM, the parametrization and optimization is done using DAKOTA, while for data analytics we use Python. Hereafter, we present several optimization examples in order to demonstrate the capability of the proposed framework in performing complex design exploration and optimization studies. We start with an academic test case where we briefly demonstrate the capabilities of the optimizer. Then we move towards practical applications, where we couple the optimizer with the multi-physics solver and the geometry generation and mesh generation tools. As an advanced proof-of-the-concept test case, we present an optimization study of a sailing yacht daggerboard, where we conduct a surrogate-based optimization and a multi-objective optimization study using 12 design variables, 2 objective functions and 1 non-linear constraint.

Joel Guerrero

Department of Environmental Engineering  
University of Genoa  
joegi.geo@yahoo.com

#### MS264

### Beyond Parameter Scans, Using Dakota for Real World Problems

The ALEGRA code developed by Sandia National Laboratories is used for applications including armor mechanics, component design and evaluation, and other shock and multiphysics applications. Users of this code have leveraged Dakota's algorithms for sensitivity analysis, calibration, and uncertainty quantification (UQ) in both a traditional black-box mode and an embedded mode with Dakota linked as a third party library for ease of use. Examples of three application areas will be discussed. For 3D exploding wire simulations, resistive MHD physics requires the use of a multigrid solver with many problem-dependent solver parameters, some of which can drastically affect computational performance. By optimizing six parameters using insight from Dakota parameter scans, performance improved by a factor of two. In another application, we want to know whether the results obtained in a simulation are attributable to uncertainties in the equation of state (EOS) and conductivity model. Recent work at Sandia has provided a self-consistent EOS/electrical conductivity table for aluminum which is delivered to the user as a random field with a compact uncertain representation. This allows the computation of associated probability density functions of output quantities of interest using Dakota. A final example will discuss the usefulness of Dakota algorithms in component evaluation.

Sharon Petney, John H. Carpenter

Sandia National Laboratories  
svpetne@sandia.gov, jhcarpe@sandia.gov

Robert L. Doney, III

Army Research Laboratory  
robert.l.doney4.civ@mail.mil

#### MS264

### Dakota: Algorithms for Design Exploration and Simulation Credibility

The Dakota toolkit provides a flexible interface between simulations and iterative analysis methods including optimization, uncertainty quantification, parameter estimation, and sensitivity analysis. Methods may be used individually or as components within advanced strategies to address questions like "What is the best design?", "How safe is it?", and "How much confidence do I have in my answer?". In this talk, we present an overview of Dakota history, capabilities, and future direction. A series of use cases also will be presented to illustrate how Dakota can inform engineering decisions.

Adam Stephens

Sandia National Laboratories  
jasteph@sandia.gov

#### MS265

### Comparison of BLR and HSS Low-Rank Formats in Multifrontal Solvers: Theory and Practice

Matrices coming from elliptic Partial Differential Equations have been shown to have a low-rank property: conveniently defined off-diagonal blocks of their Schur complements can be approximated by low-rank products and this property can be efficiently exploited in multifrontal solvers to provide a substantial reduction of their complexity. Several matrix representations have been proposed to exploit this property within direct multifrontal solvers. The simplicity and flexibility of the Block Low-Rank (BLR) format make it easy to use in a general purpose, algebraic multifrontal solver. While its theoretical complexity has been shown to be significantly lower than in full-rank, it is higher than the complexity achieved by hierarchical formats, a more complex class of low-rank matrix representations. Among the hierarchical formats, the Hierarchically Semi-Separable (HSS) one has been used to design the most advanced to date low-rank multifrontal solvers. In this talk, we compare the BLR and HSS formats, both from a theoretical and experimental standpoint. In our numerical experiments, we use the BLR-based MUMPS and HSS-based STRUMPACK multifrontal solvers on a large set of matrices coming from a variety of real-life applications.

Patrick Amestoy  
INP(ENSEEIH)-IRIT  
amestoy@enseeiht.fr

Alfredo Buttari  
CNRS-IRIT-Université de Toulouse, France  
alfredo.buttari@enseeiht.fr

Pieter Ghysels  
Lawrence Berkeley National Laboratory  
Computational Research Division  
pghysels@lbl.gov

Jean-Yves L'Excellent  
INRIA-LIP-ENS Lyon  
jean-yves.l.excellent@ens-lyon.fr

Xiaoye Sherry Li  
Computational Research Division  
Lawrence Berkeley National Laboratory  
xsli@lbl.gov

Theo Mary  
Universite de Toulouse, UPS-IRIT  
theo.mary@enseiht.fr

François-Henry Rouet  
Lawrence Berkeley National Laboratory  
frouet@lbl.gov

### MS265

#### On the Solution Phase of Sparse Direct Solvers with Many Right-Hand Sides

The cost of the solution phase of sparse direct solvers is sometimes critical. It can be larger than the cost of the factorization in applications where systems of linear equations with thousands of right-hand sides must be solved. Depending on the applications, the right-hand sides may be known all at once or may depend on previous solutions. In the context of multifrontal solvers, we consider some important characteristics of the solve phase in terms of parallelism, and present algorithms to help reduce its computational cost. By considering the critical path in assembly trees associated with 2D and 3D problems, we first show that the complexity of the solve phase offers interesting properties in terms of parallelism, compared to the factorization, both in the case of full-rank and in the case of low-rank factorizations. We then consider the case of multiple sparse right hand sides on distributed-memory systems and propose new algorithms that exploit the sparsity of their structure to reduce the computational cost (numbers of operations) while still exploiting tree parallelism efficiently. We finally give experimental results using the MUMPS solver.

Gilles Moreau  
ENS Lyon  
gilles.moreau@ens-lyon.fr

Jean-Yves L'Excellent  
INRIA-LIP-ENS Lyon  
jean-yves.l.excellent@ens-lyon.fr

Patrick Amestoy  
INP(ENSEEIH)-IRIT  
amestoy@enseiht.fr

### MS265

#### Scalable Sparse Direct Solver for Hybrid Architectures

We present new algorithmic techniques for improving the scalability and parallel efficiency of sparse direct solvers, specifically for distributed memory systems comprised of hybrid multicore CPU systems and hardware co-processors. This work extends the algorithm in SuperLU\_Dist, which is both right-looking and statically pivoted. To improve the scalability, we propose new algorithmic techniques to reduce the effects of relatively high communication cost at large number of processing cores.

First, we describe a 3-dimensional (3D) sparse LU factorization algorithm that reduces per process communication volume, albeit at expense of some additional storage and computations. We call it 3D algorithm as it uses a 3D logical arrangement of processors, in contrast to existing SuperLU\_Dist that uses 2D process grid. Secondly, we aggressively hide the communication cost by exploiting the elimination tree parallelism, a new scheduling strategy and efficient hybrid OpenMP+MPI programming; all of which serves to either reduce or hide the intra- and inter-node communication. Finally, we present some preliminary results to show the efficacy of the proposed techniques.

Piyush Sao  
Georgia Institute of Technology  
psao3@mail.gatech.edu

Xiaoye S. Li  
Lawrence Berkeley National Laboratory  
xsli@lbl.gov

Richard Vuduc  
Georgia Institute of Technology  
richie@cc.gatech.edu

### MS265

#### Auto-Tuning for Multifrontal Sparse QR Factorizations on GPU

Efficient dense matrix computation kernels are critical components for achieving high performance in sparse direct methods. However, the existence of many different GPU architectures with different characteristics and features makes it challenging to write a single code that targets them all. Moreover, the approach of writing optimized kernels targeting a specific GPU architecture to leverage all of its potential requires too much effort in practice to be a viable solution. Hence, in the context of the multifrontal QR method on GPUs, we solve this issue by relying on auto-tuning techniques. Our design consists of two levels of kernels. At the low level of micro-kernels, we express the computation algorithms in a completely parametrized way. Then, an auto-tuner is used to traverse the parameters search space in order to find the best combination of parameters for any targeted GPU architecture. At the high level of macro-kernels, we rely exclusively on the micro-kernels along with batched BLAS kernels and only express the dependency between data and low level kernels calls in a very similar way to sequential CPU codes.

Wissam M. Sid-Lakhdar  
Texas A&M  
Computer Science and Engineering  
wissam@tamu.edu

Tim Davis  
Texas A&M University  
Computer Science and Engineering  
davis@tamu.edu

### MS266

#### Hierarchically Compositional Kernels for Regression and Parameter Estimation

We present a novel class of kernels to alleviate the high computational cost of large-scale nonparametric learning with kernel methods and Gaussian processes. The proposed kernel marries the Nyström approximation with a block-diagonal approximation in a hierarchical manner

and maintains positive definiteness. With  $n$  training data points and  $r$  being the sample size in Nyström, the memory and arithmetic costs for computing the Gaussian log-likelihood and for regression are approximately  $O(nr)$  and  $O(nr^2)$ , respectively. We demonstrate the effective use of the kernels for regression and parameter estimation.

Jie Chen

IBM Thomas J. Watson Research Center  
chenjie@us.ibm.com

### MS266

#### Iterated Brownian Bridge Kernels for Sequential Kriging Optimization

Sequential kriging optimization, also referred to as Gaussian process-backed Bayesian optimization, is a process whereby a kriging (Gaussian process) model is fit to existing data to facilitate an optimal addition to the experimental design with the intent of learning the global optimum of a given black-box function. An important component of any kriging model is the choice of covariance kernel, as it defines the mechanism by which we can interpolate previously observed results into predictions at unobserved locations. In the context of sequential kriging optimization, where the goal is to execute an efficient balance of exploration and exploitation, using kernels which satisfy desired boundary conditions can provide significant benefit by preventing exploration near the boundary. We present the results of numerical experimentation with iterated Brownian bridge kernels applied to functions with their optimum in various locations relative to the boundary; we identify the circumstances under which these boundary condition satisfying kernel compare favorably to other standard kernels in the literature.

Greg Fasshauer

Illinois Institute of Technology  
fasshauer@mines.edu

### MS266

#### Knowledge Gradient for Multi-Information Source Optimization

We consider Bayesian optimization of an expensive-to-evaluate black-box objective function, where we also have access to cheaper approximations of the objective, and these approximations may have varying unknown bias. Such applications arise, for example, when optimizing an engineering system that can be simulated with less computational expense using physics with simplifying assumptions, or by using a coarse mesh when solving a PDE. We present a novel algorithm that provides a rigorous mathematical treatment of the uncertainties arising from varying unknown bias (model discrepancy) and noisy observations. Its optimization decisions rely on a stringent value of information analysis, using the Knowledge Gradient method, and maximize the predicted benefit per unit cost. We conduct an experimental evaluation that demonstrates that the method consistently outperforms other state-of-the-art techniques: it finds designs of considerably higher objective value and additionally inflicts less cost in the exploration process.

Peter I. Frazier

Dept. of Operations Research and Information Engineering  
Cornell University  
pf98@cornell.edu

Jialei Wang

Cornell University  
jw865@cornell.edu

### MS266

#### Failure Region Identification Coincident to Bayesian Optimization

While conducting a Bayesian optimization, it is possible that the function to be optimized may fail: for example, certain choices of hyperparameters could cause a machine learning algorithm to time out or encounter insufficient memory. For many applications, the nature of this so-called failure region, where inputs to the function yield a failed function evaluation, may be unknown a priori. There are heuristic strategies in software and literature for dealing with this failure region as a nuisance to the optimization, but no strategies exist for developing an understanding of the failure region while simultaneously optimizing the function of interest. In this talk, we will go over some techniques for creating a positive definite kernel representation of the failure region coincident to the kernel model underlying the Bayesian optimization. We also consider the implications of trying to modify the Bayesian optimization to supplement the acquisition function with some understanding of the accuracy of the failure region.

Michael McCourt

Department of Mathematical and Statistical Sciences  
University of Colorado Denver  
mccourt@sigopt.com

Ling Leevan

Hong Kong Baptist University  
Department of Mathematics  
lling@hkbu.edu.hk

### MS267

#### Including First-Shell Response in Continuum Solvent Models Using Multiscale Boundary Conditions

Many problems in molecular science and engineering involve the behavior of a solute molecule in a solvent fluid, often water with dissolved ions. Atomically detailed simulations of solute-solvent interactions, for instance molecular dynamics, can achieve reasonable accuracy and offer precise molecular-scale insights, but at tremendous computational cost. Continuum theories based on macroscopic dielectric theory are orders of magnitude faster, but unable to reproduce essential phenomena such as temperature dependence and asymmetric response to positive and negative atomic charges. These confounding phenomena arise largely due to the solvent molecules in the first "solvation layer" at the interface, whose non-bulk-like behavior is not captured in macroscopically based continuum models. We have improved the continuum model's accuracy by replacing the macroscopic interface condition with a multiscale solvation-layer interface condition (SLIC). Predictions offer improved accuracy while using fewer parameters than competing theories, and our new SLIC model also is the first continuum theory capable of predicting temperature dependent effects as well as the behavior of solvent mixtures.

Jaydeep P. Bardhan, Amirhossein Molavi Tabrizi, Ali Mehdizadeh Rahimi, Spencer Goossens  
Northeastern University  
j.bardhan@neu.edu, a.molavitabrizi@neu.edu,

mehdizadehrahimi.a@husky.neu.edu,  
goossens.s@husky.neu.edu

Matthew G. Knepley  
Rice University  
knepley@gmail.com \*Preferred

### MS267

#### Evaporation-Induced Nanoparticle Assembly

A critical challenge in the field of polymer nanocomposites is to control the dispersion of nanoparticles in a polymer matrix. We employ large-scale molecular dynamics simulations to study the assembly of nanoparticles as the solvent evaporates from a polymer solution containing nanoparticles. Results show that we can control the organization of nanoparticles by varying the strength of the polymer-nanoparticle interactions. When the nanoparticles and polymers strongly attract, as the solvent evaporates, a concentrated polymer film forms at the surface and entraps a layer of nanoparticles, which assemble into a close-packed hexagonal lattice. If the nanoparticle-polymer interactions are weak, then as the solvent evaporates, the surface layer is almost entirely made of polymers. The nanoparticles are largely excluded from the surface and dispersed randomly in the region below the surface layer. Also of interest is the case of a nanoparticle solution in contact with a planar polymer brush. For a relatively weak nanoparticle-brush attraction, after evaporation of the solvent the nanoparticles straddle the brush surface and form an ordered lattice. For a strong nanoparticle-polymer attraction, however, the nanoparticles are engulfed inside the brush and the packing quality diminishes because the lateral diffusion of the nanoparticles is suppressed. Our recent simulations of evaporation-induced stratification phenomena in a binary nanoparticle system will also be discussed.

Shengfeng Cheng  
Department of Physics  
Virginia Tech  
chengsf@vt.edu

### MS267

#### Recent Advances in Atomistic Simulations of Biomass Dissolution

Efficient, economical, and environmentally friendly dissolution of lignocellulosic biomass is being actively pursued as a replacement for current petroleum-derived feedstocks. Modeling the dissolution process is challenging because of the large size and complex structure of the solute, and the significant changes in dissolution that result from relatively minor changes in the structure of the solvent. Thus, a wide range of approaches have been developed to model the response of biomass to different solvent systems and processing conditions with the goals of replicating experimental results and explaining their behavior at a fundamental molecular level. In this talk, we will present an overview of currently used approaches and recent results in the field, as well as discuss some of the remaining challenges to be overcome in understanding these complex phenomena.

Ahmed E. Ismail  
Department of Mechanical Engineering  
RWTH Aachen University  
ahmed.ismail@mail.wvu.edu

Brooks Rabideau  
Department of Chemical Engineering

University of South Alabama  
brabideau@southalabama.edu

Pan Chen  
Wallenberg Wood Science Center  
KTH Royal Institute of Technology  
evan.pan.chen@gmail.com

Yoshiharu Nishiyama  
CERMAV  
University Grenoble Alpes  
yoshi@cermav.cnrs.fr

### MS267

#### Towards a Molecular Level Understanding Hydration Lubrication: Structure and Dynamics of 2-Methacryloyloxyethyl Phosphorylcholine

We explore the hydration structure and dynamics of poly(2-methacryloyloxyethyl phosphorylcholine)50, or pMPC50, in solution using molecular dynamics simulations with an emphasis on isolating the contributions of the key chemical moieties. Surface bound films of this biocompatible polymer have been shown experimentally to produce tribological properties that surpass those of the human synovial joint [Chen et al., *Science*, 2009] and are being developed for use in artificial joints [Kyomoto et al., *Biomaterials*, 2012]. The mechanism by which such materials are thought to provide ultra-low friction coefficients has been termed hydration lubrication [Klein et al., *Friction*, 2013]. Here, we characterize the molecular level structure and dynamics of this hydration water around pMPC molecules that are experimentally grafted onto substrates for biomimetic lubrication. Previous results [Roussel et al., *Young Scientist*, 2016] showed that the charged choline groups of individual monomers tend to fold onto neighboring monomers yielding solvent facing, highly hydrated phosphoryl groups. Here, we further analyze the effect of the two outer chemical moieties, the choline and phosphoryl groups, by removing them and comparing the behavior of pMPC to poly(2-methacryloyloxyethyl phosphate) and poly(ethyl methacrylate). Preliminary results strengthen the previous hypothesis [4] that the choline moiety is responsible for stabilizing the polymer.

Christoph Klein  
Department of Chemical and Biomolecular Engineering  
Vanderbilt University  
christoph.klein@vanderbilt.edu

William Roussel, Christopher Iacovella, Clare McCabe  
Vanderbilt University  
leviroussel@gmail.com,  
christopher.r.iacovella@vanderbilt.edu,  
c.mccabe@vanderbilt.edu

Peter Cummings  
Oak Ridge National Laboratory and Vanderbilt University  
peter.cummings@vanderbilt.edu

### MS268

#### Randomized Matrix Decompositions for High-Dimensional Dynamic Data

The emergence of massive data-sets driven by new sensory technologies poses a computational challenge for obtaining traditional matrix decomposition. Interestingly though,

many high-dimensional data have low intrinsic rank relative to the dimension of the ambient measurement space. We present an overview of modern randomized matrix algorithms able to ease the computational challenges substantially. The key idea of these algorithms is to utilize randomization as a computational resource to construct a smaller (compressed) matrix, which approximates a high-dimensional data matrix well. Randomized algorithms are robust, and have the ability to fully exploit modern computational architectures, i.e., parallelized and distributed computing. Specifically, we present concepts and software to compute the low-rank singular value decomposition, and dynamic mode decomposition of matrices of various sizes. We also present distributed memory codes suitable for very large matrices via novel block matrix techniques. All the developed software we present is open source and freely available for use.

Benjamin Erichson  
University of St Andrews  
nbe@st-andrews.ac.uk

Sergey Voronin  
Tufts University  
Sergey.Voronin@tufts.edu

#### MS268

##### **Model Inference in Epidemiology and Biological Networks Using Sparse Identification of Nonlinear Dynamics**

Inferring the structure and dynamics of biological networks is critical to understanding and controlling their functionality. We present implicit-SINDy, a data-driven method for inferring nonlinear dynamical systems which combines a compact feature library, implicit formulation to enable discovery of rational functions, and sparsity promoting non-convex optimization. We integrate information theoretic approaches with implicit-SINDy to identify the effects of measurement error, missing variables, incomplete feature libraries, and insufficient data. We apply this rapid data-driven method to metabolic, regulatory, and epidemiological data and generate parsimonious dynamic ODE representations of each system. In contrast to other algorithms, the models constructed using implicit-SINDy are biologically interpretable in terms of the same network motifs that would be used in bottom-up mechanistic modeling frameworks. Thus, these rapidly constructed models provide mechanistic understanding that could be leveraged for therapeutic gene modulation for treatment of genetic diseases, metabolic engineering to design novel biochemical pathways and products, and intervention in infectious disease networks such as outbreak detection and response.

Niall M. Mangan  
Department of Applied Mathematics  
University of Washington, Seattle  
niallmm@uw.edu

#### MS268

##### **Data-Based Extraction of Modal Interaction Networks**

Quantifying nonlinear interactions in unsteady fluid flows remains a challenge due to their complexity and high-dimensional physics. In this study, a novel data-based approach for accessing the nonlinear interactions in unsteady fluid flows is developed using a network-theoretic perspective. The transfer of kinetic energy among coherent flow

structures is examined by uncovering a complex network of spatial modes. The dynamics on such a network can be analyzed with a network oscillator model, where phase and amplitude knowledge can describe the energy transfers and phase dynamics among the modes. In particular, we extend the networked Stuart-Landau oscillator model to capture interactions involved in unsteady fluid flows. Leveraging the understanding of nonlinear interactions, we discuss our efforts to reveal synchronization phenomena with the networked fluid flow models. Canonical examples of two-dimensional bluff body flows are considered to validate the modeling approach. The present technique captures fundamental nonlinear interactions and energy pathways in unsteady fluid flows, and offers a promising approach for designing flow control strategies.

Aditya G. Nair, Kunihiko Taira  
Florida State University  
Mechanical Engineering  
agn13@my.fsu.edu, ktaira@fsu.edu

Steven Brunton  
University of Washington  
sbrunton@uw.edu

#### MS268

##### **Nonlinear System Identification for Complex Systems**

The increasing ubiquity of complex systems that require control is a challenge for existing methodologies in characterization and controller design when the system is high-dimensional, nonlinear, and without physics-based governing equations. Equation-free methods such as the recently developed Dynamic Mode Decomposition and the resurgence of Koopman operator theory can potentially mitigate these issues. We illustrate that this equation-free approach that is being applied to measurement data from complex systems can be extended to include inputs and control. Examples are presented that extend current equation-free architectures toward the characterization and control of complex systems. These examples motivate a potentially revolutionary shift in the characterization of complex systems and subsequent design of objective-based controllers for data-driven models.

Joshua L. Proctor  
Institute for Disease Modeling  
JoshLProctor@gmail.com

#### MS269

##### **A Case for Cross-Cutting Ideas for Computational Climate Science**

A workshop to discuss bold, new computational ideas to address longer-term science needs for climate modeling was held on September 12-13, 2016 in Rockville, MD, USA. As the climate science community explores a host of critical science questions to be answered over the next 10+ years, there is growing recognition of the expanding requirements for multiscale, global, coupled Earth system models. These models need to provide much more detail and fidelity, with a much better understanding of their uncertainties, while still executing robustly and efficiently on ever larger and more complex computing systems. The DOE Office of Science BER and ASCR programs sponsored the event, titled "Advancing Cross-cutting Ideas for Computational Climate Science", or AXICCS 2016. It was a grassroots effort to generate fresh thinking that we hope



will mold the optimal science directions and motivate the latest and as yet to be uncovered computational science developments that will be required to handle the wealth of new requirements. We outline the science challenges, strategies to address them, and possible avenues for execution using new ideas in computational science brought up at the workshop. All attendees provided "ideas" papers, which helped form the content we present.

Katherine J. Evans  
Oak Ridge National Laboratory  
evanskj@ornl.gov

Esmond G. Ng  
Lawrence Berkeley National Laboratory  
eng@lbl.gov

**MS269**  
**Quantifying Uncertainties in Climate Science**

Projected changes in climate will potentially have a large impact on society through increasing weather extremes, a rapid rise in sea level, and diminished water availability. The talk will summarize community input concerning the challenges and promising directions that could be taken for quantifying projection uncertainties of these quantities. Climate presents a challenging problem in extrapolation for which there is limited data and understanding of how modeling errors affect predictions. A much more concerted effort will need to be made to develop computational and mathematical strategies that are tailored to the characteristics of each of these science targets.

Charles Jackson  
UTIG  
University of Texas at Austin  
charles@ig.utexas.edu

**MS269**  
**A Multiscale Integrated Modeling Framework for Climate Science**

Climate change is a complex problem that intersects multiple disciplines and processes across a wide range of spatial and temporal scales. Traditionally different modeling systems and approaches have been used in climate prediction, climate mitigation, and impacts, adaptation and vulnerability research. However to inform decisions for the next decades, new modeling frameworks must recognize the co-evolution of the Earth and human systems or interactions between climate, mitigation, and adaptation. Multiscale integrated modeling frameworks consisting of Earth System Models, Integrated Assessment Models, and Sectoral Models are being developed to answer a wide range of climate science questions. This presentation will highlight examples of modeling approaches and numerical experiments to advance understanding and prediction of climate change and its impacts.

Ruby Leung  
Pacific Northwest National Laboratory  
ruby.leung@pnnl.gov

**MS269**  
**Applied Mathematics and Computational Opportunities for Climate Science**

We summarize open challenges in climate science and op-

portunities for applied mathematics research as identified by a recent community-driven workshop.

Stefan Wild  
Argonne National Laboratory  
wild@mcs.anl.gov

George Ostrouchov  
Computer Science and Mathematics Division  
Oak Ridge National Laboratory  
ostrouchovg@ornl.gov

Daniel Martin  
Lawrence Berkeley National Laboratory  
dfmartin@lbl.gov

Ray S. Tuminaro  
Sandia National Laboratories  
Computational Mathematics and Algorithms  
rstumin@sandia.gov

Sam Williams  
Lawrence Berkeley National Lab  
swwilliams@lbl.gov

Kerstin Kleese Van Dam  
Brookhaven National Laboratory  
kleese@bnl.gov

**MS270**  
**Algorithms for 3D-3D Registration with Known and Unknown References: Applications to Materials Science**

Point Set Registration in three dimensions is becoming increasingly relevant as a means of extracting structure at the atomic level from data sets containing tens of millions of atomic coordinates from experiment and simulation. Two challenges are investigated, which involve extracting a distribution of structures when the underlying reference, or average structure, is known (as possible from x-ray diffraction for crystalline materials) or unknown (as in the case of highly disordered or amorphous materials). The most complete experimental data sets come from atomic probe tomography (APT), but such data sets contain significant sources of error, including missing atoms and uncertainty in the coordinates. We describe efforts to account for this noise. We present two examples: obtaining the atomic-level structure from experimental APT data of high entropy alloys and from simulated data of lithium-ions in battery anodes composed of hierarchically structured carbon composites.

David J. Keffer, Nicholas McNutt  
University of Tennessee  
dkeffer@utk.edu, nmcnutt@utk.edu

**MS270**  
**Deviation from High-Entropy Configurations in the  $Al_{1.3}CoCrCuFeNi$  Alloy**

The alloy-design strategy of combining multiple elements in near-equimolar ratios has shown great potential for producing exceptional engineering materials, often known as "high-entropy alloys". Understanding the elemental distributions, and, thus, the evolution of the configurational entropy during solidification, are the goal of the present research. The case of the  $Al_{1.3}CoCrCuFeNi$  model al-

loy is examined, using integrated theoretical and experimental techniques, such as *ab initio* molecular-dynamics simulations, in-situ neutron levitation and scattering experiments, synchrotron X-ray diffraction, high-resolution electron microscopy, and atom-probe tomography. It is shown that even when the material undergoes elemental segregation, precipitation, chemical ordering, and spinodal decomposition, a significant amount of disorder remains, due to the distributions of multiple elements in the major phases. The results suggest that the high-entropy-alloy-design strategy may be used to develop a wide range of complex materials, which are not limited to single-phase solid solutions. The integrated experimental and theoretical techniques, discussed here, are particularly well-suited to studying partially-ordered materials, produced using the high-entropy-alloy design strategy.

Peter Liaw  
University of Tennessee  
pliaw@utk.edu

#### MS270

##### **Implementation Techniques for Point Set Registration on Multicore and Heterogeneous Hardware**

The classic image processing and computer vision problem of 3D registration attempts to find the transformation that establishes correspondence between two sets of points in three dimensions. The standard test problems check the results quality and the performance of the proposed algorithms involve image scenes with up to tens of thousands of points and rarely ever would they strain a basic/reference implementation, even the ones that were done in a high level language. When it comes to image data coming from large scientific experiments, specifically Atom Probe Tomography, the computational demand increases drastically, which hinders data analysis in material science. We present our experimental techniques and optimizations that make the 3D registration viable for large scientific data sets.

Piotr Luszczyk  
University of Tennessee  
luszczyk@icl.utk.edu

#### MS270

##### **Advanced Particle Filters for Stochastic Point Set Registration**

Point Set Registration algorithms responsible for recovering the atomic structure from a defective data obtained by Atomic Probe Tomography (APT), typically, compute pertinent transformation parameters by minimizing an appropriate metric. However, sparsity and noise may cause the optimization module of the registration algorithm to get attracted in basins of local minima. An accurate representation of the underlying structure of the experimental data may not be available. In this talk, we adopt a Bayesian perspective such that the collected data clouds influence any prior (distributional) assumptions used for the synthetic data. Consequently, the problem of registration can be viewed as a posterior estimation problem. The objective is to compute the associated posterior conditional distribution of the states of some hidden process relying on these sparse noisy observations. The a posteriori estimation is a rather formidable task due to sparse information and we execute it with advanced particle filters, which employ a Markov Chain Monte Carlo procedure. The advanced particle filter aids the registration algorithm to pick after every iteration the solution that seems more promising in

an optimization sense.

Vasileios Maroulas  
University of Tennessee, Knoxville  
maroulas@math.utk.edu

#### MS271

##### **Uncertainty Quantification for Electronic-Structure-Based Chemical Kinetics Modeling**

In this talk we discuss recent work on modeling Density Functional Theory (DFT) data and error uncertainty in complex, high-dimensional chemical reaction networks. We present related mathematical tools for sensitivity analysis and uncertainty quantification, capable to handle a large number of model parameters and high-dimensional state spaces. We also discuss the impact of electronic structure-induced parameter correlations and uncertainties on model predictions.

Markos A. Katsoulakis  
University of Massachusetts, Amherst  
Dept of Mathematics and Statistics  
markos@math.umass.edu

#### MS271

##### **Quantifying Probability of Failure Based on Noisy Multi-Fidelity Data**

For more than two centuries, solutions of differential equations have been obtained either analytically or numerically based on typically well-behaved forcing and boundary conditions for well-posed problems. We are changing this paradigm in a fundamental way by establishing an interface between probabilistic machine learning and differential equations. We develop data-driven algorithms for general linear equations using Gaussian process priors tailored to the corresponding integro-differential operators. The only observables are scarce noisy multi-fidelity data for the forcing and solution that are not required to reside on the domain boundary. The resulting predictive posterior distributions quantify uncertainty and naturally lead to adaptive solution refinement via active learning. This general framework circumvents the tyranny of numerical discretization as well as the consistency and stability issues of time-integration, and is scalable to high-dimensions. Here, we provide an overview of the proposed framework, and we show how these new algorithms can be applied for the efficient computation of probabilities of failure in classical solid mechanics.

Maziar Raissi  
Brown University  
maziar\_raissi@brown.edu

Paris Perdikaris  
Massachusetts Institute of Technology  
parisp@mit.edu

George Em Karniadakis  
Brown University  
george\_karniadakis@brown.edu

#### MS271

##### **Predicting Fracture Patterns in Simulations of Brittle Materials under Variable Load and Mate-**

**rial**

Understanding material fracture and crack propagation is a major challenge in predicting critical failure of structures and machine components. Computer simulations using advanced material models, e.g., peridynamics, have been able to predict the outcome of carefully controlled experiments, however, operational conditions and material properties in the field can have significant deviation from those in the controlled laboratory environment. Reliable predictions of crack formation and propagation require rigorous uncertainty quantification (UQ), in this work, we specifically focus on study of crack branching and fracture patterns in two dimensional brittle material with varying load force and material strength. In order to address this problem, we have to face the two challenges that come from the high computational cost associated with simulations of crack propagation and the lack of UQ techniques that are applicable to problems with discontinuous behavior. To this end, we employ a novel UQ approach that is based on surrogate modeling and reconstruction of multidimensional response surfaces using reduced basis projection and optimization techniques.

Miroslav Stoyanov  
Oak Ridge National Lab  
mkstoyanov@gmail.com

**MS271****Accelerated Sampling and Sensitivity Estimation of Multiscale Reaction Networks**

Sampling of the equilibrium distribution of a multiscale stochastic dynamical system is usually computationally prohibitive due to the very long sojourn time in the metastable regions. Based on the parallel replica method, we propose an efficient algorithm for simulating continuous time Markov chains (CTMC) in the presence of metastability. Unlike most accelerated algorithms in computational chemistry, the algorithm we propose does not require the detailed balance assumption on the underlying CTMC. Furthermore, it can be shown that the theoretical speedup is linear in terms of the number of parallel processes used for simulation. Finally, the algorithm can be easily combined with the path-space information bounds for sensitivity estimation. We present numerical example from the stochastic chemical kinetics to demonstrate the consistency of the algorithm. Results from joint work with D. Aristoff on error analysis will be also presented.

Ting Wang, Petr Plechac  
University of Delaware  
tingw@udel.edu, plechac@udel.edu

**MS272****Bayesian Inversion for Subsurface Properties from Poroelastic Forward Models and Surface Deformation Data**

In this work we solve the inverse problem to characterize heterogeneous subsurface parameters using subsidence and fluid pressure observational data. We use the infinite-dimensional Bayesian inversion theory to invert for the permeability field of the domain. We use MCMC sampling methods to characterize the posterior distribution, which represents our confidence in the calibrated parameters. Upon discretizing the infinite dimensional permeability field, we end up with as many parameters as there are mesh nodes, which makes the inversion task challeng-

ing. Therefore, we rely on dimension-independent schemes for both finding the maximum a posteriori (MAP) point of the posterior distribution and sampling (e.g. Newton and Gauss-Newton method for optimization and Stochastic Newton MCMC with MAP-based Hessian for sampling). Our forward model is based on the linear poroelasticity theory (Biot model) and is being solved using FEM in a fully coupled approach. We experiment with different choices of FEM discretizations of 2-field and 3-field formulations of the system of forward PDEs. We verify our 2D solver against the known analytical solution for the Mandel problem and our 3D solver using a manufactured solution. We apply these techniques and tools to a real-world scenario using data that have been collected from an experiment in a municipal pumping well site.

Amal M. Alghamdi  
Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
amal@ices.utexas.edu

Marc A. Hesse  
University of Texas  
Department of Geological Sciences  
mhesse@jsg.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

**MS272****Hybrid First-Principles Data-Driven Dynamic Model Misspecification Correction**

In this study, we explore the interplay between first-principles and data-driven modeling. Since both approaches have unique advantages and disadvantages, which are complementary, it is sensible to develop hybrid modeling frameworks that can leverage the strengths of both, while avoiding their associated shortcomings. Following an overview of the proposed framework, its utility is exemplified in a Steam Assisted Gravity Drainage settings.

Lior Horesh, Theodore Van Kessel  
IBM Research  
lhoresh@us.ibm.com, tvk@us.ibm.com

Omolade Saliu  
IBM Canada  
osaliu@ca.ibm.com

Andrew R. Conn  
IBM T J Watson Research Ctr  
arconn@us.ibm.com

**MS272****Characterization of Poromechanical Behavior of Faulted Reservoirs During Injection-Induced Seismicity Using a Bayesian Inversion Approach**

Characterization of frictional properties of faults and their evolution during seismic slip is a fundamental problem in seismology (Southern California Earthquake Center's Six Fundamental Problems). Recent increase in induced earthquakes has further intensified the need for robust methods of estimating fault properties. We present a new computational framework that combines coupled multiphysics simulation of injection-induced seismicity in an aquifer-fault

system with adaptive PCE surrogate-based Bayesian approach to estimate the dynamic fault friction, critical slip distance, and aquifer permeability using aquifer pressure, ground displacements, and fault slip as measurements during injection monitoring. The model allows us to observe non-monotonic evolutions of shear traction and slip on the fault resulting from interaction among different physical mechanisms of injection-induced aquifer expansion, stress transfer, and slip-induced stress relaxation. We find that combining measurements sensitive to flow with those sensitive to deformation yields improved estimates of the uncertain parameters. We identify an inverse relationship between the dynamic friction and critical slip, which is in agreement with small dynamic friction and large critical slip reported during seismicity on mature faults. We conclude that both the coupled flow-geomechanical model and the surrogate-based inference model are key to a successful poromechanical inversion framework of induced seismicity.

#### Birendra Jha

Mork Family Dept. of Chemical Eng. and Materials Sci.  
University of Southern California  
bjha@usc.edu

Jayanth Jagalur-Mohan  
Massachusetts Institute of Technology  
jagalur@mit.edu

Zheng Wang  
Massachusetts Institute of Technology  
USA  
zheng\_w@mit.edu

Ruben Juanes  
MIT  
Civil and Environmental Engineering  
juanes@mit.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

#### **MS272**

##### **Hydraulic Inverse Modeling Using Total-Variation Regularization with Relaxed Variable-Splitting**

Inverse modeling seeks model parameters given a set of observed-state variables. For practical hydrogeological problems, the inversion can be ill-posed. Regularization techniques can be employed to eliminate the ill-posedness. The hydraulic tomographic analyses of aquifer heterogeneity with Tikhonov regularization tends to yield smoothed results, while the ones with TV regularization can preserve the sharp contrast between low and high permeability regions. However, hydraulic inverse modeling with the conventional TV regularization can be unstable and yield unwanted artifacts because of the non-differentiability of the TV norm. We develop a novel inverse modeling method using a TV regularization with relaxed variable-splitting scheme to preserve sharp interfaces in piecewise-constant structures and improve the accuracy of inversion. To solve the problem, we decouple the original problem into two simple sub-problems: a standard inverse modeling sub-problem with the Tikhonov regularization, and a standard L2TV sub-problem. We solve these two sub-problems using the standard linear solver and Alternating Direction Method of Multipliers iterative methods, respectively. Our new inversion algorithm is implemented in the MADS com-

putational framework (<http://mads.lanl.gov>). Our numerical examples using synthetic data show that our new methods not only preserve sharp interfaces of subsurface permeability distribution, but also significantly improve the accuracy of the inversion.

#### Youzuo Lin

Los Alamos National Laboratory  
ylin@lanl.gov

Velimir V. Vesselinov, Dan O'Malley  
Los Alamos National Lab  
vvv@lanl.gov, omalled@lanl.gov

Brendt Wohlberg  
Los Alamos National Laboratory  
brendt@lanl.gov

#### **MS273**

##### **Stochastic Simulation of Predictive Space-Time Scenarios of Wind Speed Using Observations and Physical Models**

We propose a statistical space-time model for predicting atmospheric wind speed based on deterministic numerical weather predictions and historical measurements. We consider a Gaussian multivariate space-time framework that combines multiple sources of past physical model outputs and measurements along with model predictions in order to produce a probabilistic wind speed forecast within the prediction window. Relevant to this workshop, the approach results in the ability of quantifying the model error of numerical weather systems. We illustrate this strategy on wind speed forecast during several months in 2012 for a region near the Great Lakes in the United States. The results show that the prediction is improved in the mean-squared sense relative to the numerical forecasts as well as in probabilistic scores. Moreover, the samples are shown to produce realistic wind scenarios based on sample spectra.

Julie Bessac  
Argonne National Laboratory  
jbessac@anl.gov

Emil M. Constantinescu, Mihai Anitescu  
Argonne National Laboratory  
Mathematics and Computer Science Division  
emconsta@mcs.anl.gov, anitescu@mcs.anl.gov

#### **MS273**

##### **A Hierarchical Bayesian Approach for Characterizing Model Bias**

We introduce several approaches for embedding probabilistic model discrepancy within models, which can be described using a hierarchical Bayesian model calibration framework. For some approaches, parameters of the probabilistic model discrepancy are only weakly-identifiable when considering a single model. Due to uncertainty in model structure and parameters, multiple possible models (e.g., equifinal models) can be combined into a single hierarchical Bayesian model, with a hyperprior describing the shared model parameters. Unlike Bayesian model averaging or generalized likelihood uncertainty estimation, our hierarchical Bayesian approach does not assume independence between the posterior distributions conditioned on each model, and provides a consistent likelihood measure. We also propose several methods to construct discrete and continuous probability measures over the space of compet-

ing models to determine the *a priori* uncertainty about the “best” model.

Chi Feng  
Massachusetts Institute of Technology  
chifeng@mit.edu

### MS273

#### Approximate Bayesian Inference for Intractable Likelihood Functions

Properly modeling the discrepancy for inadequacy of models yields intractable likelihood functions, which requires the evaluation of high-dimensional integrals at every Markov Chain Monte Carlo step. As a results alternative approximations need to be developed. A new approximate Bayesian inference algorithm is proposed to obtain samples from high dimensional posterior distributions using latent variable models. This is shown to be generally applicable to both physics-based models and purely data-driven processes.

Gabriel Terejanu  
University of South Carolina  
terejanu@cse.sc.edu

### MS273

#### Quantifying Uncertainty in Predictions from Inadequate Models for Porous Media Transport

A computational model for estimating the evolution of species concentration during flow through a porous media characterized by a random permeability field is considered. A reduced model is obtained by spanwise averaging of the concentration and flow velocity. We consider mathematical formulations to potentially mitigate and quantify uncertainty owing to inadequately modelling the *unclosed* term in the reduced model. In this talk, we aim to present an analysis for the uncertainty in predictions for instantaneous mean concentration at a given location, as obtained using formulations for inadequacy in the reduced model, in a Bayesian setting.

Manav Vohra  
Duke University  
Corning Incorporated  
manav@ices.utexas.edu

Damon McDougall  
Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
damon@ices.utexas.edu

Todd A. Oliver  
PECOS/ICES, The University of Texas at Austin  
oliver@ices.utexas.edu

Robert D. Moser  
University of Texas at Austin  
rmoser@ices.utexas.edu

### MS274

#### Scalable Domain Decomposition Solvers for Embedded Boundary Methods

Embedded boundary methods are very attractive, because eliminate the need to define body-fitted meshes. In particular, at large scales, the meshing step is a bottleneck of the

simulation pipeline, since mesh generators do not usually scale properly. In some other situations, like in additive manufacturing simulations, the geometry evolves in time, and the use of body-fitted meshes is not suitable. On the contrary, algorithms to create adaptive cartesian meshes are highly scalable. However, using embedded boundary methods, one can destroy the condition number of the linear systems to be solved, since cut elements can have close to zero support. As a result, these techniques require direct linear solvers, since standard preconditioned iterative solvers are not robust and scalable. In this work, we take as a starting point a balancing domain decomposition by constraints (BDDC) preconditioner. Next, we consider a recent physics-based version of the method that is robust with respect to high variations of the materials. Finally, we show to how to make these preconditioners robust also for embedded boundary methods for coercive PDEs, by a proper modification of the inter-subdomain constraints.

Santiago Badia  
International Center for Numerical Methods in Engineering  
Universitat Politècnica de Catalunya, Barcelona, Spain  
sbadia@cimne.upc.edu

Francesc Verdugo  
CIMNE  
fverdugo@cimne.upc.edu

### MS274

#### A User-Friendly Highly Scalable Amg Solver

We present AGMG, a linear system solver tailored for discrete scalar second order elliptic PDEs. It uses a method of algebraic multigrid type with coarsening based on plain aggregation. It compares favorably with some other algebraic multigrid solvers, both in term of speed and robustness. Excellent scalability is observed with both the multithreaded (OpenMP) and MPI implementations. As an extreme example, a 3D Poisson problem with more than  $10^{12}$  unknowns has been solved in less than two minutes using 23,328 processors with 16 cores each. The solver is user friendly in the sense that no expertise is required from the end user to select variant or tune parameter, all choices being made automatically by the software. AGMG can also be used to solve other problems like discrete Stokes equations or Laplacians of graphs with possibly exotic sparsity pattern.

Yvan Notay  
Universite Libre de Bruxelles  
ynotay@ulb.ac.be

### MS274

#### Scaling Finite Element Multigrid Solvers to Ten Trillion ( $10^{13}$ ) Unknowns

Geometric multigrid methods belong to the fastest solvers for elliptic problems and systems of PDE, such as the Stokes equations. On block-structured grids they can be implemented in a matrix-free style and thus can solve systems with in excess of ten trillion ( $10^{13}$ ) unknowns on Peta-Scale systems using compute times of a few minutes.

Dominik Bartuschat  
Lehrstuhl f. Simulation  
Uni Erlangen-Nuremberg  
dominik.bartuschat@fau.de

Ulrich J. Ruede  
 University of Erlangen-Nuremberg  
 Department of Computer Science (Simulation)  
 ulrich.ruede@fau.de

**MS274****An Algebraic Multigrid Approach to PDE Systems with Variable Degrees-of-Freedom Per Node**

The application of algebraic multigrid (AMG) to PDE systems can be problematic when the number of degrees-of-freedom per node varies. Currently, most monolithic AMG approaches for PDE systems assume that the number of degrees-of-freedom per node is constant. This talk considers multi-physics flow problems that contain interfaces between different materials. The incompressible Navier-Stokes equations are discretized via stabilized linear finite elements using a moving mesh that conforms to the interface. In this formulation, discontinuities are allowed at interfaces and represented by multiple degrees-of-freedom (e.g., pressure unknowns for two materials such as air and water at the same physical location). In this way, the number of degrees-of-freedom at nodes on the interface is greater than the rest of the domain. To tackle this problem with multigrid, we consider an approach that applies AMG to a scalar problem in order to generate initial grid transfer operators. This scalar problem corresponds to a distance Laplacian that is constructed automatically using coordinate information, the matrix, and a Boolean table indicating which degrees-of-freedom are defined at each nodal location. The grid transfers for this Laplace problem are then used to build interpolation/restriction operators for the original incompressible flow application. Numerical results will be presenting highlighting convergence rates and the accurate capture of interfaces on coarse meshes.

Raymond S. Tuminaro  
 Sandia National Laboratories  
 rstumin@sandia.gov

**MS275****Hybrid MultiScale Finite Element-Finite Volume Method for Poroelastic Geological Media**

Geological porous media display highly heterogeneous poromechanical properties over a wide range of length scales. High resolution characterizations are needed to obtain reliable modeling predictions and motivate the development of multiscale solution strategies to cope with the computational burden. Here, we propose a hybrid MultiScale Finite Element-Finite Volume (h-MSFE-FV) framework for the simulation of single-phase flow through deformable porous media. A coupled two-field fine-scale mixed finite element-finite volume formulation of the governing equations, namely conservation laws of linear momentum and mass, is first implemented assuming as primary unknowns the displacement vector and pressure. For the MSFE displacement stage, we introduce local basis functions for the displacement field over coarse elements that are subject to reduced boundary condition. Such MSFE stage is then coupled with the MSFV method for flow, in which a coarse and a dual-coarse grid are imposed to obtain approximate but conservative multiscale solutions. Numerical experiments are presented to demonstrate robustness and accuracy of the proposed method—both as an approximate, non-iterative solver, and a preconditioner.

Nicola Castelletto

Stanford University  
 Department of Energy Resources Engineering  
 ncastell@stanford.edu

Hadi Hajibeygi  
 TU Delft  
 h.hajibeygi@tudelft.nl

Hamdi Tchelepi  
 Stanford University  
 Energy Resources Engineering Department  
 tchelepi@stanford.edu

**MS275****Multiscale and Upscaling Finite Volume Methods in Electromagnetic Geophysics**

Simulation of electromagnetic fields in geophysical settings that include highly heterogeneous media and features varying at multiple spatial scales is computationally expensive. These settings often require very fine and large meshes to be discretized accurately. Large meshes lead to solve enormous linear systems of equations. In this talk, we introduce upscaling and multiscale finite volume methods that help to reduce the size of the system of equations to be solved. The performance of our methods is shown in the context of EM mineral exploration.

Luz Angelica Caudillo Mata  
 Earth, Ocean and Atmospheric Sciences Department  
 University of British Columbia  
 lcaudill@eos.ubc.ca

Eldad Haber  
 Department of Mathematics  
 The University of British Columbia  
 haber@math.ubc.ca

Christoph Schwarzbach  
 University of British Columbia  
 cschwarz@eos.ubc.ca

**MS275****Spectral Upscaling for Graph Laplacian Problems with Application to Reservoir Simulation**

In this talk, we will discuss a coarsening procedure for graph Laplacian problems written in a mixed saddle-point form. In that form in addition to the original vertex degrees of freedom (dofs), we also have edge dofs. We extend previously developed aggregation-based coarsening procedures applied to both sets of dofs [Vassilevski and Zikatanov, *Commun. Numer. Lin. Alg. Appl.*, 21 (2014), pp. 297–315] to allow now more than one coarse vertex dof per aggregate. Those dofs are selected as certain eigenvectors of local graph Laplacians associated with each aggregate. Additionally, we coarsen the edge dofs by using traces of the discrete gradients of the already constructed coarse vertex dofs. As a main application, by treating a finite volume discretization as a graph Laplacian formulated in mixed form, we develop consistent and accurate coarse-scale models of a fine scale finite volume problem.

Andrew Barker  
 Lawrence Livermore National Laboratory  
 atb@llnl.gov

Chak Lee

Texas A&M University  
chakshinglee@gmail.com

Panayot Vassilevski  
Lawrence Livermore National Laboratory  
vassilevski1@llnl.gov

### MS276

#### Analysis of Partitioned Methods for the Biot System

Abstract Not Available At Time Of Publication

William Layton

University of Pittsburgh  
wjl@pitt.edu

### MS276

#### Viscodynamic Operators in Poroelastic Wave Equations

Homogenized wave equations for poroelastic materials in general contain memory terms, which take into account the frequency-dependent fluid-solid interaction in the fine scale. Viscodynamic operators refer to the operators defined by the kernels of these memory terms. The dependence of these operators on the microstructure of the poroelastic composites will be mathematically quantified in this talk. It will also be explained in this talk as to how an approximation theory linked to Stieltjes functions can be utilized to develop efficient numerical schemes for approximating the kernel function from discrete permeability data and for dealing with the memory term for time-domain solvers. Numerical results will also be presented in this talk.

Yvonne Ou

University of Delaware  
mou@math.udel.edu

### MS276

#### Direct and Inverse Problems for Biot's Equations

The classic poroelastic theory of Biot, developed in 1950's, describes the propagation of elastic waves through a porous media containing a fluid. This theory has been extensively used in various fields dealing with porous media: continuum mechanics, oil/gas reservoir characterization, environmental geophysics, earthquake seismology, etc. In this work we study the propagation of elastic waves in porous media governed by the Biot equations in the low and high frequency ranges. In the low frequency case we prove the existence and uniqueness result both for the direct problem and the inverse one, which consists in identifying the unknown scalar function  $f(t)$  in the body density force  $f(t)\mathbf{g}(\mathbf{x}, t)$  acting on a poroelastic body when some additional measurement is available. In the case of the high frequency range (Biot-JKD approach) we prove the uniqueness and continuous dependence on the data of a weak solution both in unbounded and bounded time intervals and in all space dimensions. Additionally, we discuss implementation of the matrix method to derive explicit formulas for the analysis of propagation of elastic waves through a stratified 3D porous media, where the parameters of the media are characterized by piece-wise constant functions of only

one spatial variable, depth.

Viatcheslav I. Priimenko

North Fluminense State University Darcy Ribeiro  
slava@lenep.uenf.br, slava211054@gmail.com

Mikhail P. Vishnevskii

Sobolev Institute of Mathematics, Russia  
North Fluminense State University Darcy Ribeiro, Brazil  
mikhail@uenf.br

### MS276

#### Analysis of Nonlinear Poro-Visco-Elasticity with Applications to the Human Eye

We consider a nonlinear elliptic-parabolic system of PDEs that models fluid flow through poro-visco-elastic material. The ability of the fluid to flow within the solid is described by the permeability tensor which varies nonlinearly with the structural dilation. We study the existence of weak solutions in bounded domains with physical, mixed-type boundary conditions, and we account for non-zero volumetric and boundary sources. One principal aim is to investigate the influence of viscoelasticity on the qualitative properties solution. Our analysis shows that different time regularity requirements are needed for the volumetric source of linear momentum and the boundary source of traction depending on the presence of viscoelasticity. Theoretical results are further investigated via numerical simulations; when data are appropriately regular, numerical simulations show that the solutions satisfy the predicted energy estimates. Simulations also show that, in the purely elastic case, the Darcy velocity and the related fluid energy may become unbounded if the data do not enjoy the time regularity required by the theory. These results are interpreted in the context of pressure changes in lamina cribrosa (in the human eye), and the connection between these biomechanics and the development of glaucoma.

Lorena Bociu

N. C. State University  
Raleigh, NC  
lvbociu@ncsu.edu

Giovanna Guidoboni

Indiana University-Purdue University at Indianapolis  
Department of Mathematical Sciences  
gguidobo@math.iupui.edu

Riccardo Sacco

Politecnico di Milano  
riccardo.sacco@polimi.it

Justin Webster

Department of Mathematics  
College of Charleston  
websterj@cofc.edu

### MS277

#### Inverse Medium Problem for Acoustic Problems

We describe a fast, stable algorithm for the solution of the inverse acoustic scattering problem in two dimensions. Given full aperture far field measurements of the scattered field for multiple angles of incidence, we use Chen's method of recursive linearization to reconstruct an unknown sound speed at resolutions of thousands of square wavelengths in a fully nonlinear regime. Despite the fact that the under-

lying optimization problem is formally ill-posed and non-convex, recursive linearization requires only the solution of a sequence of linear least squares problems at successively higher frequencies. By seeking a suitably band-limited approximation of the sound speed profile, each least squares calculation is well-conditioned and involves the solution of a large number of forward scattering problems, for which we employ a recently developed, spectrally accurate, fast direct solver.

Carlos Borges  
Courant Institute of Mathematical Sciences  
New York University  
borges@ices.utexas.edu

Leslie Greengard  
Simons Foundation and Courant Institute  
New York University  
greengar@cims.nyu.edu

Adrianna Gillman  
Rice University  
Department of Computational and applied mathematics  
adrianna.gillman@rice.edu

#### MS277

##### Numerical Solution to a Linearized Time Fractional KdV Equation on Unbounded Domains

An efficient numerical scheme is developed to solve a linearized time fractional KdV equation on unbounded spatial domains. First, the exact absorbing boundary conditions (ABCs) are derived which reduces the pure initial value problem into an equivalent initial-boundary value problem on a finite interval that contains the compact support of the initial data and the inhomogeneous term. Second, the stability of the reduced initial-boundary value problem is studied in detail. Third, an efficient unconditionally stable finite difference scheme is constructed to solve the initial-boundary value problem where the nonlocal fractional derivative is evaluated via a sum-of-exponentials approximation for the convolution kernel. As compared with the direct method, the resulting algorithm reduces the storage requirement from  $O(MN)$  to  $O(M \log^d N)$  and the overall computational cost from  $O(MN^2)$  to  $O(MN \log^d N)$  with  $M$  the total number of spatial grid points and  $N$  the total number of time steps. Here  $d = 1$  if the final time  $T$  is much greater than 1 and  $d = 2$  if  $T \approx 1$ . Numerical examples are given to demonstrate the performance of the proposed numerical method.

Qian Zhang  
Beijing Computational Science Research Center, Beijing, China  
qianzhang@csrc.ac.cn

Jiwei Zhang  
Beijing Computational Science Research Center  
jwzhang@csrc.ac.cn

Shidong Jiang  
Department of Mathematical Sciences  
New Jersey Institute of Technology  
shidong.jiang@njit.edu

Zhimin Zhang  
Beijing Computational Science Research Center  
Beijing, China

zzhang@math.wayne.edu

#### MS277

##### Optimal Design of Metallic Surfaces for Near-Field Energy Extraction

An emitter in an inhomogeneous medium can radiate energy at rates dramatically enhanced relative to free-space spontaneous emission. Yet the highest rate enhancements are often demonstrated in structures that predominantly decay into plasmonic channels that generate heat instead of radiation. How does one optimally couple the near field to the far field? Through systematic computation and shape optimization we identify structures that start to approach global bounds on this process.

Owen D. Miller  
Department of Applied Physics  
Yale University  
owen.miller@yale.edu

Yuxiang (Larry) Liu  
Yale  
yuxiang.liu@yale.edu

#### MS277

##### Scattering of Transient Waves by Piecewise Homogeneous Obstacles

Some recent work has shown that semidiscretization in space of some time-domain boundary integral equations related to transmission problems for the wave equation is equivalent to a dynamical system where a double transmission problem (with two fields on each side of the interface) is subject to some exotic transmission conditions that reflect the trial-test characteristics of Galerkin semidiscretization. Here we consider a more complicated topological situation, where several non-overlapping domains generate a similar dynamical system, and then we pull the thread to discover an automatically stable semidiscrete system of integral equations related to the scattering of acoustic (or elastic) waves by piecewise homogeneous obstacles. We will also briefly address new analytical techniques to study full discretization of the problem using Convolution Quadrature in the time variable.

Francisco J. Sayas  
Department of Mathematical Sciences  
University of Delaware  
fjsayas@math.udel.edu

Alexander Rieder  
Technical University of Vienna, Austria  
alexander.rieder@tuwien.ac.at

#### MS278

##### Preliminary Studies of an Implicit Navier-Stokes Solver on a Many-Core Machine

In this talk, we present some preliminary studies of an implicit, unstructured grid, incompressible Navier-Stokes solver on a many-core computer, whose processors include management processing cores (MPCs) and computing processing cores (CPCs). Newton-Krylov-Schwarz is the general framework that we use to solve the large sparse nonlinear system of equations, and the focus of the talk is on the re-factorization and optimization techniques of the shared memory parallelization of the subdomain solvers.



We report some performance results obtained using a large number of MPCs and CPCs.

Rongliang Chen  
Chinese Academy of Sciences  
rl.chen@siat.ac.cn

Li Luo, Ziju Liao  
Shenzhen Institutes of Advanced Technology  
Chinese Academy of Sciences  
li.luo@siat.ac.cn, zj.liao@siat.ac.cn

Zhengzheng Yan  
Shenzhen Institutes of Advanced Technology  
zz.yan@siat.ac.cn

Xiao-Chuan Cai  
University of Colorado, Boulder  
Dept. of Computer Science  
cai@cs.colorado.edu

### MS278

#### A Study on the Impact of Extra Precision on the Unsymmetric Eigenvalue Problem

First, we discuss the optimization of extra precision calculations like double-double on recent multicore and many-core processors. Second, we provide some new insights on building a library based on extra precision, comparing to other libraries like the XBLAS. Finally, for the purpose of illustration, we show a new algorithm related to the unsymmetric eigenvalue problem which can exploit extra precision and show when it works and doesn't work, comparing favorably to a theoretical analysis of the problem.

Greg Henry  
Intel Corporation  
greg.henry@intel.com

### MS278

#### Vectorizing and Multithreading the Calculation of Electronic Integrals

In computational chemistry, the computation of electron repulsion integrals (ERI) is often considered a bottleneck in many calculations. Because computation of these integrals scales (formally) as the fourth power with the size of the system, there has been much interest in their efficient computation. We have developed an ERI library (Simint) which utilizes single instruction, multiple data (SIMD) instructions to speed up the calculation of these integrals. Higher performance is obtained by writing the computationally-expensive inner loop structure in terms of vector intrinsics, rather than relying on compiler optimization. Benchmarks display 2-4x speedup over scalar code on the Haswell microarchitecture. Speedups obtained with the Psi4 computational chemistry package, as well as benchmarks of calculations with the Intel Xeon Phi (Knights Landing) will be presented. Some new schemes of assembling the calculated data within a computational chemistry program will also be discussed.

Benjamin Pritchard  
Georgia Institute of Technology  
ben@bennyp.org

Edmond Chow  
School of Computational Science and Engineering

Georgia Institute of Technology  
echow@cc.gatech.edu

### MS278

#### MATLAB and MEX for Multicore Machines

MATLAB, as a general purpose technical computing environment, does its best to leverage the full computational power of the host machine. Often, however, toolbox authors can exploit knowledge about their particular problems and build higher performance applications by using C/C++/FORTRAN codes linked into MATLAB by using the MEX infrastructure. This talk will provide some tips and tricks for squeezing performance out of MATLAB on multicore machines including how to make MEX files that use various multicore technologies without interfering with the MATLAB environment.

Pat Quillen  
MathWorks  
Pat.Quillen@mathworks.com

### MS279

#### Topology Optimization and Mesh Adaptation

New techniques, such as topology and shape optimization, have been recently proposed and developed to design new several modern structures, essentially to have better performances with respect to certain paradigms of interest. In particular, topology optimization, starting from the so-called design domain, optimizes the distribution of mass to minimize quantities of design interest such as, e.g., mechanical or thermal stresses, compliance. Several disciplines in the aerospace, automotive, or biomedical engineering, may benefit of any new contribution in topology optimization. Our interest is in Additive Layer Manufacturing (ALM) which has been growing fast in the last decade. In this context, an optimal distribution of the material turns out to be of paramount importance. Thus, for example, 3D printing-based production can be assisted by these techniques. From a mathematical point of view, topology optimization requires solving the linear elasticity partial differential equation. For this purpose, a suitable computational grid is to be devised; the grid should be picked such that the numerical solution can be considered reliable with respect to measurement data. Aim of this presentation is to analyze the effects of mesh adaptation on topology optimization, in order to propose new competitive numerical methods which are physically reliable and computationally affordable.

Nicola Ferro  
Politecnico di Milano  
nicola.ferro@polimi.it

Simona Perotto  
MOX - Modeling and Scientific Computing  
Dipartimento di Matematica  
simona.perotto@polimi.it

Stefano Micheletti  
Politecnico di Milano  
stefano.micheletti@polimi.it

### MS279

#### A Flexible Conservative Remapping Framework for

## Exascale Computing

We present an open-source remapping framework called *Portage* that can transfer fields between general polyhedral meshes while conserving some integral quantity of interest. At the heart of Portage is a templated driver that can be constructed with any set of classes that can perform the steps required in conservative remapping - searching, intersection and interpolation. The driver is also templated on classes for the mesh and state managers that can answer a specific set of queries. Distributed memory parallelism in Portage requires that the mesh and state managers support the concept of partitions surrounded by a halo of ghost elements. On-node parallelism in Portage requires the search, intersection and interpolation classes to be written such that that can operate on each element without any side effects. This allows Portage to use NVIDIA's Thrust interface to enable on-node parallelism using OpenMP or Intel TBB backends. We have tested Portage in 2D/3D for remapping between general polygonal and polyhedral meshes. The code can remap cell-centered or node-centered variables using 1st or 2nd order accurate remapping with Barth-Jespersen limiting of gradients to enforce monotonicity. We will present convergence and scaling results.

Ondrej Certik, Charles R. Ferenbaugh, Rao V. Garimella, Angela Herring, Brian Jean, Christopher Malone, Christopher Sewell  
Los Alamos National Laboratory  
certik@lanl.gov, cferenba@lanl.gov, rao@lanl.gov,  
angelah@lanl.gov, baj@lanl.gov, cmalone@lanl.gov,  
csewell@lanl.gov

## MS279

### Portably Performant and Conservative Mesh Adaptivity for Lagrangian Shock Dynamics

A system is presented that performs general mesh adaptivity to support a Lagrangian shock simulation. Our first priority is portability across all new DOE systems, including NVidia GPUs, Intel Xeon Phis. We present a general strategy to keep all data and most work on the accelerator hardware. We make use of ghost (halo) element layers, and select sets of cavities that can be modified independently. We also ensure termination when no modifications are needed. Our second priority is strict conservation of mass, energy, and momentum during adaptivity. Effective local algorithms for such conservation are presented, as well as challenges they present to parallelism and multi-material settings. Conservation is based on a local common refinement approach which considers intersections of old and new elements for a local cavity modification. Quantities such as mass and energy which are discontinuous between elements are conserved via sums of intersection volumes. Piecewise linear continuous quantities such as velocity require buffer layers of elements around the modification such that cavity boundary values are not modified in conflict. We are able to implement this in the same unified independent set framework. The implementation of complex operations in a GPU kernel will be discussed, including porting of the R3D library for element intersection and the solution of local linear systems.

Dan A. Ibanez  
Rensselaer Polytechnic Institute  
SCOREC  
ibaned@rpi.edu

Mark S. Shephard

Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

Thomas Voth, Edward Love  
Sandia National Laboratory  
tevoth@sandia.gov, elove@sandia.gov

James Overfelt, Glen Hansen  
Sandia National Laboratories  
jroverf@sandia.gov, gahanse@sandia.gov

## MS279

### A Global Optimization and Adaptivity-Based Algorithm for Automated Edge Grid Generation

Digital curves can be found in fields as diverse as computer-aided design, image processing, and kinematics. Typically discretizations (i.e., edge grids) of digital curves are required for use in computational simulations. Such edge grids must accurately represent the underlying physical object or phenomenon in order to obtain accurate results in a predictive computational simulation. In particular, an edge grid should do a good job of approximating the arc length of the curve. In other words, the discretization should minimize the arc length deficit. The edge grids should also be generated without manual intervention. In earlier work by McLaurin and Shontz, we developed a method for automatically generating optimal discretizations of curves through local minimization of the arc-length deficit. The method also incorporated adaptivity in that it added nodes to the existing grid until the arc-length deficit was less than a prescribed tolerance. In this talk, we propose a method for automatic edge grid generation through global minimization of the arc-length deficit via an underlying pattern search method. The method is also adaptive. We present results from our numerical experiments which demonstrate the robustness and accuracy of our global optimal discretization technique for automatic edge grid generation. We also compare the performance of our global optimization method to that of our earlier local optimization method.

Suzanne M. Shontz  
University of Kansas  
shontz@ku.edu

David McLaurin  
Cd-Adapco  
david.mclaurin@gmail.com

## MS280

### Band Structure Inverse Problem

In mean-field models, the electronic structure of a perfect crystal (i.e. no defects) is characterized by a periodic Schrodinger operator

$$H = -\frac{1}{2}\Delta + V_{\text{per}}$$

acting on  $L^2(\mathbb{R}^d)$  with domain  $L^2(\mathbb{R}^d)$  where  $V_{\text{per}}$  is a real valued periodic potential. The Bloch-Floquet theory allows one to characterize the spectrum of  $H$  through an infinite set of functions  $q \in \Gamma^* \mapsto \varepsilon_n(q)$  defined on the First Brillouin zone  $\Gamma^*$ . These functions are called the energy bands. The inverse problem of band structure consists in finding an admissible periodic potential  $V_{\text{per}}$  from the knowledge of partial spectral information. In our case, from the knowledge of only a few energy bands. One way of talking the

question is to formulate it as an optimization problem. In this talk, some well known results on the spectra of periodic Schrodinger operators will be recalled as well as some generalizations. A numerical method based on gradient descent will be proposed to solve the inverse problem. The proof of existence of a minimizer will be sketched for the one dimensional case and one target band.

Athmane Bakhta  
Ecole des Ponts Paristech  
bakhtaa@cermics.enpc.fr

#### MS280

##### **Electronic Transport in Incommensurate Van Der Waals 2D Heterostructures**

Electronic Transport in Incommensurate Van Der Waals 2D Heterostructures Two-dimensional materials have generated a lot of interest in the past decade as a new toolbox for electronics. This family includes insulators (boron nitride), semiconductors (transition metal dichalcogenides), and conductors (graphene). Vertical stacks of a few layers of such materials, interacting through van der Waals forces, create a venue to explore and tune desired electronic properties. However, the generically incommensurate character of these systems represents a significant hurdle for theoretical and numerical advances. Due to the lack of periodicity, many mathematical problems remain open. In this talk, we present some progress towards multi-scale calculations aimed at rigorously predicting macroscopic properties of heterostructures. First, we recall some basic mathematical notions of electronic structure calculations, as well as some properties of single-layer 2D materials. We then present an application of non-commutative geometry for the rigorous modelling, definition and calculation of density of states and transport properties in heterostructures. These results have been obtained in the framework of a collaboration with the groups of M. Luskin (Math, UofM), Efthimios Kaxiras (Physics, Harvard) and Eric Cancs (Math, Ecole des Ponts).

Paul Cazeaux  
University of Minnesota  
pcazeaux@umn.edu

#### MS280

##### **Advances in Real Space Methods to Solve the Kohn-Sham Equation**

We will discuss advances in solving the Kohn-Sham equation using pseudopotentials implemented in real space. A solution within density functional theory is often limited by the high computational demand in solving an eigenvalue problem at each self-consistent-field (SCF). Our code replaces the explicit eigenvalue calculations by an approximation of the wanted invariant subspace, obtained with well-selected Chebyshev polynomial filters. Only the initial SCF iteration requires solving an eigenvalue problem, in order to provide a good initial subspace. In the remaining SCF iterations, no iterative eigensolvers are involved. Instead, Chebyshev polynomials are used to refine the subspace. The subspace iteration at each step is easily five to ten times faster than solving a corresponding eigenproblem by the most efficient eigen-algorithms. Moreover, the subspace iteration reaches self-consistency within roughly the same number of steps as an eigensolver-based approach. This results in a significantly faster SCF iteration. We will discuss recent advances in data partitioning. These advances include improvements to sparse matrix vector mul-

tiplication, which makes up the large computational component of the filtering process. We demonstrate an efficient scaling of the Laplacian component to more than  $10^8$  grid points and of the Rayleigh-Ritz component to  $10^5$  electronic states. As an application, we consider nanocrystals of silicon containing over 10,000 atoms.

James R. Chelikowsky  
Institute for Computational Engineering and Sciences  
University of Texas at Austin  
jrc@utexas.edu

Charles Lena  
University of Texas  
charles@ices.utexas.edu

Ariel Biller, Leeor Kronik  
Weizmann Institute  
billerel@gmail.com, leeor.kronik@weizmann.ac.il

#### MS280

##### **Title Not Available At Time Of Publication**

Abstract Not Available At Time Of Publication.

Xiaoying Dai  
Institute of Computational Mathematics  
Chinese Academy of Sciences, Beijing, China  
daixy@lsec.cc.ac.cn

#### MS281

##### **Simultaneous Parallel-in-Time Optimization with Unsteady PDEs Using a Multigrid Reduction in Time Software Library**

In this talk we present a simultaneous optimization framework for unsteady PDEs using the parallel-in-time software library XBraid. XBraid provides a non-intrusive approach for parallelization in the time domain by applying an iterative multigrid reduction in time algorithm to existing spatially parallel time-stepping codes. We develop a non-intrusive adjoint solver for XBraid that enhances these primal iterations by an iteration for computing adjoint sensitivities parallel in time. The adjoint sensitivities are then used in a simultaneous optimization method, namely the One-shot method which solves the optimization problem in the full space. We validate the time-parallel simultaneous optimization approach by applying it to an inverse design problem with unsteady PDEs that mimics the behavior of separated flows past bluff bodies: the near wake is mimicked by a nonlinear ODE, namely the Lorenz attractor, which exhibits self-excited oscillations. The far wake is modeled by an advection-diffusion equation whose upstream boundary condition is determined by the ODE mimicking the near wake.

Stefanie Günther, Nicolas R. Gauger  
TU Kaiserslautern  
stefanie.guenther@scicomp.uni-kl.de,  
nicolas.gauger@scicomp.uni-kl.de

Jacob B. Schroder  
Lawrence Livermore National Laboratory  
schroder2@llnl.gov

#### MS281

##### **Towards a Posteriori Error Estimation for Least Squares Shadowing Sensitivity Analysis of Chaotic**

## Systems

Many important phenomena in engineering are chaotic, such as problems involving turbulent flow or fluid-structure interaction. For these chaotic systems, we are often interested in long-time averaged objective functions rather than instantaneous quantities. Furthermore, we want to perform sensitivity analysis, i.e. to know how changing the design parameters of the systems would affect those long-time averaged objectives. The Least Square Shadowing (LSS) method can perform such task. In LSS, we first solve a least square problem to approximate the true shadowing direction of a trajectory, then use this approximate shadowing direction to compute sensitivity. It has been estimated that the approximation in shadowing direction can result in an  $O(\frac{1}{T})$  error in the sensitivity result, where  $T$  is the length of the trajectory. However, the error in the approximate shadowing direction is also influenced by many other factors, such as the angle between stable and unstable subspaces, the Lyapunov exponents, and the amplitude of the shadowing direction itself. In this talk, we develop a posteriori estimation of the error in LSS due to approximating the shadowing direction, by taking all the above factors into consideration.

Angxiu Ni  
MIT  
niangxiu@mit.edu

Qiqi Wang  
Massachusetts Institute of Technology  
qiqi@mit.edu

## MS281

### Uncertainty Estimates for Statistics Computed from Simulations of Chaotic Systems

Systematic assessment of uncertainty is essential in scientific computation. Developing accurate uncertainty estimates is particularly important and challenging in the context of simulations of chaotic systems, where the computed quantities of interest (QoIs) are often statistical quantities. Such predictions are contaminated by both discretization error and finite sample size effects, and thus, it is necessary to estimate both of these errors in order to characterize the uncertainty in the QoIs. This talk describes a statistical methodology for estimating both of these errors. The sampling error estimate is formulated based on the extended central limit theorem that allows for correlated random variables and can be understood as estimating the effective sample size, accounting for correlation, for the given sample. The discretization error is estimated using an extension of Richardson extrapolation that accounts for the estimated sampling error by reformulating Richardson extrapolation as a Bayesian inference problem. These techniques have previously been used to investigate errors in direct numerical simulation of turbulent channel flow. We will review these results and show applications to other examples in fluid mechanics.

Todd A. Oliver  
PECOS/ICES, The University of Texas at Austin  
oliver@ices.utexas.edu

Sigfried Haering  
The University of Texas at Austin  
shaering@utexas.edu

Nicholas Malaya, Robert D. Moser  
University of Texas at Austin

nick@ices.utexas.edu, rmoser@ices.utexas.edu

## MS281

### Mathematical and Numerical Challenges in Shadowing Methods for Sensitivity Analysis of Chaotic Simulations

Chaotic simulations have many applications in fluid dynamics, combustion, molecular dynamics, climate, and astrophysics. Their solutions exhibit sensitive dependence on initial conditions and simulation parameters, but statistical quantities of these solutions, e.g., long time averages, can depend smoothly on simulation parameters. The shadowing method aims at computing the derivative of these statistical quantities, sometimes called the climate of the simulation, with respect to simulation parameters. The shadowing method has been applied to a range of chaotic simulations, ranging from the Lorenz attractor to direct numerical simulation of turbulent flows. This talk will introduce the mathematical formulation and computational algorithm of the shadowing method. It will also expose the challenges current facing researchers. These include extending the theory to non-uniformly hyperbolic systems, and systems with weaker ergodic properties. Also challenging are reducing the cost of shadowing methods, particularly for simulations whose strange attractors have hundreds or more dimensions.

Qiqi Wang  
Massachusetts Institute of Technology  
qiqi@mit.edu

## MS282

### Parallel Versions of the Southwell Iterative Method with Low Communication for Preconditioning and Smoothing

The Southwell iterative method for solving linear systems is similar to the Gauss-Seidel method, relaxing one row at a time, but choosing the row that has the highest residual. We present parallel versions of this method, including one that is appropriate for distributed memory computers. We study the use of these methods as preconditioners and multigrid smoothers. The new methods have the potential of reducing computation and communication compared to the Gauss-Seidel method.

Edmond Chow, Jordi Wolfson-Pou  
School of Computational Science and Engineering  
Georgia Institute of Technology  
echow@cc.gatech.edu, jwp3@gatech.edu

## MS282

### Estimating of the 2-Norm Forward Error for SYMMLQ and CG

For positive definite linear systems (or semidefinite consistent systems), we use Gauss-Radau quadrature to obtain a cheaply computable upper bound on the 2-norm error of SYMMLQ iterates. The close relationship between SYMMLQ and CG iterates lets us construct an upper bound on the 2-norm error for CG. For indefinite systems, the upper bound becomes an estimate of the 2-norm SYMMLQ error. Numerical experiments demonstrate that the bounds and estimates are remarkably tight.

Ron Estrin  
ICME  
Stanford University

restrin@stanford.edu

Dominique Orban  
GERAD and Dept. Mathematics and Industrial  
Engineering  
Ecole Polytechnique de Montreal  
dominique.orban@gerad.ca

Michael A. Saunders  
Systems Optimization Laboratory (SOL)  
Dept of Management Sci and Eng, Stanford  
saunders@stanford.edu

### MS282

#### A Minimum Residual Method for Saddle Point Systems

We introduce SPMR, a new family of methods for iteratively solving saddle-point systems using a minimum or quasi-minimum residual approach. No symmetry assumptions are made. The basic mechanism underlying the method is a novel simultaneous bidiagonalization procedure that yields a simplified saddle-point matrix on a projected Krylov-like subspace, and allows for a monotonic short-recurrence iterative scheme. We develop a few variants, demonstrate the advantages of our approach, derive optimality conditions, and discuss connections to existing methods. Numerical experiments illustrate the merits of this new family of methods.

#### Chen Greif

Department of Computer Science  
The University of British Columbia  
greif@cs.ubc.ca

Ron Estrin  
ICME  
Stanford University  
restrin@stanford.edu

### MS282

#### Matrix-Equation-Based Strategies for Certain Structured Algebraic Linear Systems

The mathematical formulation of various application problems naturally gives rise to *systems* of linear matrix equations. This is the case for instance in regulator equations in robust control, in mixed finite element formulations of certain stochastic Galerkin diffusion problems, and in discretized Stokes and Navier-Stokes equations. We describe some numerical strategies for solving these systems, both in the small and large scale cases. Numerical experiments illustrate the performance of the new methods.

#### Valeria Simoncini

Universita' di Bologna  
valeria.simoncini@unibo.it

### MS283

#### An Explicit Partitioned Elastodynamics Method Based on Lagrange Multipliers

Traditional partitioned methods for explicit elastodynamics analysis exchange boundary condition information between subdomains and then advance the solution to the next time step. In so doing these algorithms treat the interface nodes similarly to boundary nodes in finite element simulations, i.e., they couple the subdomains problems by

specifying Neumann, Dirichlet or Robin type boundary conditions. In contrast, our approach connects the subdomains by exchanging forces and masses across the material interface in a manner that resembles the treatment of interior nodes in finite elements. The exchanged mass and force quantities at each interface node are estimated using an auxiliary monolithic Lagrange multiplier formulation of the interface problem. In so doing our algorithm avoids some accuracy and/or stability issues inherent to traditional partitioned solvers, yet it does not require the solution of the monolithic system.

#### Pavel Bochev

Sandia National Laboratories  
Computational Math and Algorithms  
pbboche@sandia.gov

Paul A. Kuberry  
Clemson University  
pkuberr@clemson.edu

Kara Peterson  
Sandia Natl. Labs  
kjpeter@sandia.gov

### MS283

#### Non-Iterative Multi-Physics Domain Decomposition Method for Coupled Free Flow and Porous Media Flow Problem

The Stokes-Darcy and Navier-Stokes-Darcy model have attracted significant attention in the past ten years since they arise in many applications involving with coupled free flow and porous media flow such as surface water flows, ground-water flows in karst aquifers, petroleum extraction and industrial filtration. They have higher fidelity than either the Darcy or Navier-Stokes systems on their own, but coupling the two constituent models leads to a very complex system. This presentation discusses a non-iterative multi-physics domain decomposition method for efficiently solving this type of problems in parallel. The key idea is to decouple the free and porous media flow through Robin type boundary conditions which arise from the three interface conditions. Optimal convergence is proved and computational results are presented to illustrate their features.

#### Xiaoming He

Department of Mathematics and Statistics  
Missouri University of Science and Technology  
hex@mst.edu

### MS283

#### Coupling Hyperbolic PDEs Over Non-Coincident Interfaces

When subdomains are decomposed and discretized independently, it frequently occurs that the interfaces between two subdomains have mismatched nodes and even non-coincident interfaces that are described as having gaps and overlaps. Constructing algorithms to couple the problems in a way that will satisfy the interface conditions on the subdomains' boundaries and also preserve accuracy and stability is a challenging problem. We introduce the construction of a virtual interface and operators with which to map solutions to and from subdomain interfaces. Additionally, we develop operators useful for mapping a single function to two distinct subdomain interfaces in a manner that preserves global flux conservation. Posing the problem as a linear system, we will generate the coupling algorithm

through a Schur complement and provide the results of numerical experiments demonstrating successfully passing a pass test, optimal convergence rates, and global flux conservation.

Paul Kuberry  
Sandia National Laboratories  
pakuber@sandia.gov

Kara Peterson  
Sandia Natl. Labs  
kjpeter@sandia.gov

Pavel Bochev  
Sandia National Laboratories  
Computational Math and Algorithms  
pbboche@sandia.gov

### MS283

#### Numerical Study of Viscoelastic Fluid-Structure Interaction

Simulating viscoelastic fluid-structure interactions is challenging not only due to the coupling between solid and fluid substructures, but also the complexity of fluid model in a moving domain. In this talk, both monolithic and decoupling approaches are considered for the numerical study of fluid-structure interaction problems, where the fluid is governed by a viscoelastic model. Monolithic and partitioned algorithms are presented, where the viscoelastic fluid is stabilized by the streamline upwind Petrov-Galerkin (SUPG) approximation. We will discuss some issues with the stress boundary condition on the interface and present simulation results with and without interface stress boundary conditions.

Hyesuk Lee  
Clemson University  
Dep. of Mathematical Sciences  
hklee@clemson.edu

### MS284

#### Title Not Available At Time Of Publication

Abstract Not Available At Time Of Publication.

Andreas H. Hielscher  
Departments of Biomedical Engineering and Radiology  
Columbia University  
ahh2004@columbia.edu

### MS284

#### Curvature Correction and Corner Layer of the Transport Equation

In this talk we study the Knudsen layer for the transport equation in its parabolic regime. In particular, we study if curvature correction is needed for non-flat geometry. In 2015 Wu-Guo proved an unexpected result which showed that the classical boundary layer analysis fails over the unit disk  $B(0,1)$  in 2D due to the curvature, and we show in this talk that such inconsistency does not propagate into the interior, and the classical half-space equation still provides a valid boundary condition for the interior Laplace equation. We validate the result numerically at the end of the talk too. It is a joint work with Jianfeng Lu and Weiran Sun.

Qin Li

University of Wisconsin  
qinli@math.wisc.edu

Jianfeng Lu  
Mathematics Department  
Duke University  
jianfeng@math.duke.edu

Weiran Sun  
Simon Fraser University  
weirans@sfu.ca

### MS284

#### Void-Compatible Second-Order Forms of the Transport Equation

Second-order forms of the linear Boltzmann equation have several advantages over the first-order construction. Nevertheless, these second-order forms have issues modeling the transport of particles in voids. In this work we look at two approaches to ameliorating these issues. One takes the least-squares form of the transport equation and a self-adjoint and conservative equation by modifying the weak-form of the equation. The other approach uses first-order upwinding to join heterogeneous regions of the problem. This approach also allows for varying angular resolution between domains. Both of these approaches do sacrifice the symmetry of the system, but do allow for the use of continuous finite elements in space, whereas first-order forms typically require discontinuous finite elements, and therefore, more degrees of freedom. Numerical results demonstrate that our fixes enable second-order forms to solve eigenvalue problems accurately and efficiently.

Ryan McClarren  
Texas A&M University  
USA  
rgm@tamu.edu

Weixiong Zheng, Vincent Laboure  
Texas A&M University  
zwxne2010@tamu.edu, vincent.laboure@neo.tamu.edu

### MS284

#### Title Not Available At Time Of Publication

Abstract Not Available At Time Of Publication.

Justin Ponders  
Bettis Atomic Power Laboratory  
justin.ponders@unnpp.gov

### MS285

#### Iterative Lanczos Methods for Principal Components Analysis of Genomics Data

Finding the largest few principal components of a matrix of genetic data is a common task in genome-wide association studies (GWASs), both for dimensionality reduction and for identifying unwanted factors of variation. We describe a simple random matrix model for matrices that arise in GWASs, showing that the singular values have a bulk behavior that obeys a Marchenko-Pastur law with a handful of large outliers. We also implement Golub-Kahan-Lanczos (GKL) bidiagonalization in the Julia programming language, providing thick restarting and a choice between full and partial reorthogonalization strategies to control numerical roundoff. Our implementation of GKL bidiagonalization

nalization is up to 36 times faster than software tools used commonly in genomics data analysis for computing principal components, such as EIGENSOFT and FlashPCA, which use dense LAPACK routines and randomized subspace iteration respectively.

Jiahao Chen  
MIT  
jiahao@mit.edu

#### MS285

##### **Applying the Feast Sparse Eigensolver to the Singular Value Decomposition**

The FEAST eigenvalue solver is a contour integration-based algorithm for solving large, sparse eigenvalue problems. By mathematically separating the spectrum of a matrix into multiple non-intersecting intervals, FEAST makes it possible solve an eigenvalue problem in parallel by finding the eigenvalues and eigenvectors in each interval independently. In this talk we discuss the application of the FEAST algorithm to finding the singular values and singular vectors of a matrix by solving the corresponding eigenvalue problem.

Brendan Gavin  
ECE Department  
University of Massachusetts, Amherst  
bgavin@ecs.umass.edu

Eric Polizzi  
University of Massachusetts, Amherst, USA  
polizzi@ecs.umass.edu

#### MS285

##### **PRIMME\_SVDS: A High-Performance Preconditioned SVD Solver for Accurate Large-Scale Computations**

The increasing number of applications requiring the solution of large scale singular value problems have rekindled interest in iterative methods for the SVD. Some promising recent advances in large scale iterative methods are still plagued by slow convergence and accuracy limitations for computing smallest singular triplets. Furthermore, their current implementations in MATLAB cannot address the required large problems. Recently, we presented a preconditioned, two-stage method to effectively and accurately compute a small number of extreme singular triplets. In this research, we present a high-performance software, PRIMME\_SVDS, that implements our hybrid method based on the state-of-the-art eigensolver package PRIMME for both largest and smallest singular values. PRIMME\_SVDS fills a gap in production level software for computing the partial SVD, especially with preconditioning. The numerical experiments demonstrate its superior performance compared to other state-of-the-art software and its good parallel performance under strong and weak scaling.

Lingfei Wu  
IBM T. J. Watson Research Center  
lwu@email.wm.edu

Andreas Stathopoulos  
College of William & Mary  
Department of Computer Science  
andreas@cs.wm.edu

Eloy Romero  
College of William and Mary  
eloy@cs.wm.edu

#### MS285

##### **Randomized Algorithm for Computing Or Updating SVD**

Randomized algorithms have emerged as promising algorithms to compute truncated SVD for analyzing large data sets. In this talk, we present our recent studies in using randomized algorithms for computing truncated SVD on the current computers.

Ichitaro Yamazaki, Jakub Kurzak  
University of Tennessee, Knoxville  
iyamazak@icl.utk.edu, kurzak@icl.utk.edu

Piotr Luszczek  
University of Tennessee  
luszczek@icl.utk.edu

Jack J. Dongarra  
University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

#### MS286

##### **Lagrangian Particle Methods for Euler Equations: Stability, High Order Convergence, and Limiters**

A new Lagrangian particle method for solving Euler equations for compressible inviscid fluid or gas flows is proposed. Similar to smoothed particle hydrodynamics (SPH), the method represents fluid cells with Lagrangian particles and is suitable for the simulation of complex free surface and multiphase flows. The main contributions of our method, which is different from SPH in all other aspects, are as follows. (a) Significant improvement of approximation of differential operators based on a polynomial fit via weighted least squares approximation and the convergence of prescribed order. (b) An upwind first order and unsplit second-order particle-based algorithms providing accuracy and long-term stability. A limiter method that combines the second and the first order methods with proper weights has been developed for the reduction of dispersive behavior of the second order scheme. (c) Accurate resolution of states at free interfaces using a ghost particle method. Numerical verification tests demonstrating accuracy the convergence order will be presented. In particular, wave propagation problems and the Gresho vortex, known as a difficult test for SPH, demonstrated the expected convergence order. Examples of complex free surface flows as well as applications of the method to high energy density problems will also be discussed.

Roman Samulyak  
Stony Brook University  
Brookhaven National Laboratory  
rosamu@bnl.gov or roman.samulyak@stonybrook.edu

Xingyu Wang  
Stony Brook University  
xingyu.wang@stonybrook.edu

#### MS286

##### **Adaptive Particle-in-Cloud Method for Optimal**

### Solutions of Vlasov-Poisson Equations

A new adaptive Particle-in-Cloud (AP-Cloud) method for obtaining optimal numerical solutions to the Vlasov-Poisson equation has been developed. Unlike the traditional particle-in-cell (PIC) method, which is commonly used for solving this problem, the AP-Cloud adaptively selects computational nodes or particles to deliver higher accuracy and efficiency when the particle distribution is highly non-uniform. Unlike other adaptive techniques for PIC, our method balances the errors in PDE discretization and Monte Carlo integration, and discretizes the differential operators using a generalized finite difference (GFD) method based on a weighted least square formulation. As a result, AP-Cloud is independent of the geometric shapes of computational domains and is free of artificial parameters. Efficient and robust implementation is achieved through an octree data structure with 2:1 balance. We analyze the accuracy and convergence order of AP-Cloud theoretically, and verify the method using an electrostatic problem of a particle beam with halo. Simulation results show that the AP-Cloud method is substantially more accurate and faster than the traditional PIC, and it is free of artificial forces that are typical for some adaptive PIC techniques. AP-Cloud has been applied to the simulation of modulation problem in the concept of Coherent Electron Cooling. It has achieved a good agreement with theoretical verification tests and predicted the modulation for real particle beams.

Xingyu Wang  
Stony Brook University  
xingyu.wang@stonybrook.edu

Roman Samulyak  
Stony Brook University  
Brookhaven National Laboratory  
rosamu@bnl.gov or roman.samulyak@stonybrook.edu

Xiangmin Jiao  
Stony Brook University  
xiangmin.jiao@stonybrook.edu

Kwangmin Yu  
Brookhaven National Laboratory  
yukwangmin@gmail.com

### MS286

#### Scalable Mesh Infrastructure for Particle-in-Cell Magnetic Fusion Simulations

Particle-in-cell (PIC) method has been widely adopted to solve 5D Gyrokinetic (GK) Vlasov-Poisson (or Maxwell) PDE system for electrostatic (or electromagnetic) plasma turbulence simulations at first principle level using supercomputers in magnetic fusion community. In actual implementation, one of popular ways to manage the mesh for scalability concern is that each compute node has a copy of entire mesh and its 3D fields even with domain and particle decomposition techniques in order to avoid communicating field information for particle push with partitioned mesh. However, considering increasing demands on larger number of elements and particles such as multi-scale turbulence simulation for a whole device, scalability of that algorithm is questionable for the larger scale of simulations. In this presentation, we introduce a new unstructured mesh distribution method with associated mesh-based particle data structure and particle search algorithm using topological adjacency. This method is built on the parallel unstruc-

tured mesh infrastructure (PUMI) [D.A. Ibanez, E.S. Seol, C.W. Smith, M.S. Shephard, PUMI: Parallel Unstructured Mesh Infrastructure, ACM Transactions on Mathematical Software, 42(3), Article No. 17 (2016)] considering physics of particle orbits in magnetic fusion device. The detailed algorithm and performance comparison with the conventional algorithm will be presented.

Eisung Yoon  
Rensselaer Polytechnic Institute  
yoone@rpi.edu

Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

Kaushik Kalyanaraman  
Rensselaer Polytechnic Institute  
kalyak@rpi.edu

E. Seegyoung Seol  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
seols@rpi.edu

Dan A. Ibanez  
Rensselaer Polytechnic Institute  
SCOREC  
ibaned@rpi.edu

### MS287

#### Data-Driven Model Reduction and the Preservation of DAE Structure

In this talk we will consider data-driven model reduction in the Loewner framework. In general this approach does not preserve the DAE structure (for instance, a system frequency response which tends to infinity with increasing frequency). We will address this issue in the case that the defining system matrices do not have structure. An important tool in this respect is the Moore-Penrose pseudoinverse. Several examples will illustrate the proposed method.

Athanasios C. Antoulas  
Dept. of Elec. and Comp. Eng.  
Rice University  
aca@rice.edu

### MS287

#### Preserving Passivity with Data-Driven Models

A data-driven model reduction strategy is introduced that preserves both stability and passivity by recovering and optimizing port-Hamiltonian (PH) structure consistent with the observed system response. The approach is based on a successful strategy for reducing PH systems with (explicit) interpolatory projections that are iteratively refined to produce good  $\mathcal{H}_\infty$  performance. The data-driven reformulation developed here involves semidefinite programming problems of small size. The effect of measurement noise is considered.

Christopher A. Beattie  
Virginia Polytechnic Institute and State University



beattie@vt.edu

### MS287

#### A Comparison of Three MOR Methods for LSS: The Loewner Framework, Balanced Truncation and Optimal H2

We discuss model order reduction (MOR) of linear switched systems (LSS). Such a system switches among a finite number of linear subsystems (or modes). For the clarity of exposition, we analyze solely the case with two modes. First contribution is generalizing the Loewner framework to the class of LSS. We devise a moment matching method that constructs reduced order models whose generalized transfer functions match those of the original system on a grid of selected frequencies, or interpolation points. This requires the appropriate definition of transfer functions for LSS. The next method represents a novel approach towards balanced reduction for linear switched systems. We compute the controllability and observability gramians  $P$  and  $Q$  corresponding to each active mode by solving systems of coupled generalized Sylvester equations. As for bilinear systems, there are infinitely many time domain Volterra kernels that are used to define an  $H_2$ -type of norm. We show how to directly evaluate this norm using the system matrices and the gramians. Finally, the third method is an adaptation of the BIRKA method for LSS. The main goal is constructing locally optimal reduced order models for which the approximation error is minimal w.r.t. the new introduced norm. We provide a recursive algorithm (named SIRKA) that, upon convergence, produces the desired reduced system. We illustrate the practical applicability of the proposed methods by means of several numerical examples.

Ion Victor Gosea

Jacobs University Bremen  
i.gosea@jacobs-university.de

Athanasios C. Antoulas  
Dept. of Elec. and Comp. Eng.  
Rice University  
aca@rice.edu

### MS287

#### $\mathcal{H}_2$ -Optimal Model Approximation by Structured Time-Delay Reduced Order Models

Dynamical systems are the basic framework used for modeling, controlling and analyzing a large variety of systems and phenomena. Due to the increasing use of dedicated computer-based modeling design software, numerical simulation turns to be more and more used to simulate a complex system or phenomenon and shorten both development time and cost. However, the need of an enhanced model accuracy inevitably leads to an increasing number of variables and resources to manage at the price of a high numerical cost. This counterpart is the justification for model reduction and several approaches have been effectively developed since the 60's. Among these, interpolation-based methods stand out due to their flexibility and low computational cost, making them a predestined candidate in the reduction of truly large-scale systems. In this contribution, we address the model reduction problem for reduced order models with delay-structure. Some recent developments on the  $\mathcal{H}_2$  model reduction problem and interpolation based framework are presented for the case of Time-Delay Systems (TDS). Firstly, some new results on  $\mathcal{H}_2$ -optimal reduction for input/output time-delay models

will be presented as interpolation conditions. Secondly, the  $\mathcal{H}_2$  model reduction problem will be addressed in the case the reduced-order model is a linear combination of simple coprime structures. All theoretical results are illustrated with some numerical examples.

Igor Pontes Duff, Charles Poussot-Vassal, Cédric Seren  
Onera - Toulouse

igorp89@hotmail.com, charles.poussot-vassal@onera.fr,  
cedric.seren@onera.fr

### MS288

#### Towards an Exascale Hyperbolic PDE Engine: High-Order ADER-DG on Tree-Structured Cartesian Meshes

The ExaHyPE project develops an engine for the solution of hyperbolic PDEs at exascale. Our approach is based on a space-tree discretization of the computational domain, a numerical scheme with high-order Discontinuous Galerkin space time discretization and a posteriori subcell limiting. This setup is envisaged to combine highest numerical accuracy with the exploitation of massive parallelism and an intense level of adaptive mesh refinement. The ultimate applications of the ExaHyPE engine will be grand challenge simulations stemming from the fields of astro-physics and geo-sciences. In this talk, we present the current status of the ExaHyPE project in general and in particular the parallel adaptive mesh refinement capabilities of the engine.

Vasco Varduhn  
TU München  
varduhn@tum.de

Jean-Matthieu Gallard  
Technical University of Munich  
jean-matthieu.gallard@tum.de

Dominic E. Charrier, Tobias Weinzierl  
Durham University  
dominic.e.charrier@durham.ac.uk,  
tobias.weinzierl@durham.ac.uk

Michael Bader

Technical University of Munich  
bader@in.tum.de

### MS288

#### Adaptive Methods in Uintah for Mesh Refinement and Resilience in Time-Dependent PDEs

Two non-standard approaches involving the use of adaptive methods are considered in the Uintah software framework. Modeling thermal radiation is both important for real applications and computationally challenging in parallel due to its all-to-all physical and resulting computational connectivity. As a direct all-to-all treatment of radiation is prohibitively expensive on petascale computers, We show that scalability can be achieved through a combination of reverse Monte Carlo ray tracing (RMCRT) techniques combined with AMR, to reduce the amount of global communication leading to excellent strong scaling results to 16384 GPUs on Titan. An AMR approach to resilience is considered using task replication at a coarse level as a task may be replicated on a coarser mesh at 1/8 of the cost of its parent. If the parent task fails then interpolation is used to reconstruct the parent task data. This requires high-accuracy interpolation routines that preserve the physical

characteristics of the underlying solution, such as positivity without introducing errors that persist. Experiments using Uintah illustrate the success of these two approaches.

Martin Berzins, Alan Humphrey  
Scientific Computing and Imaging Institute  
University of Utah  
mb@sci.utah.edu, ahumphrey@sci.utah.edu

Todd Harman  
Department of Mechanical Engineering  
University of Utah  
t.harman@utah.edu

Aditya Pakki, Sahithi Chaganti, Damodar Sahasrabudhe  
Scientific Computing and Imaging Institute  
pakkiadi@sci.utah.edu, sahithe@cs.utah.edu,  
damodars@mail.sci.utah.edu

### MS288

#### The Resiliency of Multilevel Methods on Next Generation Computing Platforms: Probabilistic Model and Its Analysis

With the the advent of exascale computing expected within the next few years, the number of components in a system will continue to grow. The error rate per individual component is unlikely to improve, however, meaning that future high performance computing will be faced with faults occurring at significantly higher rates than present day installations. Therefore, the resilience properties of numerical methods will become important factors in both the choice of algorithm and in its analysis. In this talk we present a framework for the analysis of linear iterative methods in a fault-prone environment. The effects of random node failures and bit flips are taken into account through a probabilistic model involving random diagonal matrices. Using this model, we analyze the behavior of two- and multigrid methods as well as domain decomposition solvers. Both analytic convergence estimates for these methods and simulation results will be discussed.

Mark Ainsworth  
Division of Applied Mathematics  
Brown University  
Mark\_Ainsworth@brown.edu

Christian Glusa  
Brown University  
christian\_glusa@brown.edu

### MS288

#### Multilevel Inference with PETSc

We present an example of the inference of uncertain parameters in multiphysics simulation using PETSc, the Portable Extensible Toolkit for Scientific Computing. For our simulations, we construct high resolution, 3D regional thermal structures for mantle convection simulations using the SubductionGenerator program. In the presented work, we seek to improve the 3D thermal model and test different model geometries and dynamically driven flow fields using constraints from observed seismic velocities and plate motions. We incorporate a variety of data sources, including sea floor ages and geographically referenced 3D slab locations based on seismic observations. Through a formal adjoint analysis, we construct the primal-dual version of the multi-objective PDE-constrained optimization problem for the plate motions and seismic misfit. We have efficient,

scalable preconditioners for both the forward and adjoint problems based upon a block preconditioning strategy. The full optimal control problem is formulated on a nested hierarchy of grids, allowing a nonlinear multigrid method to accelerate the solution.

Matthew G. Knepley  
Rice University  
knepley@gmail.com \*Preferred

Margarete A. Jadamec  
Brown University  
Department of Geological Sciences  
mjadamec@central.uh.edu

Tobin Issac  
University of Chicago  
tisaac@uchicago.edu

### MS289

#### Exploring On-Node Programming Models for Irregular Algorithms

Lately, computing systems that are of heterogeneous and manycore in nature have become mainstream delivering a thousand-fold increase in processing capabilities. Although modern architectures provide several opportunities to achieve good performance, there are still numerous software challenges that need to be addressed. Especially with algorithms that are irregular such as sparse FFT, the software challenges are multifold since the memory access patterns in these kinds of applications are irregular and they remain unknown until runtime and often vary across the computation, dynamically. This talk will present parallelization and optimization techniques used by both high-level directive-based and low-level languages along with runtime transformation techniques to improve cache utilization and increase temporal and spatial locality in order to port irregular applications to heterogeneous platforms.

Sunita Chandrasekaran  
University of Delaware  
schandra@udel.edu

### MS289

#### Dense Linear Algebra at Exascale: Challenges and Opportunities

A very wide range of science and engineering applications depend on the functionality of dense linear algebra libraries. Our current libraries have historically been in high demand because of their inherent properties, namely, high performance, coverage, efficiency, scalability, reliability, portability, and reasonable ease of use. Exascale-ready applications has one clear implication: they must be able to make effective and efficient use of distributed memory at scale. But with the end of Dennard scaling and the consequent rise of processor architectures that are heterogeneous, multicore, bandwidth deficient, and energy constrained, this already formidable challenge has only gotten harder. On all the systems projected along the road to Exascale, the problem of computing with distributed memory has become the problem of computing with heterogeneous distributed memories; executing at scale means the coordinated use of various different kinds of memory, on various different kinds of devices, that have to be accessed through various different hardware/software interfaces. We present techniques to implement a DLA library that integrates seamlessly into the Exascale ecosystem. To

redesign and deliver a state-of-the art implementation of the foundational DLA package that many Exascale applications will require we combine the strengths of existing technologies and engineering them for the complexities of Exascale.

Jakub Kurzak

University of Tennessee, Knoxville  
kurzak@icl.utk.edu

#### MS289

##### **Interoperability, Portability and Productivity Challenges for Transforming Legacy Applications**

We need the programming environments to step up to meet with the rapidly changing hardware architectures. The programming environments is expected to consider a number of factors such as performance, portability, productivity/programmability, interoperability along with correctness and reproducibility and strive to achieve a balance between these factors. To that end, this talk focusses on Legion, a data-centric model that facilitates co-designing programming models with programming systems that includes efficient placement and movement of data. The talk will also discuss how Legion has been evaluated using real applications such as S3D, a production combustion simulation, on Titan and Piz-Daint supercomputers demonstrating better scheduling strategies and scalability analysis.

Patrick McCormick

Los Alamos National Laboratory  
pat@lanl.gov

#### MS289

##### **Preparing Sandia's Application Portfolio for the Future Using Kokkos**

Recognizing the challenges laying ahead with radically changing processor design for High Performance Computing, Sandia started the development of a programming model for performance portability in 2011. The outcome of this effort is Kokkos, a programming model which provides on-node abstractions for complex memory hierarchies and heterogeneous execution resources. Its implementation through C++ template meta-programming is now ready for production use, and Sandia embarked on the journey to port its application portfolio to Kokkos. In this talk the six core abstractions of Kokkos (execution spaces, execution patterns, execution policies, memory spaces, memory layout and memory traits) will be briefly introduced, before discussing their use in applications. Additionally the particular challenges of moving a programming model from research state to production use will be highlighted.

Christian Trott, Daniel Sunderland

Sandia National Laboratories  
crtrott@sandia.gov, dsunder@sandia.gov

Simon D. Hammond

Scalable Computer Architectures  
Sandia National Laboratories  
sdhammo@sandia.gov

H. Carter Edwards

Sandia National Laboratories

hcedwar@sandia.gov

#### MS290

##### **Enabling Low Mach Fluid Simulations Using Trilinos**

In this talk I will provide a high level overview of Sandia National Labs' Trilinos project: current capabilities, application impact, and future directions. Special attention will be given to linear solver capabilities and fundamental underlying libraries. I will then discuss the impact of Trilinos parallel solver technology on simulations in the open source low Mach turbulent flow code Nalu. In particular, I will focus on the solution methods applied to the pressure system. These methods involve Krylov iterative solvers preconditioned by algebraic multigrid (AMG). The AMG itself depends on iterative smoothers, direct solvers, and load-balancing, all supplied by other Trilinos libraries. The talk will conclude with parallel numerical results.

Jonathan J. Hu

Sandia National Laboratories  
Livermore, CA 94551  
jhu@sandia.gov

Andrey Prokopenko, Mark Hoemmen, Paul Lin,  
Christopher Siefert, Christian Trott  
Sandia National Laboratories  
prokopenkoav@ornl.gov, mhoemme@sandia.gov,  
ptlin@sandia.gov, csiefer@sandia.gov, crtrott@sandia.gov

Stefan P. Domino

Sandia National Laboratories  
spdamin@sandia.gov

#### MS290

##### **A Parallel Approach for the Multi-Level HP-Adaptive Finite Cell Method**

The multi-level hp-refinement scheme [ N. Zander *et al.*, Multi-level hp-adaptivity, *Comput. Mech.*, vol. 55, no 3, pp. 499-517, 2015 ] is an extension of the finite element method that allows local mesh adaptation without the trouble of constraining hanging nodes by performing spatial refinement based on superposition. This contribution presents an MPI-parallel implementation of this scheme, adapted to its hierarchical structure. Our implementation utilizes routines from the Trilinos project to achieve better performance for different problem classes. Furthermore, we combine the hp-scheme with the finite cell method [ A. Düster *et al.*, The finite cell method for three-dimensional problems of solid mechanics, *Comput. Method. Appl. M.*, vol. 197, no 4548, pp. 3768-3782, 2008 ], a high-order fictitious domain method, in order to perform computations on domains of complex shape.

John Jomo, Nils Zander, Mohamed Elhaddad, Ali

Oezcan, Stefan Kollmannsberger  
Technical University of Munich  
john.jomo@tum.de, nils.zander@tum.de,  
mohamed.elhaddad@tum.de, ali.oezcan@tum.de,  
kollmannsberger@mytum.de

Ralf-Peter Mundani

TUM, Faculty of Civil, Geo and Environmental  
Engineering  
Chair for Computation in Engineering  
mundani@tum.de

Ernst Rank  
 Computation in Engineering  
 Technische Universität München  
 ernst.rank@tum.de

### MS290

#### Large Scale Computational Inelasticity Using the Trilinos Framework

Several manufacturing and processing applications involve the use of materials beyond their elastic limit, and their simulation requires the development and implementation of complex nonlinear inelastic constitutive models. Within this broad class of problems, we focus on a rate-objective constitutive model that combines J2 plasticity with the Norton creep law for a thermoelastic material. We discretize the resulting equations using a stabilized finite element method in space. We apply the backward Euler method in time, and at every time step solve the resulting set of nonlinear equations through the use of a return-mapping algorithm, where we construct an elastic predictor and an efficient inelastic corrector, which requires the solution to (just) two scalar equations. We have implemented this algorithm within Albany, a multi-physics code that heavily leverages the Trilinos framework. The components that are key to our application are the automatic differentiation capability to calculate the consistent tangent of a complex constitutive model, and the use of Krylov-subspace methods with efficient multigrid preconditioners to solve the resulting system of linear equations. We describe the optimal choices of parameters within these algorithms, and the simulation of large-scale problems with upwards of billions of degrees of freedom. We also discuss the strong parallel scaling of the overall scheme, and other applications and extensions of this approach.

Assad Oberai  
 Department of Mechanical, Aerospace and Nuclear  
 Engineering  
 Rensselaer Polytechnic Institute  
 oberaa@rpi.edu

Zhen Li, Brian Granzow  
 Rensselaer Polytechnic Institute  
 lzsijia@gmail.com, granzb@rpi.edu

Dan A. Ibanez  
 Rensselaer Polytechnic Institute  
 SCOREC  
 ibaned@rpi.edu

Max Bloomfield  
 RPI  
 bloomm2@rpi.edu

Mark S. Shephard  
 Rensselaer Polytechnic Institute  
 Scientific Computation Research Center  
 shephard@rpi.edu

Glen Hansen, Christopher Siefert  
 Sandia National Laboratories  
 gahanse@sandia.gov, csiefer@sandia.gov

### MS290

#### Using Trilinos Components to Construct Reduced Order Models for Wind Turbine Wake Develop-

#### ment from Lidar Measurements

The focus of our research is the real-time characterization of the wake behind wind turbines in the setting of a wind farm. We accomplish this through simulations based on a reduced-order model (ROM) derived from a snapshot matrix. The data that comprises the snapshot matrix comes from a library of simulations. These simulations are solved using a fully coupled, fully implicit, second-order accurate method for the primitive variable form of the Navier-Stokes Equations. This method was formed by Elman, Howle, Shadid, and Tuminaro in their 2003 paper. We use the iterative linear solve package Belos with block preconditioning handled via Playa and multi-level preconditioning handled via ML to implement this solving strategy.

We consider several different approaches for constructing the ROM from the proper-orthogonal decomposition (POD), which we acquire through use of the Anasazi package. These ROM techniques share two main principles in common. First, the methods do not do not require the velocity field to be discretely divergent-free in the resulting discretization of the Navier-Stokes Equation. Second, these methods incorporate modes for the pressure fields into the ROM.

Simon Rush, Kevin Long  
 Texas Tech University  
 simon.rush@ttu.edu, kevin.long@ttu.edu

### MS291

#### Convergence Between Big Data and HPC from the Application Point of View

Convergence between Big Data and HPC is a major concern since many years now and has a strong influence on supercomputer design. But this only a small part of the problem. We have now more and more scientific cases to deal with needing both huge data management (storage, treatment, analysis, distribution ...) and huge computing power. The data are coming both from numerical simulations and experiments and the treatment associated are getting more and more complex from different point of views: worldwide data distribution, volume of data, complexity and cost of data treatment... We will illustrate in this talk some of these problematics using different scientific use cases (astronomy, genomics, climate...) and we will try to identify some major challenges.

Christophe Calvin  
 CEA/DSM  
 christophe.calvin@cea.fr

### MS291

#### Application Development Framework for Many-core Architectures in Post-K/Post-Moore Era

Increasing compute intensity is essential for extracting potential performance of modern HPC systems efficiently. On the contrast, we should pay more attention to data movement intensity on the Post Moore Systems (PMS), where the features of the PMS are (1) high bandwidth in memory and network, (2) large capacity of memory and cache, and (3) large and heterogeneous latency due to deep hierarchy in memory and network. This presentation describes the strategy of research and development of numerical algorithms and methods for applications, such as linear solvers, and automatic-tuning (AT) framework, on Exas-

cale Systems and the PMS. Actually, problems on large and heterogeneous latency due to deep hierarchy in memory and network are common issues for Exascale Systems and the PMS. Therefore, we are focusing on the methods, which are suitable for hiding network latency, such as pipelined algorithms, and parallel-in-space/time (PiST) algorithms. Since 2011, we have been developing ppOpen-HPC which is an open source infrastructure for development and execution of optimized and reliable simulation codes with automatic-tuning (AT) on post-peta-scale (pp) parallel computers based on manycore architectures. Finally, ppOpen-HPC will be extended for the Exascale Systems and the PMS by integrating these newly developed libraries for linear solvers considering hierarchy of network, and new AT framework based on strategy for data movement intensity.

Kengo Nakajima  
The University of Tokyo  
Information Technology Center  
nakajima@cc.u-tokyo.ac.jp

#### MS291

##### **ShyLU: A Collection of Node-Scalable Sparse Linear Solvers**

As node level parallelism increases manifold new ideas are needed for sparse subdomain solvers on the node. These ideas are both at algorithmic level and at the parallel implementation level. This talk will focus on both these aspects. At the algorithmic level we will focus algorithms on asynchronous algorithms and algorithms for two-dimensional data layouts. At the implementation level we will focus on task-based implementations and point-to-point communication between threads. Performance portability of these ideas in various architectures will also be discussed. All the work that is discussed here will be based on the ShyLU package in Trilinos. This talk will also provide an overview of the ShyLU software framework and how it fits into the larger system-level solvers such as domain decomposition or multigrid solvers.

Siva Rajamanickam, Kyungjoo Kim, Andrew Bradley  
Sandia National Laboratories  
srajama@sandia.gov, kyukim@sandia.gov,  
ambradl@sandia.gov

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

#### MS291

##### **CHiS: An Efficient CPU/GPU Direct Linear System Solver for 3D FDFD Photonic Device Analysis**

Finite-difference frequency-domain method is an important tool for photonics analysis such as resonance, sensor, and spectrum-sensitive applications. Discretizing the 3D Maxwell equations results in sparse complex unsymmetric ill-conditioned linear systems. An efficient preconditioner is absent for an iterative solver, and the vector field computing leads to high computation cost and memory requirement for a direct solver. We propose Compressed Hierarchical Schur method (CHiS) to overcome the difficulties. The CHiS concept is derived from domain decomposition and nested dissection, while computation redundancy is identified among subdomains and separators based on the physical problem, discretization, and relations among the

hierarchy. The interfaces between subdomains and separators can be greatly reduced by defining of unit face and edge interfaces. For photonic devices with large homogeneous bulks or periodic structures, a large portion of redundancy can be identified and removed during numerical factorization. By using the CHiS, computation time and memory requirement can be significantly reduced. The efficiency of the CHiS is further enhanced by modern processors and accelerators. The proposed CHiS is tested on multicore-CPU with or without multiple-GPUs for performance analysis. These new techniques can greatly reduce design timeframe in the future development of photonic devices and circuits.

Weichung Wang  
National Taiwan University  
Institute of Applied Mathematical Sciences  
wwang@ntu.edu.tw

Cheng-Han Du  
Department of Mathematics, National Taiwan University  
b92006@csie.ntu.edu.tw

#### MS292

##### **Backward SDE Methods for Nonlinear Filtering Problems**

A nonlinear filtering problem can be classified as a stochastic Bayesian optimization problem of identifying the random state of a stochastic dynamical system given noisy observations of the system. Well known numerical simulation methods include unscented Kalman filters and particle filters. In this talk, we attempt to construct efficient numerical methods using forward backward stochastic differential equations. The backward SDEs for the nonlinear filtering problems are the counter parts of Fokker-Planck equations for SDEs with no observation constraints. In this talk, we will describe the process of deriving such backward SDEs as well as the corresponding high order numerical algorithms for nonlinear filtering problems.

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

#### MS292

##### **Effective Computational Methods for Optimal Control Problems of Random PDEs**

In this talk, we consider an optimal control problem for partial differential equations with random inputs. To determine an applicable deterministic control  $\hat{f}(x)$ , we consider the four cases which we compare for efficiency and feasibility. We prove the existence of optimal states, adjoint states and optimality conditions for each cases. We also derive the optimality systems for the four cases. The optimality system is then discretized by a standard finite element method and sparse grid collocation method for physical space and probability space, respectively. The numerical experiments are performed for their efficiency and feasibility.

Hyung-Chun Lee  
Ajou University  
Department of Mathematics  
hclee@ajou.ac.kr

Max Gunzburger

Florida State University  
Department of Scientific Computing  
mgunzburger@fsu.edu

**MS292****Multi-Level Monte Carlo Finite Element Method for Stochastic Optimal Control Problem**

We consider the implementation of multi-level Monte Carlo finite element method (MLMC-FEM) to an elliptic optimal control problem with uncertain coefficients and its surrogate model. Sample size formulas for each level of MLMC from the perspective of optimization, i.e., minimizing the computational error/cost with given computational cost/error, were derived. A gradient-based optimization algorithm using MLMC-FEM was proposed and compared to the results obtained by classical Monte Carlo method that employs many more degrees of freedom. These comparisons show the effectiveness and feasibility of the use of MLMC for obtaining accurate optimal solutions, which are required to construct statistical moments associated with the quantity of interest (QoI), of the stochastic control problem at low cost.

Ju Ming

Beijing Computational Science Research Center  
jming@csrc.ac.cn

Qiang Du

Columbia University  
Department of Applied Physics and Applied Mathematics  
qd2125@columbia.edu

Qi Sun

Beijing Computational Science Research Center  
sunqi@csrc.ac.cn

**MS292****Inverse Problems Matching Statistical Moments in PDEs with Random Data**

We present a parameter identification problem constrained by PDEs with random input data. Several identification objectives are discussed that either minimize the expectation of a tracking cost functional or minimize the difference of desired statistical quantities in the appropriate  $L^p$  norm. The stochastic parameter identification algorithm integrates an adjoint-based deterministic algorithm with the sparse grid stochastic collocation FEM approach. The proof of the error estimates uses a Fink-Rheinboldt theory for parametrized nonlinear equations.

Catalin S. Trenchea

University of Pittsburgh  
trenchea@pitt.edu

**MS293****Optimal Regularized Inverse Matrices for Inverse Problems**

We consider optimal low-rank regularized inverse matrix approximations and their applications to inverse problems. We give an explicit solution to a generalized rank-constrained regularized inverse approximation problem. We propose an efficient rank-update approach that decomposes the problem into a sequence of smaller rank problems. Using various examples, we demonstrate that more accurate solutions to inverse problems can be achieved by

using rank-updates to existing regularized inverse approximations.

Matthias Chung

Department of Mathematics  
Virginia Tech  
mcchung@vt.edu

**MS293****Efficient UQ Strategies in High-Dimensional Inverse Problems**

It is important to separate the representations of the inverse problem and solution from those of the forward problem and its solution. This prevents accuracy requirements on one from influencing the other and avoids unnecessary mappings across multiple meshes and its associated problems. This in turn provides several advantages for efficient sampling strategies for high dimensional inverse problems.

Eric De Sturler

Virginia Tech  
sturler@vt.edu

Misha E. Kilmer

Mathematics Department  
Tufts University  
misha.kilmer@tufts.edu

Arvind Saibaba

North Carolina State University  
asaibab@ncsu.edu

Eric L. Miller

Tufts University  
elmiller@ece.tufts.edu

**MS293****Uncertainty Quantification in the Accelerated Climate Modeling for Energy**

Climate models contain a large number of empirical parameters that are inherently uncertain. The standard practice of parameter estimation is to sequentially conduct a series of model simulations, subjectively evaluate the models predictive skill after each simulation, and manually adjust the parameter values based on experts interpretation of the models response to parameter changes. Neither the current parameter estimation practice nor the standard UQ methods will be affordable for the climate models that are typically very time-costly. We apply the compressive sensing based techniques to build surrogate model for the climate simulations. Therefore the uncertainty quantification and the parameter estimate of the climate model can be performed based on the surrogate model. Our method requires fewer data than the standard compressive sensing method to reach the same accuracy level as it enhances the sparse of the surrogate model. We will use 100km resolution Accelerated Climate Modeling for Energy (ACME) simulations to demonstrate the efficacy of our method.

Hui Wan

Pacific Northwest National Laboratory  
Richland, WA 99354  
hui.wan@pnnl.gov

Xiu Yang

Pacific Northwest National Laboratory

xiu.yang@pnnl.gov

Alexandre M. Tartakovsky  
PNNL  
alexandre.tartakovsky@pnnl.gov

Yun Qian  
Pacific Northwest National Laboratory  
yun.qian@pnnl.gov

Philip J. Rasch  
Pacific Northwest Nat'l Lab  
philip.rasch@pnnl.gov

### MS293

#### A Data-Driven Approach to Quantify and Reduce Model-Form Uncertainty in Turbulent Flow Simulations

Model-form uncertainty in Reynolds-Averaged Navier-Stokes (RANS) simulations is a major obstacle for predictive applications, especially when they are used to support high-consequence decisions related to turbulent flows. In this work, we present a data-driven, physics-informed Bayesian approach to quantify and reduce the model-form uncertainties in RANS simulations. Perturbations are directly introduced into the RANS modeled Reynolds stresses and are represented with compact parameterization accounting for prior knowledge and physical constraints. These physics-based perturbations are compared to those obtained by random matrix theoretic approach with maximum entropy, which can be utilized to assess the specification of physical priors. With sparse available online data, the model-form uncertainties can be reduced based on an iterative ensemble Kalman method by incorporating the data into the prior. Two representative cases, flow over periodic hills and flow in a square duct, are tested. Simulation results suggest that the propagated mean quantities of interest (e.g., mean velocity, wall shear stresses) from Reynolds stress posterior are significantly improved over the baseline prediction.

Jianxun Wang, Jinlong Wu, Heng Xiao  
Dept. of Aerospace and Ocean Engineering, Virginia Tech  
vtwjx@vt.edu, jinlong@vt.edu, hengxiao@vt.edu

### MS294

#### A Compatible Finite Element Dynamical Core

I will describe our research on numerical methods for atmospheric dynamical cores based on compatible finite element methods. These methods extend the properties of the Arakawa C-grid to finite element methods by using compatible finite element spaces that respect the elementary identities of vector-calculus. These identities are crucial in demonstrating basic stability properties that are necessary to prevent the spurious numerical degradation of geophysical balances that would otherwise make numerical discretisations unusable for weather and climate prediction without the introduction of undesirable numerical dissipation. The extension to finite element methods allow these properties to be enjoyed on non-orthogonal grids, unstructured multiresolution grids, and with higher-order discretisations. In addition to these linear properties, for the shallow water equations, the compatible finite element structure can also be used to build numerical discretisations that respect conservation of energy, potential vorticity and enstrophy; I will survey these properties. We

are currently developing a discretisation of the 3D compressible Euler equations based on this framework in the UK Dynamical Core project (nicknamed "Gung Ho"). The challenge is to design discretisation of the nonlinear operators that remain stable and accurate within the compatible finite element framework. I will survey our progress on this work to date and present some numerical results.

Colin J. Cotter  
Imperial College London  
Department of Mathematics  
colin.cotter@imperial.ac.uk

Jemma Shipton, Hiroe Yamazaki, David Ham  
Department of Mathematics  
Imperial College London  
j.shipton, h.yamazaki@imperial.ac.uk,  
david.ham@imperial.ac.uk

### MS294

#### The Numa Framework for Nonhydrostatic Atmosphere and Ocean Modeling

We give a progress report on the development of the NUMA framework which is a modeling infrastructure for high-order nonhydrostatic modeling of both the atmosphere and ocean. High-order is achieved via the use of spectral element and/or discontinuous Galerkin methods on tensor-product grids. In addition, NUMA is highly scalable and is being retro-fitted to handle non-conforming adaptive grids. NUMA is the original atmospheric component which is used within the U.S Navy's NEPTUNE next-generation weather prediction system. The ocean component is known as NUMO (O for Ocean). Both NUMA and NUMO share the same: 1) high-order numerics, 2) time-integrators, 3) iterative solvers, 3) parallel communication stencil, and 4) grid generation. In this talk, we shall describe the NUMA/NUMO components and show the latest results for both atmospheric and ocean applications in fully three-dimensional domains.

Frank Giraldo  
Naval Postgraduate School  
fxgiral@nps.edu

Michal Kopera  
University of California - Santa Cruz  
makopera@ucsc.edu

Jeremy E. Kozdon  
Department of Applied Mathematics  
Naval Postgraduate School  
jekozdon@nps.edu

Daniel Abdi  
Naval Postgraduate School  
dsabdi@nps.edu

### MS294

#### Multirate Time Stepping for Accelerating Explicit Discontinuous Galerkin Computations with Application to Geophysical Flows

This talk deals with transient adaptive high order methods for the resolution of complex environmental fluid mechanics problems. In a simple presentation, we will present the most important challenges for the spatial and time discretizations on an one-dimensional simplified example, il-

lustrating the fact that complex concepts can be explained with simple equations and simple simulations. Adaptive schemes is a research field that is likely to expand rapidly in the near future. Global spatial and time high order adaptive methods are attractive for ocean modeling because of their accuracy and their geometrical flexibility for representing irregular boundaries and for local mesh refinement. Running globally well-resolved general circulation ocean models for climatic timescales is still an elusive goal. A serious difficulty to reaching this goal is the likely need to fully, or at least partially, resolve relatively small topological and dynamical features as mesoscale eddies, narrow sills, mid-ocean ridge, western boundary current and equatorial circulation. This work is done in the framework of the SLIM project (Second-generation Louvain-la-Neuve Ice-ocean Model), which aims to build a numerical ice-ocean model using adaptive finite elements. Three-dimensional and parallel computations will also be shown in order to demonstrate that we can also handle very large computations :-)

Vincent Legat  
 Université catholique de Louvain  
 Belgium  
 vincent.legat@uclouvain.be

#### MS294

##### **High Level Implementation of Geometric Multigrid Solvers for Finite Element Problems: Applications in Atmospheric Modelling**

In recent years, finite element methods have gained popularity in the numerical weather prediction community due to their ability to deal with non-orthogonal meshes and high order discretisations. Compatible (or mimetic) spaces offer some appealing features for models such as stable discretisations of pressure gradients and geostrophic balance. However, they present challenges in solving the elliptic system resulting from implicit timestepping schemes. The high aspect ratio domains typical in these problems require us to treat the "vertical" direction specially when designing a solver if we want to achieve mesh-independent iteration counts. In this talk I will discuss work in the Firedrake finite element framework on developing custom geometric multigrid solvers for these problems and present an analysis of the resulting performance, solving problems with up to two billion degrees of freedom.

Lawrence Mitchell  
 Department of Computing  
 Imperial College London  
 lawrence.mitchell@imperial.ac.uk

Eike H. Mueller  
 University of Bath  
 e.mueller@bath.ac.uk

#### MS295

##### **Sharp Interface Immersed Boundary Methods for Heart Valve Fluid-Structure Interaction**

The immersed boundary (IB) method is a mathematical formulation and numerical approach to fluid-structure interaction (FSI). The typical IB formulation of FSI considers the case of an elastic structure immersed in a viscous incompressible fluid. The deformations and stresses of the structure are described using Lagrangian variables, and the momentum, viscosity, and incompressibility of the fluid-solid system are described using Eulerian variables. Inte-

gral transforms with delta function kernels connect the Lagrangian and Eulerian frames. Conventional IB methods approximate these Lagrangian-Eulerian interaction equations using regularized delta functions, but doing so limits the formal order of accuracy of the IB method. This talk will describe two approaches to obtaining higher-order accuracy in the IB method for FSI problems involving thick (codimension-0) immersed bodies. In one approach, we use a version of the immersed interface method to treat traction-like forces present at fluid-solid interfaces. In the other, we split discontinuous field variables into the sum of several smooth fields, which allows us to use standard discretization technology while still sharply resolving flow features at fluid-solid interfaces. We also apply these methods to model cardiovascular FSI, including in models of prosthetic heart valves.

Boyce E. Griffith  
 Departments of Mathematics and Biomedical Engineering  
 University of North Carolina at Chapel Hill  
 boyceg@email.unc.edu

Amneet P.S. Bhalla, Ebrahim Kolahdouz  
 University of North Carolina at Chapel Hill  
 mail2amneet@gmail.com, ebrahimk@email.unc.edu

#### MS295

##### **Confining Active Matter**

Abstract Not Available At Time Of Publication.

David Saintillan  
 UCSD  
 dsaintillan@ucsd.edu

#### MS295

##### **Soft Matter and Fluids: An Overview and Challenges**

The modeling of soft matter system in fluids is a rapidly evolving area, with the development of new modeling techniques and the application of existing techniques to novel systems. The purpose of this talk is to introduce the typical physical systems under consideration, in addition to providing an overview of the various modeling methods from different researchers.

David Salac  
 University at Buffalo - SUNY  
 davidsal@buffalo.edu

#### MS295

##### **The Role of Cytoplasmic Rheology in Blebbing Dynamics**

Blebs are pressure-driven protrusions that play an important role in cell migration, particularly in 3D environments. A bleb is initiated when the cytoskeleton detaches from the cell membrane, resulting in the pressure-driven flow of cytosol towards the area of detachment and local expansion of the cell membrane. Recent experiments involving blebbing cells have led to conflicting hypotheses regarding intracellular pressure dynamics. A dynamic computational model of the cell is presented to simulate these experiments, and results show that complex rheology of cytoplasm is essential to explain experimental observations. The model is also used to quantify intracellular pressure dynamics dur-



ing different mechanisms of bleb initiation.

Wanda Strychalski  
Department of Mathematics, Applied Mathematics, &  
Statistics  
Case Western Reserve University  
wis6@case.edu

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

#### MS296

##### **A Framework for Data Assimilation and Identifiability Issues in Computational Hemodynamics**

Three-dimensional models for hemodynamics have been developed for more than twenty years. Tremendous progress has been made: more and more scales and coupled physical phenomena are taken into account, including control and regulation mechanisms. But to make them useful to clinical trials or surgical planning, these models have to be patient-specific. This talk will be devoted to a general framework to assimilate clinical measurements into hemodynamics models. In this approach, state and parameter estimation is performed by nonlinear filtering in lumped parameter models, which act as reduced order surrogates. A special emphasis is put on the issue of identifiability, which is addressed with a generalized sensitivity analysis. The approach is applied to a patient specific aortic coarctation case. Less than ten percents errors between the clinically measured quantities and three-dimensional simulation results for rest and stress are obtained. Lastly, it is demonstrated that the proposed approach is capable of dealing with a wide variety of measurements and cases where the pressure and flow clinical measurements are not taken simultaneously. Another application to liver surgery simulation will be also presented.

Jean-Frederic Gerbeau  
INRIA Paris-Rocquencourt  
jean-frederic.gerbeau@inria.fr

#### MS296

##### **Fluid-Structure Interaction in Abdominal Aortic Aneurysms: Multiscale Modeling of Tissue Mechanics and a Novel Wall Shear Stress Risk Assessment**

Fluid-structure interaction (FSI) in arterial vessels is modeled by considering an anisotropic non-linear elastic tissue response. Tissue constitutive description accounts for arterial wall features, in which collagen hierarchical organization is modeled by integrating histological and chemico-physical data within a multiscale structurally-motivated approach [D. Bianchi et al. 2016, An integrated computational approach for aortic mechanics including geometric, histological and chemico-physical data, *J Biomech* S0021-9290:30136]. Pulsatile blood flow is described via the Arbitrary-Lagrangian-Eulerian method and by involving a 3D-0D coupling for treating adsorbing boundary conditions on the vessel segment. Nonlinear finite-element schemes are adopted to solve the 3D multiscale FSI problem and thrombus deposition risk assessment is conducted via wall shear stress analysis. Classical and novel risk indicators based on the Three-Band Decomposition method [Nestola et al. 2016 Three-band decomposition analysis in multiscale FSI models of abdominal aortic aneurysms Int

*J Mod Phys C* 27:1650017] are compared and contrasted with the expected clinical evidence. The proposed computational framework allows to identify accurate stress and strain patterns at the tissue level and depending on tissue microstructure, opening to clarify the etiology of some cardiovascular diseases and physiopathological remodeling mechanisms.

Alessio Gizzi  
University Campus Bio-medico of Rome  
a.gizzi@unicampus.it

Daniele Bianchi, Elisabetta Monaldo  
University of Rome Tor Vergata  
daniele.bianchi@uniroma2.it, elisabettamonaldo@libero.it

Michele Marino  
Leibniz Universität Hannover  
marino@ikm.uni-hannover.de

Giuseppe Vairo  
University of Rome Tor Vergata  
vairo@ing.uniroma2.it

Simonetta Filippi  
University Campus Bio-Medico of Rome, Italy  
ICRA University of Rome, La Sapienza, Italy  
s.filippi@unicampus.it

#### MS296

##### **A Comprehensive Framework for Thoracic Aortic Endograft Simulations: from Virtual Deployment to Computational Fluid-Dynamics**

Thoracic endovascular aortic repair (TEVAR) is nowadays a consolidated alternative to open surgery to treat aortic diseases such as aneurysms and dissections. A correct pre-operative planning is essential for positive long-term outcomes of the procedure and to avoid drawbacks such as endoleaks. Diseased thoracic aorta is characterized by complex and tortuous geometries with high inter-individual variability. In particular, angulated arches are more prone to suffer post-operative complications. Thus, even the selection of simple geometrical variables of the endograft such as length and diameter could be cumbersome. Motivated by such considerations, in this work we propose a framework to simulate the stent deployment within an aorta creating then the computational mesh to perform fluid-dynamics simulations. The final result is a well fitted geometry representing the fluid dynamics scenario as it would be after the surgery. Prospective endograft apposition, blood dynamics and drag force are compared with the real post operative simulation showing good agreement. We conclude that this framework is a key step to simulate TEVAR in order to predict stent malappositions and to evaluate post-surgery complications.

Rodrigo M. Romarowski, Elena Faggiano, Michele Conti, Simone Morganti  
University of Pavia  
rodrigo.romarowski@unipv.it, elena.faggiano@unipv.it, michele.conti@unipv.it, simone.morganti@unipv.it

Santi Trimarchi  
IRCCS Policlinico San Donato  
santi.trimarchi@grupposandonato.it

Ferdinando Auricchio  
DICAr, University of Pavia

via Ferrata 3, 27100 Pavia, Italy  
auricchio@unipv.it

**MS296****Numerical Assessment of Hemodynamics in Retro-Fontan Connections**

The Fontan operation is the third and final step of a multi-stage surgery designed to separate systemic and pulmonary circulations in patients born with single-ventricle heart defects. The Fontan operation results in a surgical anastomosis called the total cavopulmonary connection (TCPC) that allows deoxygenated blood to bypass the heart and continue to the lungs. A surgeon can construct variations of the TCPC depending on the location of the heart and other organs. This study examines a cohort of ten patients with retro-Fontan connections. These connections include an extended conduit routing more laterally and anteriorly than a traditional extra-cardiac baffle. This unique approach provides the means for patients with anatomical abnormalities of the heart and visceral organs to undergo these life-saving surgeries, despite obstructions for the traditional lateral tunnel or extra-cardiac conduit. This study evaluates the hemodynamics of the retro-Fontan connection using computational fluid dynamics to assess hepatic blood flow distribution and energy losses in the TCPC. Alternative hypothetical traditional Fontan surgical options will be modeled on the same patient group to compare the performance of the retro-Fontan option. This study will provide insight into the effectiveness of this retro-Fontan connection as an alternative to other surgical means.

Camille Johnson, Zhenglun Wei  
Department of Biomedical Engineering  
Georgia Institute of Technology  
cjohnson96@gatech.edu, alan.wei@bme.gatech.edu

Ritchie Sharon, Mahesh Kappanayil  
Amrita Institute of Medical Sciences and Research Centre  
Department of Pediatric Cardiology  
maheshpeds@yahoo.co.in, maheshpeds@yahoo.co.in

Ajit P. Yoganathan  
Department of Biomedical Engineering  
Georgia Institute of Technology  
ajit.yoganathan@bme.gatech.edu

**MS297****Phase Transition in a Model for Gang Territorial Development**

One of the most important mechanisms for gang territorial development involves putting down graffiti to give the gang a claim to the marked area, a job which is often relegated to the youngest members of the gang. Here, we examine an agent-based model for two gangs, where each agent puts down graffiti markings and moves on the lattice to preferentially avoid areas marked by the other gang. A similar spin-model was shown to have a phase transition by Barbaro, Chayes, and D'Orsogna, using methods from statistical physics. Here, we take an agent-based approach and observe that the systems undergoes a phase transition as the intensity of the avoidance of the other gang's graffiti is varied. We then derive a system of macroscopic PDEs from the model and examine the stability of its stationary states.

Alethea Barbaro  
Dept. of Mathematics, Applied Mathematics & Statistics

Case Western Reserve University  
alethea.barbaro@case.edu

Abdulaziz Alsenafi  
Case Western Reserve University  
asa83@case.edu

**MS297****A Transport Model of Data Flows in HPC Environments**

Scientific computing in high performance computing (HPC) environments is characterized by distributing over many processors a computational task which is then performed by iteratively performing some subtasks on each processor and then communicating relevant information across the network of processors. Our goal is to develop a quantitative predictive model for the performance of computations performed on large-scale systems. To this end, we describe a discrete model of inter-processor communications. We then study the asymptotic limit of this discrete model; this results in a multi-dimensional transport equation presenting several theoretical and numerical challenges which we discuss.

Richard C. Barnard  
RWTH Aachen University  
Center of Computational Engineering Sciences  
barnardrc@ornl.gov

Cory Hauck  
Oak Ridge National Laboratory  
hauckc@ornl.gov

**MS297****Inhomogeneous Boltzmann-Type Equations Modelling Opinion Leadership and Political Segregation**

In recent years different kinetic models to describe the process of opinion formation have been proposed. Such models successfully use mathematical tools from statistical mechanics to describe the behaviour of a large number of interacting individuals in a society. This leads to generalizations of the classical Boltzmann equation for gas dynamics. Most approaches in the literature assume a homogeneous society. To model additional sociologic effects in real societies, e.g. the influence opinion leaders, one needs to consider inhomogeneous models. In this talk we discuss such kinetic models for opinion formation, where the opinion formation process depends on an additional independent variable, e.g. a leadership or a spatial variable. More specifically, we consider: (i) opinion dynamics under the effect of opinion leadership, where each individual is characterised not only by its opinion, but also by another independent variable which quantifies leadership qualities; (ii) opinion dynamics modelling political segregation in 'The Big Sort', a phenomenon that US citizens increasingly prefer to live in neighbourhoods with politically like-minded individuals. Based on microscopic opinion consensus dynamics such models lead to inhomogeneous Boltzmann-type equations for the opinion distribution. We derive macroscopic Fokker-Planck-type equations in a quasi-invariant opinion limit and present results of numerical experiments.

Bertram Düring  
University of Sussex

b.during@sussex.ac.uk

### MS297

#### Uncertainty Quantification for Kinetic Equations in Socio-Economic Sciences

In this talk we will focus on numerical methods for kinetic equations of collective behavior with random input. These models play a relevant role in for the description of complex systems in social sciences, biology and robotics. In such cases uncertainty exists in most modeling parameters. The evolution of mesoscopic quantities for the mentioned phenomena is usually given in terms of the so called aggregation diffusion equations, for which steady state solutions are known in some special cases. We design steady-state preserving numerical methods in the UQ setting.

Mattia Zanella

University of Ferrara  
znlmtt@unife.it

### MS298

#### Geometry-Dependent Viscosity Reduction in Sheared Active Fluids

We investigate flow pattern formation and viscosity reduction mechanisms in active fluids by studying a generalized Navier-Stokes model that captures the experimentally observed bulk vortex dynamics in microbial suspensions. We present exact analytical solutions including stress-free vortex lattices and introduce a computational framework that allows the efficient treatment of previously intractable higher-order shear boundary conditions. Large-scale parameter scans identify the conditions for spontaneous flow symmetry breaking, geometry-dependent viscosity reduction and negative-viscosity states amenable to energy harvesting in confined suspensions. The theory uses only generic assumptions about the symmetries and long-wavelength structure of active stress tensors, suggesting that inviscid phases may be achievable in a broad class of non-equilibrium fluids by tuning confinement geometry and pattern scale selection.

Jorn Dunkel

Mathematics, MIT  
dunkel@math.mit.edu

Jonasz Slomka

MIT  
jslomka@mit.edu

### MS298

#### Structure and Dynamics of Dense Active Suspension Under Confinement

We analyze what is perhaps the simplest active fluid under rigid confinement with complex dynamics: a suspension of non-motile, but mobile, “extensor” rods that exert active dipolar stresses on a fluid. We first describe the system through a kinetic theory based on microscopic modeling. The stresses produced by particle activity produces long-ranged hydrodynamic coupling, and for extensors can lead to complex time-dependent flows and, depending upon flow geometry, to a form of singularity dynamics through disclination defect pair production, propagation, and annihilation. We then study various useful closures of the kinetic theory, particularly the “Q-tensor” Bingham closure that has been used to study suspensions of passive microscopic

rods. We demonstrate analytically and numerically that the Bingham closure gives a very good statistical accounting, at much reduced computational cost, of the complex dynamics particularly when alignment forces are predominant.

Tong Gao

Michigan State University  
gaotong@egr.msu.edu

### MS298

#### Individual and Collective Dynamics of Bacteria and Micro-Algae in Confinement

Solid boundaries alter both motion and spatial distribution of microorganisms in ways that are currently not completely understood. We present novel micro-swimmer models and simulations able to display correct features seen in experiments such as bacteria circling near surfaces or micro-algae scattering from them. For pushers like bacteria we show that the correct flow singularity is more complex than a force dipole. For bi-flagellates like micro-algae we show that their behavior at surfaces results from a nuanced interplay of flagellar contact, hydrodynamics, noise and cell spinning, with the swimmer geometry being a crucial component. Our results compare well with the most recent experimental data and suggest ways of designing multi-swimmer simulations that capture the correct physics.

Enkeleida Lushi

School of Engineering, Brown U.  
lushi@cims.nyu.edu

### MS298

#### Transitions to Spontaneous Flows in Confined Active Suspensions

Recent experimental studies have shown that confinement can profoundly affect self-organization in semi-dilute active suspensions, leading to striking features such as the formation of steady and spontaneous vortices in circular domains and the emergence of unidirectional pumping motions in periodic racetrack geometries. Motivated by these findings, we analyze the two-dimensional dynamics in confined suspensions of active self-propelled swimmers using a mean-field kinetic theory where conservation equations for the particle configurations are coupled to the forced Navier-Stokes equations for the self-generated fluid flow. In circular domains, a systematic exploration of the parameter space casts light on three distinct states: equilibrium with no flow, stable vortex, and chaotic motion, and the transitions between these are explained and predicted quantitatively using a linearized theory. In periodic racetracks, similar transitions from equilibrium to net pumping to traveling waves to chaos are observed in agreement with experimental observations and are also explained theoretically. Our results underscore the subtle effects of geometry on the morphology and dynamics of emerging patterns in active suspensions and pave the way for the control of active collective motion in microfluidic devices.

David Saintillan

Mechanical and Aerospace Engineering  
UCSD  
dsaintillan@eng.ucsd.edu

### MS299

#### Adjoint-Based Methods for Design of Optimal Ex-

## periments for PDE Models

Mathematical models are of great importance for natural sciences and engineering. Besides providing scientific insight into processes, mathematical models are fundamental for process simulation, optimization and control. However, the results from simulation and optimization can be only trusted as a basis for decision and control if an underlying model describes a given process quantitatively and qualitatively correctly. This implies a model which is validated by experimental data with sufficiently good estimates for model parameters. The development and quantitative validation of complex nonlinear differential equation models is a difficult task that requires the support by numerical methods for sensitivity analysis, parameter estimation, and the optimal design of experiments. The talk will address new developments in optimization methods for validation of models, in particular design of optimal experiments for PDE models based on the adjoint approach for the computation of necessary reduced gradients and Hessians.

Ekaterina Kostina, Gregor Kriwet  
Heidelberg University  
ekaterina.kostina@iwr.uni-heidelberg.de,  
gregor.kriwet@iwr.uni-heidelberg.de

## MS299

### Scalable Methods for Optimal Experimental Design for Inverse Scattering

We consider the problem of finding the sensor locations that maximize expected information gain in the recovered medium in Bayesian inverse problems governed by PDEs. The inverse problem is cast in the framework of Bayesian inference, the solution of which is given by the posterior probability density. The optimization criterion of maximizing expected information gain is expressed by maximizing the Kullback-Leibler divergence from the prior to the posterior. Computing this KL divergence in high parameter dimensions and for expensive forward models is prohibitive. Thus we turn to approximation of the integrand via a quadratic Taylor series, which is made tractable by fast randomized linear algebra methods. This leads to large-scale optimization problems governed by multiple PDEs. Examples are given for optimizing receiver locations for inverse scattering problems governed by Helmholtz forward models.

Jeonghun Lee  
Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
jeonghun@ices.utexas.edu

Umberto Villa  
University of Texas at Austin  
uvilla@ices.utexas.edu

Peng Chen  
UT Austin  
peng@ices.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

## MS299

### Approximate Experimental Design in High Dimen-

## sional Spaces

Calibration and validation of models are inherently data-driven processes. In order to maximize the amount of information provided by experimental data, data collection needs to be judiciously performed, especially for those with time consuming and costly operations. However, experimental design in high dimensional spaces is always challenging as directly comparing Cartesian combinations is usually impractical. In this work we propose an approximate Bayesian experimental design strategy which is shown to be effective working in high dimensional spaces.

Xiao Lin  
University of South Carolina  
University of South Carolina  
lin65@email.sc.edu

Gabriel Terejanu  
University of South Carolina  
terejanu@cse.sc.edu

## MS299

### Challenges in Computing Bayesian Designs for Complex Models

The estimation of empirical and physical models is often performed using data collected via experimentation. The design of the experiment is crucial in determining the quality of the results. For complex models, an optimal design often depends on features, particularly model parameters, which are uncertain prior to experimentation. This dependence leads naturally to a Bayesian approach which can (a) make use of any prior information on these features, and (b) be tailored to the reduction of posterior uncertainty. Bayesian design for most realistic models is complicated by the need to approximate an analytically intractable expected utility; for example, the expected gain in Shannon information from the prior to posterior distribution. For models which are nonlinear in the uncertain parameters, this expected gain must be approximated numerically. The standard approach employs “double-loop” Monte Carlo integration using nested sampling from the prior distribution. Although this method is easy to implement, it produces biased approximations and is computationally expensive. In this talk, we will describe, assess and compare some recent alternatives to simple Monte Carlo sampling from the prior for the approximation of expected utilities. The presented methods include combinations of features from importance sampling and Laplace approximations. Assessments will include both computational cost and the statistical qualities of the resulting approximations.

David Woods, Yiolanda Englezou  
University of Southampton  
d.woods@southampton.ac.uk, ye1g14@soton.ac.uk

Tim Waite  
University of Manchester  
timothy.waite@manchester.ac.uk

## MS300

### Machine Learning Methods for Pairwise Scoring of Protein Models

Generating accurate three dimensional conformations of proteins from an amino acid sequence has been widely studied and has many applications. Recent advances in protein conformational pool generation have allowed for greater

focus in finding a satisfactory ranking of structure quality. The goal is to generate an accurate and computationally tractable global energy function for conformations. Previous attempts in the area have utilized GDT-TS - a recursive sub-structure overlap algorithm that computes similarity to the native conformation - as the response variable for training a variety of machine learning models. Unfortunately, this approach does not take into account the relative nature of the scoring function for each protein. The predictor variables used are protein features based on physics and statistical calculations, possible variations in the importance of each predictor variable for different types of proteins may also decrease previous models accuracy. Since GDT-TS is relative within a single protein, our work tries to identify a way around this by utilizing pairwise comparisons to train multiple scoring functions which are combined based on similarity to the test protein. The new model was tested extensively and is competitive with the current state-of-the-art, a result that may be further improved by investigating alternative similarity and ensemble functions.

Conrad Czejdo  
University of North Carolina, Chapel Hill  
conradcze@gmail.com

Silvia N. Crivelli  
Lawrence Berkeley National Laboratory  
University of California, Davis  
sncrivelli@lbl.gov

### MS300

#### Protein Structure Scoring Using Kernel Ridge Regression and Support Vector Machine

The human body is composed of proteins whose functions are determined by their 3D structure. It is very complex and expensive to experimentally determine protein structures; therefore computational methods have become highly necessary. Computational methods generate a large number of models from a given sequence of amino acids and then select the models that most closely resemble protein structures found in nature. Machine learning algorithms offer great potential to quickly identify the best models and thus aide biologists in designing new drugs and experiments. The objective of this summer research was to implement two different mathematical algorithms, kernel ridge regression (KRR) and support vector machine (SVM) into a machine learning environment to accurately select the best predicted protein structures from a large set of protein models. This study comparatively examined the results of KRR and SVM with each other and with other methods conducted during previous Critical Assessment of protein Structure Prediction (CASP) experiments, which gauges progress in the field. The goal was to minimize the loss and error metrics by exploring these computational techniques. This was done by manipulating the parameters in both algorithms and employing different feature selection approaches. This talk will discuss the performance achieved by these procedures and compare them to those presented in the literature.

Itzhel Dimas, Austin Chung, Shokoufeh Mirzaei  
Cal Poly Pomona  
itzhelndimas@gmail.com, aus71nchung@gmail.com, shokoufeh.mirzaei@gmail.com

### MS300

#### Experiments in Protein Folding Research Using a

#### Highly Dynamic and Extendable Computational Framework

Extensive research aims to discover how amino acids form a given protein. The experimental process to determine these structures is expensive and time-consuming. Computer modeling could reduce all that but in the modelling process many wrong models are produced. Consequently, the best models must be identified. Numerous teams around the globe are using various features, machine learning (ML) algorithms, and hardware to get the most protein-like models. The goal of this project is to create a computational research framework to accelerate such experiments. It requires generalization of the use of ML in protein model scoring and creation of typical research workflows. We identified main workflow components, designed the supporting software and integrated it with the existing software libraries. Currently our framework supports workflows for three fundamental ML algorithms: Support Vector Machines, Random Forest and Neural Networks. Each workflow has a sequence of processing steps or components. Each component can be executed separately and includes required software, sample datasets (training data), and description of component properties: function (ML algorithm), hardware requirements (GPU/supercomputer) and time limitations. Using this framework we were able to quickly duplicate the results of leading researchers in scoring protein models. The framework is expandable and we are currently testing it at Fayetteville State University by involving a group of student researchers.

James R. Inscoe, Bogdan Czejdo  
Fayetteville State University  
skywise314@gmail.com, bczejdo@uncfsu.edu

Silvia N. Crivelli  
Lawrence Berkeley National Laboratory  
University of California, Davis  
sncrivelli@lbl.gov

### MS301

#### Performance of Scalable AMG-based Preconditioners for MHD and Multifluid Plasma Simulations

Scientific understanding through computational simulation of problems governed by the resistive magnetohydrodynamics (MHD) equations often requires high resolution simulations. Fully-coupled Newton-Krylov approaches can be advantageous because of their robustness for complex multiphysics problems. Block preconditioners can often be employed for situations where fully-coupled approaches cannot be employed. Both approaches require the scalable solution of very large sparse linear systems, either for the entire system or for sub-block systems. One approach that offers the potential of scalable solutions is a multi-level or multigrid approach. The smoothers can have a significant impact on the performance and robustness of a multigrid preconditioner. We examine the performance of various smoothers for our fully-coupled algebraic multigrid preconditioned Newton-Krylov solution approach for a finite element variational multiscale turbulence model for incompressible MHD. Our focus will be on large-scale, transient MHD simulations on unstructured meshes. We present scaling results for resistive MHD test cases, including results for over 500,000 cores on an IBM BG/Q.

Paul Lin  
Sandia National Laboratories  
ptlin@sandia.gov

John Shadid  
Sandia National Laboratories  
Albuquerque, NM  
jshadi@sandia.gov

Edward G. Phillips  
Sandia National Laboratories  
egphill@sandia.gov

Jonathan J. Hu  
Sandia National Laboratories  
Livermore, CA 94551  
jhu@sandia.gov

Andrey Prokopenko  
Sandia National Laboratories  
prokopenkoav@ornl.gov

Paul Tsuji  
Lawrence Livermore National Laboratory  
tsuji1@llnl.gov

Eric C. Cyr  
Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

Roger Pawlowski  
Multiphysics Simulation Technologies Dept.  
Sandia National Laboratories  
rppawlo@sandia.gov

### MS301

#### A Highly Scalable Implementation of the BDDC Preconditioner

Paving the road to exascale will likely imply reformulating science problems and redesigning or reinventing their solution algorithms and software for exascale systems in order to face challenges such as energy consumption, resilience, and extreme concurrency. With the extreme concurrency challenge in mind, in this talk we will summarize our recent efforts towards fast sparse linear solver codes tailored for the Finite Element approximation of elliptic PDEs (e.g., Laplacian or Linear Elasticity equations). At the kernel of these codes lies the so-called Balancing Domain Decomposition by Constraints (BDDC) preconditioner. In the first part of the talk, we will introduce this preconditioner and some of the salient properties that make it highly suitable for extreme scale solver design. Then, in the second part of the talk, we will cover how we progressively took profit of these properties to end up with a fully-distributed, asynchronous, communicator-aware, recursive and interlevel-overlapped MPI-parallel implementation of the Multilevel BDDC (MLBDDC) preconditioner. A comprehensive set of numerical experiments reveals remarkable weak scalability up to 458,752 IBM BG/Q cores and 1.8M MPI tasks (full JUQUEEN) for the 3/4-level BDDC preconditioner in the solution of 3D Laplacian and Linear Elasticity discrete problems with up to several dozens of billions of unknowns.

#### Alberto Martin

International Center for Numerical Methods in Engineering,  
Universitat Politècnica de Catalunya, Barcelona, Spain  
amartin@cimne.upc.edu

Santiago Badia, Javier Principe  
International Center for Numerical Methods in Engineering  
Universitat Politècnica de Catalunya, Barcelona, Spain  
sbadia@cimne.upc.edu, principe@cimne.upc.edu

### MS301

#### Scalable Spectral-Geometric-Algebraic Multigrid and Schur Complement Preconditioning for Non-linear, Multiscale, Heterogeneous Flow in Earth's Mantle

We target the simulation of earth's instantaneous mantle convection at high-resolutions and at global scales. While being a fundamental geophysical process, enormous knowledge gaps about mantle convection remain. A reason for the gaps is that realistic mantle models pose computational challenges due to highly nonlinear rheologies, severe heterogeneities, anisotropies, and wide ranges of spatial scales. We present new advances in implicit solvers that include: heterogeneity-robust Schur complement preconditioning (weighted BFBT) and hybrid spectral-geometric-algebraic multigrid methods (HMG). These methods operate within an inexact Newton-Krylov method, on aggressively adapted meshes and mixed continuous-discontinuous discretizations with high-order accuracy. The implicit solver maximizes accuracy and minimizes runtime, while exhibiting optimal algorithmic performance for mantle flow problems. These solver features, however, present enormous challenges for extreme scalability. But due to the mentioned HMG and weighted BFBT methods, which constitute the core of the solver, we demonstrate that algorithmic optimality can be combined with parallel scalability up to 1.6 million cores. This capability for extreme scaling allows the solution of severely nonlinear, ill-conditioned, heterogeneous, and anisotropic PDEs, such as stemming from earth's mantle convection.

#### Johann Rudi

Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
johann@ices.utexas.edu

Georg Stadler  
Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

A. Cristiano I. Malossi  
IBM Research - Zurich  
acm@zurich.ibm.com

Tobin Isaac  
The University of Chicago  
tisaac@ices.utexas.edu

Michael Gurnis  
Caltech  
gurnis@caltech.edu

Peter W. J. Staar, Yves Ineichen, Costas Bekas  
IBM Research - Zurich  
taa@zurich.ibm.com, zrlyin@ch.ibm.com,  
BEK@zurich.ibm.com

Alessandro Curioni  
IBM Research-Zurich  
cur@zurich.ibm.com

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

### MS301

#### A Performance Study of the Use of GPUs in Hypre

The hypre software library provides parallel high performance preconditioners and solvers for the solution of large sparse linear systems, with focus on both structured and unstructured algebraic multigrid methods. While it has shown good performance on many massively performance computers, it has not been able to take advantage of GPUs, however work is underway to change this. In this presentation, we describe our efforts to prepare hypre for emerging architectures with both CPUs and GPUs.

Ulrike Meier Yang  
Lawrence Livermore National Laboratory  
yang11@llnl.gov

### MS302

#### A Multiscale Mimetic Method for Maxwell's Equations

Abstract Not Available At Time Of Publication.

Eldad Haber  
Department of Mathematics  
The University of British Columbia  
haber@math.ubc.ca

### MS302

#### A General, Efficient and Self-Correcting Enrichment Strategy for Multiscale Methods in Reservoir Simulation

We present a general, efficient and self-correcting approach for enriching the coarse space of Algebraic Multiscale Solver: Enriched Algebraic Multiscale Solver (EAMS). EAMS is motivated by ideas developed to enhance the convergence of algebraic multigrid, where the coarse-space enrichment is adaptive to both the underlying problem characteristics and the solver settings. For a given multiscale-method configuration, EAMS enriches the coarse space with local basis functions specifically aimed at the largest error components in the solution space. For this purpose, the homogeneous version (i.e., with 0 right-hand side) of the discrete error equation is used to identify the solution modes that are missing from the two-level multiscale operator. The identified error modes, which are complex combinations of a spectrum of wave numbers, are then localized (truncated) and added to the prolongation operator. The enrichment process is repeated iteratively until the desired convergence rate is reached. The identification and enrichment processes are algebraic, and they are performed adaptively during the iterative solution process. Using challenging test cases from the literature, we show that EAMS leads to great improvements in the robustness and efficiency of existing state-of-the-art multiscale methods. The algorithm is general, and can be used to enhance the convergence of any existing multiscale implementation.

Abdulrahman M. Manea  
Stanford University  
Energy Resources Engineering Department  
amanaea@stanford.edu

Hadi Hajibeygi  
TU Delft  
h.hajibeygi@tudelft.nl

Panayot Vassilevski  
Lawrence Livermore National Laboratory  
vassilevski1@llnl.gov

Hamdi Tchelepi  
Stanford University  
Energy Resources Engineering Department  
tchelepi@stanford.edu

### MS302

#### PDE-Constrained Optimization with Multiscale Methods

PDE parameter estimation problems arise frequently in many applications such as geophysical and medical imaging. Commonly, the inverse problem is formulated as an optimization problem whose constraints are given by the underlying PDEs. We consider the so-called reduced formulation in which the PDEs are eliminated. The resulting unconstrained optimization problem is computationally expensive to solve because they require solving the underlying PDEs numerous times until the reconstruction of the parameter is sufficiently accurate. We consider the case in which many measurements are available leading to a large number of PDE constraints. To reduce the costs of the PDE solvers, we discuss multiscale reduced order modeling (ROM) scheme that projects the discretized forward problems onto a lower dimensional subspace. This is done through an operator-induced interpolation by solving PDEs locally on coarse cells. By design, this interpolation accounts for the parameter changes and thus the basis adapts to the current estimate of the parameters. We outline an optimization method that includes the derivatives of the adaptive multiscale basis. We demonstrate the potential of the method using examples from geophysical imaging.

Samy Wu Fung  
Emory University  
400 Dowman Dr, 30322 Atlanta, GA, USA  
samy.wu@emory.edu

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

Eldad Haber  
Department of Mathematics  
The University of British Columbia  
haber@math.ubc.ca

### MS302

#### Tensor T-Product Model Reduction

In this study, we introduce t-product based tensor model reduction frameworks, namely t-POD and t-DEIM. Unlike their matrix based counterparts, the tensor formulations account for intrinsic multi-dimensional correlations, and thereby provide superior reduced models.

Jiani Zhang  
Department of Mathematics  
Tufts University

jiani.zhang@tufts.edu

Lior Horesh  
IBM Research  
lhoresh@us.ibm.com

Haim Avron  
Tel Aviv University  
haimav@post.tau.ac.il

Misha E. Kilmer  
Mathematics Department  
Tufts University  
misha.kilmer@tufts.edu

### MS303

#### **A Lagrange Multiplier Method for a Stokes-Biot Model of Flow in Fractured Poroelastic Media**

We study a finite element computational model for solving the coupled problem arising in the interaction between fluid in a poroelastic material and fluid in a fracture. The fluid flow in the fracture is governed by the Stokes equations, while the poroelastic material is modeled using the Biot system. Equilibrium and kinematic conditions are imposed on the interface via the Lagrange multiplier method. A complete stability and error analysis is performed for both semidiscrete continuous-in-time and fully discrete formulations. A series of numerical experiments are presented in order to study the convergence rates, the applicability of the method to modeling physical phenomena and the sensitivity of the model with respect to its parameters.

Ilona Ambartsumyan, Eldar Khattatov  
University of Pittsburgh  
Department of Mathematics  
ila6@pitt.edu, elk58@pitt.edu

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

Ivan Yotov  
University of Pittsburgh  
Department of Mathematics  
yotov@math.pitt.edu

### MS303

#### **Finite Element Approximation of Nonlinear Poroelasticity**

Abstract Not Available At Time Of Publication.

Song Chen  
University of Wisconsin at La Crosse  
schen@uwlax.edu

### MS303

#### **Efficient Numerical Method for Fluid Flows in Poroelastic Media**

Abstract Not Available At Time Of Publication.

Xiaobing H. Feng  
The University of Tennessee

xfeng@math.utk.edu

### MS303

#### **Analysis and Numerical Approximation of Nonlinear Poroelasticity**

Poromechanics is the science of energy, motion, and forces, and their effect on porous material and in particular the swelling and shrinking of fluid-saturated porous media. Modeling and predicting the mechanical behavior of fluid-infiltrated porous media is significant since many natural substances, for example, rocks, soils, clays, shales, biological tissues, and bones, as well as man-made materials, such as, foams, gels, concrete, water-solute drug carriers, and ceramics are all elastic porous materials (hence poroelastic). After a brief overview I will describe some nonlinear problems in poroelasticity and their mathematical analysis. I will also describe finite element based numerical methods for efficiently and accurately approximating solutions of (nonlinear) model problems in poroelasticity, and the available a-priori error estimates.

Amnon J. Meir  
Auburn University  
ajmeir@smu.edu

### MS304

#### **Wave-Structure Interaction Problems with a Piezoelectric Solid**

We present a novel time-domain analysis of a wave-structure interaction problem involving an acoustic wave propagating through a fluid in an unbounded domain interacting with an elastic solid with piezoelectric properties. The solid is initially at rest, and when the acoustic wave interacts with the solid it initiates an elastic wave coupled with an electric potential. While previous analysis has relied on the use of transforms to move the equations to the frequency domain, our analysis allows us to do everything in the time domain. Our theoretical results allow us to use data with less regularity in time than previous results. In addition to analyzing the abstract problem, we also consider a discretized version of the problem which uses Finite Elements to discretize the solid and Boundary Elements to discretize the fluid. We will also present numerical experiments which we hope will illustrate the validity of our analysis and regularity results.

Thomas S. Brown  
Department of Mathematical Sciences  
University of Delaware  
tsbrown@udel.edu

Tonatiuh Sanchez-Vizuet  
New York University  
tonatiuh@cims.nyu.edu

Francisco J. Sayas  
Department of Mathematical Sciences  
University of Delaware  
fjsayas@math.udel.edu

### MS304

#### **A Tensor Train Acceleration for the ICVSIE for 3D High-Contrast Scatterers**

Electromagnetic (EM) scattering analysis of highly-heterogeneous objects is essential for many applications



ranging from communication and satellite systems to implantable and non-invasive electrical stimulation devices. Oftentimes, EM analysis is conducted using volume and volume-surface integral equation formulations (VIE and VSIE). These formulations, however, suffer from ill-conditioning when there is high permittivity contrast, as well as from low-frequency breakdown. In the recent work of Gomez, Yucel and Michielssen, a novel internally combined VSIE formulation has been proposed to address these issues.

The tensor train (TT) decomposition is an extremely effective tool for tensor approximation. It has been employed as an extremely memory-efficient framework to accelerate the solution of volume and surface integral equations in three dimensions for non-oscillatory kernels. In this talk, we will show how these results extend to oscillatory kernels (Helmholtz), and we will showcase performance and rank behavior for the VIE and ICVSIE formulations.

Eduardo Corona

Department of Mathematics  
University of Michigan  
coronae@umich.edu

#### MS304

##### **Robust Integral Formulations for Electromagnetic Scattering from Three-Dimensional Cavities**

Scattering from large, open cavity structures is of importance in a variety of electromagnetic applications. In this paper, we propose a new well conditioned integral equation for scattering from general open cavities embedded in an infinite, perfectly conducting half-space. The integral representation permits the stable evaluation of both the electric and magnetic field, even in the low-frequency regime, using the continuity equation in a post-processing step. We establish existence and uniqueness results, and demonstrate the performance of the scheme in the cavity-of-revolution case. High-order accuracy is obtained using a Nyström discretization with generalized Gaussian quadratures.

Jun Lai

Courant Institute of Mathematical Science  
lai@cims.nyu.edu

Leslie Greengard

Simons Foundation and Courant Institute  
New York University  
greengar@cims.nyu.edu

Michael O'Neil

New York University  
oneil@cims.nyu.edu

#### MS304

##### **Simulating Surface Plasmons on Periodic Gratings Coated with Graphene: A High Order Perturbation of Surfaces Approach**

In this talk we describe a High Order Perturbation of Surfaces (HOPS) method for simulating the scattering of electromagnetic waves by periodic grating structures coated with graphene. Graphene is one example of a two-dimensional crystal with extraordinary optical properties which present new opportunities for novel nano photonics and plasmonics. In contrast with volumetric approaches, HOPS algorithms are inexpensive interfacial methods which rapidly and recursively estimate scattering

returns by perturbation of the interface shape. In comparison with Boundary Integral/Element Methods, the stable HOPS algorithm we describe here does not require specialized quadrature rules, periodization strategies, or the solution of dense nonsymmetric positive definite linear systems. With numerical experiments we show the remarkable efficiency, fidelity, and accuracy one can achieve with an implementation of this algorithm.

David P. Nicholls

University of Illinois at Chicago  
davidn@uic.edu

#### MS305

##### **Opportunities and Challenges in Sparse Linear Algebra on Many-Core Processors with High-Bandwidth Memory**

High-performance computing systems are increasingly using many cores backed by high bandwidth memory, which is well represented by Intel's Knights Landing with up to 68 cores with 3 double-precision teraflops and MCDRAM with 4x more bandwidth than DDR4. Efficiently harnessing the high compute density and high bandwidth poses interesting challenges particularly to sparse linear algebra algorithms with input-dependent behavior and high-bandwidth demand. We have been looking at various aspect like architecture-aware optimizations, programming system, and processor architecture. In this talk, I will share our experience with examples like the following: Architecture-aware optimizations: trading-off parallelism, locality, and convergence of pre-conditioners. Avoiding global synchronizations like barriers and making algorithm more asynchronous. Fusing operations and reducing memory sweeps to save memory bandwidth, the primary bottleneck in sparse linear algebra algorithms. Judiciously using privatization depending on input size and distribution. Programming system: we should not expect algorithm developers have skill and time to manually apply the above mentioned optimizations. I will present automatic wave-front parallelization of pre-conditioners, and automatic cross-kernel/matrix optimizations implemented in a Julia prototype. Processor architecture: we have been identifying bottlenecks in current processor design and would like to hear wish-list.

Jongsoo Park

Intel  
jongsoo.park@intel.com

#### MS305

##### **Exploiting Modern Manycore Architecture in Sparse Direct Solver with Runtime Systems**

Sparse direct solvers is a time-consuming operation required by many scientific applications to simulate physical problems. By its important overall cost, many studies tried to optimize the time to solution of those solvers on multi-core and distributed architectures. More recently, many works have addressed heterogeneous architectures to exploit accelerators such as GPUs or Intel Xeon Phi with interesting speedup. Despite researches towards generic solutions to efficiently exploit those accelerators, their hardware evolution requires continual adaptation of the kernels running on those architectures. The recent Nvidia architectures, as Kepler, present a larger number of parallel units thus requiring more data to feed every computational unit. A solution considered to supply enough computation has been to study problems with a large number of small com-

putations. The batched BLAS libraries proposed by Intel, Nvidia, or the University of Tennessee are examples of this solution. We discuss in this talk the use of the variable size batched matrix-matrix multiply to improve the performance of a the PaStiX sparse direct solver. Indeed, this kernel suits the supernodal method of the solver, and the multiple updates of variable sizes that occur during the numerical factorization. Performance results on a spectrum of matrices with different properties will be presented.

Pierre Ramet  
Bordeaux University - INRIA  
pierre.ramet@inria.fr

Mathieu Faverge  
Bordeaux INP - Labri - Inria  
mathieu.faverge@inria.fr

Gregoire Pichon  
INRIA  
gregoire.pichon@inria.fr

### MS305

#### Optimizing the Earthquake Simulation Code SeisSol for Heterogeneous Xeon Phi Supercomputers

SeisSol uses the high-order ADER-DG method to solve seismic wave propagation and dynamic rupture problems, as they occur in earthquake simulation. It relies on a recently extended code generation approach, which is fundamental to achieve high node-level performance on current CPU architectures. This presentation will focus on issues to optimize SeisSol for heterogeneous supercomputers. For platforms based on Xeon-Phi coprocessors, we will discuss recent work on MPI parallelization in "symmetric mode" (i.e., use one MPI rank per CPU or co-processor). We will also present optimizations of large-scale I/O on machines where (co-)processors may have different I/O-capabilities.

Leonhard Rannabauer  
Technical University of Munich  
rannabau@in.tum.de

Sebastian Rettenberger  
Technische Universität München  
Department of Informatics, Scientific Computing  
rettenbs@in.tum.de

Carsten Uphoff  
Technische Universität München  
uphoff@in.tum.de

Michael Bader  
Technical University of Munich  
bader@in.tum.de

### MS305

#### Refactoring Atmospheric Models on Sunway TaihuLight Many-Core Supercomputer

The Sunway TaihuLight supercomputer is the world first system with a peak performance greater than 100 PFlops. In this talk, we present our experience on refactoring and tuning different atmospheric models on TaihuLight. (1) we will give an introduction to TaihuLight from the programmer perspective, focusing on the architecture of the many-core SW26010 CPU and the customized Sunway OpenACC

tool. SW26010 empowers the computing part of TaihuLight, which includes both 4 management-processing-elements (MPEs) and 256 computing-processing-elements (CPEs) in one chip, achieving a peak performance of 3.06 TFlops. (2) We report our efforts on refactoring the Community Atmosphere Model (CAM). When comparing the original version using only MPEs and the refactored version using both MPE and CPEs, we achieve up to 22x speedup for the compute-intensive kernels. For the 25km resolution CAM global model, we manage to scale to 24,000 MPEs, and 1,536,000 CPEs, and achieve a simulation speed of 2.81 simulated-years-per-day. (3) We demonstrate a successful co-design case on TaihuLight, a fully-implicit solver for non-hydrostatic atmospheric dynamics. The fully-implicit solver can scale to the entire system of TaihuLight with over 10.5M cores, sustaining an aggregate performance of 7.95 PFLOPS in double-precision, and enables fast and accurate atmospheric simulations at the 488-m horizontal resolution (over 770 billion unknowns) with 0.07 simulated-years-per-day.

Wei Xue, Haohuan Fu  
Tsinghua University  
xuewei@tsinghua.edu.cn, haohuan@tsinghua.edu.cn

Chao Yang  
Software Institute, Chinese Academy of Sciences  
P. R. China  
yangchao@iscas.ac.cn

Lanning Wang  
Beijing Normal University  
wangln@bnu.edu.cn

Dexun Chen  
NRCPC and Tsinghua University  
adch@263.net

### MS306

#### A Unified Finite Element Method for Fluid-Structure Interaction

We present a new unified finite element method (UFEM) for simulation of general Fluid-Structure interaction (FSI) based upon the interaction of an unstructured ALE (Arbitrary Lagrangian-Eulerian) mesh, for the deforming solid, with an h-adaptive mesh covering the entire domain (i.e. occupied by both fluid and solid). Our proposed approach has similarities with classical immersed finite element methods (IFEMs), by approximating a single velocity and pressure field in the entire domain, but differs by treating the corrections due to the solid deformation on the left-hand side of the modified fluid flow equations (i.e. implicitly). The talk will provide an explanation of the computational technique, followed by the presentation of a broad set of computational examples in two and three space dimensions in order to validate the proposed technique across a wide range of fluid and solid parameters and interactions. It will be demonstrated that the UFEM has the same generality and robustness as traditional monolithic methods (which solve the fluid and solid equations simultaneously in a fully-couple manner) but is significantly more computationally efficient and easier to implement.

Yongxing Wang, Peter K. Jimack, Mark A. Walkley  
School of Computing  
University of Leeds  
scywa@leeds.ac.uk, p.k.jimack@leeds.ac.uk,

m.a.walkley@leeds.ac.uk

### MS306

#### A Parallel Algorithm for Variational Mesh Quality Improvement Method

Parallel mesh quality improvement is needed whenever meshes of low quality arise in computationally-intensive simulations. In particular, mesh quality improvement increases the accuracy of the associated numerical PDE algorithm and maintains its stability. Although there are numerous geometric mesh quality improvement algorithms in the literature, there are a few parallel variational methods for mesh quality improvement. In this talk, we propose a parallel algorithm designed for distributed memory machines based on the sequential variational mesh quality improvement method of Huang and Kamenski. Our parallel implementation employs a partitioned mesh, which is generated using METIS, and then solves an ordinary differential equation on each interior node of each region. For the nodes that belong to the boundaries between regions, only a partial ODE solution is calculated, as these nodes are shared among cores. The ODE solution at the shared nodes is completed by adding the partial solutions of each processor using non-blocking MPI collective communication instructions to overlap communication with computation. We present results from several numerical experiments which examine the strong and weak scalability of our method.

Maurin A. Lopez Varilla, Suzanne M. Shontz  
University of Kansas  
maurin@ku.edu, shontz@ku.edu

Weizhang Huang  
Department of Mathematics  
University of Kansas, Lawrence, KS  
huang@math.ku.edu

### MS306

#### Advances in Grid-Based All-Hex Mesh Generation for High Performance Computing

Automatically generating an all-hex mesh for arbitrary geometry is an enormously complex challenge. Adding the difficulty of scaling to billions of elements and thousands of processors can significantly compound the problem. Large scale meshing problems must necessarily remove the user from any practical user interaction making semi-automated methods such as block-structured and sweeping tools unrealistic. Grid-based technologies have, however provided an avenue to overcome many of these challenges eliminating the need to perform complex geometric reasoning for defeaturing and decomposition. In this talk we introduce a highly scalable algorithm for generating all-hex meshes based on an adaptive grid-based technology currently implemented in Sandias Sculpt tool. We describe new parallel algorithms for capturing geometric features in the mesh through dynamic conformal all-hex mesh adaption along with methods for topologically resolving curves and surfaces. We also illustrate the combination of distributed and shared memory parallelism while maintaining a parallel-consistent mesh regardless of the number of processors and configuration of the domain decomposition. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energys National Nuclear Security Admin-

istration under contract DE-AC04-94AL85000.

Steve J. Owen  
Sandia National Laboratories  
Albuquerque, NM  
sjowen@sandia.gov

### MS306

#### Parallel Geometry and Meshing Adaptation with Application to Problems with Evolving Domains

Evolving domains are a feature of a variety of physics problems such as metal forming, additive manufacturing, ballistics and crack propagation. Effective simulation of these problems requires that the model geometry evolves as the simulation progresses, and that the mesh faithfully evolves and adapts to changes in the geometry, as well as to other analysis needs such as error estimation results. Resolution of the often complex physics of these problems requires simulating them on parallel computers, and for scaling it becomes necessary that the evolving geometry and meshing is also supported in a parallel distributed environment. In problems where the model geometry is complex with many thousands of entities, for example moving propellant grains in a ballistics simulation, maintaining the entire geometric model on all processors does not scale in parallel, and it becomes necessary to distribute the geometric model itself among the processors involved in the simulation. Simmetrix has developed tools and methods to support parallel distributed geometry, mesh generation and mesh adaptation to simulate evolving geometry problems. These methods will be described with particular emphasis on the support of evolving geometries on distributed memory parallel computers. Several applications of these methods to problems in which the geometry and mesh evolve will be demonstrated.

Saurabh Tendulkar  
Simmetrix Inc.  
saurabh@simmetrix.com

Ottmar Klaas, Mark Beall  
Simmetrix, Inc.  
oklaas@simmetrix.com, mbeall@simmetrix.com

Mario Juha, Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
juham@rpi.edu, shephard@rpi.edu

### MS307

#### A Posteriori Error Estimation for Nonlinear Schrödinger Equations

The determination of the electronic structure of molecular systems is very costly, especially for large systems. For example, the use of a model based on Density Functional Theory (DFT) requires to solve a nonlinear eigenvalue problem, the computational cost being at least  $O(N^3)$ , where  $N$  is the number of valence electrons of the system. Such problems are solved numerically using a discretization (planewaves, wavelets, finite elements, etc.), and a given (generally iterative) algorithm. The a priori analysis for the discretization of these problems has been performed a few years ago (see e.g. [1]), proving the convergence and the optimality of the method used. Then an a posteriori analysis, when available, provides a guaranteed upper bound of the total error. In this talk, we shall present an a posteriori analysis for the 3D Kohn-Sham model in a

periodic setting using a planewave discretization, leading to computable bounds on the error between the exact and approximate solutions. Then, we shall explain how the a posteriori bound can be split into two components, each of them depending mainly on one approximation parameter. Numerical simulations will be presented to illustrate the possibility of balancing the different error components and to check the efficiency of the approach. Reference: [1] E.Cances, R.Chakir and Y. Maday, Numerical analysis of the planewave discretization of orbital-free and Kohn-Sham models, M2AN (highlight article) 46 (2012) 341-388.

Eric Cances  
Ecole des Ponts and INRIA, France  
cances@cermics.enpc.fr

Geneviève Dusson  
Laboratoire Jacques-Louis Lions - Institut du Calcul et de l'Analyse  
Université Pierre et Marie Curie  
dusson@ljl.math.upmc.fr

Yvon Maday  
Universite Pierre et Marie Curie  
and Brown university  
maday@ann.jussieu.fr

Benjamin Stamm  
Center for Computational Engineering Science, RWTH Aachen  
stamm@ann.jussieu.fr

Martin Vohralik  
UPMC Univ. Paris 06, France  
vohralik@ann.jussieu.fr

### MS307 Robust Computation of Wannier Functions

We propose an algorithm to determine localized Wannier functions. This algorithm, based on recent theoretical developments, does not require any physical input such as initial guesses for the Wannier functions, in contrast to previous algorithms. We demonstrate that such previous approaches based on the projection method can fail to yield the continuous Bloch gauge needed to initialize the Marzari-Vanderbilt algorithm used to construct Maximally Localized Wannier Functions (MLWFs). We show however that our algorithm is able to find localized Wannier functions, and present tests on two-dimensional systems with arbitrary periodic potentials and on Silicon. Joint work with E. Cances, G. Panati and G. Stoltz.

Antoine Levitt  
Inria and Ecole des Ponts ParisTech  
antoine.levitt@inria.fr

Eric Cances  
Ecole des Ponts and INRIA, France  
cances@cermics.enpc.fr

Gabriel Stoltz  
CERMICS  
Ecole des Ponts ParisTech  
stoltz@cermics.enpc.fr

Gianluca Panati  
Dipartimento di Matematica

Universita di Roma "La Sapienza"  
panati@mat.uniroma1.it

### MS307 Second-Order Optimization of the Cascf Wavefunction in a Relativistic Framework

A correct description of heavy transition-metal chemistry is challenging for computational chemistry, as it requires taking into account both the possible multireference character of such systems and relativistic effects. Four component approaches starting from the Dirac equation offer a rigorous tool to describe relativistic effects but introduce a significant computational overhead. When spin-orbit interactions are expected to be small compared to scalar relativistic effects, spin separation can be used to define a spin-free Dirac-Coulomb Hamiltonian, which allows one to exploit fully spin and spatial symmetry as well as real algebra, with significant computational savings. In this contribution we present a quadratically convergent implementation of the CAS-SCF method tailored for the spin-free Dirac-Coulomb Hamiltonian. The focus will be on the wave function optimization, a second-order procedure called norm-extended optimization (NEO). Such a strategy is a reformulation of the trust-region Newton method which is compatible with an efficient, direct implementation. The NEO algorithm defines the step as an eigenvector of an augmented Hessian which, for relativistic problems, that correspond to a high-order saddle point optimization, is not the one associated to the lowest eigenvalue. The strategies to solve such an eigenvalue problems will be presented.

Filippo Lipparini, Jürgen Gauss  
Johannes Gutenberg Universität Mainz  
flippari@uni-mainz.de, gauss@uni-mainz.de

### MS307 Locality of Electronic Structure Models

I will give an overview of locality results for electronic structure models and discuss their applications in materials simulations. These include the construction of QM/MM coupling methods and linear scaling algorithms. I will also introduce a lattice relaxation problem, which considers the rearrangement of a crystal lattice after the introduction of a defect as a variational problem. This treats both point defects and dislocations. Our main result establishes the far-field decay of minimising lattice displacements, which applies to several electronic structure models including: Tight-Binding and DFT models, such as the restricted Hartree-Fock Yukawa model and a Thomas-Fermi type model with full Coulomb interaction. This talk is based on joint work with Huajie Chen and Christoph Ortner.

Faizan Nazar  
MMath  
University of Warwick  
f.q.nazar@warwick.ac.uk

### MS308 Controlling Electroconvection with Surface Patterning

Electroconvection is a chaotic microscale fluid dynamic phenomena which displays many of the multiscale properties characteristic of classical turbulence but in a regime where inertial effects are negligible. Instead, nonlinear interactions between ion-transport and electric fields drive

electroconvection in electrolytes when large electric potentials are applied across ion-selective interfaces. Electroconvection greatly enhances ion transport rates, and offers a promising potential for improving otherwise diffusion-limited electrochemical processes. Under natural conditions and near flat and homogeneous surfaces, electroconvective vortices are induced intermittently and are governed by a highly chaotic dynamics. Therefore under these conditions, at any given time, a small subdomain of the system is benefitting from enhanced transport by electroconvection. We have investigated whether this inefficiency can be ameliorated through passive control of electroconvection by patterning selective surfaces with impermeable stripes. Using direct numerical simulation of the governing equations, we show that the system is highly sensitive to this type of surface modification exhibiting  $O(1)$  changes to bulk properties over a broad range of parameters. The surface modifications induce additional flows which regularize the chaotic electroconvective flows enhancing transport and generating an up to 80% increase in current density relative to homogeneous surfaces.

Scott Davidson, Ali Mani  
Stanford University  
davscott@stanford.edu, alimani@stanford.edu

### MS308

#### Numerically Computing the Sensitivity Derivatives for Chaotic High Fidelity Simulation

The Smale horseshoes are used to study the sensitivity derivatives of chaotic system. We numerically calculate the sensitivity derivatives of the Smale horseshoes over a range of parameters, including the intermittent and chaotic regimes, using the direct method, the sampling method and the least square method. It is found that the results from the sampling method are divergent and the ones from the least square method are consistent with the direct method. The least square method is further used to the coupled map lattices and provides the reasonable estimations of their sensitivity derivatives. The applications of sensitivity derivatives to the optimization of flexible bodies are also discussed.

Guowei He  
Chinese Academy of Sciences  
hgw@lnm.imech.ac.cn

### MS308

#### Scaling of Lyapunov Exponents in Homogeneous, Isotropic Turbulence

One of the defining characteristics of turbulence is that it is unstable, with small perturbations to the velocity growing exponentially on average. Indeed, turbulent flows in closed domains appear to be chaotic dynamical systems. The result is that the evolution of the detailed turbulent fluctuations can only be predicted for a finite time into the future, due to the exponential growth of errors. This prediction horizon is inversely proportional to the largest Lyapunov exponent of the system, which is the average exponential growth rate of typical linear perturbations. This work focuses on how the maximum Lyapunov exponent and hence the predictability time horizon scales with Reynolds number and computational domain size of a numerically simulated homogeneous isotropic turbulence. We will also discuss results from a short-time Lyapunov exponent analysis to characterize the nature of the instabilities underly-

ing the turbulence.

Prakash Mohan, Robert D. Moser  
University of Texas at Austin  
prakash@ices.utexas.edu, rmoser@ices.utexas.edu

### MS309

#### Fast Iterative Solvers for Cahn-Hilliard Problems

The Cahn-Hilliard equation models the motion of interfaces between several phases. The underlying energy functional includes a potential for which different types were proposed in the literature. We consider smooth and non-smooth potentials with a focus on the latter. In the non-smooth case, we apply a function space-based algorithm, which combines a Moreau-Yosida regularization technique with a semismooth Newton method. We apply classical finite element methods to discretize the problems in space. At the heart of our method lies the solution of large and sparse fully discrete systems of linear equations. Block preconditioners using effective Schur complement approximations are presented. For the smooth systems, we derive optimal preconditioners, which are proven to be robust with respect to crucial model parameters. Further, we prove that the use of the same preconditioners give poor approximations for the nonsmooth formulations. The preconditioners we present for the nonsmooth problems incorporate the regularization terms. Extensive numerical experiments show an outstanding behavior of our developed preconditioners. Our strategy applies to different Cahn-Hilliard problems including phase separation and coarsening processes, image inpainting, and two-phase flows.

Jessica Bosch  
Max Planck Institute Magdeburg  
jbosch@cs.ubc.ca

Martin Stoll  
Max Planck Institute, Magdeburg  
stollm@mpi-magdeburg.mpg.de

### MS309

#### Efficient Solvers for Stochastic Galerkin Linear Systems

Many challenging problems in computational science and engineering are characterized by uncertainty, nonlinearity and high dimensionality. Existing numerical techniques for the models of these problems would typically require considerable computational resources. This is the case, for instance, for optimization problems governed by certain time-dependent PDEs with stochastic coefficients. In particular, using the stochastic Galerkin finite element method to discretize this class of models often leads to a prohibitively high dimensional saddle-point system with tensor product structure. Crucially, we present a numerically efficient and elegant low-rank approach that makes the solution of the resulting linear systems tractable. Our numerical experiments confirm that our proposed approach reduces the computational complexity of the solution by two-three orders of magnitude, depending on discretization parameters.

Akwum Onwunta  
Max Planck Institute, Magdeburg, Germany  
onwunta@mpi-magdeburg.mpg.de

Sergey Dolgov  
University of Bath

Department of Mathematical Sciences  
s.dolgov@bath.ac.uk

Martin Stoll  
Max Planck Institute, Magdeburg  
stollm@mpi-magdeburg.mpg.de

Peter Benner  
Max Planck Institute for Dynamics of Complex Technical  
Systems, Magdeburg, Germany  
benner@mpi-magdeburg.mpg.de

### MS309

#### Asynchronous Optimized Schwarz Methods for the Solution of PDEs

An asynchronous version of the optimized Schwarz method is presented for the solution of differential equations on a parallel computational environment. In a one-way subdivision of the computational domain, with overlap, the method is shown to converge when the optimal artificial interface conditions are used. Convergence is also proved under very mild conditions on the size of the subdomains, when approximate (non-optimal) interface conditions are utilized. Numerical results are presented on large three-dimensional problems illustrating the efficiency of the proposed asynchronous parallel implementation of the method. The main application shown is the calculation of the gravitational potential in the area around the Chicxulub crater, in Yucatan, where an asteroid is believed to have landed 66 million years ago contributing to the extinction of the dinosaurs.

Daniel B. Szyld  
Temple University  
Department of Mathematics  
szyld@temple.edu

Frédéric Magoules, Cédric Venet  
CentraleSupélec, Université Paris-Saclay  
Châtenay Malabry, France  
frederic.magoules@hotmail.com,  
eric.magoules@hotmail.com

fred-

### MS309

#### Robust and Fast Preconditioners for Porous Media Flow

We consider systems originating from the simulation of multi-phase flow through porous media. The spatially discretized coupled nonlinear equations are solved with a two-stage preconditioner. We consider the first step of this procedure, i.e. solving of the pressure equation. An aspect of porous media flow is the small number of wells which are positioned quite far apart. This results in pressure solutions that are dominated by the near-well pressure behaviour in addition to the geological spatial features. In this paper we develop preconditioners based on deflation and a selection of deflation vectors motivated by proper orthogonal decomposition (POD) to a number of pre-computed solutions. We investigate alternatives using different varieties of reduced-order modelling. Furthermore we explore the connection between POD-based preconditioning and deflation methods. One of the difficulties for deflation methods is to find the right deflation vectors for general problems. The combination of deflation with the POD methods looks very promising in this respect. Some numerical experiments are given to illustrate the theory.

We start with a simple layered problem and also give results for the SPE10 benchmark.

Kees Vuik  
T. U. Delft  
c.vuik@math.tudelft.nl

Gabriela Diaz  
TU Delft  
g.b.diazcortes@tudelft.nl

Jan Dirk Jansen  
Delft University of Technology  
Department of Geoscience and Engineering  
j.d.jansen@tudelft.nl

### MS310

#### A Generalized Approach For Coupling Elliptic Problems With Material Discontinuities with Spatially Noncoincident Interfaces Through Linear Extensions

In this work, we consider coupling second order elliptic PDE together through spatially non-coincident meshes. This is an important problem to consider when two different codes are to be coupled together. While there are many methods in the literature that allow models to be coupled together across non-coincident interfaces, few achieve the desirable property of achieving optimal accuracy while also passing the linear consistency test. The main idea of our approach is to prescribe approximate coupling conditions on the non-coincident interfaces so that the optimal accuracy of the finite element method is achieved, while also allowing the resulting set of equations to pass the linear consistency test. This is done by introducing linear extension operators into the approximate interface conditions. We apply these approximate coupling conditions in the Steklov–Poincaré, Mortar Element, and Optimization Based coupling frameworks for both monolithic and iterative methods. We demonstrate that, for each of the methods presented in the talk, we achieve first order  $H^1(\Omega)$  and second order  $L^2(\Omega)$  convergence rates. In addition, we demonstrate that the linear consistency test is always passed by the set of methods presented.

James Cheung  
Florida State University  
Department of Scientific Computing  
jc07g@my.fsu.edu

### MS310

#### Inf-Sup Stability of Geometrically Unfitted Stokes Finite Elements

We demonstrate inf-sup stability property for several well-known 2D and 3D Stokes elements on triangulations which are not fitted to a given smooth or polygonal domain. The property implies stability and optimal error estimates for a class of unfitted finite element methods for the Stokes and Stokes interface problems, such as Nitsche-XFEM or cut-FEM. The error analysis is presented for the Stokes problem. All assumptions made in the paper are satisfied once the background mesh is shape-regular and fine enough.

Johnny Guzman  
Brown University  
johnny\_guzman@brown.edu

Maxim A. Olshanskii

Department of Mathematics  
University of Houston  
molshan@math.uh.edu

**MS310****An Explicit Mortar Flux Recovery Approach for Interface Coupling**

Multiphysics applications on complex geometries often require the coupling of different codes or separately meshed regions through non-matching interfaces. In this context we seek to define fluxes along the interface that can be used as boundary conditions to close the system of equations in each subdomain. We present a new explicit Lagrange multiplier based interface coupling method for this problem in the case where the subdomain governing equations are time-dependent scalar conservation equations. By differentiating the continuity constraint in the Lagrange multiplier system we are able to solve directly for the flux along the interface that satisfies the interface conditions. We refer to this approach as the Mortar Flux Recovery (MFR) method. Numerical results for advection-diffusion equations demonstrate second-order convergence in both the advection and diffusion dominated regimes and when the interface is matching the single domain solution is recovered exactly.

Kara Peterson  
Sandia Natl. Labs  
kjpeter@sandia.gov

Pavel Bochev  
Sandia National Laboratories  
Computational Math and Algorithms  
pbboche@sandia.gov

Paul Kuberry  
Sandia National Laboratories  
pakuber@sandia.gov

**MS310****Concurrent Multiscale Coupling in Solid Mechanics via the Schwarz Alternating Method**

This talk presents some recent developments in advancing the Schwarz alternating method as a means for concurrent multiscale coupling in finite deformation solid mechanics. First, we describe the algorithm's variational formulation and convergence properties: the method's convergence rate is geometric provided each of the subdomain problems is well-posed and the overlap size is non-zero. Next, we show that the use of a Newton-type method for the solution of the resultant nonlinear system leads to two kinds of block linearized systems, depending on the treatment of the Dirichlet boundary conditions. The first kind is a symmetric block-diagonal linear system in which each diagonal block is the tangent stiffness of each subdomain, so that the coupling is only through the right-hand side. The second kind is a nonsymmetric block system with off-diagonal coupling terms. We present several variants of the Schwarz alternating method that we have developed for the first kind of linear system, including one in which the Schwarz and Newton iterations are combined into a single scheme. An upshot of this version of the method is that it lends itself to a minimally intrusive implementation into existing finite element codes. We describe our implementation of the Schwarz alternating method in Albany, an open-source multiphysics research platform. Finally, we demonstrate the accuracy, convergence and scalability of the proposed

Schwarz variants on several solid mechanics examples.

Irina K. Tezaur, Alejandro Mota, Coleman Alleman  
Sandia National Laboratories  
ikalash@sandia.gov, amota@sandia.gov,  
callemas@sandia.gov

**MS311****Uncertainty Quantification for Multiscale Transport Equations**

In this talk we will study the generalized polynomial chaos-stochastic Galerkin (gPC-SG) approach to transport equations with uncertain coefficients/inputs, and multiple time or space scales, and show that they can be made asymptotic-preserving, in the sense that the gPC-SG scheme preserves various asymptotic limits in the discrete space. This allows the implementation of the gPC methods for these problems without numerically resolving (spatially, temporally or by gPC modes) the small scales. Rigorous analysis will be provided to prove that these schemes are stochastically asymptotic preserving. Examples to be discussed include the linear neutron transport equation and radiative heat transfer equations.

Shi Jin  
Shanghai Jiao Tong University, China and the  
University of Wisconsin-Madison  
sjin@wisc.edu

**MS311****Inverse Transport and Acousto-Optic Imaging**

We consider the inverse problem of recovering the optical properties of a highly-scattering medium from acousto-optic measurements. Using such measurements, we show that the scattering and absorption coefficients of the radiative transport equation can be reconstructed with Lipschitz stability by means of algebraic inversion formulas.

John Schotland  
University of Michigan  
schotland@umich.edu

**MS311****Title Not Available At Time Of Publication**

Abstract Not Available At Time Of Publication.

Guangwei Yuan  
Institute of Applied Physics and Computational  
Mathematics  
Laboratory of Computational Physics  
ygw8009@sina.com

**MS311****Preconditioning the Discrete Ordinate Equations in Optically Thick Media**

We present here an efficient preconditioner for the solution of the discrete ordinate equations in highly scattering regime. This simple method is based on the Sherman-Morrison-Woodbury identity for inverse matrices and the fact that the factorization of the scattering matrix can be done efficiently with the help of its tensor product structure. This preconditioner is efficient and results in small and almost scattering-independent iteration counts with

tuned parameters. Numerical simulations are presented to demonstrate the performance of the methods.

Rongting Zhang, Yimin Zhong, Kui Ren  
University of Texas at Austin  
rzhang@math.utexas.edu, yzhong@math.utexas.edu,  
ren@math.utexas.edu

### MS312

#### QCD with Eigenvectors

Lattice Quantum Chromodynamics is one of the areas in physics where efficient solvers of sparse matrices play a critical role. Recently there has been a tremendous amount of progress in techniques in using exact or approximate low-lying eigenvectors, as well as those to generate them. Generating thousands of low-lying eigenvectors for a matrix of rank on the order of  $10^9$  is being done regularly, which in turn enables a significant savings in subsequent calculations via techniques such as All-Mode Averaging and sparsening. We describe the algorithms used for eigenvector generation as well as techniques for using those. Challenges in scaling those algorithms to exascale machines are also discussed.

Chulwoo Jung  
Brookhaven National Laboratory  
chulwoo@quark.phy.bnl.gov

### MS312

#### Domain Decomposition Approaches for Accelerating Contour Integration Eigenvalue Solvers for Symmetric Eigenvalue Problems

In this talk we discuss techniques for computing a few selected eigenvalue-eigenvector pairs of large and sparse symmetric matrices. A recently developed powerful class of techniques to solve this type of problems is based on integrating the matrix resolvent operator along a complex contour that encloses the interval containing the eigenvalues of interest. This paper considers such contour integration techniques from a domain decomposition viewpoint, and proposes two schemes. The first scheme can be seen as an extension of domain decomposition linear system solvers in the framework of contour integration methods for eigenvalue problems, such as FEAST. The second scheme focuses on integrating the resolvent operator primarily along the interface region defined by adjacent subdomains. A parallel implementation of the proposed schemes is described and results on distributed computing environments are reported. These results show that domain decomposition approaches can lead to reduced runtimes and improved scalability.

Vassilis Kalantzis  
University of Minnesota  
kalan019@umn.edu

### MS312

#### Accelerating Nuclear Configuration Interaction Calculations through a Preconditioned Block Iterative Eigensolver

We describe a number of recently developed techniques for improving the performance of large-scale nuclear configuration interaction calculations on high performance parallel computers. We show the benefit of using a preconditioned block iterative method to replace the Lanczos algorithm that has traditionally been used to perform this type of computation. The rapid convergence of the block iterative

method is achieved by a proper choice of starting guesses of the eigenvectors and the construction of an effective preconditioner. These acceleration techniques take advantage of special structure of the nuclear configuration interaction problem which we discuss in detail. The use of block method also allows us to improve the concurrency of the computation, and take advantage of the memory hierarchy of modern microprocessors to increase the arithmetic intensity of the computation relative to data movement. We also discuss implementation details that are critical to achieving high performance on massively parallel multi-core supercomputers, and demonstrate that the new block iterative solver is two to three times faster than the Lanczos based algorithm for problems of moderate sizes on a Cray XC30 system.

Meiyue Shao  
Lawrence Berkeley National Laboratory  
myshao@lbl.gov

H. Metin Aktulga, Chao Yang  
Lawrence Berkeley National Lab  
hma@cse.msu.edu, cyang@lbl.gov

Esmond G. Ng  
Lawrence Berkeley National Laboratory  
egng@lbl.gov

Pieter Maris, James Vary  
Iowa State University  
pieter.maris@gmail.com, jvary@iastate.edu

### MS312

#### Fast Computation of Spectral Projectors of Banded Matrices

We consider the approximate computation of spectral projectors for symmetric banded matrices. While this problem has received considerable attention, especially in the context of linear scaling electronic structure methods, the presence of small relative spectral gaps challenges existing methods based on approximate sparsity. In this work, we show how a data-sparse approximation based on hierarchical matrices can be used to overcome this problem. We prove a priori bounds on the approximation error and propose a fast algorithm based on the QDWH algorithm, along the works by Nakatsukasa et al. Numerical experiments demonstrate that the performance of our algorithm is robust with respect to the spectral gap. A preliminary MATLAB implementation becomes faster than `eig` already for matrix sizes of a few thousand.

Ana Susnjara  
EPF Lausanne  
ana.susnjara@epfl.ch

Daniel Kressner  
EPFL Lausanne  
Mathicse  
daniel.kressner@epfl.ch

### MS313

#### Patient-Specific Blood Flow Simulations at Reduced Computational Cost: The Transversally Enriched Pipe Element Methodology (TEPEM)

In computational hemodynamics, it is well known the existing trade-off between accurate but very costly simulations



and cheap but sometimes insufficient over-simplified models. Taking advantage of the branching nature of arterial geometries, and inspired in the hierarchical model reduction [S. Perotto and A. Ern and A. Veneziani. Hierarchical local model reduction for elliptic problems: a domain decomposition approach. *Multiscale Model. Simul.*, 8 (2010), 4, 1102-1127], the Transversally Enriched Pipe Element Method (TEPEM) was recently proposed in [Mansilla Alvarez L. and Blanco P. and Bulant C. and Dari E. and Veneziani A. and Feijóo R (2016). Transversally enriched pipe element method (TEPEM): An effective numerical approach for blood flow modeling. *International Journal for Numerical Methods in Biomedical Engineering.*], in which the split of axial (dominant direction) and transversal dependence of physical fields yields a very adaptable methodology capable to render accurate solutions retaining a low computational cost in comparison to traditional 3D blood flow models. In this work, we explore the capabilities of this methodology on patient-specific coronary geometries featuring bifurcations, and focusing on accuracy and computational cost.

Alonso M. Alvarez, Pablo Blanco, Raúl A. Feijóo  
Laboratorio Nacional de Computacao Cientifica  
Brasil  
lalvarez@lncc.br, pjblanco@lncc.br, feij@lncc.br

### MS313

#### Reduced Basis Element for Solid Mechanics: Nitsche-Based Reduced Basis Element

The requirement for rapid solution in accurate computational domain pose significant challenges for the biomedical applications. To overcome this issue the recent ROM techniques takes advantage of the domain decomposition to decouple the domain into smaller lego blocks that are geometrically parametrized with respect to few reference shapes and the construction of local reduced basis. In the RBE "[Maday, Y. and E. M. Ronquist, A reduced-basis element method, *Journal of scientific computing* 17.1-4 (2002): 447-459]", the global system is obtained by gluing the individual reduced basis functions via finite element tearing and interconnecting (FETI). In spite of introducing new degree of freedom, we present an algorithm that exploits the potentiality of Nitsche-based domain decomposition "[Hansbo P., Nitsche's method for interface problems in computational mechanics, *GAMM-Mitteilungen* 28.2 (2005): 183-206]" for online gluing in combination with RBe approach. This combination allows to tackle both the geometric reduction of each legos block and the algebraic dimension of the full linear system. In particular, we focus on the possibility of handling in online gluing both conforming and non-conforming subdomains, and also overlapping one. This latter case is of particular interest since it allows to consider moving object without the deterioration of the mesh quality. We present some examples of Reduced basis element-Nitsche based applied to steady biomechanics problem.

Davide Baroli  
University of Luxembourg Faculty of Science, Technology and Communication  
davide.baroli@uni.lu

### MS313

#### Hierarchical Model Reduction: Theory and Practice

Customization of methods to exploit all the possible fea-

tures of the phenomenon of interest may lead to a significant improvement in terms of computational efficiency. With this respect to, Hierarchical Model (HiMod) reduction methods are ideal to describe phenomena with a dominant dynamics, locally enriched via transverse components. The HiMod procedure follows the idea of combining separation of variables with a diverse numerical approximation. This separate description of dynamics leads to construct psychological 1D models, yet able of switching to a locally higher fidelity. Thus, along the mainstream we consider a classical one-dimensional finite element approximation to exploit easiness and versatility of this method. The transverse components are tackled by a modal approximation. In particular, the number of modes can be locally tuned along the mainstream, according to the meaningfulness of the transverse information. Relatively few modes are expected to capture the transverse dynamics with an overall reduction of computational costs. In this presentation, we focus on the most recent advances in Hi-Mod reduction, after introducing the basics of such an approach. Work in collaboration with A. Veneziani, D. Baroli, C.M. Cova, S. Guzzetti, M. Lupo Pasini, L. Sala. Supported by NSF DMS 1412963.

Simona Perotto  
MOX - Modeling and Scientific Computing  
Dipartimento di Matematica  
simona.perotto@polimi.it

### MS313

#### Perspectives and Advances in Reduced Order Methods for Parametrized Viscous Flows

We present the state of the art of our research in methodological and numerical developments for reduced order methods, focused on parametric flows (shape of domain and physics/data). Current efforts are devoted to the detection of instabilities and bifurcations, in the increase of Reynolds number, as well as in data assimilation, parameter estimation, uncertainty quantification. Extension to multi physics setting is undergoing (fluid-structure interactions, porous media), as well as to more efficient techniques in dealing with parametric shapes and interfaces. Special attention is devoted in the development of open source libraries to guarantee competitive computational performances, thanks to offline-online computing splitting.

Gianluigi Rozza  
SISSA, International School for Advanced Studies  
Trieste, Italy  
gianluigi.rozza@sissa.it

Francesco Ballarin  
SISSA, International School for Advanced Studies  
francesco.ballarin@sissa.it

Giovanni Stabile  
SISSA mathLab, Trieste, Italy  
giovanni.stabile@sissa.it

### MS314

#### Discovering Governing Equations by Sparse Identification of Nonlinear Dynamics

This work develops a general framework to discovering the governing equations underlying a dynamical system simply from data measurements, leveraging advances in sparsity techniques and machine learning. The resulting models are parsimonious, balancing model complexity with descriptive

ability while avoiding overfitting. The only assumption about the structure of the model is that there are only a few important terms that govern the dynamics, so that the equations are sparse in the space of possible functions. This perspective, combining dynamical systems with machine learning and sparse sensing, is explored with the overarching goal of real-time closed-loop feedback control of complex systems. There are many more critical data-driven problems, such as understanding cognition from neural recordings, inferring patterns in climate, determining stability of financial markets, predicted and suppressing the spread of disease, and controlling turbulence for greener transportation and energy. With abundant data and elusive laws, data-driven discovery of dynamics will continue to play an increasingly important role in these efforts.

Steven Brunton

University of Washington  
sbrunton@uw.edu

Joshua L. Proctor  
Institute for Disease Modeling  
JoshLProctor@gmail.com

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

**MS314**

**Interpolatory Decompositions in an Instrumented Building**

Virginia Tech's Goodwin Hall is instrumented with 212 accelerometers welded to the steel frame of the building. Each of these sensors records vibrations at the rate of 25,600 samples per second. From this massive data set, the Virginia Tech Smart Infrastructure Lab seeks to learn about many aspects of building occupancy, ranging from day-to-day operating conditions to the detection of an active shooter. This talk will describe the use of interpolatory matrix factorizations (e.g., CUR decompositions) to identify a minimal set of sensors that capture a pedestrian walking down a heavily-instrumented hallway, toward establishing a minimal sensor configuration for instrumenting other buildings.

Mark Embree

Department of Mathematics  
Virginia Tech  
embree@vt.edu

Dustin Bales  
GAiTE, LLC  
dbales@gaitellc.com

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

Pablo Tarazaga  
Virginia Tech  
Department of Mechanical Engineering  
ptarazag@vt.edu

**MS314**

**Fast Data-Driven System Identification from Im-**

**pulse Response Measurements**

With vast amounts of data, the cost of identifying a state-space system model via nonlinear least squares becomes dominated by the dimension of the data. Our approach reduces the effective data dimension by constructing a series of inexpensive surrogates where the data is projected onto a low dimensional subspace. Although a naive implementation would still require an  $\mathcal{O}(n)$  operation inner-product to minimize each surrogate, where  $n$  is the data dimension, we can exploit the structure of this problem to evaluate this inner-product in  $\mathcal{O}(1)$  operations, reducing the cost of optimization. When there is noise in the data, each projected surrogate problem yields a different system estimate from the original problem. We present a bound on this mismatch and use it to control the error by updating the subspace during the course of optimization. If perturbations to the measurements sample a Gaussian distribution, the mismatch can be interpreted in statistical sense and the covariance of the surrogate estimate obtains the Cramér-Rao bound to within an arbitrary tolerance. The net result is an algorithm that obtains an estimate that is 98% as accurate as the conventional nonlinear least squares solution while being 100× faster. This approach is also competitive with subspace based system identification approaches.

Jeffrey M. Hokanson  
Colorado School of Mines  
hokanson@mines.edu

**MS314**

**Data-Driven Model Order Reduction for Linear Parametric Systems**

Data-driven model order reduction in the Loewner framework has attracted increasingly more research attentions. Since parameters are often present in industrial applications and parametric studies are crucial in system design and analysis, developing parametric model order reduction in the Loewner framework is important to further improve its applicability. Interpolation methods in the Loewner framework have been proposed to reduce systems with one or two parameters. Although these methods can be extended to accommodate more parameters theoretically, this faces several difficulties. First, the mathematical formulation needs to be specially derived for each number of parameters. In addition, the complexity grows fast with the number of parameters and the formulation is already complicated for two parameters. In this talk, we propose another type of interpolation method under the Loewner framework. This method builds parametric reduced models by interpolating nonparametric reduced models built by the Loewner approach. To interpolate these nonparametric reduced models, our method first transforms all of them under a consistent set of generalized coordinates. Then, these transformed reduced models are interpolated on their appropriate matrix manifold. The proposed method provides a uniform framework for building parametric reduced models in the Loewner framework because it has the same formulation and can be efficiently implemented for different numbers of parameters.

Yao Yue

Max Planck Institute for Dynamics of Complex  
Technical Systems, Magdeburg  
yue@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

### MS315

#### Modeling Laser Absorption in Low-Speed Reacting Flows using Adaptive Finite Elements

This work presents modeling of laser absorption in gaseous reacting flows using adaptive finite elements driven by dual-weighted residual error estimates. Laser absorption spectroscopy is a valuable tool that allows experimentalists to study the characteristics of gaseous flows without directly affecting the flow itself. However, this technique can be limited when the flow characteristics vary significantly along the optical path, as the data obtained represents spatially-averaged quantities. These difficulties can be circumvented by using detailed computational models in conjunction with the laser absorption data. Integration of the Beer-Lambert law is performed using the GRINS multiphysics finite element framework by supplying the relevant quantity-of-interest (QoI) module whereby the laser intensity is integrated along the optical path using quadrature. The libMesh FEMSystem framework, upon which GRINS is built, facilitates the element-wise QoI integration across the mesh as well as provides a framework for computing dual-weighted residual error estimates that are used to drive adaptive mesh refinement (AMR) in areas with high laser absorption. The talk discusses the formulation of the problem, details of the software implementation and libraries used, and numerical examples.

Timothy Adowski, Paul Baumann  
University of Buffalo  
tradowsk@buffalo.edu, pbauman@buffalo.edu

### MS315

#### Parallel Goal-Oriented Adaptive Solid Mechanics

We present a software framework for goal-oriented error estimation and mesh adaptation in quasi-steady Lagrangian solid mechanics applications. This framework is being developed using the Trilinos software environment taking advantage of the PUMI unstructured mesh adaptation components. Key features being developed include the capability to: (1) Execute all steps of the adaptive loop workflow in parallel to support the scalable execution of large problems. (2) Effectively implement the primal problem and dual problems via using automatic differentiation. (3) Investigate alternative quantities of interest relevant to solid mechanics applications. We will discuss the software components used to construct the framework, the steps performed for an adaptive goal-oriented analysis, and how the framework generalizes goal-oriented error estimation for a class of solid mechanics problems. Work has been done to coordinate a variety of software packages from Trilinos and PUMI to efficiently solve primal problems, solve dual problems, estimate the error in a functional quantity, and perform mesh adaptation to reduce the functional error. Based on a mixed displacement-pressure formulation, the primal problem, dual problem, and error contribution terms have been implemented in a generic form to allow the analysis of variety of different types of solid mechanics applications. Results for linearly elastic, Neo-Hookean, and

plasticity models will be presented.

Brian Granzow  
Rensselaer Polytechnic Institute  
granzb@rpi.edu

Assad Oberai  
Department of Mechanical, Aerospace and Nuclear  
Engineering  
Rensselaer Polytechnic Institute  
oberaa@rpi.edu

Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

### MS315

#### Adaptivity in High-Order Finite Element ALE Simulations

The Arbitrary Lagrangian-Eulerian (ALE) framework forms the basis of many large-scale multi-physics codes, and in particular those centered around radiation diffusion and shock hydrodynamics. We are developing general high-order finite element discretization framework that aims to improve the quality of current ALE simulations, while also improving their performance on modern data-centric computing architectures. We use the de Rham complex to guide the discretization of different physics components. In particular, kinematic quantities (e.g. velocity, position) are discretized with continuous (H1) finite elements, thermodynamic quantities (e.g. internal energy) use continuous (L2) elements, while H(div)-conforming finite elements are used for the fluxes in radiation diffusion. To address adaptivity in these settings, we have developed a general unstructured nonconforming mesh refinement capability that support adaptive mesh refinement on triangular, quadrilateral and hexahedral curvilinear meshes, at arbitrarily high order, for any finite element space. Our approach can handle complex 3D anisotropic refinements, unlimited refinement levels of adjacent elements, and MPI parallelism with load balancing. In this talk we present the high-order AMR algorithms and demonstrate their performance on model compressible hydrodynamics and computational electromagnetics problems. We also discuss our ongoing efforts to develop practical error estimators in these settings.

Jakub Cerveny, Veselin Dobrev  
Lawrence Livermore National Laboratory  
cerveny1@llnl.gov, dobrev1@llnl.gov

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

Robert Rieben, Vladimir Tomov  
Lawrence Livermore National Laboratory  
rieiben1@llnl.gov, tomov2@llnl.gov

### MS315

#### Mesh Motion and Adaptation for Two-Phase Flow Problems with Moving Objects

Reliable simulation of fluid flow problems with fluid-structure and multiphase interactions are challenging because of the inherent difficulty in capturing and/or keeping

track of the evolving interfaces as well as moving objects (e.g., a floating bar or a moving ship on water). In this talk, we will focus on an adaptive unstructured approach for two-phase problems with moving objects. Our adaptive approach is based on local error indicators and estimators of both explicit and implicit types that are used to construct a mesh sizefield which accounts for anisotropy (e.g., near the interface or boundary). The anisotropic mesh sizefield is used to drive the overall mesh adaptation process, which is based on a set of local mesh modification operations (e.g., edge split, collapse and swap, and compound operators) and operates in parallel on distributed/partitioned meshes. To account for moving objects, an arbitrary Lagrangian-Eulerian (ALE) description is employed and mesh motion is used. In summary, the current adaptive approach (based on local error indicator/estimator and mesh modification) supports evolving interfaces, moving objects involving mesh motion, and high anisotropy, and operates in parallel. This adaptive approach will be demonstrated on different problems of interest such as a bar plunging into a pool of liquid, a floating bridge deck, etc.

Onkar Sahni, Alvin Zhang  
Rensselaer Polytechnic Institute  
sahni@rpi.edu, zhanga@rpi.edu

Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

Chris Kees  
U.S. Army Engineer Research and Development Center  
Coastal and Hydraulics Laboratory  
christopher.e.kees@usace.army.mil

### MS316

#### Energy Conserving, Linear Scaling, Real Space Ab Initio Molecular Dynamics

Born-Oppenheimer molecular dynamics (BOMD) calculates forces on the nuclei via an iterative electronic optimization at each time step. A tight convergence of this optimization is typically required in order to prevent a systematic drift in the total energy over time. Using extended Lagrangian BOMD (XL-BOMD), where the initial guess to this optimization procedure is propagated alongside the nuclei, energy conservation can be achieved without requiring an expensive tightly-converged electronic solution. We have investigated the application of XL-BOMD for the scalable real-space BOMD strategies developed by Fattbert and Osei-Kuffuor. The fusion of these two methods requires the propagation of both the initial guess for the Kohn-Sham subspace as well as the centers of the localized orbitals. XL-BOMD has shown remarkable energy conservation when the localized orbitals are not truncated. We are currently exploring the magnitude and nature of the energy drift caused by truncating the localized orbitals.

Ian Dunn  
Columbia University  
iansdunn@gmail.com

### MS316

#### O(N) Density Functional Theory Calculations: Beyond Ground State of Insulators

Numerous approaches have been proposed in recent years to reduce computational complexity from  $O(N^3)$  to  $O(N)$

in Density Functional Theory simulations. Practical and accurate methods exist for ground state calculations of insulators. Extending these methods to more difficult problems such as molecular dynamics and metallic systems present some challenges, such as energy conservation, fast time-to-solution, critical problem size where  $O(N)$  become competitive. These issues will be discussed in light of our recent experience with a Finite Difference discretization of the Kohn-Sham equations, a localized orbitals representation of the electronic structure, and distributed approximate solvers for the matrices associated with the operators projected onto the localized orbitals basis set. Parallel scalability will also be discussed. This work was performed under the auspices of the U.S Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344.

Jean-Luc Fattbert  
Lawrence Livermore National Lab.  
fattbert1@llnl.gov

Daniel Osei-Kuffuor, Tadashi Ogitsu  
Lawrence Livermore National Laboratory  
oseikuffuor1@llnl.gov, ogitsu1@llnl.gov

### MS316

#### Next Generation Quantum Molecular Dynamics

We are developing a modern framework for quantum based molecular dynamics simulations that combines some of the best features of regular Born-Oppenheimer and Car-Parrinello molecular dynamics. The new framework is based on an extended Lagrangian formulation of Born-Oppenheimer molecular dynamics that allows, for the first time, efficient energy conserving Born-Oppenheimer molecular dynamics simulations with a computational complexity that scales only linearly with the system size. Only a single diagonalization per time step is required. The efficiency and accuracy of the new dynamics can be understood from a backward analysis. Instead of integrating an underlying exact dynamics with approximate forces, exact forces that do not rely on the fulfillment of the Hellmann-Feynman theorem are used to integrate the equations of motion for an approximate shadow Hamiltonian that closely follows the exact solution. In this way properties such as the total energy can be controlled. This geometric approach to integration is widely used in classical molecular dynamics, e.g. in the velocity Verlet algorithm. Our new framework allows this geometric technique to be applied also to self-consistent field theory. Extended Lagrangian Born-Oppenheimer molecular dynamics represents a general approach that can be applied to a broad variety of quantum based molecular dynamics applications.

Anders Niklasson  
Los Alamos National Laboratory  
amn@lanl.gov

### MS316

#### Spectral Quadrature Method for Large-Scale, High-Temperature Quantum Molecular Dynamics

The computational cost of conventional diagonalization-based Density Functional Theory (DFT) calculations grows rapidly with increasing temperature due to a large number of states becoming partially occupied. Since most diagonalization algorithms scale quadratically with the number of states to be computed, and scale poorly in the context of scalable high performance computing due to the

orthogonality constraint on the states, high temperature calculations for even small to moderately sized systems are currently intractable. In this talk, the recently developed spectral quadrature DFT (SQDFT) method will be discussed. Specifically, SQDFT formulates Kohn-Sham densities, energies, and forces as spectral integrals, which are then evaluated by quadrature rules. This approach scales linearly with respect to the number of states, is identically applicable to insulating and metallic systems, and has decreasing computational cost with increasing temperature. In addition, the SQDFT approach admits an efficient decomposition of full sparse Hamiltonian operations to local dense sub-Hamiltonian operations, thus facilitating parallel scalability. This decomposition is exploited to maximize parallel scalability, while retaining strict equivalence to the full-Hamiltonian formulation, and so systematic convergence to the exact diagonalization result. Results using SQDFT for high-temperature quantum molecular dynamics simulations will be presented.

Phanish Suryanarayana  
Georgia Institute of Technology  
phanish.suryanarayana@ce.gatech.edu

### MS317

#### **Panzer: A Finite Element Assembly Engine within the Trilinos Framework**

Panzer is a general finite element assembly engine for multiphysics simulation, which leverages a variety of Trilinos packages. The library provides abstractions that accelerate the implementation of weak forms by automating the local to global assembly. A wide variety of finite element bases are possible, including physics-compatible discretizations. To specify the problem, users divide the domain into element blocks, and then associates physics blocks (equation sets and boundary conditions) with them. Fully coupled systems composed of different equation sets in different element blocks are possible, along with mixed bases for degrees of freedom within an element block. The assembly process utilizes a directed acyclic graph-based equation description, in which each graph node defines a computational kernel that computes a term of the weak form. Executing these kernels according to the graph dependency results in the complete computation of the residual. The power of this model is the interchangeability of kernels to build up different variational forms. Panzer utilizes embedded template based generic programming to assemble quantities of interest (residuals, Jacobians, Hessians) from the same code. Users specify their equations once and automatic differentiation determines the rest to machine precision. It therefore naturally supports adjoint-based error analysis, stability and bifurcation analysis, optimization, and uncertainty quantification.

Jason Gates  
Sandia National Laboratories  
jmgate@sandia.gov

Roger Pawlowski  
Multiphysics Simulation Technologies Dept.  
Sandia National Laboratories  
rppawlo@sandia.gov

Eric C. Cyr  
Computational Mathematics Department  
Sandia National Laboratories

eccyr@sandia.gov

### MS317

#### **Porting Spectral/hp Element Framework Nektar++ to Kokkos**

As high-performance computers with distributed memory are becoming increasingly heterogeneous, parallel execution based exclusively on MPI is no longer sufficient. Recent years have seen a rapid development of multi-core architectures and associated software toolkits that enable intra-node acceleration, but present also a significant challenge for numerical software. Established algorithms often need to be redesigned in order to fit the underlying hardware and given the variety of available chips, it is desirable to maintain portability and sufficient level of abstraction without sacrificing performance. Finite element methods are one of the most important applications of high-performance computing and their high-order variants are amenable to execution on multi-core processors due to their data locality and high arithmetic intensity. We are interested in performance-portable implementation of continuous and discontinuous Galerkin methods that are currently available in spectral element solver Nektar++. The Trilinos library Kokkos is a suitable candidate for making the existing code base multi-core aware and we present our experience with refactoring the code to Kokkos programming model. Nektar++ also contains a high-order mesh generator called NekMesh and we discuss its acceleration of NekMesh through Kokkos and associated challenges.

Mike Kirby  
University of Utah  
School of Computing  
kirby@cs.utah.edu

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

Martin Vymazal  
Imperial College  
martin.vymazal@imperial.ac.uk

David Moxey, Jan Eichstaedt, Michael Turner  
Department of Aeronautics  
Imperial College London  
d.moxey@imperial.ac.uk, eichstaedt13@imperial.ac.uk,  
m.turner14@imperial.ac.uk

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

### MS317

#### **A Novel Algebraic Multigrid Method for Q2-Q1 Discretizations of the Navier-Stokes equations**

Algebraic multigrid methods are well established for PDE systems with co-located unknowns. In this talk, we propose a novel multigrid method for mixed discretizations of the incompressible Navier-Stokes equations for the case where unknowns are not co-located at mesh points such as  $Q_2-Q_1$  discretizations. The main idea of the new approach is to first define coarse pressures in a somewhat standard AMG fashion and then to automatically choose coarse velocity unknowns so that the spatial location relationship between pressure and velocity dofs resembles that on the

finest grid. To define coefficients for the transfer operators, an energy minimization multigrid is used. We present numerical results for a variety of Stokes and Navier-Stokes problems.

Andrey Prokopenko  
Sandia National Laboratories  
prokopenkoav@ornl.gov

Ray S. Tuminaro  
Sandia National Laboratories  
Computational Mathematics and Algorithms  
rstumin@sandia.gov

### MS317

#### Multiphysics Preconditioning with the MueLu Multigrid Library

Computer simulations of multiphysics problems not only grow in size but also in complexity including more and more types of physics and coupling equations. Multiphysics systems typically characterized by multi-scale temporal and spatial mechanisms often demand implicit time integration. For certain classes of single-field problems multigrid methods are known to be efficient for solving the linear systems arising from the time integrator. However, the design and handling of fully coupled physics-based multigrid preconditioners for multiphysics problems remains a challenge. MueLu provides a flexible multigrid framework in Trilinos which allows one to flexibly build and extend application-specific block preconditioners based on algebraic multigrid methods. This drastically simplifies the implementation of block multigrid preconditioners for new types of applications. Additionally, it allows one to define simple user interfaces for application-specific preconditioners that reduce the number of parameters to minimum. This is of interest for end users who want to apply the preconditioners in a production environment. We demonstrate the design principles and usage of the MueLu framework for multiphysics problems using several examples including Navier-Stokes and MHD. Finally, we compare the performance of the multigrid preconditioners with alternative preconditioning methods.

Tobias A. Wiesner  
Sandia National Laboratories  
tawiesn@sandia.gov

Ray S. Tuminaro  
Sandia National Laboratories  
Computational Mathematics and Algorithms  
rstumin@sandia.gov

Eric C. Cyr  
Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

John Shadid  
Sandia National Laboratories  
Albuquerque, NM  
jnshadi@sandia.gov

Jonathan J. Hu  
Sandia National Laboratories  
Livermore, CA 94551

jhu@sandia.gov

### MS318

#### Stochastic and Hybrid (Stochastic/Deterministic) Methods for Data and Compute Intensive Problems

Current and emerging extreme-scale systems require novel scientific algorithms to hide network and memory latency, have very high computation/communication overlap, have minimal communication, have no synchronization points. Scientific algorithms for multi-petaflop and exa-flop systems also need to be fault tolerant and fault resilient, since the probability of faults increases with scale. Therefore, key science applications require novel mathematical methods and algorithms as well as system software that address the challenges of current- and future generation extreme-scale HPC systems. Note also that with the advent of Big Data in the past few years the need of such scalable mathematical methods and algorithms able to handle data and compute intensive applications at scale becomes even more important. In this talk the author will focus on novel parallel stochastic and hybrid (stochastic/deterministic) methods for solving Linear Algebra Problems as one possible approach in addressing the above challenges. Stochastic approaches usually produce high level of parallelism with minimal communication in the parallel case. In particular Monte Carlo and quasi-Monte Carlo scalable hybrid methods for matrix inversion and solving systems of linear algebraic equations will be presented. In addition various approaches how these can be used as efficient Sparse Approximate Inverse Preconditioners will be given. Examples will be drawn from a set of diverse applications.

Vassil Alexandrov  
Barcelona Supercomputing Centre, Spain  
vassil.alexandrov@bsc.es

### MS318

#### Hierarchical Matrices and Low-Rank Methods for Extreme-Scale Solvers

Solving large linear systems is a key component of many computations in CS&E. Recently, hierarchical matrices and low-rank methods have attracted much attention. We give an overview of the main types of algorithms in this area, and discuss what makes them attractive for extreme-scale computing. Finally, we show some results from a parallel hierarchical solver under development that is based on  $H^2$  matrices and the inverse FMM (IFMM).

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

Siva Rajamanickam  
Sandia National Laboratories  
srajama@sandia.gov

Ray S. Tuminaro  
Sandia National Laboratories  
Computational Mathematics and Algorithms  
rstumin@sandia.gov

Chao Chen  
Stanford University  
cchen10@stanford.edu

Eric F. Darve  
Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

### MS318

#### Automatic Mapping Operator Construction for a Subspace Correction Method Applied to a Series of Linear Systems

We focus on solving a series of linear systems with an identical or similar coefficient matrix. The linear systems are sequentially processed due to the right hand side vector depending on the solution vector of the prior linear system. For the problem, we investigate the subspace correction method and its implicit version which accelerate the convergence of an iterative solver. The key to the correction method is how effective mapping operators are constructed. When the range of the mapping operator covers a subspace including slow convergent error vectors, the correction method works well. We have developed a new mapping operator construction method for the problem above in which the information obtained in the prior solution steps is used. The subspace including slow convergent vectors is automatically identified. Numerical tests on electromagnetic field analyses confirm that the correction method with the generated mapping operator can accelerate the convergence of the iterative solver.

Takeshi Iwashita  
Hokkaido University  
iwashita@iic.hokudai.ac.jp

Shigeru Kawaguchi  
Kyoto University  
s.kawaguchi@fem.kuee.kyoto-u.ac.jp

Takeshi Mifune  
Department of Electrical Engineering  
Kyoto University  
mifune@fem.kuee.kyoto-u.ac.jp

Tetsuji Matsuo  
Kyoto University  
tmatsuo@kuee.kyoto-u.ac.jp

### MS318

#### Convergence Between Big Data and HPC

In general, High Performance Computing and Big Data have different programming and computing paradigms. Big Data is more usually concerns by I/O optimization and storage, while HPC is mainly often about computation optimizations. It is of course a very easy comparison of two very complex domains and a larger number of applications have to address all these problems. Recent evolutions generate large applications with a convergence between these scientific domains. The increasing size of problems, new methods, potential computing power of future supercomputers, associated with the shared problems of minimizing the energy consumption and multi-core processor evolutions, would merge computational and data sciences into a general scientific programming and computing paradigm for a large class of applications. New classes of methods, algorithms and data representations would emerge and applications would be adapted. In this talk, we discuss on differences between HPC and Big Data programming and computing. For large scale application associated with very

large data modified during the computation, we propose ideas to mix those approaches and we discuss on several other existing ideas and experiments for HPC and Big Data programming and computing. As an example, we propose a solution based on using graphs of components for HPC and DAG for Big Data applications (respectively YML and TEZ). We conclude on the necessity to deploy experimental platforms.

Serge G. Petiton  
University Lille 1, Science and Technologies - CNRS  
serge.petiton@lifl.fr

Nahid Emad  
PRiSM laboratory and Maison de la Simulation,  
University of  
Nahid.Emad@prism.uvsq.fr

Laurent Bobelin  
LI, universit e Franois Rabelais  
Tours, France  
laurent.bobelin@univ-tours.fr

### MS319

#### Estimation of Cardiac Conductivities by a Variational Data Assimilation Approach: Analysis and Validation

Computational electrocardiology is becoming a formidable tool to deeply understand cardiac electrophysiology. Mathematical and numerical modeling have been continuously refined and made closer and closer to clinical application so that it can be used to improve diagnosis and prognosis of cardiac arrhythmia. However, many popular models have been shown to be strongly sensitive to the cardiac conductivity parameters. The accurate quantification of cardiac conductivities is crucial for extending computational electrocardiology from medical research to clinical practice. With this motivation, we consider a variational data assimilation approach for the estimation of the cardiac conductivities which allows us to combine available patient-specific measures with mathematical models. We regard the parameters as control variables to minimize the mismatch between available measures and computed potentials. By resorting to a derivative-based optimization method, we significantly improve the numerical approaches present in literature. We demonstrate the reliability of the conductivity estimation approach in presence of noise. A sensitivity analysis with respect to different possible experimental set ups is provided. In order to extend our procedure to an operative stage, an extensive validation using in vitro animal measures will be discussed. Finally, a Bayesian approach to the inverse problem aiming at assessing the uncertainty about the parameter estimation will be considered.

Alessandro Barone  
Department of Mathematics and Computer Science  
Emory University, Atlanta  
alessandro.barone@emory.edu

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

Huanhuan Yang  
Florida State University  
Department of Scientific Computing  
hyang3@fsu.edu

Flavio H. Fenton  
Georgia Institute of Technology  
flavio.fenton@physics.gatech.edu

Alessio Gizzi  
University Campus Bio-medico of Rome  
a.gizzi@unicampus.it

Simonetta Filippi  
University Campus Bio-Medico of Rome, Italy  
ICRA University of Rome, La Sapienza, Italy  
s.filippi@unicampus.it

### MS319

#### Taylor Approximation for PDE-Constrained Optimal Control Problems Under High-Dimensional Uncertainty: Application to a Turbulence Model

In this talk, we present an efficient method based on Taylor approximation for PDE-constrained optimal control problems under high-dimensional uncertainty. The computational complexity of the method does not depend on the nominal but only on the intrinsic dimension of the uncertain parameter, thus the curse of dimensionality is broken for intrinsically low-dimensional problems. Further correction for the Taylor approximation is proposed, which leads to an unbiased evaluation of the statistical moments in the objective function. We apply our method for a turbulence model with infinite-dimensional random viscosity.

Peng Chen  
UT Austin  
peng@ices.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

Umberto Villa  
University of Texas at Austin  
uvilla@ices.utexas.edu

### MS319

#### Simulation Based Optimal Experimental Design: A Measure-Theoretic Perspective

The typical work-flow in the end-to-end quantification of uncertainties in model input parameters requires first formulating and solving stochastic inverse problems (SIPs) using output data on available quantities of interest (QoI). The solution to a SIP is often written in terms of a probability measure, or density, on the space of model inputs. Then, we can formulate and solve a stochastic forward problem (SFP) where the uncertainty on model inputs is propagated through the model to make quantitative predictions on either unobservable or future QoI data. As the fidelity of models increases to include behavior at a wide range of scales, the number of uncertain input parameters often increases significantly, while the number of additional experimental data, which are often costly to obtain, are typically limited. It is therefore becoming increasingly important that we optimally design data collection networks so that the data collected on available QoI leads to improved predictive capabilities of the computational model. In this talk, we use a measure-theoretic framework to formulate and solve both SIPs and SFPs. From this perspective, we quantify the geometric characteristics of inverse images using hypothetical sets of QoI. This leads to a nat-

ural definition of the optimal experimental design, i.e., the optimal configuration of a finite set of sensors is in some space-time domain. Several numerical examples and applications are discussed.

Scott Walsh  
University of Colorado Denver  
scott.walsh@ucdenver.edu

### MS319

#### Stochastic Optimization for Turbofan Noise Reduction Using Parallel Reduced-Order Modeling

Simulation-based optimization of acoustic liner design in a turbofan engine nacelle for noise reduction purposes can dramatically reduce the cost and time needed for experimental designs. Because uncertainties are inevitable in the design process, a stochastic optimization algorithm is posed in this talk based on the conditional value-at-risk measure so that an ideal acoustic liner impedance is determined that is robust in the presence of uncertainties. A parallel reduced-order modeling framework is presented that dramatically improves the computational efficiency of the stochastic optimization solver for a realistic nacelle geometry. The reduced stochastic optimization solver takes less than 500 seconds to execute.

Huanhuan Yang, Max Gunzburger  
Florida State University  
Department of Scientific Computing  
hyang3@fsu.edu, mgunzburger@fsu.edu

### MS320

#### A Hydrostatic Dynamical Core using Higher-order Structure-Preserving Finite Elements

It is well known that the inviscid, adiabatic equations of atmospheric motion constitute a non-canonical Hamiltonian system, and therefore possess many important conserved quantities such as mass, potential vorticity and total energy. In addition, there are also key mimetic properties (such as  $\text{curl grad} = 0$ ) of the underlying continuous vector calculus. Ideally, a dynamical core should have similar properties. A general approach to deriving such structure-preserving numerical schemes has been developed under the frameworks of Hamiltonian methods and mimetic discretizations, and over the past decade, there has been a great deal of work on the development of atmospheric dynamical cores using these techniques. An important example is *Dynamico*, which conserves mass, potential vorticity and total energy; and possesses additional mimetic properties such as a curl-free pressure gradient. Unfortunately, the underlying finite-difference discretization scheme used in *Dynamico* has been shown to be inconsistent on general grids. To resolve these accuracy issues, a scheme based on mimetic Galerkin discretizations has been developed that achieves higher-order accuracy while retaining the structure-preserving properties of the existing discretization. This presentation will discuss the new dynamical core, termed *Dynamico-FE*, along with a more general discussion on mimetic methods as used in atmospheric science.

Chris Eldred  
University of Paris 13  
chris.eldred@gmail.com

Thomas Dubos  
LMD



Ecole Polytechnique  
dubos@lmd.polytechnique.fr

Evangelos Kritsikis  
LAGA  
wkrits@gmail.com

### MS320

#### A Computational Algebra of Hybridization Methods for Ocean and Atmosphere Discretizations

Recently developed finite element discretizations maintain the exact force balances which are essential for accurate ocean and atmosphere simulations. However, efficient solution of geophysical flow equations depends on transforming the mixed velocity-pressure system into an elliptic system for the pressure. Implementing this transformation in the finite element method requires the exploitation of “hybridization” techniques in which the discretized systems are algebraically manipulated during the equation assembly process to produce the desired elliptic system. Using conventional model development techniques, hybridization of complex discretizations requires manual intervention in intricate numerical code, and this intervention must be repeated every time the model is modified, extended, or debugged. In contrast, the Firedrake project at Imperial College takes the discretized equations in symbolic form as input, and automatically generates high performance parallel code from this mathematical specification. By introducing symbolic operations for hybridization operations, and generating code from them, we aim to successfully extend this system for automated simulation to this important class of problem.

Thomas Gibson, David Ham  
Department of Mathematics  
Imperial College London  
t.gibson15@imperial.ac.uk, david.ham@imperial.ac.uk

Colin J. Cotter  
Imperial College London  
Department of Mathematics  
colin.cotter@imperial.ac.uk

### MS320

#### R-Adaptive Mesh Generation for Global Weather Prediction

We consider the generation of meshes adapted to a scalar monitor function through equidistribution. Together with an optimal transport condition, this leads to a nonlinear equation of Monge-Ampere type for generating meshes on the plane. As our interest is in numerical methods for global weather prediction, we must formulate the equivalent equation for performing this task on the sphere. Finally, we present robust iterative numerical methods for generating appropriate meshes, using finite element methods to solve the nonlinear equations.

Andrew McRae  
Department of Mathematics  
Imperial College London  
a.t.t.mcrae@bath.ac.uk

Colin J. Cotter  
Imperial College London  
Department of Mathematics  
colin.cotter@imperial.ac.uk

Chris Budd  
Dept. of Mathematical Sciences  
University of Bath, UK  
mascjb@bath.ac.uk

### MS320

#### The Acme Spectral Finite Element Non-Hydrostatic Dynamical Core

We will describe the development of a non-hydrostatic dynamical core for the Department of Energy’s (DOE) Accelerated Climate Model for Energy (ACME) program. ACME is developing a high-resolution Earth System Model designed to run efficiently on DOE’s upcoming pre-exascale and exascale computers. ACME’s current hydrostatic dynamical core is from the High Order Method Modeling Environment (HOMME). Our non-hydrostatic core is being developed in HOMME and shares many of the underlying computational kernels including full support for variable resolution meshes. When running at hydrostatic scales, it is designed to be as similar as possible to the hydrostatic dycore. We use a terrain-following, vertically-Lagrangian pressure coordinate. The discretization relies on mimetic methods: spectral finite elements in the horizontal and finite differences in the vertical. We use a horizontally explicit vertically implicit IMEX approach for the time discretization. We introduce a moist potential temperature as a prognostic variable, and formulate the equations so that moist total energy are conserved through the use of mimetic discretizations. Initial results will be presented from several of the DCMIP idealized test cases.

Mark A. Taylor  
Sandia National Laboratories, Albuquerque, NM  
mataylo@sandia.gov

David M. Hall  
University of Colorado at Boulder  
Department of Computer Science  
david.hall@colorado.edu

Paul Ullrich  
University of California, Davis  
paulullrich@ucdavis.edu

Oksana Guba  
Sandia National Laboratories  
NM, USA.  
onguba@sandia.gov

### MS321

#### Complex Optimization for Big Computational and Experimental Neutron Datasets

We present a framework to use high performance computing to determine accurate solutions to the inverse optimization problem of big experimental data against computational models. We demonstrate how image processing, mathematical regularization, and hierarchical modeling can be used to solve complex optimization problems on big data. We also demonstrate how both model and data information can be used to further increase solution accuracy of optimization by providing confidence regions for the processing and regularization algorithms. We use the framework in conjunction with the software package SIMPHONIES to analyze results from neutron scattering experiments on silicon single crystals, and refine first principles calculations to better describe the experimental data.

Feng Bao  
Oak Ridge National Laboratory  
baof@ornl.gov

**MS321****Stochastic Dynamics for Cross-Scale Uncertainty Propagation in Molecular Systems**

Modeling uncertainty propagation with respect to time is important for time-dependent multiscale systems. In this work, the evolution of uncertainty associated with thermodynamics properties as the quantities of interests (QoIs) in molecular dynamics (MD) simulation is modeled at two time scales. At short time scale, the probability density distributions (PDFs) for QoIs are evolved at each step of MD simulation. At long time scale, a meta-model of uncertainty dynamics is proposed to propagate PDFs with much less computational time. An example of MD simulation with liquid argon with Lennard-Jones interatomic potential is used to demonstrate the feasibility of the meta-modeling approach.

Anh Tran, Yan Wang  
Georgia Institute of Technology  
anh.vt2@gatech.edu, yan.wang@me.gatech.edu

**MS321****A Multiscale Data-Driven Stochastic Method for Elliptic PDEs with Random Coefficients**

We propose a multiscale data-driven stochastic method (MsDSM) to study stochastic partial differential equations (SPDEs) with multiscale features in a multiquery setting. This method combines the advantages of the recently developed multiscale model reduction method and the data driven stochastic method (DSM). The MsDSM reduces both the stochastic and the physical dimensions of the solution. Numerical results are presented to demonstrate the accuracy and efficiency of the proposed method for several multiscale stochastic problems without scale separation.

Maolin Ci  
Caltech  
mci@cms.caltech.edu

Thomas Hou  
California Institute of Technology  
hou@cms.caltech.edu

Zhiwen Zhang  
Department of Mathematics,  
The University of Hong Kong,  
zhangzw@maths.hku.hk

**MS321****Conex Splitting Method for Transition State in Phase Field Model**

Convex splitting method is a useful numerical strategy to maintain unconditional energy stability for large time step size in solving the partial differential equations arising from steepest descent dynamics of energy functionals. In this work, we show how to use this strategy to help solve the problem of finding transition states, i.e., index-1 saddle points of the same energy functionals. Based on the previous work of iterative minimization formulation (IMF)

for saddle point (SIAM J. Numer. Anal., vol. 53, p1786, 2015), we introduce the convex splitting idea to minimize the auxiliary functionals at each cycle of the IMF. For any given convex splitting forms of the energy functional in the model, we present the general principle of constructing convex splitting forms for these auxiliary functionals and show how to avoid solving nonlinear equations by using two convex splitting forms. The new numerical scheme allows a much larger time step size than many traditional methods. The numerical results for the Ginzburg-Landau energy functional demonstrate that our new method not only helps to solve the subproblems with a less number of iterations but also considerably improves the overall performance of searching transition states.

Xiang Zhou  
Department of Mathematics  
City University of Hong Kong  
xiang.zhou@cityu.edu.hk

Shuting Gu  
City University of Hong Kong  
shutinggu2-c@my.cityu.edu.hk

**MS322****An Implementation of a Variational Immersed Boundary Method for Fluid-Structure Interaction using GRINS and libMesh**

The Immersed Boundary Method has been an important tool for studying fluid-structure interaction (FSI) problems, but is traditionally implemented through Dirac delta forcing terms in the fluid equations. In this presentation, we investigate an existing variational formulation which allows for a natural coupling into existing finite element (FEM) formulations. We implement this variational formulation in the GRINS multiphysics framework, built on the libMesh finite element library. Using GRINS allows for designing the immersed boundary kernel separately from other weak forms thereby providing a large amount of flexibility in both the fluid and solid properties such that they may be specified at runtime. The talk focuses on the architectural framework, implemented design strategies, and results relevant to many classes of problems in both soft matter and other application areas.

Boris Boutkov  
SUNY AT BUFFALO  
borisbou@buffalo.edu

Paul Bauman  
Mechanical and Aerospace Engineering  
University at Buffalo, State University of New York  
pbauman@buffalo.edu

**MS322****Dynamics of a Multicomponent Vesicle in Shear Flow**

Using a pseudo-spectral boundary integral method, we investigate dynamical patterns induced by inhomogeneous bending. Numerical results reveal that there exist novel phase treading and tumbling mechanisms with oscillatory inclination angle. We observe that tumbling dynamics can be triggered even for very low shear rate, and the excess arc-length required for tumbling is significantly smaller than the value required for a homogeneous case. These findings highlight the level of complexity of vesicle dynam-

ics due to inhomogeneous elasticity.

Kai Liu  
Mathematics  
University of California, Irvine  
liuk10@uci.edu

Gary Marple  
Department of Mathematics  
University of Michigan  
gmarple@umich.edu

Shuwang Li  
Department of Applied Mathematics  
Illinois Institute of Technology  
sli@math.iit.edu

John Lowengrub  
Department of Mathematics  
University of California at Irvine  
lowengrb@math.uci.edu

Shravan Veerapaneni  
Department of Mathematics  
University of Michigan  
shravan@umich.edu

### MS322

#### Stokes Flows for Non-Intact Moving Internal Boundaries

Calculations for streaming flows that include a deformable surface are of interest to a number of areas including physiology and chemical engineering. These surfaces need not be intact, but can possess a pore allowing the passage of once impermeant molecules through the membrane, as occurs in electroporation or cell lysis, for example. In order to make progress in modeling these phenomena, and eventually use the models to guide the development of medical applications, it is useful to know the details of the flow patterns and quantify the rate at which energy is dissipated by viscous losses. Recently, a fairly complete analytical description was obtained for motions of the spherical cap—a hollow sphere with a single hole—characterizing the streaming flow for pore closure and for leak out. Apart from possible applications in physiology and chemical engineering, these determinations have mathematical interest because of their role in the numerical analysis of mixed boundary-value problems where the internal boundary possesses an edge. In particular, an analysis of pressure and stress singularities near the edge reveals how the degradation in accuracy of the finite element method can be avoided through the construction of suitable local solutions.

Rolf Ryham  
Fordham University  
rryham@fordham.edu

James von Albade  
Department of Mathematics  
Fordham University  
jvonalbade@fordham.edu

### MS322

#### Modeling of the Interaction Between a Lipid Bilayer Membrane and a Solid Particle

The spreading dynamics of a thin layer of viscous New-

tonian fluid between an elastic sheet and a wetting solid substrate is examined using lubrication theory. On the wetting substrate an ultra thin film (precursor film) develops as a result of an intermolecular force between the fluid and the wetting solid substrate. Such a precursor film prevents the stress singularity associated with a moving contact line. Following the methodology by Glasner (2003), the effects of elasticity on the macroscopic contact line structure in the quasistatic limit are elucidated by an ordinary differential equation derived from an analysis of the energy and its dissipation. Similar to the case of a regular fluid interface with surface tension (capillary spreading), the elasto-capillary thin film profile also consists of a core at the center, an ultra thin film in the far field, and a contact line region where the core film profile connects smoothly to the precursor film. For capillarity-dominated spreading, the precursor film transitions monotonically to the core film. For elasticity-dominated spreading, a spatial oscillation of film height in the contact line region is found instead. In addition, it is found that elasticity causes a sliding motion of the thin film: the contact angle is close to zero due to elasticity and the contact line moves at a finite speed.

Yuan-Nan Young  
Department of Mathematical Sciences  
NJIT  
yyoung@njit.edu

Howard Stone  
Princeton University  
hastone@princeton.edu

### MS323

#### On Optimization Approaches to Boundary Conditions Setting in Computational Hemodynamics

The so-called Navier-Stokes equations for incompressible flow have now been used extensively to model hemodynamics problems. However, the required amount of data to close the system is in practice likely not to be attained. In a clinical setting, for example, we in general only have access to average fluxes from Echo-Doppler or due to measurement errors the data may not satisfy the incompressibility constraint, etc. In a direct approach, the researcher calls upon his experience to infer on the missing data to close the system. An alternative approach, the so-called optimization approach, consists in expressing the primal problem as a subproblem to one that minimize the distance between the available data and the data consistent with the model at hand. However, such a technique is in general way too expensive computationally. We will show how simpler surrogates to the Navier-Stokes model may indeed provide better solutions to a direct approach while being computationally efficient.

Adrien Lefieux, Alex Viguerie  
Emory University  
adrien.lefieux@emory.edu, aviguer@emory.edu

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

### MS323

#### Software Tools for the Assessment of Cardiovascular Diseases: The Emory Cardiac Toolbox Experi-

ence

Myocardial Perfusion Imaging (MPI) modalities, such as Single Photon Emission Computed Tomography (SPECT) and Positron Emission Tomography (PET), have attained widespread clinical acceptance as a standard of care for patients with known or suspected coronary artery disease. An important contribution to this success has been provided by the development and use in clinical environments of computerized techniques for the objective and robust quantification and interpretation of these imaging studies. The Emory Cardiac Toolbox (ECTb) has been implemented over the years as a pipeline to distribute the software tools that our team and other have researched, developed and validated to be clinically useful so that clinicians everywhere can benefit from this work. Fundamental attributes of the ECTb are: an extensive number of normal perfusion databases for both SPECT and PET studies; analysis of myocardial regional thickening from gated acquisitions to quantify left ventricular function and dyssynchrony; development of tools for the PET-derived quantification of myocardial hibernation and viability; development of a platform for the creation of three-dimensional displays that allow the integration of functional images with anatomical ones such as coronary computed tomography angiography (cCTA) by means of multimodality image fusion; extraction of measurements such as absolute myocardial blood flow and myocardial flow reserve from list-mode PET acquisitions.

Marina Piccinelli

Department of Radiology, Emory University  
Atlanta, USA  
mpiccin@emory.edu

### MS323

#### **Sickle Cell Anemia and Pediatric Strokes: Computational Fluid Dynamics Analysis in the Middle Cerebral Artery**

Abstract Not Available At Time Of Publication.

Manu Platt

Department of Biomedical Engineering, GA Tech  
Atlanta, USA  
manu.platt@bme.gatech.edu

### MS323

#### **Patient-Specific Simulations of the Ascending Aorta in Aortopathies**

Abstract Not Available At Time Of Publication.

Adelia Sequeira

Department of Mathematics  
Instituto Superior Tecnico  
adelia.sequeira@math.ist.utl.pt

### MS324

#### **Mean-Field Optimal Control Hierarchy in Consensus Models**

In this talk we will discuss the role of a policy maker on a large interacting multi-agent system as a mean field optimal control problem. Such control problems are constrained by a PDE of continuity-type, governing the dynamics of the probability distribution of the agent population, example will be presented in consensus dynamics. We will first present well-posedness results for the mean

field optimal controls in the stochastic and deterministic setting, and we will derive the first order optimality conditions. In the second part we will introduce a novel approximating hierarchy of sub-optimal controls based on a Boltzmann approach, whose computation requires a very moderate numerical complexity. We will compare the behavior of the control hierarchy with respect to the solution of the optimal control problem, providing numerical experiments for models in consensus formation.

Giacomo Albi

Dipartimento di Matematica e Informatica  
Università di Ferrara  
giacomo.albi@ma.tum.de

Young-Pil Choi

Technische Universität München  
young-pil.choi@ma.tum.de

Massimo Fornasier

Technical University Munich  
massimo.fornasier@ma.tum.de

Dante Kalise

Radon Institute for Computational  
and Applied Mathematics (RICAM)  
dante.kalise@oeaw.ac.at

### MS324

#### **Meanfield Games and Model Predictive Control**

Mean-Field Games are games with a continuum of players that incorporate the time-dimension through a control-theoretic approach. Recently, simpler approaches relying on the Best Reply Strategy have been proposed. They assume that the agents navigate their strategies towards their goal by taking the direction of steepest descent of their cost function (i.e. the opposite of the utility function). In this paper, we explore the link between Mean-Field Games and the Best Reply Strategy approach. This is done by introducing a Model Predictive Control framework, which consists of setting the Mean-Field Game over a short time interval which recedes as time moves on. We show that the Model Predictive Control offers a compromise between a possibly unrealistic Mean-Field Game approach and the sub-optimal Best Reply Strategy.

Michael Herty

RWTH Aachen University  
herty@igpm.rwth-aachen.de

### MS324

#### **Kinetic Models of Conservative Economies with Welfare Thresholds**

As a subset of the field of econophysics, kinetic exchange models of markets have been used to recover and analyze characteristics of wealth distributions. We develop two kinetic models of wealth redistribution that aim to bring individuals with wealth below a fixed welfare threshold up to it. The first model uses need-based transfers where welfare occurs via binary gifts from individuals above threshold to individuals below. The second model centrally collects and redistributes wealth. We discuss and compare features of each of these models.

Kirk Kayser, Dieter Armbruster

Arizona State University  
School of Mathematical and Statistical Sciences

kirk.kayser@asu.edu, dieter@asu.edu

Michael Herty  
RWTH Aachen University  
herty@igpm.rwth-aachen.de

Christian Ringhofer  
Arizona State University  
School of Mathematical and Statistical Sciences  
ringhofe@asu.edu

### MS324

#### Econophysical Financial Market Models: A Mean-Field Game Approach

In the recent years there has been a growing number of financial market models, described as large interacting particle systems, which are part of the research field econophysics. The starting point of our investigations is the well-known econophysical Levy-Levy-Solomon [Levy, Levy Solomon, A microscopic model of the stock market: cycles, booms, and crashes, 1994] model, where each financial agent is faced with the question of allocating capital between a risky and risk-free asset. Inspired by classical work of portfolio optimization by Markowitz [Markowitz, Portfolio selection, 1952] each agent is faced with an optimization problem, which needs to be solved in a game theoretic setting. Starting at the microscopic level, we formally derive a mean-field game model. Mean-field game theory has been mainly developed by Lions and collaborators [Lasry, Lions, Mean field games, 2007] and has been already applied to several economic models. The novelty of our work is the complex behavior of our state equations compared to several prototype examples. The mathematical model is described by a coupled system of two PDEs: A Fokker-Planck-Kolmogorov equation models the time evolution of the wealth distribution, and is coupled with a classical Hamilton-Jacobi-Bellmann equation, which describes the behavior of the value function. Finally, we discuss numerical examples and further perspectives

Torsten Trimborn  
RWTH Aachen University  
trimborn@mathcces.rwth-aachen.de

Martin Frank  
RWTH Aachen University  
Center for Computational Engineering Science  
frank@mathcces.rwth-aachen.de

Michael Herty  
RWTH Aachen University  
herty@igpm.rwth-aachen.de

### MS325

#### The Genesis of Critters: A Multi-part Story

Self-assembly in colloidal systems often requires finely tuning the interactions between particles. When colloids are active, or moving due to an external drive, the assembly is even harder to achieve. In this talk we will show that long-lived compact motile structures, called ‘critters’, can be formed just with hydrodynamic interactions with a wall. They naturally emerge from a fingering instability recently discovered in an experimental system of microrollers near a floor. We will focus on the genesis of this phenomenon which involves a cascade of instabilities: the formation of a ‘shock’ front due to the nonlocal nature of hydrodynamic

interactions, the destabilization of this front which leads to the formation of fingers, which then detach to form moving persistent autonomous clusters called ‘critters’. Combining nonlocal continuum models, linear stability analysis, 3D large scale simulations, and comparison with experiments, we will explain each step in detail and show how the distance of the particles to the nearby floor controls the dynamics of the system. These critters are a persistent state of the system, move much faster than individual rollers, and quickly respond to a changing drive. The formation of critters is robust to any initial conditions and our experiments suggest that similar structures are formed in a thermal colloidal system. We believe the critters are a promising tool for microscopic transport, flow, aggregation and mixing.

Blaise Delmotte  
Courant Institute  
NYU  
delmotte@cims.nyu.edu

Michelle Driscoll  
New York University  
md180@nyu.edu

Aleksandar Donev  
Courant Institute of Mathematical Sciences  
New York University  
donev@courant.nyu.edu

Paul Chaikin  
CSMR-NYU  
chaikin@nyu.edu

### MS325

#### Theory of Meiotic Spindle Assembly

The meiotic spindle is a biological structure that self-assembles from the intracellular medium to separate chromosomes during meiosis. It consists of filamentous microtubule (MT) proteins that interact through the fluid in which they are suspended and via the associated molecules that orchestrate their behavior. We aim to understand how the interplay between fluid medium, MTs, and regulatory proteins allows this material to self-organize into the spindle’s highly stereotyped shape. To this end we develop a continuum model that treats the spindle as an active liquid crystal with MT turnover. In this active material, molecular motors, such as dyneins which collect MT minus ends and kinesins which slide MTs past each other, generate active fluid and material stresses. Moreover nucleator proteins that are advected with and transported along MTs control the nucleation and depolymerization of MTs. This theory captures the growth process of meiotic spindles, their shapes, and the essential features of many perturbation experiments. It thus provides a framework to think about the physics of this complex biological suspension.

Sebastian Fuerthauer  
Simons Foundation and New York University  
sebastian.fuerthauer@gmail.com

Peter Foster, Daniel Needleman  
Harvard University  
pfoster@fas.harvard.edu, dan.needleman@gmail.com

Michael Shelley  
Courant Institute of Mathematical Sciences

New York University  
shelley@courant.nyu.edu

**MS325****A Fast Platform to Simulating Cellular Assemblies of Semi-Flexible Filaments and Motor Proteins**

We present a novel platform for the large-scale simulation of fibrous structures immersed in a Stokesian fluid and evolving under confinement or in free-space. One of the main motivations for this work is to study the dynamics of fiber assemblies within biological cells. For this, we also incorporate the key biophysical elements that determine the dynamics of these assemblies, which include the polymerization and depolymerization kinetics of fibers, their interactions with molecular motors and other objects, their flexibility, and hydrodynamic coupling. This work, to our knowledge, is the first technique to include many-body hydrodynamic interactions (HIs), and the resulting fluid flows, in cellular fiber assemblies. We use our method to investigate two important questions in the mechanics of cell division: (i) the effect of confinement on the hydrodynamic mobility of microtubule asters; and (ii) the dynamics of the positioning of mitotic spindle in complex cell geometries.

Ehssan Nazockdast, Abtin Rahimian  
Courant Institute of Mathematical Sciences  
New York University  
ehssan@cims.nyu.edu, arahimian@acm.org

Denis Zorin  
Computer Science Department  
Courant Institute, New York University  
dzorin@cs.nyu.edu

Michael J. Shelley  
New York University  
Courant Inst of Math Sciences  
shelley@cims.nyu.edu

**MS325****High-Order Adaptive Time Stepping in Vesicle Suspensions**

Accurate time stepping in interfacial dynamics, such as vesicle suspensions, is crucial for obtaining stable and accurate simulations. I will describe two algorithms that can be used to accomplish this task. First, a deferred correction method uses a low-order semi-implicit time integrator to construct high-order results. Second, physically conserved quantities are used to adaptively select appropriate time step sizes.

Bryan Quaife  
Florida State U.  
bquaife@gmail.com

**MS326****Optimal Sequential Experimental Design using Adaptive Transport Maps**

The *optimal* design of a sequence of experiments accounts for (1) feedback between experiments, and (2) future consequences of each design decision. Common design practices such as batch/open-loop design (choose all experiments simultaneously without feedback) and greedy/myopic design (optimally select the next experiment without accounting for future effects) are thus suboptimal. We formulate the

optimal sequential experimental design problem using dynamic programming, and maximize expected information gain under continuous parameter, design, and observation spaces. The dynamic program can only be solved approximately and numerically, and we find a near-optimal policy via approximate value iteration. To make this iteration tractable, we employ *transport maps* to represent non-Gaussian posteriors for continuous parameters, and to enable fast approximate Bayesian inference. We construct a Knothe-Rosenblatt map that couples a standard Gaussian to the joint distribution of designs, observations, and parameters, so that posterior maps can easily be obtained by conditioning on the design and observation values. The maps are constructed using sample trajectories from exploration and exploitation, favoring accuracy over state regions that are more likely to be visited. The exploitation measure is updated as better policies become available, thus improving the map accuracy. The overall method is demonstrated on a sensor placement problem for source inversion.

Xun Huan  
Sandia National Laboratories  
xhuan@sandia.gov

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

Andrew D. Davis  
MIT  
davisad@mit.edu

**MS326****Sub-Modularity Based Optimal Experimental Design of Kernelized Mis-Specified Dynamic Model**

In this study we aim to infer a functional correction for a mis-specified dynamic model, where the systems dynamics is only approximately known. Our primary objective is to faithfully represent the dynamics for unobserved initial conditions. Assuming the availability of observations of the system for various initial conditions, we propose a formulation for estimating correction terms in a Reproducing Kernel Hilbert Space (RKHS) of candidate correction functions. Further, we analyze the problem of performing efficient experimental design to find an optimal correction term under a limited experimental budget and experimental constraints. The problem is computationally intractable when formulated in the D-Bayes optimality framework. However, by introducing a suitable approximate proxy, the problem can be cast as a submodular optimization problem, for which there are computationally efficient and approximately optimal solvers. Our numerical experiments exemplify the efficiency of these techniques.

Gal Shulkind  
Massachusetts Institute of Technology  
shulkind@mit.edu

Lior Horesh  
IBM Research  
lhoresh@us.ibm.com

Haim Avron  
Tel Aviv University

haimav@post.tau.ac.il

### MS326

#### A Measure Space Approach to Optimal Experimental Design

In this talk we present a new approach for the optimal placement of measurement sensors in the context of optimal design of experiments for the estimation of finitely dimensional parameters describing a system of PDEs. Hereby the possible distributions of measurement sensors in an experimental domain  $\Omega$  is modelled by the set of positive Borel-measures  $M^+(\Omega)$ . The costs of the experiment are modelled by the total variation norm. This approach leads to optimal control problems of the form

$$\text{minimize}_{\omega \in M^+(\Omega)} \phi(C(\omega)) + \beta \|\omega\|_{M(\Omega)}$$

where  $\phi$  is a (smooth) optimality criterion acting on the eigenvalues of the Fisher-Information matrix  $C(\omega)$ ,  $C(\omega)_{ij} = \langle \partial_i S(\bar{q}) \partial_j S(\bar{q}), \omega \rangle$ , e.g.  $\phi(C(\omega)) = \text{trace}(C(\omega)^{-1})$ .  $\partial_i S(\bar{q})$  denotes the sensitivity of the parameter-to-state operator  $S$  with respect to the  $i$ -th parameter at an a priori guess  $\bar{q}$ . We examine the solvability of such problems and state optimality conditions. Furthermore we consider a variational discretization of this problem, i.e. we only discretize the sensitivities, but not the design space  $M^+(\Omega)$  and prove some convergence results. For the algorithmic solution we consider a generalized conditional gradient method combined with a semismooth newton method and a postprocessing strategy. The talk is completed by numerical experiments.

#### Daniel Walter

Technische Universität München  
walter@ma.tum.de

Boris Vexler  
Technical University of Munich  
Centre for Mathematical Sciences, M17  
vexler@ma.tum.de

Ira Neitzel  
Technische Universitaet Muenchen,  
Department of Mathematics  
neitzel@ma.tum.de

Konstantin Pieper  
Florida State University  
Dept of Scientific Computing  
kpieper@fsu.edu

### MS327

#### Two-Scale Simulation of Low-Alloyed Trip Steels based on Phase Transformation at the Microscale

Low-alloyed TRIP steels are characterized by a high ductility and strength at moderate production costs. These steels possess a heterogeneous multiphase microstructure, initially consisting of ferrite, bainite and retained austenite which is responsible for the mechanical properties. Upon deformation, a diffusionless, stress-induced, martensitic phase transformation from face-centered cubic austenite to body-centered cubic martensite is observed, enhancing ductility and strength. To incorporate this process in sheet metal forming simulations we focus on a multi-scale method in the sense of FE<sup>2</sup>, where additional microscopic boundary value problems are solved taking into account suitable representative volume elements. At the mi-

cro-scale, the material behavior of the matrix, consisting of ferrite and bainite, is modeled by an isotropic hyperelastic-plastic constitutive model. The inelastic processes in the austenitic inclusions involve the phase transformation from austenite to martensite and the inelastic deformation of these two phases. To capture this material behavior, an analytic homogenization approach describing the effective viscoplastic response of the two-phase (austenitic-martensitic) material is combined with a transformation model representing both the stress-dependent evolution of retained austenite and the deformations associated with the transformation. Numerical two-scale calculations will be presented to illustrate the performance of the approach.

#### Daniel Balzani

TU Dresden  
Institut of Mechanics and Shell Structures  
daniel.balzani@tu-dresden.de

Ashutosh Gandhi, Stefan Prueger  
TU Dresden  
Institute of Mechanics and Shell Structures  
"gandhi, ashutosh rajiv" jashutosh\_rajiv.gandhi@tu,  
stefan.prueger@tu-dresden.de

### MS327

#### Highly Scalable Implicit Solvers for Elasticity Problems on More Than Half a Million Parallel Tasks

The numerical solution of partial differential equations as, e.g., linear or nonlinear elasticity problems, on modern and future supercomputers requires fast and highly scalable parallel solvers. In this talk, we will focus on the different implicit nonlinear solvers and their performance we consider in the EXASTEEL project for the solution of the microscopic problems inside our scale bridging method FE2TI. Domain decomposition methods such as FETI-DP (Finite Element Tearing and Interconnecting - Dual Primal) and BDDC (Balancing Domain Decomposition by Constraints) are robust and efficient solvers for implicit problems arising in structural mechanics as, e.g., in a simulation of deformation processes. In this talk, we will present our highly efficient FETI-DP and BDDC implementations and show scalability to half a million parallel tasks on Mira (Argonne National Laboratory, USA) and JUQUEEN (Jülich Supercomputing Center, Germany) supercomputers. We will discuss several recent developments in FETI-DP and BDDC, i.e., nonlinear domain decomposition or nonlinear preconditioning, and the combination of FETI-DP with an algebraic multigrid method tailored for elasticity problems. Here we used the local neighborhood (LN) approach implemented in BoomerAMG from the hyper package, which provides an exact interpolation of the rigid body modes.

#### Martin Lanser

Universitaet zu Koeln  
Mathematisches Institut  
martin.lanser@uni-koeln.de

Axel Klawonn  
Universitaet zu Koeln  
axel.klawonn@uni-koeln.de

Oliver Rheinbach  
Technische Universität Bergakademie Freiberg

oliver.rheinbach@math.tu-freiberg.de

### MS327

#### FE2TI - Parallel Computational Scale Bridging for Dual-Phase Steels

Modern high-strength steels are simulated using the FE<sup>2</sup> multiscale method combined with efficient parallel FETI-DP domain decomposition/multigrid solvers. In this multiscale approach, in each Gaussian integration point of the macroscopic problem, a microscopic problem is solved on a representative volume element (RVE). A material law on the macroscale is not needed. The problems on the RVEs are coupled only through the macroscopic problem and thus can be solved in parallel. The implementation uses PETSc and efficient solver packages including Boomer-AMG, MUMPS and UMFPACK. Weak scalability results are presented, filling the complete JUQUEEN at FZ Jülich (458,752 BG/Q cores) and the complete Mira at Argonne National Laboratory (786,432 BG/Q cores).

#### Oliver Rheinbach

Technische Universität Bergakademie Freiberg  
oliver.rheinbach@math.tu-freiberg.de

#### Axel Klawonn

Universitaet zu Koeln  
axel.klawonn@uni-koeln.de

#### Martin Lanser

Universitaet zu Koeln  
Mathematisches Institut  
martin.lanser@uni-koeln.de

### MS327

#### Modeling of Polycrystals Using a Fcc/bcc Crystal Visco-Plasticity Theory Applying a Robust Complex-Step Derivative Approximation for the Tangent Moduli

The modeling of polycrystals using crystal visco-plasticity models lead to complicated formulations of the exact consistent tangent modulus tensor and therefore its derivation and implementation might be elaborate. In this case numerical approximations of the moduli can provide an alternative in particular for scientific development purposes, especially for the implementation of different complex hardening laws. Here, a robust numerical approximation of the fourth-order algorithmic consistent tangent tensor that can be implemented within FE-software is proposed using an innovative numerical derivative scheme. The role of consistent tangent moduli in the implicit nonlinear finite element method is important in order to achieve quadratic convergence within the Newton-Raphson iteration scheme. Furthermore, these moduli are needed to detect material instabilities especially in localization analysis. However, numerical tangent moduli of large-strain constitutive models using finite difference schemes typically suffer from round-off errors in floating point arithmetics and they are quite sensitive with respect to the perturbation especially in highly nonlinear problems. The complex-step derivative approximation scheme, which uses complex numbers in computation and can provide first derivatives with high accuracy that are almost identical to the analytical ones.

#### Lisa Scheunemann

Universitaet Duisburg-Essen  
Institut fuer Mechanik  
lisa.scheunemann@uni-due.de

#### Joerg Schroeder

Uni Duisburg Essen  
Institut f. Mechanik, Abteilung Bauwissenschaften  
j.schroeder@uni-due.de

#### Daniel Balzani

TU Dresden  
Institut of Mechanics and Shell Structures  
daniel.balzani@tu-dresden.de

#### Masato Tanaka

Toyota Central R&D Labs, Inc  
tanamasa@mosk.tytlabs.co.jp

### PP1

#### Preconditioning for Incompressible Two-Phase Flow

Two-phase flows arise in many coastal and hydraulic engineering applications such as the study of coastal waves, dam breaking scenarios, and the design of coastal structures. However, modelling two-phase incompressible flow with a level set formulation results in a variable coefficient Navier-Stokes system that is challenging to solve computationally. Here we consider preconditioners for such systems, looking to adapt efficient preconditioners for single-phase flows. In particular we consider systems arising from the application of finite element methodology and preconditioners based on approximate block factorisation. A crucial ingredient is a good approximation to the Schur complement which can be computed efficiently.

#### Niall Bootland

University of Oxford  
bootland@maths.ox.ac.uk

### PP1

#### Nonlinear Model Reduction in Computational Fluid Dynamics

”Developing effective reduced-order models (ROMs) for compressible, turbulent fluid flows is a challenging problem whose solution can enable an important breakthrough: the routine use of large eddy simulation (LES)—or even direct numerical simulation (DNS)—models in time-critical applications such as aircraft design. We present a number of advances to making this goal a reality, including 1) least-squares Petrov–Galerkin projection (which exhibits superior accuracy to Galerkin projection), 2) the sample-mesh concept (which enables small-footprint ROM simulations), and 3) structure preservation (which ensures the ROM conserves mass, momentum, and energy globally).”

#### Kevin T. Carlberg

Sandia National Laboratories  
ktcarlb@sandia.gov

### PP1

#### Scikit-Shape: A Python Package for Shape Optimization and Analysis

Many problems in science and engineering are expressed as shape optimization problems, in which the variable is a shape such as a curve in 2d or a surface in 3d. Examples are optimization of the shape of an airplane wing, and delineation of boundaries of biological structures in medical scans. We typically express such problems as energies with data (or target) mismatch and geometric regularization



tion 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 will be freely available as an open source package for the research community at: <http://scikit-shape.org>

Gunay Dogan  
Theiss Research, NIST  
gunay.dogan@nist.gov

### PP1

#### **JInv - A Flexible Julia Package for Parallel PDE Parameter Estimation**

jInv is a Julia framework for the solution of large-scale PDE constrained optimization problems. It supports linear and nonlinear PDE models and provides many tools commonly used in parameter estimation problems such as different misfit functions, regularizers, and efficient methods for numerical optimization. Also, it provides easy access to both iterative and direct linear solvers for solving linear PDEs. A main feature of jInv is the provided easy access to parallel and distributed computation supporting a variety of computational architectures: from a single laptop to large clusters of cloud computing engines. Being written in the high-level dynamic language Julia, it is easily extendable and yet fast. This poster gives an overview about the package, recent applications, and outlines future extensions.

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

Samy Wu Fung  
Emory University  
400 Dowman Dr, 30322 Atlanta, GA, USA  
samy.wu@emory.edu

Eran Treister  
University of British Columbia  
erant@bgu.ac.il

Eldad Haber  
Department of Mathematics  
The University of British Columbia  
haber@math.ubc.ca

### PP1

#### **Offline-Enhanced Reduced Basis Method Through Adaptive Construction of the Surrogate Parameter Domain**

Classical Reduced Basis Method (RBM) is a popular certified model reduction approach for solving parametrized partial differential equations. However, the large size or

high dimension of the parameter domain leads to prohibitively high computational costs in the offline stage. In this work we propose and test effective strategies to mitigate this difficulty by performing greedy algorithms on surrogate parameter domains that are adaptively constructed. These domains are much smaller in size yet accurate enough to induce the solution manifold of interest at the current step. In fact, we propose two ways to construct the surrogate parameter domain, one through an Inverse Cumulative Distribution Function (ICDF) and the other based on the Cholesky Decomposition of an error correlation matrix. The algorithm is capable of speeding up RBM by effectively alleviating the computational burden in offline stage without degrading accuracy. We demonstrate the algorithms effectiveness through numerical experiments.

Jiahua Jiang  
University of Massachusetts, Dartmouth  
jjiang@umassd.edu

### PP1

#### **Minisymposium: Solitary Waves and Shock Waves in Periodic and Random Media**

First-order hyperbolic systems do not describe dispersive waves, but dispersive effects can arise in such systems as a result of periodically- or randomly-varying coefficients or ambient state. This e-poster allows you to explore examples of this in the context of shallow water waves over variable bathymetry. Depending on the amplitude of the initial waves and the kind of variation of the bottom, a variety of effects can be observed, including wave breaking, formation of solitary waves, and transitions between the two. The poster includes interactive demonstrations of simulation results computed with Clawpack. This e-poster is part of the Minisymposium: Clawpack and GeoClaw - Software Developments and Applications.

David I. Ketcheson  
CEMSE Division  
King Abdullah University of Science & Technology  
david.ketcheson@kaust.edu.sa

### PP1

#### **Scalable Time-Stepping for PDEs Through Componentwise Approximation of Matrix Functions**

Krylov subspace spectral (KSS) methods are high-order accurate, explicit time-stepping methods with stability characteristic of implicit methods. This “best-of-both-worlds” compromise is achieved by computing each Fourier coefficient of the solution using an individualized approximation, based on techniques from “matrices, moments and quadrature” for computing bilinear forms involving matrix functions. This poster will present an overview of their derivation and essential properties, and also highlight ongoing projects aimed at enhancing their performance and applicability.

James V. Lambers  
University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

### PP1

#### **Adaptive Turbulence Simulations with Moving Domains and Multi-Phase Flow, with Applications in**

## Biomechanics and Renewable Energy

Although turbulent flows play a major role in many areas of natural sciences and engineering, numerical simulations are currently able to fully resolve a limited range of flows of practical interest, and a turbulence model is often required to make the simulation feasible. We present the latest developments in our implicit turbulence modeling framework based on adaptive finite elements. Compared to competing methods, ours has an advantage in that it features dual-based goal-oriented adaptivity to approximate a given functional while optimally employing the available computational resources [J. Hoffman et al. (2016). Computability and Adaptivity in CFD. in: *Encycl. Comput. Mech.* ]. Challenging unsteady configurations such as a full aircraft are now computable in acceptable times, estimating aerodynamic forces with an error of a few percentage points. Here we focus on initial results on how our methodology can be extended to include simulations featuring moving meshes and multi-phase flows. This greatly increases the number of solvable problems by including, for example, flows over objects with moving parts with a reduced need for re-meshing, notoriously a costly operation. In this regard, we plan to present left ventricle blood flow simulations, as a part of a collaboration aimed at clinical applications, and vertical axis turbines for renewable energy production, an international collaboration involving also industrial partners.

### Massimiliano Leoni

KTH Royal Institute of Technology  
Basque Center for Applied Mathematics  
leoni.massimiliano1@gmail.com

### Johan Hoffman

Royal Institute of Technology KTH  
jhoffman@kth.se

### Johan Jansson

School of Computer Science and Communication  
Royal Institute of Technology KTH  
jjan@kth.se

### Van Dang Nguyen, Jeannette Spuhler

KTH Royal Institute of Technology  
vdnguyen@kth.se, spuhler@csc.kth.se

## PP1

### Mixed-Integer PDE-Constrained Optimization

Many complex applications can be formulated as optimization problems constrained by partial differential equations (PDEs) with integer decision variables. Examples include the remediation of contaminated sites and the maximization of oil recovery; the design of next generation solar cells; the layout design of wind-farms; the design and control of gas networks; disaster recovery; and topology optimization. We will present emerging applications of mixed-integer PDE-constrained optimization, review existing approaches to solve these problems, and highlight their computational and mathematical challenges. We introduce a new set of benchmark set for this challenging class of problems, and present some early numerical experience using both mixed-integer nonlinear solvers and heuristic techniques.

### Sven Leyffer

Argonne National Laboratory

leyffer@mcs.anl.gov

## PP1

### Factorization Based Sparse Solvers and Preconditioners – Recent Developments for Superlu and Strumpack

Many extreme-scale simulation codes encompass multiphysics components in multiple spatial and length scales. The resulting discretized sparse linear systems can be highly indefinite, nonsymmetric and extremely ill-conditioned. For such problems, factorization based algorithms are often the most robust algorithmic choices among many alternatives, either being used as direct solvers, or as coarse-grid solvers in multigrid, or as preconditioners for iterative solvers which otherwise rarely converge. We present our recent work on the two software packages SuperLU and STRUMPACK. The first is based on a supernodal method and the second is based on a multifrontal method. Moreover, in STRUMPACK, we incorporate the data-sparse low-rank approximate factorization techniques such as HSS to achieve lower arithmetic and communication complexity as well as robust preconditioner. We will illustrate both theoretical and practical aspects of the methods, and demonstrate their performance on newer parallel machines, such as those with GPU and Intel Xeon Phi, using a variety of real world problems.

### Xiaoye Sherry Li

Computational Research Division  
Lawrence Berkeley National Laboratory  
xsli@lbl.gov

### Pieter Ghysels

Lawrence Berkeley National Laboratory  
pghysels@lbl.gov

### Chris Gorman

UC Santa Barbara  
gorman@math.ucsb.edu

### Artem Napov

Université Libre de Bruxelles  
anapov@ulb.ac.be

### Francois-Henry Rouet

Livermore Software Technology Corp.  
frouet@lstc.com

### Piyush Sao

Georgia Institute of Technology  
piyush3@gatech.edu

## PP1

### SpinDoctor: A Simulation Tool for Diffusion MRI

I will showcase SpinDoctor, a Matlab based simulation tool that generates the MRI signal arising from heterogeneous sample geometries. The user defines the geometry by specifying the shape and positions of biological cells and the diffusion MRI sequence to be simulated, the simulation tool generates a finite elements mesh of a computational volume containing the biological cells and solves the underlying partial differential equation governing the water proton magnetization during the diffusion MRI sequence.

### Jing-Rebecca Li

INRIA Saclay

jingrebecca.li@inria.fr

### PP1

#### Physically Constrained Kalman Filter

We present an augmented version of the standard Kalman filter in the case where a physical constraint exists in the system. By enforcing this constraint in the analysis step, a more accurate solution can be obtained for little additional computational cost. We illustrate the effectiveness of our method by comparing its results on several systems with conserved quantities to the results of the standard Kalman filter.

Erin Linebarger  
University of Utah  
Scientific Computing and Imaging Institute  
erinline@sci.utah.edu

Akil Narayan  
University of Utah  
akil@sci.utah.edu

Dongbin Xiu  
Ohio State University  
xiu.16@osu.edu

### PP1

#### Computing Reduced Order Models Using Randomization

In partial differential equations-based (PDE-based) inverse problems with many measurements, we have to solve many large-scale discretized PDEs for each evaluation of the misfit or objective function. In the nonlinear case, each time the Jacobian is evaluated an additional set of systems must be solved. This leads to a tremendous computational cost and this is by far the dominant computational cost for these problems. One can use reduced order models to reduce this cost. However, computing the reduced order model still requires solving many systems. In this talk, we discuss how to generate reduced order models using randomized techniques at a greatly reduced cost.

Selin Sariaydin, Eric De Sturler  
Virginia Tech  
selin@vt.edu, sturler@vt.edu

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

### PP1

#### A Fully-Coupled Discontinuous Galerkin Spectral Element Method for Two-Phase Flow in Petroleum Reservoirs

We present a discontinuous Galerkin spectral element method (DGSEM) to simulate two-phase incompressible fluid flow in petroleum reservoirs. Petroleum reservoirs are porous media with heterogeneous geologic features and permeability, and the flow of two immiscible phases involves sharp, moving interfaces. Fluid flow in porous media is governed by a generalized Darcy's law with relative permeability functions for each phase. The relative permeability introduces a strongly non-linear coupling between fluid velocity,  $\mathbf{u}$  and local saturation,  $s$ . The governing equations

of motion involve an incompressible pressure equation and a convective-diffusive transport equation for saturation,  $s$  with strong non-linearity associated with the  $(\mathbf{u}, s)$  coupling. Such coupled PDE systems present various computational challenges owing to sharp, moving fronts, and spatial discontinuities in material properties such as permeability, porosity, fluid viscosity. We develop a fully coupled numerical scheme using discontinuous Galerkin method with spectral element basis functions and implicit Rosenbrock time-stepping schemes. Stable numerical simulations are achieved without the use of slope limiters. High-order time accuracy and  $hp$ -convergence in spatial variables of DGSEM is demonstrated. We present the performance capabilities of DGSEM for simulating fluid flow in heterogeneous reservoirs with low-permeable regions and impermeable shale barriers.

Ankur Taneja, Jonathan Higdon  
University of Illinois at Urbana-Champaign  
taneja2@illinois.edu, jhigdon@illinois.edu

### PP1

#### Exahype - An Exascale Engine for Solving Hyperbolic PDEs: Various Equations and Quick Adaption to the Users Needs Becoming Open Source

The ExaHyPE project develops an exascale simulation engine for the solution of hyperbolic partial differential equations. We tackle these challenges with an approach based on the Discontinuous Galerkin method with high-order accuracy in space and time and a discretization based on structured space-tree meshes with adaptive refinement. ExaHyPE is an open-source engine and a key focus is the fast adaption to the research needs of other groups. An E-poster will be a great way to convey the philosophy of the engine character of ExaHyPE. We will prepare different problem setups for interactive presentation covering standard cases of hyperbolic conservation laws such as Euler flow, wave equations or special relativistic magneto hydrodynamics. Beyond that, there will be the possibility for the audience to experience ExaHyPE: We will together with and driven by the users exchange the PDE formulation, solve multiple PDEs at the same time, adapt the mesh and show the potential and flexibility of the engine. In a live demonstration and interaction with the visitors we will extend the simulation by additional user-specific parameters and investigate the implications to the whole simulation process. This interactive approach will be very useful concerning user input to us as the project consortium, will greatly convey the mission statement of the project and is only feasible with approaches such as an E-poster.

Dominic E. Charrier  
Durham University  
dominic.e.charrier@durham.ac.uk

Jean-Matthieu Gallard  
Technical University of Munich  
jean-matthieu.gallard@tum.de

Sven Köppel  
Goethe University Frankfurt  
koeppel@th.physik.uni-frankfurt.de

Vasco Varduhn  
TU München  
varduhn@tum.de

Tobias Weinzierl

Durham University  
tobias.weinzierl@durham.ac.uk

Michael Bader  
Technical University of Munich  
bader@in.tum.de

### PP1

#### Optimized Least-Squares Rational Filters for Interior Eigenvalue Problems

FEAST-like solvers have gained popularity for interior eigenproblems. Originally derived from the Cauchy's integration formula, a spectral projector is approximated using numerical integration. This results in a matrix-valued rational transfer function that is used in a linear system solve to dampen the unwanted eigenvalues. The poles of the transfer function depend on the integration rule that is used. Recent approaches interpret the spectral projector not as a numerical contour integration, but directly as a rational transfer function. We propose a non-linear Least-Squares (LS) optimization of the rational function. Such an approach can generate filter functions with varying properties, via appropriate constraints. For example, if the poles of the rational function have small imaginary parts the convergence of the resulting linear system solve may suffer. An appropriate constraint mitigates the problem, at the cost of a slightly worse convergence ratio. We focus on symmetric filter functions for hermitian eigenproblems, although our approach allows for generalization to asymmetric filter function, as well as, non-hermitian eigenproblems. We present: (1) a formulation of the optimization problem that allows fast solves; (2) a means to compare filter functions for a large sampling of "interesting" intervals; (3) filter functions that outperform the state-of-the-art for most desirable properties.

Jan Winkelmann  
AICES  
RWTH Aachen  
winkelmann@aices.rwth-aachen.de

Edoardo A. Di Napoli  
Juelich Supercomputing Centre  
e.di.napoli@fz-juelich.de

### PP2

#### A Numerical Method for High Energy X-Ray Source Shape Reconstruction

Quantifying the source shape of a pulsed power X-ray imaging system is essential in determining a lower bound for the overall resolution of the system. A method of reconstructing the source shape involves solving an inverse problem with an operator matrix of size  $n \times n$ , where  $n$  is the number of pixels in the image. Building and inverting such a large matrix quickly becomes intractable as  $n$  increases due to both memory and computational constraints, even in well-posed problems. The ill-posed nature of this problem adds further complications, making direct inversion unstable. Taking advantage of the structure of the operator matrix and utilizing the Projected Newton (PN) and Conjugate Gradient (CG) algorithms, we are able to perform reconstructions of images that are orders of magnitude larger than previously possible.

Jesse Adams  
The University of Arizona

jadams@math.arizona.edu

### PP2

#### Electrical Impedance Tomography for Damage Detection in Concrete

In this poster, the mathematical formulation of the inverse problem in electrical impedance tomography for damage detection in concrete is presented. The appropriate function spaces and regularization required to solve this ill-posed inverse problem is described. Both the deterministic and the statistical approaches for the inversion with preliminary results are discussed.

Sanwar Ahmad  
Clemson University  
1982  
sanwara@g.clemson.edu

Tao Ruan  
Department of Civil Engineering  
Clemson University  
truan@clemson.edu

Thilo Strauss  
Department of Pediatrics  
University of Washington  
thilos@uw.edu

### PP2

#### Simulation of Cerebrospinal Fluid Flow Using Finite Pointset Method

Cerebrospinal fluid (CSF) fills a system of cavities at the centre of the brain, surrounds the brain and spinal cord and is in free communication with the interstitial fluid of the nervous tissue. For problems involving CSF flow, the nervous tissue is thus often modelled as a poroelastic material. Further interesting aspects arise due to the rigid enclosure of the brain by the skull and the pressure oscillations originating in the large arteries of the brain. The finite pointset method (FPM) is a generalized finite difference method for the numerical solution of partial differential equations on an unstructured, meshfree point cloud. Since no neighborhood information is necessary between the points, it is particularly well suited for applications with changing geometries, such as free surface problems or fluid structure interactions. At the Fraunhofer Institute for Industrial Mathematics ITWM, FPM has been developed into a simulation tool for several types of industrial applications over the last two decades. The aim of this project is to examine in a simplified geometry, how the CSF-brain interactions can be simulated with the FPM package in its current form or where new development would first be necessary.

Emmanuel O. Akinlabi  
African In  
akinlabi.e.olutayo@aims-senegal.org

### PP2

#### Measuring and Modeling Bipartite Graphs with Community Structure

Network science is a powerful tool for analyzing complex systems in fields ranging from sociology to engineering to biology. Here, we focus on generative models of bipartite graphs. We propose two generative models that can

be easily tuned to reproduce the characteristics of real-world networks, not just qualitatively, but quantitatively. The measurements we consider are the degree distributions and the bipartite clustering coefficient, which we refer to as the metamorphosis coefficient. We define edge, node, and degreewise metamorphosis coefficients, enabling a more detailed understanding of the bipartite community structure. Our proposed bipartite Chung-Lu model is able to reproduce real-world degree distributions, and our proposed bipartite "BTER" model reproduces both the degree distributions as well as the degreewise metamorphosis coefficients. We demonstrate the effectiveness of these models on several real-world data sets.

Sinan G. Aksoy  
Department of Mathematics  
UC San Diego  
saksoy@ucsd.edu

Tamara G. Kolda  
Sandia National Laboratories  
tgkolda@sandia.gov

Ali Pinar  
Sandia National Labs  
apinar@sandia.gov

## PP2

### Implicit Unstructured Aerodynamics on Emerging Multi and Many-Core HPC Architectures

Shared memory parallelization of the flux kernel of PETSc-FUN3D, an unstructured tetrahedral mesh Euler code previously characterized for distributed memory SPMD for thousands of nodes, is hybridized with shared memory SIMD for hundreds of threads per node. We explore thread-level performance optimizations on state-of-the-art multi- and many-core Intel processors, including the first and second generation Intel Xeon Phi, Knights Corner and Knights Landing, respectively. While linear algebraic kernels are bottlenecked by memory bandwidth for even modest numbers of cores sharing a common memory, the flux kernel, which arises in the control volume discretization of the conservation law residuals and in the formation of the preconditioner for the Jacobian, is compute-intensive and effectively exploits contemporary multi-core hardware. We study its performance on the Xeon Phi in three thread affinity modes, namely scatter, compact, and balanced, with different configurations of memory and cluster modes on Knights Landing, with various code optimizations to improve alignment and reduce cache coherency penalties. The optimizations employed to reduce the data motion and cache coherency protocol penalties are expected to be of value other unstructured applications as many-core architecture evolves. We explore large-scale distributed-shared memory performance on the Haswell-based Cray XC40 supercomputer, and demonstrate that optimizations employed on Phi are beneficial in this context.

Mohammed A. Al Farhan  
KAUST  
mohammed.farhan@kaust.edu.sa

Dinesh Kaushik  
Qatar Foundation  
dkaushik@qf.org.qa

David E. Keyes  
KAUST

david.keyes@kaust.edu.sa

## PP2

### Padé Time-Stepping Methods for Solving Partial Differential Equations

One of the difficulties of the numerical solution of time-dependent ordinary and partial differential equations is the severe restrictions on time-step sizes for stability in explicit methods. Otherwise, challenging, generally nonlinear systems of equations in implicit schemes would be imposed to solve such problem. A class of explicit methods, based on use of Padé approximation, is introduced. They are inexpensive per time-step and they possess stability restrictions similar to the one offered by implicit schemes. Due to the rationality form of PTS, some numerical error occurs and then some a kind of control is imposed. We find that the Padé time stepping (PTS) showed a preferable behaviour when reaction diffusion equations are considered. We also notice that the PTS schemes have less computational time than the compared ones. Finally, numerical runs are conducted to obtain the optimal local error control threshold.

Said Algarni  
Assistant Professor - Department of Mathematics  
King Fahd University of Petroleum and Minerals  
garnis@kfupm.edu.sa

## PP2

### Asynchronous Task-Based Parallelization of Algebraic Multigrid

Exascale architectures will exacerbate the challenges of increased SIMT concurrency, reduced memory capacity and bandwidth per core, and vulnerability to global synchronization that already complicate programming for performance in today's petascale environment. Algebraic multigrid, the solver of choice in many large-scale simulations, scales well in the weak sense, with fixed problem size per node. However, its strong scaling within a node is challenging. Reducing synchronization and increasing concurrency are essential for AMG to exploit hybrid architectures. Recent communication-reducing improvements to classical additive AMG by Vassilevski and Yang increase communication-computation overlap, but remain bulk synchronous. We extend additive AMG with asynchronous task-based parallelism within a node, using OmpSs (from the Barcelona Supercomputer Center), along with MPI for internode communications. We implement a tiling approach to decompose the grid hierarchy into parallel units within task containers, gaining more than a factor of 2 in time to solution on small thread counts and still a factor of 1.5 on 32 threads within a Haswell node of the Cray XC40. We profile the performance of our approach compare to standard alternatives within the hypre library for the 3D Laplacian and the Auxiliary space Maxwell Solver (AMS).

Amani Alonazi, David E. Keyes  
KAUST  
amani.alonazi@kaust.edu.sa, david.keyes@kaust.edu.sa

## PP2

### Methodology for Optimized Generation of Virtual Environments Based on Hydroelectric Power Plants

Based on the possibility of real-time interaction with three-

dimensional environments through an advanced interface, Virtual Reality consist in the main technology of this work, used in the design of virtual environments based on real Hydroelectric Plants. Prior to the process of deploying a Virtual Reality System for operation, three-dimensional modeling and interactive scenes settings are very important steps. However, due to its magnitude and complexity, the generation of virtual power plants demands a high computing cost. This work aims to present a methodology to optimize the production process of virtual environments associated to real hydroelectric power plants. In partnership with electric utility CEMIG, several real hydroelectric power plants are present in the scope of this work. During the modeling of each environment, the authors applied techniques within the proposed methodology. After the evaluation of the computational techniques presented here, it is possible to confirm a reduction in the time required to complete each virtual hydroelectrical power plant. Thus, this work presents the current scenario about the development of virtual hydroelectric power plants and discusses the proposed methodology that seeks to optimize this process in the electricity generation sector.

Ígor A. Andrade Moraes, Alexandre Cardoso, Edgard Afonso Lamounier Jr, Milton Miranda Neto  
Federal University of Uberlândia  
igorandradesystem@gmail.com,  
alex.cardoso.ufu@gmail.com, elamounier@gmail.com,  
voidmmn@gmail.com

## PP2

### Explicit Continuous Finite Element Methods on Triangles

Performing explicit time advancement using continuous Finite Element Methods (FEMs) requires a mass matrix inversion. On quadrilaterals using the spectral element method, the mass matrix is diagonal and there is no difficulty using explicit methods, but on triangles this is not the case. There is no Gauss-Lobatto (GL) rule on triangles similar to that on quadrilaterals that allows a nodal basis and diagonal mass matrix. It has been suggested that an approximate diagonal mass matrix could exist for triangles that, although not associated with a nodal basis, allows fast and accurate explicit time advancement [Helenbrook, 2009]. In this work we define a linear transformation that maps between sets of basis functions suitable for continuous FEMs. The transformation is then used to prove that no GL equivalent approximate diagonal mass matrix exists for triangular elements. A lower-triangular mass matrix that could be used for explicit FEMs is presented as an alternative. We then discuss how the transformation is used in identifying the elements of this lower-triangular mass matrix and its associated basis functions.

Jay M. Appleton, Brian Helenbrook  
Clarkson University  
appletj@clarkson.edu, bhelenbr@clarkson.edu

## PP2

### A Two Phase Approach for Modeling Sedimentation of Pigments in Liquid Coating Formulations

Liquid coating formulations such as latex paints may form irreversible sediments during storage or transportation, causing significant economic losses to manufacturers or end-users. We develop a two phase mathematical model which takes into consideration the surface chemistry at the fluid-particle interface and vibration of the container to elu-

cidate the sedimentation and consolidation process. Our computer simulations indicate that formulating a slightly flocculated dispersion could significantly improve product shelf life.

Emmanuel O. Asante-Asamani  
Department of Mathematical Sciences  
University of Wisconsin-Milwaukee  
emmanuel.asante-asamani@drake.edu

Bruce Wade  
Department of Mathematical Sciences, UW-Milwaukee  
Milwaukee, Wisconsin 53201-0413  
wade@uwm.edu

Lei Wang  
University of Wisconsin, Milwaukee  
wang256@uwm.edu

## PP2

### Numerical Analysis of Set and Fuzzy Integral Equations and Its Applications

We study linear and nonlinear integral equations of functions with values in so-called L-spaces. We prove theorems of existence and uniqueness for solutions of such equations and provide algorithms to find approximate solutions. To develop the algorithms we use methods of approximation of functions with values in L-spaces, and the methods of approximate calculation of integrals of such functions. Obtained results include corresponding results for integral equations with set-valued and fuzzy-valued functions. The proposed algorithms can be used to solve some of the problems arising in numerical analysis (algorithms for the solution of IVP and BVP for differential equations with derivatives of the Hukuhara-type), in the theory of optimal control (calculation of reachable sets), medicine (prediction of the spread of infections), and engineering (design of feedback systems). Another area of application of the results is associated with propagation of uncertainty.

Vira Babenko  
Ithaca College  
vbabenko@ithaca.edu

## PP2

### Compressed Sensing with Corrupted Measurements and Its Application to Uncertainty Quantification

A major problem in signal and image processing is to recover a signal or image from undetermined sets of measurements. Rather than the traditional sampling at a high rate and then compress the sampling data, compressed sensing (CS) allows us to directly sense the data at a lower sampling rate. This research focuses on a natural generalization of compressed sensing- CS with corruptions, which is to address the signal recovery problem when the measurements are corrupted and become unreliable. We compare and contrast several algorithms for this problem and illustrate both analytically and empirically the optimal algorithm. Practical applications of CS with corruption include clinical MRI scans, transmission failures in sensor networks, and processor failures in large-scale computational modeling. In particular, we demonstrate the application of this technique to the problem quantifying uncertainty in physical systems by computing polynomial chaos expansions of high-dimensional functions from corrupted

data.

Casie Bao, Ben Adcock  
Simon Fraser University  
cbao@sfu.ca, ben\_adcock@sfu.ca

## PP2

### Numerical MOOC: Collaborating in Open Education for CSE

The course "Practical Numerical Methods with Python" launched in 2014 as an "indie" massive open online course (MOOC), simultaneous to on-campus courses at three universities. Prof. Barba spearheaded the effort, creating infrastructure and content in the form of modules to teach the foundation of scientific computing with Python. In 2015, a new on-campus course adopted the curriculum and methods of "Numerical MOOC," expanded with a new course module. We wish to share the experiences of collaborating openly in teaching CSE, like we collaborate in research. We have all the course materials on GitHub, and contributions have come in from course instructors as well as avid students. Our basic curriculum can be expanded, modified, remixed, etc., in the open-source model. The course's stacked learning modules are somewhat self-contained, each one motivated by a problem modeled by a differential equation (or system of Des), building new concepts in numerical computing, new coding skills and ideas about analysis of numerical solutions. They cover methods for time integration of simple dynamical systems (systems of ordinary differential equations); finite-difference solutions of various types of partial differential equations (hyperbolic, parabolic or elliptic); assessing the accuracy and convergence of numerical solutions; and using the scientific Python libraries to write these numerical solutions. (Course repository: <https://github.com/numerical-mooc/numerical-mooc>)

Lorena A. Barba  
Department of Mechanical and Aerospace Engineering  
George Washington University  
labarba@gwu.edu

Ian Hawke  
Southampton University  
i.hawke@soton.ac.uk

Bernard Knaepen  
Université Libre de Bruxelles  
bknaepen@ulb.ac.be

## PP2

### Krylov Subspace Spectral Methods for Navier-Stokes in Cylindrical Geometries

Existing time-stepping methods for PDEs such as Navier-Stokes equations are not as efficient or scalable as they need to be for high-resolution simulation due to stiffness. The failure of existing time-stepping methods to adapt to changes in technology presents a dilemma that is becoming even more problematic over time. By rethinking approaches to time-stepping, dramatic gains in efficiency of simulation methods can be achieved. Krylov subspace spectral (KSS) methods have proven to be effective for solving time-dependent, variable-coefficient PDEs. The objective of this research is to continue the development of KSS methods to provide numerical solution methods that are far more efficient and scalable to high resolution for Navier-Stokes equations. So far, KSS methods have been applied

only on 1-dimensional, 2-dimensional, and 3-dimensional rectangular domains, but current work is extending them to polar domains using polar coordinates and expansions in Legendre polynomials instead of Fourier series. We will utilize these techniques for Navier-Stokes equations on rectangular domains as well as polar and cylindrical domains.

Brianna Bingham  
The University of Southern Mississippi  
brianna.bingham@usm.edu

James V. Lambers  
University of Southern Mississippi  
Department of Mathematics  
James.Lambers@usm.edu

## PP2

### Shape-Memory Polymers: A 3D Phenomenological Constitutive Model for Engineering Applications

Shape-memory polymers (SMPs) are smart materials able to store a temporary shape and to recover the original shape upon an external stimulus; the most common are thermo-responsive SMPs. Thanks to this peculiar behavior, SMPs can be used in a variety of advanced applications, such as self-adjustable utensils, biomedical devices, deployable structures, microsystems, smart sensors and actuators. The increasing interest for SMPs, together with the vast range of the available polymer types, motivates a deep investigation of the material behavior as well as the introduction of effective constitutive models. The present work proposes a user-friendly constitutive model for thermo-responsive SMPs, formulated within a three-dimensional finite-strain and thermomechanically consistent mathematical framework. The goal of the model is an easy applicability to engineering purposes, therefore it has a phenomenological character and it can reproduce the main constitutive behavior of SMPs using a low number of intuitive material parameters. Several numerical simulations have been performed to demonstrate the ability of the model to reproduce the SMP shape-memory cycle, both heating-stretching-cooling and cold drawing shape-fixing procedures, and the non-ideal behavior of real SMPs. Comparisons with experimental data from literature have been presented for model validation. The application of the proposed model to the simulation of realistic biomedical devices is also reported.

Elisa Boatti  
John A. Paulson School of Engineering and Applied Sciences  
Harvard University  
elisa.boatti@unipv.it

Giulia Scalet, Ferdinando Auricchio  
DICAr, University of Pavia  
via Ferrata 3, 27100 Pavia, Italy  
giulia.scalet@unipv.it, auricchio@unipv.it

## PP2

### The Stability Region for Feedback Control of the Wake Behind Twin Oscillating Cylinders

Linear feedback control has the ability to stabilize vortex shedding behind twin cylinders where cylinder rotation is the actuation mechanism. Complete elimination of the wake is only possible for certain Reynolds numbers and cylinder spacing. This is related to the presence of asymmetric unstable modes in the linearized system. We

investigate this region of parameter space using a number of closed-loop simulations that bound this region. We then consider the practical issue of designing feedback controls based on limited state measurements by building a nonlinear compensator using linear robust control theory with and incorporating the nonlinear terms in the compensator (e.g., using the extended Kalman filter). Interpolatory model reduction methods are applied to the large discretized, linearized Navier-Stokes system and used for computing the control laws and compensators. Preliminary closed-loop simulations of a three-dimensional version of this problem will also be presented.

Jeff Borggaard, Serkan Gugercin  
Virginia Tech  
Department of Mathematics  
jborggaard@vt.edu, gugercin@vt.edu

Lizette Zietsman  
Virginia Tech  
Interdisciplinary Center for Applied Mathematics  
lzietsma@vt.edu

## PP2

### Iterative Hybrid Methods with Wavelet Denoising

Solving ill-posed inverse problems is difficult due to noise amplification during inversion. Regularization is employed to dampen the most severe components of these effects, but an inevitable consequence is that features of the desired solution are filtered out during the process. A two-step process called Fourier-Wavelet Regularized Deconvolution (ForWaRD) has been proposed in the signal processing literature where Fourier inversion is followed by wavelet-based denoising in order to alleviate inverted noise and obtain denoised images. In this work, we extend this framework to a wider variety of inverse problems by considering iterative hybrid methods followed by post-wavelet denoising. Particular attention is paid to selecting proper choices of parameters for regularization and wavelet denoising. Numerical examples include deblurring of spatially varying blur and tomographic reconstruction problems.

Matthew Brown  
Department of Mathematics  
Virginia Polytechnic Institute and State University  
brownm12@vt.edu

Julianne Chung  
Department of Mathematics  
Virginia Tech  
jmchung@vt.edu

## PP2

### Computational Advances by Variational Formulation for Breast Cancer Detection Through Electrical Impedance Tomography

Electrical Impedance Tomography (EIT) is a rapidly developing non-invasive medical imaging technique recently gaining popularity in various medical applications including breast screening and cancer detection. Despite its obvious benefits (no harm to the tissues, technologically easier and much cheaper), nowadays it is still considered as an alternative or complimentary technique to x-rays mammography and magnetic resonance imaging. A higher level of reliability for EIT techniques could be achieved by improving the mathematical theory by adding a thorough research on the variational formulation for the inverse problem of

cancer detection. The improved mathematical foundations will give a rise to a new generation of highly sophisticated computational algorithms which comprise the methodology of solving the inverse problem and dealing with multiple complications caused by its ill-posedness. Increased performance and reliability of the new engineering design procedure will be achieved by adding proven numerical techniques for proper regularization, parameter space upscaling and re-parameterization.

Vladislav Bukshtynov  
Florida Institute of Technology  
Department of Mathematical Sciences  
vbukshtynov@fit.edu

Ugur G. Abdulla  
Florida Institute of Technology  
abdulla@fit.edu

## PP2

### A Strategy to Improve 3D Arrangement Modeling on Virtual Electrical Substations

Virtual Reality Environments in Power Substations provides a new supervisory control paradigm. The existence of a virtual environment, geometrically equal to the real substation, reduces the difference between the mental models built by the field operators and by the System Operation Center employees, improving their communication. These systems may be used as visualization interfaces for power systems simulators and training environments for students, teachers, field operators and other professionals. These applications developing process is complex, involving several activities such as programming, 3D modeling, virtual scenes construction and usability studies. This work presents a GUI strategy for the creation of field arrangements based on Scene Graphs, to: speed up the Virtual Environments of Electrical Substations' development process, avoid mistakes from the VE construction, reduce development time and improve quality and visual fidelity of the scene. These tools are evaluated during the development of Virtual Environments related to 51 real Power Substations, in Brazil. As case study and tools evaluation, two different sized virtual substations were built by two different methods: the manual, where every virtual models are inserted manually, and the semi-automatic, where most of the virtual models are inserted with the help of automatic tools. The obtained results present a considerably large gain in relation to the semi-automatic and manual methods of construction.

Camilo Lellis Barreto Junior  
Federal University of Uberlândia  
barretojunior@gmail.com

Alexandre Cardoso  
Federal University of Uberlandia  
Faculty of Electrical Engineer  
alexandre@ufu.br

Edgard Lamounier, Nicholas Duboc  
Federal University of Uberlandia  
lamounier@ufu.br, nick\_cduboc@hotmail.com

Paulo Prado  
Energy Company of Minas Gerais  
prmprado@cemig.com.br



prmprado@cemig.com.br

## PP2

### Project Research Model Canvas An Auxiliary Tool to Create and Manage Research Projects

Canvas is a tool for strategic and direct form of information visualization of a project elements, and is based on the concept of a screen subdivided into blocks that can be used simultaneously by multiple people over a brainstorm or presentation already conceived ideas in concepts. The Model Canvas Project Research is based on successful administrative examples, especially in the innovation sector. It approaches in a visual way the essential components of a project, facilitating the identification of steps, team management, time and resources. It can be used as an aid tool to document what is proposed more clearly and concisely. Since during the text of traditionally written projects the elements are arranged in a straight stream without further concerns with the connection between the ideas presented, leaving the essence of the idea on each fragmented part of the text, one after another. The proposed Canvas intend to contribute with the reduction of time to develop an efficiency project in a composition design as a whole. Its application is essential at the beginning of the project, even if it wasnt written, becoming also a tool to track the progress according to what was previously established.

Hiago A. Silva  
Federal University of Uberlandia  
hiago13@gmail.com

Alexandre Cardoso  
Federal University of Uberlandia  
Faculty of Electrical Engineer  
alexandre@ufu.br

Edgard Lamounier  
Federal University of Uberlandia  
lamounier@ufu.br

## PP2

### A Virtual Reality Based Approach to Improve Human Performance and to Minimize Safety Risks When Operat-Ing Power Electric Systems

Power systems require continuous operation for reasons of public safety, emergency management and business continuity. Companies today control power systems by 2D line diagrams, whereas a substation in the field is a 3D space. Many control center operators have never been into a real substation environment. This fact reduces human performance when it comes to operating the power system, since the operator has to associate 2D and 3D worlds. This association isnt intuitive, since a 2D line diagram doesnt represent the details of a real substation. Therefore, it is important to seek new alternatives to ensure that systems are designed in a way to optimize human performance and minimize risks. Virtual Reality (VR) is known to provide the feeling of being there and this paper proposes a Virtual Reality approach for the simulation, training and control of power substations. In this approach, a substation is replicated according to its dimensions, using pictures, videos and floor plans. Next, by web service, data from a supervisory system is allocated to each component in the virtual substation, so the operator can attain access to all the information required for possible intervention. Cemig, a power company, which has over 50 transmission substations spread across the state of Minas Gerais, Brazil, is a

research partner to this project. All the features explored in this work have the capacity to increase human performance when operating a power electric substation.

Alexandre C. Silva  
Federal University of Uberlandia  
acs.carvalho10@gmail.com

Alexandre Cardoso  
Federal University of Uberlandia  
Faculty of Electrical Engineer  
alexandre@ufu.br

Edgard Lamounier  
Federal University of Uberlandia  
lamounier@ufu.br

Paulo Prado  
Energy Company of Minas Gerais  
prmprado@cemig.com.br  
prmprado@cemig.com.br

Milton Miranda, Gerson Lima, Camilo Barreto, Nicholas Duboc  
Federal University of Uberlandia  
voidmmn@gmail.com, gfm169@gmail.com, barretojunior-mail@gmail.com, nick\_cduboc@hotmail.com

## PP2

### Second-Order Langevin Markov Chain Monte Carlo

Bayesian methods for system identification and estimation are critical for quantifying uncertainty and for making robust predictions. Dynamical systems approaches to Markov Chain Monte Carlo (MCMC) have become a popular choice for solving these Bayesian inference problems as they can handle high dimensional and complex posterior distributions. Approaches like Hamiltonian Monte Carlo (HMC) allow samplers to achieve high acceptance rates and effective sample sizes by exploiting the structure of the posterior distribution to explore areas of high probability. We present our algorithm, Second-Order Langevin MCMC (SOL-MC), a stochastic dynamical system based MCMC algorithm which uses a damped second-order Langevin stochastic differential equation (SDE) to sample a desired posterior distribution. This SDE combines Hamiltonian dynamics using a Euclidean or Riemannian metric with an OrnsteinUhlenbeck process. The Hamiltonian dynamics exploit the structure of the posterior to encourage exploring high probability regions, while the OU process introduces diffusion to reduce correlation. Since this method is based upon an auxiliary dynamical system, we can apply ideas from dynamics and control theory to inform the choice of parameters to best tradeoff the drift and diffusion processes. We apply SOL-MC to a challenging system identification problem to illustrate the effectiveness of this method and gain insight into its characteristics.

Thomas A. Catanach  
California Institute of Technology  
tcatanac@caltech.edu

James Beck  
Division of Engineering and Applied Science  
California Institute of Technology

jimbeck@caltech.edu

## PP2

### A Splitting Discontinuous Galerkin Method for the Coupling of Flow and Geomechanics

A numerical method is proposed and analyzed for solving the Biot problem where a splitting technique is used. The interior penalty discontinuous Galerkin method is employed for the spatial discretization and the backward Euler method is employed for the time discretization. Unlike other splitting techniques, this method is not iterative allowing a speed up of the computations. We provide a convergence analysis of the scheme along with numerical simulation to validate the theoretical results.

Nabil Chaabane  
Virginia Tech  
nabilch@vt.edu

## PP2

### Pricing and Hedging in a Stochastic Volatility Model with Finite Element Method

We consider computing numerical solutions of Heston partial differential equation (PDE), which is frequently used in computational finance for pricing problems of financial derivatives in a stochastic volatility environment. The Heston PDE is a time-dependent, two-dimensional convection-diffusion-reaction equation. The finite difference method is commonly used to obtain numerical solutions of the PDE but it has its limitation on the shape of domains and smoothness of terminal and boundary conditions. We utilize finite element methods to solve the Heston PDE for European call options. Since its closed-form solution is available for a European call option with a numerical approximation of an integral, relative errors of the results are obtained. Using the numerical solutions of the Heston PDE, we can compute the Greeks of European call options, which are partial derivatives used to hedge in financial markets. We also consider the Heston PDE for digital options, whose initial condition has a jump discontinuity. The finite element method is applied to obtain the numerical solutions and the results show the accuracy and efficiency of the method.

Jungmin Choi  
East Carolina University  
choiju@ecu.edu

## PP2

### Fast and Efficient Hyperbolic Embeddings for Large Graphs

Because hyperbolic space has properties that make it amenable to graph representations, there is significant interest in developing scalable embedding methods for this purpose. Such embeddings are often used for constant-time approximate distance calculations, which are much more effective than distances calculated via breadth-first shortest path methods. In this paper we introduce an improved landmark-based hyperbolic embedding algorithm for large-scale graphs. Whereas previous algorithms compute the embedding by using the derivative-free simplex optimization method, this algorithm uses a limited-memory quasi-Newton optimization method (L-BFGS) with analytic gradients. The L-BFGS method is not only significantly faster, but it also produces higher-quality embed-

dings. Moreover, we are able to include the hyperbolic curvature as a variable in the optimization. For comparison, we use Rigel, a the gradient-free Nelder-Mead simplex algorithm, with fixed curvature. Our method shows significantly consistent improvements in the accuracy of the shortest path distance calculations. Tests are performed on community data sets available on the SNAP (Stanford Network Analysis Project) database.

Kenny Chowdhary  
Sandia National Laboratories  
kchowdh@sandia.gov

## PP2

### Computational Modeling of Shock Wave Reflections Over a Wedge

Computational Fluid Dynamics (CFD) modeling is used to efficiently simulate the supersonic flow field of shock-wave reflections in compressible flow with alternated Mach number, viscosity, and altitude. The research investigated the reflection of shock-waves in a computational shock tube problem. The research can be split into two categories of study: To find the computational range where flow features are accurately resolved in an acceptable time frame. To compare and validate the shock-wave reflections of computational results to experimental data of the Applied Numerical Algorithm Group.

Uthman Clark  
Tuskegee University  
uclark6030@mytu.tuskegee.edu

## PP2

### Direct Solution of High-Dimension Sylvester Equations

Sylvester equations occur frequently in several areas of applied mathematics including control theory and solution of PDEs. The most frequent case involves matrices as the unknown. In the general case, the unknown objects are high-dimensional arrays, or multiarrays. The direct solution of these general equations can be formulated in terms of Kronecker sums of matrices which adds several Kronecker products of the various matrices with appropriate identity matrices. However, when the matrices are large, alternative computations are needed to avoid computing Kronecker products. Our approach extends the standard approach available for two-dimensional arrays to the general case of higher-dimensional multiarrays by using diagonalization of the matrices or Schur decompositions for each of the non-diagonalizable matrices. By splitting off the diagonals from the strictly triangular part, the inverse of the Kronecker sum of triangular matrices will result in a set of recursive matrix multiplication with levels bounded by the sum of the nilpotency of the strictly triangular parts minus the dimension of the multiarray (where null matrices have nilpotency of one). The complexity of the proposed approach will grow linearly with the dimension of the multiarray. Numerical examples are included in the paper, e.g. a high-accuracy finite difference solution of a three dimensional Poisson equation approximated by large number of grids.

Tomas Co  
Department of Chemical Engineering  
Michigan Technological University

tbco@mtu.edu

**PP2****Dislocation Dynamics Simulations of Void and Precipitate Strengthening**

The strength of crystalline material is controlled, to a large extent, by the motion of dislocations. In materials with a high density of microstructural features, the motion of dislocations is restricted, resulting in increased strength. The small scale plasticity occurring in the vicinity of microstructure is a fundamentally multiscale phenomena. Bridging the characterization of individual dislocations at the atomistic scale with the macroscopic plastic response from cooperative dislocation motion at the continuum scale remains an open challenge. A method well suited for bridging this gap is discrete dislocation dynamics (DDD), where plasticity is explicitly captured by the motion of individual dislocations. In this work, we present a scalable algorithm for modeling DDD with microstructural effects. The method involves coupling a highly parallel DDD simulator for bulk materials (ParaDiS) with a highly parallel finite element method (FEM) solver to capture microstructural effects. We use the coupled DDD-FEM simulator to model the strengthening effects of nanoscale voids and impenetrable precipitates. We also perform large scale simulations to investigate the collaborative effects of dislocation and obstacle hardening mechanisms.

Joshua C. CroneUS Army Research Laboratory  
joshua.crone.civ@mail.mil

James J. Ramsey

Army Research Laboratory  
james.j.ramsey14.civ@mail.mil

Jaroslaw Knap

U.S. Army Research Laboratory  
jaroslaw.knap.civ@mail.mil**PP2****A Multilevel Low-Rank Preconditioner for Indefinite Linear Systems**

Sparse nonsymmetric indefinite linear systems have long been difficult to solve via iterative methods. In this work we extend the Multilevel Schur Low-Rank preconditioner of Xi, Li, and Saad to this difficult case. This method generates a tree structure  $\mathcal{T}$  that represents a hierarchical decomposition of the original matrix. This decomposition gives rise to a block structured matrix  $A_l$  at each level of  $\mathcal{T}$ . An approximate inverse based on the block  $LU$  factorization of  $A_l$  is computed at each level where the inverse of the Schur Complement  $S_l = C_l - E_l B_l^{-1} F_l$  is approximated by  $C_l^{-1}$  plus a low-rank correction term. The low-rank correction term is computed by several steps of the Arnoldi process. Numerical results will show that this preconditioner is robust with respect to indefiniteness and symmetry for discretized PDEs and publicly available test problems.

Geoffrey DillonVirginia Tech University  
gdillon@umn.eduYuanzhe Xi, Yousef Saad  
University of Minnesota

yxi@cs.umn.edu, saad@umn.edu

**PP2****Multilevel Stochastic Collocation in Linear Gyrokinetics**

The analysis of uncertainty in plasma fusion problems is a challenging task, both mathematically as well as from a computational perspective. In this contribution, we employ the established plasma microturbulence code GENE (<http://genecode.org/>) and we focus on 5D linear gyrokinetic eigenvalue problems to understand the impact of microinstabilities in magnetically confined plasma. This is accomplished by analyzing the real and imaginary parts of the dominant eigenvalue, which represent the growth rate and frequency of the microinstabilities. Since input parameters such as temperature gradients of ions and electrons are inherently uncertain, these simulations need to be performed within the framework of uncertainty quantification. We address the challenges posed by the resulting stochastic problem using multilevel stochastic collocation and generalized polynomial chaos expansion. We perform a hierarchical decomposition in both phase space and parameter domains. In the former, we consider a hierarchy of discretization grids with a resolution increasing by a constant factor, whereas in the parameter domain, we employ sparse grids. In particular, we use the sparse grid combination technique, constructed using non-uniform, nested Leja points to create a surrogate for the calculation of polynomial chaos expansion's coefficients. Furthermore, we exploit the non-intrusiveness of our stochastic approach and we simulate the underlying problem using two layers of parallelism.

Ionut-Gabriel FarcasTechnical University of Munich  
Department of Scientific Computing in Computer Science  
farcasi@in.tum.de

Hans-Joachim Bungartz

Technical University of Munich, Department of Informatics  
Chair of Scientific Computing in Computer Science  
bungartz@in.tum.de

Tobias Goerler

Max Planck Institute for Plasma Physics  
Group of Turbulence in Laboratory and Astrophysical Plasmas  
tbg@ipp.mpg.de

Denis Jarema

Technische Universitaet Muenchen  
Department of Scientific Computing  
jarema@in.tum.de

Tobias Neckel

Technische Universität München  
neckel@in.tum.de**PP2****Identification of Gene Regulatory Models Using Data-Driven Reduction of the Chemical Master Equation**

The biological processes in a cell are induced and regulated by an intricate network of reacting chemical species such as DNA, RNA and proteins. Building predictive mod-

els for these cellular interactions requires parameter values calibrated to experimental observations. The intrinsically stochastic nature of cellular interactions is well captured by the modeling framework of the chemical master equation (CME). The CME describes the time evolution of a discrete-state, continuous-time Markov process, where the state space is defined by combinations of molecule counts of the cellular species. Since the size of the problem grows exponentially in the number of species considered, parameter identification using CME solutions is computationally costly and only practical for small networks. To overcome this limitation, we first project the state space of the CME onto a smaller subspace defined by a wavelet decomposition of experimentally measured data. After this reduction, we use the Metropolis Hasting algorithm to explore parameter uncertainties of the model and we compare them to results of the full CME analysis. We tested our approach on the two-species genetic toggle switch, and we achieved similar accuracy in parameter identification, but with a ten-fold speedup in computation time.

Zachary Fox  
School of Biomedical Engineering  
Colorado State University, Fort Collins  
zrfox@rams.colostate.edu

Ania-Ariadna Baetica  
Controls and Dynamical Systems  
California Institute of Technology  
abaetica@caltech.edu

Huy D. Vo  
University of Alabama  
hvo@crimson.ua.edu

Brian Munsky  
Department of Chemical and Biological Engineering  
Colorado State University, Fort Collins  
munsky@colostate.edu

## PP2

### Setting Tunable Solver Parameters with Performance Models

As machine architectures get more complex, so does the importance of correctly tuning solver parameters to ensure good performance. This can create large spaces that are time-consuming to explore, which is a real burden to users who are focused on getting results for their simulations and would ideally not have to spend time tuning them. In this poster, we present our experiences using performance models to tune these parameters. We show how performance models enable the straightforward selection of these parameters so that good performance is obtained with minimal user effort for two solvers, the algebraic multigrid solver BoomerAMG and the multigrid-in-time solver XBraid.

Hormozd Gahvari  
LLNL  
gahvari1@llnl.gov

William D. Gropp  
University of Illinois at Urbana-Champaign  
Dept of Computer Science  
wgropp@illinois.edu

Kirk E. Jordan  
IBM T.J. Watson Research  
kjordan@us.ibm.com

Jacob B. Schroder, Martin Schulz, Ulrike M. Yang  
Lawrence Livermore National Laboratory  
schroder2@llnl.gov, schulzm@llnl.gov, umyang@llnl.gov

## PP2

### Schur Complement Domain Decomposition for Large Scale Kinetic Transport Problems

Implicit linear kinetic equations for neutral particle transport can be solved explicitly via a procedure termed sweeps. Sweeps induce a directed acyclic graph of dependency between calculations which makes the algorithm partly parallelizable, but not completely. For large scale computations, sweep algorithms eventually scale poorly. We propose using a Schur complement domain decomposition to make a trivially parallelizable algorithm. The trade-off is the necessity to iterate until convergence is achieved. We compare the sweep algorithm to the domain decomposed algorithm at large scale.

Charles K. Garrett  
Oak Ridge National Laboratory  
ckgarrett@lanl.gov

Kevin Procopio  
University of Pennsylvania  
procopio@seas.upenn.edu

## PP2

### A New Goal-Oriented A Posteriori Error Estimation for 2D and 3D Saddle Point Problems in $hp$ Adaptive Fem

In this research we introduce a new approach on goal-oriented a posteriori error estimation for an automatic  $hp$ -Adaptive Finite Element Method. The presented method is based on the classical dual-weighted algorithm. The proposed goal-oriented error estimation requires the solution of local dual problems on patches. The idea of having local patch problems is to apply Clément and Scott–Zhang type interpolation operators to estimate point values with the finite element polynomials. The reliability and also the efficiency of the proposed a posteriori error estimator which are the upper and the lower bounds of the functional error, have been proved. Moreover, the performance of the proposed goal-oriented a posteriori estimator for both  $h$ - and  $hp$ -Adaptive FEM has been investigated in numerical examples for Saddle point problem.

Arezou Ghesmati, Wolfgang Bangerth, Bruno Turcksin  
Texas A&M University  
aghesmati@math.tamu.edu, bangerth@colostate.edu, turcksinbr@ornl.gov

## PP2

### A Subspace Pursuit Method to Invert the Refractivity Profile within the Marine Atmospheric Boundary Layer

Inferring electromagnetic (EM) propagation characteristics (e.g. EM ducting) within the marine atmosphere boundary layer (MABL) from data in real time is crucial for maritime navigation and communications. The propagation of EM waves is well modeled by a partial differential equation: a Helmholtz equation. A natural way to address the current MABL characterization inverse problem is to minimize an objective function which models the difference between the EM observations and the analogous propagation predicted

by the PDE. However, the optimization of the resulting objective function is intractable due to its extreme multimodality. We propose an alternative solution that relies on the properties of the PDE, but does not involve solving the full forward model. Ducted environments induce a field which can be decomposed into a few propagating, trapped modes. These modes are a subset of the solutions to a Sturm-Liouville eigenvalue problem. We design a new objective function that performs a projection of the data onto the low dimensional subspace spanned by these eigenvectors. Our method identifies the properties of the environment through an optimization over parameters defining a MABL refractivity profile of that objective function. We show that the resulting optimization problem is much easier to solve than the one that arises in the classical method, and describe how to solve the associated non-linear eigenvalue problem efficiently; leading to a faster than real-time method.

Marc Aurele T. Gilles

Cornell University, Center for Applied Mathematics  
mtg79@cornell.edu

Christopher J. Earls, David Bindel  
Cornell University  
cje23@cornell.edu, bindel@cs.cornell.edu

## PP2

### A Resistive Magneto-Hydrodynamic Numerical Model in the Context of Cell-Averaged Adaptive Multiresolution Methods: Verification Tests

Magnetohydrodynamics enables us to understand the intrinsic behavior of macroscopic conducting fluid phenomena. To study non-ideal fluids, i. e., those where there is no conservation of magnetic flux, we introduce dissipative effects into the equations. In this work, we are interested in resistivity, obtained from Ohm's law, when the resistivity does not vanish over time, and there is no conservation of energy density anymore. The ohmic term can trigger interesting resistive instabilities in the physical problem. For the numerical simulation, we present a resistive magnetohydrodynamic model with Harten's adaptive multiresolution analysis for cell-averages. In simulations of high computational cost, the adaptivity can reduce considerably the number of cells needed for the simulation, by creating a computational mesh that adjusts according to the local regularity of the numerical solution. This work presents verification tests and shows speed-up obtained with this methodology.

Anna Karina F. Gomes

Brazilian Institute for Space Research  
São José dos Campos, Brazil  
annakfg@gmail.com

Margarete Domingues  
Instituto Nacional de Pesquisas Espaciais  
margarete.domingues@inpe.br

Odim Mendes  
INPE, Sao Jose dos Campos, Brazil  
odim.mendes@inpe.br

Kai Schneider  
Aix-Marseille University  
I2M-CNRS, Centre de Mathematiques et Informatique

kai.schneider@univ-amu.fr

## PP2

### A Simplified Human Birth Model: Translation of a Rigid Cylinder Through a Passive Elastic Tube

This work uses a simplified numerical model to explore the forces on an infant during human birth. Numerical results are compared with the results of a physical model which represents the fetus moving through the birth canal using a rigid cylinder (fetus) that moves at a constant velocity through the center of a passive elastic tube (birth canal). The entire system is immersed in a highly viscous fluid; low Reynolds number allows the Stokes equations to approximate fluid behavior. The pulling force necessary to move the rigid inner cylinder at a constant velocity through the tube is measured, and considered along with the time-evolving behavior of the elastic tube. The discrete tube through which the rigid cylinder passes has macroscopic elasticity matched to the tube used in the physical experiment. The buckling behavior of the elastic tube is explored by varying velocity, length, and diameter of the rigid cylinder, and length of the elastic tube. More complex geometries as well as peristaltic activation of the elastic tube can be added to the model to provide more insight into the relationship between force and velocity during human birth.

Roseanna Gossmann

Tulane University  
rpealate@tulane.edu

Alexa Baumer  
George Washington University  
abaumer@gwmail.gwu.edu

Lisa J. Fauci  
Tulane University  
Department of Mathematics  
fauci@tulane.edu

Megan C. Leftwich  
George Washington University  
mleftwich@gwu.edu

## PP2

### Validation of Large Fluid Dynamic Simulations of Complex Geometries with 3D Printing

Computational fluid dynamics (CFD) plays an increasing important role in circulatory disease modeling. Advances in high performance computing have made it possible to simulate fluid flow through complex geometries at high resolutions. However, for simulations to impact clinical care, detailed validation must be conducted with realistic data sets. For a CFD application focus on simulating blood flow in vascular geometries, standard benchmark flows may not fully provide a comprehensive assessment of the CFD application's stability and validity. For instance, a simulation of flow through a coarcted aorta must accurately resolve the narrow branches of the aortic arch and handle the high Reynolds numbers in the coarctation. A comparison with experimental results for a physiologically relevant flow in a vascular geometry resolves this deficiency. 3D-printing allows the same geometry to be used for computation and experiment. HARVEY is a parallel hemodynamics application based on the lattice Boltzmann method. We compare the results from HARVEY with particle image velocimetry

(PIV) results for flow in patient-specific femoral arteries and aortas with coarctations. For this comparison, HARVEY simulations used as many as 32768 cores of a Blue Gene/Q supercomputer.

John Gounley  
Old Dominion University  
john.gounley@duke.edu

Rafeed Chaudhury, Priya Nair, Girish Pathangey, Kevin Winarta  
Arizona State University  
rafeed@asu.edu, pnair3@asu.edu, gpathang@asu.edu, kwinarta@asu.edu

Justin Ryan  
Phoenix Children's Hospital  
jrryan@asu.edu

Erik W. Draeger  
Lawrence Livermore Nat. lab.  
draeger1@llnl.gov

David Frakes  
Arizona State University  
dfrakes@asu.edu

Amanda Randles  
Duke University  
amanda.randles@duke.edu

## PP2

### New Clustering Algorithms to Identify Nonlinear Behavior During a Car Crash Simulation

During the design phase of a new car, several thousand of model variants need to be analyzed for crash safety since a late change in the design is not feasible. The vast amount of simulation runs needed to assess and improve the crash behavior calls for the use of model order reduction to speed up the response time for the designing engineer. Since a car crash has highly nonlinear behavior like large deformations or complex contact scenarios, nonlinear model reduction needs to be applied. Unfortunately, nonlinear methods are not yet as evolved as linear ones. We therefore suggest a combined approach by only reducing linear behaving parts with linear model reduction methods and apply nonlinear or no model reduction to the remaining part of the car model. A speed-up is only of practical use if it is applicable to commercial crash software that is used widely in the industry. Unfortunately, only heuristic approaches remain for the identification of nonlinearly behaving areas in this setting. We first define measures that quantify the degree of nonlinear behavior, e.g., based on the scatter of nodal positions, stress or strain among several simulation runs with only little modification in the input parameters. Then, self-developed clustering algorithms are applied that use these measures as input in order to identify areas with probably nonlinear behavior. These algorithms are faster and more reliable than state of the art approaches for this separation problem.

Dennis Grunert, Joerg Fehr  
Institute of Engineering and Computational Mechanics  
University of Stuttgart  
dennis.grunert@itm.uni-stuttgart.de, joerg.fehr@itm.uni-

stuttgart.de

## PP2

### Nonconforming Immersed Finite Element Spaces For Elliptic Interface Problems

We shall present a unified framework to study two classes of nonconforming immersed finite element (IFE) spaces constructed using the average integral value degrees of freedom over each edge on either triangular or rectangular mesh. On every interface element, the shape functions are piecewise polynomials defined on sub-elements separated either by the actual interface or its line approximation. We construct the Sherman-Morison systems to show the existence and uniqueness of those shape functions. Besides we establish the multi-edge expansion by the standard multi-point Taylor expansion and derive a group of identities which enable us to show that these IFE spaces have optimal approximation capability.

Ruchi Guo  
Virginia Polytechnic Institute and State University  
The Department of Mathematics  
ruchi91@vt.edu

Tao Lin  
Virginia Tech  
tlin@vt.edu

## PP2

### A Generalized Constitutive Relation Error

A generalized constitutive relation error (GCRE) is proposed in an analogous form to Fenchel-Young inequality on the basis of the key idea of Legendre-Fenchel duality theory. In this work, Legendre-Fenchel duality is reviewed with application to the establishment of dual variational formulation of the potential energy principle for hyperelastic problems, which are typical cases of the constitutive relations expressed by mutually convex conjugate energy functionals. Then the constitutive relation error (CRE) defined by the duality for hyperelasticity is generalized in an analogous form to Fenchel-Young inequality. The GCRE is linked with the global errors of some admissible solutions for the problem in question, and is of wide applicability, especially in *a posteriori* error estimations of numerical methods. A class of elliptic variational inequalities is also examined using the proposed approach and a strict upper bound of global energy errors of admissible solutions is obtained.

Mengwu Guo  
Department of Civil Engineering  
Tsinghua University  
gmw13@mails.tsinghua.edu.cn

Weimin Han  
Department of Mathematics  
University of Iowa  
weimin-han@uiowa.edu

Hongzhi Zhong  
Department of Civil Engineering  
Tsinghua University  
hzz@tsinghua.edu.cn

## PP2

### Efficient Computation of Sobol' Indices for

## Stochastic Models

Sobol indices are used extensively for global sensitivity analysis of deterministic models. However, stochastic models are replacing their deterministic counterparts in many applications. Generalizing Sobol indices to stochastic models poses both theoretical and computational challenges. We provide a theoretical framework to define the stochastic Sobol indices and present an efficient method for their computation. Preliminary error analysis is given and numerical results presented showing the efficiency and value of our proposed method.

Joseph L. Hart  
North Carolina State University  
jlhart3@ncsu.edu

Pierre Gremaud  
Department of Mathematics  
North Carolina State University  
gremaud@ncsu.edu

Alen Alexanderian  
NC State University  
alexanderian@ncsu.edu

## PP2

### Locally Adaptive Discriminant Analysis

Supervised image segmentation techniques utilize user-defined training data (parts of the image where the user is certain of each class) to predict a class label for each region in the image. Many classical methods rely on sharp changes in luminance, chromaticity, or texture and images without these features are difficult to segment correctly. We present a supervised segmentation method named Locally Adaptive Discriminant Analysis (LADA) that efficiently handles images containing classes with spatially varying intensities. This local method trains the algorithm with spatially local training data, decreasing the effect of noise and intensity gradients on the segmentation.

Margaret C. Hock  
The University of Alabama in Huntsville  
mch0021@uah.edu

Marylesa Howard  
National Security Technologies, LLC  
howardmm@nv.doe.gov

Tim Meehan  
National Security Technologies LLC  
meehanbt@nv.doe.gov

## PP2

### Prestructuring to Eliminate Dense Rows from Sparse Matrices

The presence of a dense row in an otherwise sparse matrix can significantly affect the performance of several direct solvers. Through an unconventional application of a null space method and the right choice of null basis, dense rows can be eliminated from the matrix efficiently and may result in a much faster direct solve. This user-level modification of a linear system prior to solution by a direct solver can be thought of as a prestructuring technique, an analogue to preconditioning techniques for iterative solvers.

Jason Howell

College of Charleston  
howelljs@cofc.edu

## PP2

### Galerkin Differences: Very High-Order Accurate and Energy Stable PDE Discretizations

We discuss the Galerkin Difference method for providing high-order accurate, energy stable schemes for hyperbolic PDE systems. As in classical FEM the method is built using the framework of Galerkin projection, but unlike traditional FEM the scheme achieves high-order accuracy through inclusion of neighboring degrees of freedom. We discuss implementation and results for the acoustic and elastic wave equations in two space dimensions.

John Jacangelo, Jeffrey W. Banks  
Rensselaer Polytechnic Institute  
jjack12@gmail.com, banksj3@rpi.edu

## PP2

### A Weno-Based Method of Lines Transpose Approach for Vlasov Simulations

We will show a high order implicit Method of Lines Transpose (MOLT) method based on a weighted essentially non-oscillatory (WENO) methodology is developed for one-dimensional linear transport equations and further applied to the Vlasov-Poisson (VP) simulations via dimensional splitting. In the MOLT framework, the time variable is first discretized by a diagonally implicit strong-stability-preserving Runge-Kutta method, resulting in a boundary value problem (BVP) at the discrete time levels. Then an integral formulation coupled with a high order WENO methodology is employed to solve the BVP. As a result, the proposed scheme is high order accurate in both space and time and free of oscillations even though the solution is discontinuous or has sharp gradients. Moreover, the scheme is able to take larger time step evolution compared with an explicit MOL WENO scheme with the same order of accuracy. The desired positivity-preserving (PP) property of the scheme is further attained by incorporating a newly proposed high order PP limiter. We perform numerical experiments on several benchmarks including linear advection, solid body rotation problem; and on the Landau damping, two-stream instabilities, bump-on-tail, and plasma sheath by solving the VP system. The efficacy and efficiency of the proposed scheme is numerically verified.

Yan Jiang  
Michigan State University  
jiangyan@math.msu.edu

## PP2

### MultiMesh: Fem on Arbitrarily Many Intersecting Meshes

We will present a technology called MultiMesh, which allows the finite element method to be used on arbitrarily many intersecting meshes. Using this technique, a computational domain  $\Omega$  in  $\mathbb{R}^d$  consisting of many parts  $\Omega = \cup_i \Omega_i$  can be described by individual meshes on each part  $\Omega_i$ . The authors suggest that this technique can alleviate geometrical modeling, since one is not restricted to describing a single domain using a single mesh. By allowing for separate meshes of arbitrary individual parts of a computational domain one avoids the sometimes costly and difficult construction of a single mesh. The interface conditions, such

as continuity, between the meshes of  $\Omega_i$  are enforced weakly using the Nitsche method. The key challenge in setting up the procedure is to create appropriate quadrature rules for performing the necessary volume and boundary integrals. We will in this talk present a novel procedure to systematically construct quadrature rules with appropriate positive and negative weights using a basic result from combinatorics. The procedure is quite general: it is valid for arbitrary dimensions, both conforming and non-conforming discretizations as well as for high order approximations. Its implementation in the open source FEniCS framework allows for easy use on a wide variety of problems.

August Johansson  
Simula Research Laboratory  
Norway  
august@simula.no

Benjamin Kehlet  
Simula Research Laboratory  
benjamik@simula.no

Mats G. Larson  
Department of Mathematics  
Umea University  
mats.larson@math.umu.se

Anders Logg  
Chalmers University of Technology  
logg@chalmers.se

## PP2

### Validated Computation of Topological Entropy in Hybrid Dynamical Systems

A hybrid dynamical system is a discrete dynamical system generated by composing the time- $\tau$  map of a continuous dynamical system with a continuous map. Hybrid maps are becoming more popular in biological and physical models due to their ability to model variables which change continuously in time (modeled by the ODE), yet also undergo periodic discontinuous changes called “kicks” (modeled by a continuous map on phase space). Global dynamics for these systems can be simulated using numerical methods to perform the integration. However, sharp rigorous error estimates for approximate solutions of ODEs are difficult to obtain. This makes it difficult to apply existing techniques for computing validated global dynamics in classical discrete systems. We present a rigorous integration method for ODEs which combines with the well developed machinery for obtaining outer approximations for discrete dynamical systems. The resulting combination allows analysis of global dynamics for hybrid dynamical systems. We illustrate the method by proving the existence of chaos in a particular hybrid map and computing validated bounds on the topological entropy.

Shane D. Kepley  
Florida Atlantic University  
skepley@fau.edu

William D. Kalies  
Florida Atlantic University  
Department of Mathematical Sciences  
wkalies@fau.edu

## PP2

### Finite Volume Methods for Visco-Plastic Flow and

## Dispersive Waves

Submarine landslides can impose great danger to the underwater structures and generate destructive tsunamis. Submarine debris flows often behave like visco-plastic materials, and the Herschel-Bulkley rheological model is known to be appropriate for describing the motion. In this work, we present numerical schemes for the visco-plastic debris flows using finite volume methods in Eulerian coordinates with two horizontal dimensions. The landslide deformation is coupled to the water column and simulated in the Clawpack framework. For the propagation of the tsunamis, the shallow water equations and the Boussinesq-type equations are employed to observe how important the wave dispersion is. We present the newly developed Bouss-Claw which solves the Boussinesq equations with the finite volume method. Finally, two cases in central Norway, i.e. the subaerial quick clay landslide at Byneset in 2012, and the submerged tsunamigenic Rissa landslide in 1979, are both presented for validation. The research leading to these results has received funding from the Research Council of Norway under grant number 231252 (Project TsunamiLand) and the European Unions Seventh Framework Programme (FP7/2007-2013) under grant agreement 603839 (Project ASTARTE).

Jihwan Kim  
University of Oslo  
jhkim2@uw.edu

## PP2

### Simulation of Rarefied Gases Using Hyperbolic Moment Models

The simulation of rarefied gases is challenging due to the high computational effort of particle methods like DSMC on the one hand and the lack of accuracy of standard models like Euler or Navier-Stokes on the other hand. Moment models close the gap between both approaches and yield good accuracy in the transition regime for moderate to large Knudsen numbers. Several new hyperbolic moment models have been derived recently, see [J. Koellermeier, R. P. Schaerer and M. Torrilhon. A Framework for Hyperbolic Approximation of Kinetic Equations Using Quadrature-Based Projection Methods, *Kinet. Relat. Mod.* 7(3), 531–549, 2014], [Y. Fan, J. Koellermeier, J. Li, R. Li and M. Torrilhon. Model Reduction of Kinetic Equations by Operator Projection, *J. Stat. Phys.*, 162(2), p. 457–486, 2016]. We compare the different models with respect to their accuracy in some 1D and 2D test cases to show the capability of moment models to be used for the simulation of rarefied gases.

Julian Koellermeier, Manuel Torrilhon  
RWTH Aachen University  
koellermeier@mathcces.rwth-aachen.de,  
mt@mathcces.rwth-aachen.de

## PP2

### Condition Assessment and Prognosis Using Fluid-Structure Interaction Model Updating Within a Stochastic Inversion Framework

It is essential to identify damage within a structure as early as possible in order to propose corrective measures to prevent mechanical failure ultimately extending service life. Such damage detection can be effected through non-destructive means that employ the updating of a physics model describing the system of interest. The resulting in-



verse problem is concerned with locating and characterizing the damage, by considering the structural dynamic response, before and after the onset of damage. Such model-based approaches also afford prognosis potential; as a calibrated physics model is subsequently available for the damaged system. In this work, a partitioned fluid-structure forward modeling capability is developed from combining an open source CFD tool (OpenFOAM) with an open source CSD solver (CU-BEN) in tightly coupled framework that affords stable solutions. The forward model is then employed within a Bayesian inversion framework, based on sampling strategies, to arrive at the inverse solution for the damage assessment. The focus of this work is propellers operating within high Reynolds number incompressible flows.

Justyna Kosianka, Christopher Earls  
Cornell University  
jwk239@cornell.edu, earls@cornell.edu

## PP2

### Melting-Refreezing Cycles of Sea Water Ice An Enthalpy-Based Fixed Grid Approach

For water ice undergoing melting and re-freezing cycles, a natural question concerns the ice thickness, hence the evolution of the water ice interface and its spatio-temporal temperature evolution. In the absence of convection and forces, this question is classically known as the Stefan problem. It consists of two free-boundary second order partial differential equations for each phase. Closure is provided by considering local energy conservation at the interface, a jump condition imposed on the heat flux, which is commonly referred to as the Stefan condition. In reality, however, the situation often is more complex and e.g. includes natural water convection due to density variations. In this contribution, we present an extension to the so-called enthalpy-porosity approach that is capable of simultaneously tracking the water ice boundary, while taking forced convection, as well as natural convection due to salinity and temperature driven buoyancy into account. It is implemented into the open source finite element library dealii. Local mesh refinement and a Non-linear Newton iteration account for the strong nonlinearities at the solid-liquid interface. We describe the mathematical model, as well as our numerical finite element approach. Next, we present convergence studies for problems of reduced complexity. We conclude by showing and discussing simulation results for prototype problems in physical regimes that resemble real world situations.

Julia Kowalski, Marco Schoos, Alexander Gary Zimmermann  
AICES  
RWTH Aachen University  
kowalski@ices.rwth-aachen.de, marco.schoos@rwth-aachen.de, zimmermann@ices.rwth-aachen.de

## PP2

### Resolution Analysis of Pod-Based Imaging Using Tikhonov Regularized Geophysical Inversion

Nonlinear resolution analysis is investigated using proper orthogonal decomposition (POD) for basis-constrained inversion of sparse geophysical data. Resolution analysis is applied to the inverted models from an adaptive POD-constrained inversion method to quantify the uncertainty in representing models of hydrogeophysical process-driven systems when using physics-based regularization. Model resolution from the POD-constrained approach is analyzed

in its ability to reconstruct representative parameters of an electrical conductivity model from electrical resistivity data of a subsurface contaminant plume. In addition, the resolution density and Backus-Gilbert resolution kernels are computed from the model resolution matrix to further validate the uniqueness and stability of the predicted electrical conductivity measurements.

Shyla R. Kupis, Stephen Moysey  
Clemson University  
Department of Environmental Engineering and Earth Sciences  
skupis@g.clemson.edu, smoysey@clemson.edu

Taufiqar Khan  
Clemson University  
Department of Mathematical Sciences  
khan@g.clemson.edu

## PP2

### Bayesian Optimization with a Finite Budget: An Approximate Dynamic Programming Approach

Bayesian Optimization is a technique adapted to the optimization of expensive objective functions. It sequentially constructs a statistical model for the objective function and leverages this surrogate to select the next design to evaluate. One of the current limitation of Bayesian optimization is the greedy selection of such design which leads to suboptimal optimization performances. To address this limitation, we propose a lookahead approach to make decisions that maximize a long-term reward. The proposed approach uses rollout, a closed-loop approximate dynamic programming technique to define this long-term reward. We numerically demonstrate the performance increase of such lookahead approach on several test functions.

Remi Lam, Karen E. Willcox  
Massachusetts Institute of Technology  
rlam@mit.edu, kwillcox@MIT.EDU

David Wolpert  
Santa Fe Institute  
dhw@santafe.edu

## PP2

### Alternative Algebraic Structures for Modelling and Computation

Various fields of application can benefit from using operations in non standard algebras rather than classical frames for computation (linear algebra, real and complex numbers). In fact, these alternative algebraic structures allow to really fit the physical and computational aspects as already demonstrated in robotics, dynamic rotation control, signal and image processing or continuous mechanics. Many other physical problems could take advantage from these alternative algebras such as quantum physics, electromagnetism, waves, ... Largely based on the specific possible properties of multiplication (non commutativity, ring structure with or without division), these structures introduce a path to tackle non linearity. A preliminary exploration of this novel path will be presented.

Jean-Baptiste Latre  
Cerfacs  
latre@cerfacs.fr

Philippe Ricoux

TOTAL SA  
philippe.ricoux@total.com

Françoise Chatelin  
Cerfacs  
chatelin@cerfacs.fr

## PP2

### Stochastic Simulation of Multilevel Monte Carlo on Multi-GPU Systems

We explore solving systems of stochastic differential equations by implementing the Milstein method using antithetic Multilevel Monte Carlo on three different architectures. First, a baseline is established using a single threaded CPU implementation. Next the program is converted to run on an NVIDIA GPU with CUDA. Performance results are obtained for both the Quadro 5000 and the Quadro K3100M. Finally, the GPU code is extended to support multiple GPU's within the same host. Our GPU implementation of the antithetic multilevel Monte Carlo displays a major speedup in computation when compared with many commonly used approaches in the literature. While our work is focused on the simulation of the stochastic volatility and interest rate model from Finance it is easily extendable to other stochastic systems, and it is of particular interest to those with non-diagonal, non-commutative noise including large scale models from Biology and Chemistry.

Jj Lay  
Thompson Machinery Commerce Corporation and  
Middle Tennessee State University  
jj.lay@tmcat.com

## PP2

### Computing Particle Trajectories with Full Phase Information in a Magnetic Field by a Multi-scale Hybridization Technique Derived from Time-Parallel Computing Methodology

A multiscale hybridization method is presented that carefully couples coarse scale gyrokinetic models with exact charged particle solution trajectories (that is, with full phase information) in a magnetic field. The approach is based on the careful approximation of a sum, generally employed for time-parallel (TP) computing applications. While the hybridization method presented is highly parallelizable, a computational efficiency gain is seen from considering serial computation only. While a complete numerical method is only presented for the aforementioned charged particle application, the general approach depicted likely has relevance to a wide swath of challenging multi-scale/multiphysics problems. Additionally, the approach has obvious relevance to TP computing applications (such as variable selection on which to perform TP calculations and fine scale sampling strategies) Distribution A Approved for public release; PA Clearance Number 16386

Carl D. Lederman  
ERC Inc.  
carl.lederman.ctr@edwards.af.mil

David Bilyeu  
Air Force Research Laboratory  
david.bilyeu.1@us.af.mil

## PP2

### Acceleration of a Multiscale Model of An Energetic

### Material Through Speculative Computation

We have recently developed a model of an energetic material using a concurrent multiscale approach. The multiscale model is composed of a continuum finite-element multi-physics simulation at the macroscale and a particle-based dissipative particle dynamics (DPD) model at the microscale. The equation of state of 1,3,5-trinitroperhydro-1,3,5-triazine (RDX) is evaluated by the microscale model to inform the constitutive response of the continuum simulation. The multiscale model was constructed using a general adaptive computational framework for scale-bridging. We employ the model for the simulation of a Taylor anvil impact experiment. The computational cost is staggering due to the execution of a parallel DPD simulation at each finite-element at each time step of the continuum simulation. We will discuss the development of a dynamic surrogate modeling approach using Gaussian process regression to reduce the computational cost of the multiscale model by several orders of magnitude. The dynamic nature of the surrogate modeling approach makes computational resource utilization a challenge. To overcome this, we employ speculative computation to fully utilize available computational resources and accelerate evaluation of the overall model, through evaluation of the microscale model at state points the macroscale model is expected to require at future times in the simulation.

Kenneth Leiter, Jaroslaw Knap, Brian Barnes, Claire Eisner, Richard Becker  
U.S. Army Research Laboratory  
kenneth.w.leiter2.civ@mail.mil,  
jaroslaw.knap.civ@mail.mil,  
brian.c.barnes11.civ@mail.mil,  
claire.g.eisner.ctr@mail.mil, richard.c.becker.civ@mail.mil

## PP2

### Nonnegative Sparse Tensor Decomposition on Distributed Memory Systems

Tensor decomposition is one of the popular methods in data analysis and mining applications, especially for large-scale data. Nonnegative CANDECOMP/PARAFAC alternating Poisson regression (CP-APR) model was proposed to analyze sparse count data which occurs in many real-world applications, such as social network analytics, healthcare analytics, and machine learning area. This paper presents coarse-grained and medium-grained distribution strategies to parallelize CP-APR algorithm on a distributed memory system. We also employ an input-aware method to allocate nonzero entries to each process for a good load balance. Our distributed CP-APR achieves up to 94 times speedup over sequential CP-APR on a Power 8 cluster with 320 cores. Our experiments show the distributed CP-APR is well load-balanced and has good strong scalability.

Jiajia Li  
College of Computing  
Georgia Institute of Technology  
jiajiali@gatech.edu

Jee Choi  
IBM TJ Watson  
jwchoi@us.ibm.com

Xing Liu  
IBM  
xing.research@gmail.com

Richard Vuduc

Georgia Institute of Technology  
richie@gatech.edu

## PP2

### Diffusion Mri in the Aplysia Neuronal Network: Experiments and Numerical Simulations

We simulated the diffusion MRI signal by solving the Bloch-Torrey equation using the geometrical description of Aplysia ganglia obtained from high resolution T2-weighted images and matched the simulations with experimental dMRI data.

Jing-Rebecca Li  
INRIA Saclay  
jingrebecca.li@inria.fr

Khieu Van Nguyen, Luisa Ciobanu  
Neurospin, CEA, Saclay, France  
nvkhieu89@gmail.com, luisa.ciobanu@gmail.com

Denis Le Bihan  
CEA-Neurospin, France  
denis.lebihan@gmail.com

## PP2

### A Real-Time Automatic Characterization of Fractures in Enhanced Geothermal Systems Using Machine-Learning Technique

Fractures play an essential role in the characterization of a geothermal reservoir, since they provide crucial information about the fluid-flow pathways. Conventional seismic techniques are often used to understand the subsurface structure; however, limited data coverage, low computational efficiency, and limited resolution relative to imaging fractures hinder their widespread application in fractured geothermal systems. We propose to employ a novel approach based on machine-learning techniques to characterize the fractures. We build data-driven strategies to effectively identify latent aspects of the measurements that are crucial in predicting the fractures. The feature-learning module is coupled with a regression module that learns a direct functional map for fracture parameters from the measurements. The parameters for both the feature-learning and regression modules are obtained in an offline-training phase, thereby enabling real-time prediction for new measurements during the online-testing phase. To validate the quality and efficiency of our approach, we carried out evaluations with synthetic data. We explored the use of kernel feature design and deep neural networks for learning the features, along with effective regression schemes such as linear and gradient-boosted regressors. We show that our method yields both high accuracy and efficiency in the detection of fractures, thereby demonstrating its value in real situations.

Youzuo Lin  
Los Alamos National Laboratory  
ylin@lanl.gov

Jayaraman Thiagarajan  
Lawrence Livermore National Labs  
jayaramanthi1@llnl.gov

George Guthrie  
Earth and Environment Science Division

Los Alamos National Laboratory, Los Alamos  
geo@lanl.gov

## PP2

### Convergence Study for Stochastic Inversion Framework to Monitor Evolving Surface Ship Mass Properties During Arctic Operations

An accurate assessment of a ship's mass properties during operation is an important concern for ships traveling in adverse conditions. Specifically, in Arctic conditions, the risk of ice accumulation on the topside of the ship is heightened, affecting the seakeeping and stability, thus putting the ship crew and mission at risk. We present a stochastic inversion framework through which we can recover a second moment mass property, the roll gyradius, using an on-board inertial measurement unit in conjunction with existing seakeeping software. We use a Markov chain Monte Carlo (MCMC) inversion scheme implemented in Python, using Ship Motion Prediction, software made available through the U.S. Navy, as the forward model. The inversion scheme requires a known ground truth for the roll period and associated signal-to-noise ratio for the roll period in order to recover a posterior distribution for the likely roll gyradius of the ship. These parameters are obtained experimentally at both full- and model-scale. We conduct an independent sensitivity study of the framework parameters to ensure convergence. Specifically, we seek to characterize the behavior of the MCMC chain at various lengths and identify the length at which convergence is achieved by comparing multiple independent simulations. The framework is then demonstrated at full- and model-scale, without any ice accumulation, and again at model-scale with additional hypothetical icing scenarios.

Yolanda C. Lin, Christopher Earls  
Cornell University  
ycl4@cornell.edu, earls@cornell.edu

## PP3

### Analytical Stacked Gaussian Process Model

A probabilistic model is proposed by stacking a set of independently trained Gaussian processes to obtain prediction of quantities of interests that require composition of functions. Analytical derivations are provided for first and second-order moments of the stacked Gaussian process using RBF and polynomial kernels. The StackedGP model can be extended to any number of layers and nodes per layer, and it provides flexibility in kernel selection for each node. The proposed nonparametric stacked model is validated using different synthetic datasets and its performance is measured in two real-world applications.

Kareem Abdelfatah  
PhD Student  
University of South Carolina  
krabea@email.sc.edu

Junshu Bao  
Graduate Teaching Assistant Department of Statistics  
University of South Carolina  
bao3@email.sc.edu

Gabriel Terejanu  
University of South Carolina

terejanu@cse.sc.edu

alemazk2@illinois.edu, meidani@illinois.edu

### PP3

#### Uncertainty Quantification in Multi-Scale Materials Modeling

In this work we present a Bayesian framework to include uncertainty quantification in predictions for the case of alloy modeling using the cluster expansion (CE), which is a surrogate model that allows to compute physical properties which are a function of the configuration. Using a Bayesian framework, the expansion coefficients are described by a multivariate probability distribution. Initially, we select a sparsity inducing prior distribution for the coefficients. In particular, we use a relevance vector machine, which provides automatic relevance determination of the expansion coefficients. We use data obtained from first principles density functional theory calculations to obtain the joint posterior probability distribution of the CE coefficients as well as the predictive distribution for the computed properties. The resulting probability distributions for the predicted material properties reflect the uncertainty arising from a limited training data set and the model error from using a truncated expansion. We also show how to augment the Bayesian CE with a Gaussian Process in order to model the temperature dependence of the material properties. We show the performance of the framework for a range of material properties, including electrical (conductivity), optical (dielectric function), elastic (elastic constants) and thermodynamic (phase diagrams), with their corresponding predictive uncertainties.

Manuel Aldegunde  
Department of Aerospace and Mechanical Engineering  
University of Notre Dame  
maldegun@nd.edu

Nicholas Zabarar  
University of Notre Dame  
nzabarar@gmail.com

### PP3

#### Near-Optimal Sampling Approach for Estimating Sparse Polynomial Chaos Expansions

Recently, compressive sampling has been widely used to estimate the coefficients of Polynomial Chaos (PC) expansions, when the number of available samples is limited. In such cases, the location of sample points can significantly impact the accuracy of constructed PCE. We aim to determine where to sample the underlying system in advance of any (computationally) expensive sampling. We present a near-optimal sampling strategy for compressive sampling approach for PCE construction. We provide theoretical motivation of near-optimal sample set along with several numerical examples in which the proposed sampling strategy is implemented. The results show that using near-optimal sample set rather than a random sample set substantially improves the accuracy of the approximated PCE using compressive sampling.

Negin Alemazkooor, Hadi Meidani  
University of Illinois at Urbana-Champaign

### PP3

#### Mathematical Analysis of Sludge Disintegration

We investigate a model for an activated sludge process connected to a sludge disintegration unit. We formulate a model for the sludge disintegration unit in which the disintegration processes occur at a finite rate. We then take an appropriate asymptotic limit to obtain a model when these processes occur infinitely quick. We use our limiting model to show that an infinite-rate model proposed earlier in the literature is incorrectly formulated. Our principle aim is to investigate how the disintegration rate in the sludge disintegration unit affects the formation of sludge in the reactor. In the limiting case of infinite disintegration rate we show that there is a critical value of the sludge disintegration factor, above which the reactor system is guaranteed to be in a state of negative excess sludge production. For the case of finite rate processes we show that if the disintegration rate is sufficiently high then the error in assuming an infinite rate is less than 10% of the exact value using finite rate. In such cases, the behaviour of the reactor can be estimated, within experimental error, by assuming an infinite rate. We further show that if the reaction rate in the sludge disintegration unit is sufficiently small then there is no longer a critical value of the sludge disintegration factor above which the reactor operates in a state of negative excess sludge production for all residence times.

Rubayyi Alqahtani  
Imam University - Saudi Arabia  
rtaa648@uowmail.edu.au

### PP3

#### Selective Time Step Adaptivity for Non-Linear Reactive Transport Problems

Numerical simulation of reactive, flow and transport in porous media requires that the increasingly non-linear models be solved in an accurate and computationally efficient manner. However, the time step sizes associated with each of these processes differ greatly. For example, kinetic reaction rates varying over order of magnitudes or layered reservoirs with fast and slow concentration changes in high and low permeability layers, respectively. A few texts can be found in literature to address this area: Mortar element technique in both space and time domain, Two-ways methods (Offline/Online), Optimized Schwarz wave-form relaxation (OSWR) method. We present an adaptive time-stepping approach with different time-step sizes in different spatial domains. In summary, we proposed a mass conservative, adaptive time-stepping scheme with different time step sizes in different spatial domains, a fully implicit solution algorithm was developed for coupled flow and transport problems, using block Jacobi for improved convergence rates, numerical results using DG0 in time and RT0 Mixed FEM in space (finite difference), and preliminary results are in good agreement with reduced error in the fine time domain.

Yerlan Amanbek  
The Institute for Computational Engineering and Sciences  
yerlan@ices.utexas.edu

Gurpreet Singh  
The University of Texas at Austin  
Center for Subsurface Modeling  
gurpreet@ices.utexas.edu

Mary F. Wheeler  
Center for Subsurface Modeling, ICES  
University of Texas at Austin  
mfw@ices.utexas.edu

### PP3

#### Upwind Dissipation for Stability on Composite Grids

Upwind schemes for the second-order wave equation have been developed and are known to be remarkably stable. Here we present results for a centered finite-difference scheme incorporating artificial dissipation in the same form of dissipation appearing naturally in the upwind scheme. The stability of the scheme is studied using the theory of normal modes and results are confirmed numerically.

Jordan B. Angel  
Rensselaer Polytechnic Institute  
angelj2@rpi.edu

### PP3

#### Fully Bayesian Deep Gaussian Processes for Uncertainty Quantification

Gaussian processes have been established as attractive models for developing efficient surrogates for uncertainty quantification when the system of interest is computationally expensive to study directly. However, such approaches struggle when the stochastic dimensionality of the problem is high or when the response surface is strongly nonstationary. We develop a fully Bayesian deep Gaussian process (DGP) model that allows for a natural handling of uncertainty propagation and demonstrate how it can be applied to address both of these issues, using the Darcy flow problem for illustration. First, we use an unsupervised DGP to perform nonlinear model reduction on high-dimensional stochastic input data corresponding to measurements of the random permeability field. The hierarchical structure of DGPs allows for flexible nonlinear dimensionality reduction that outperforms standard methods including the well-known Karhunen-Löve expansion (KLE) method. Second, we use a supervised DGP to construct a surrogate for the corresponding multi-output response surface. We are able to model sharp discontinuities in the response surface without needing to resort to specialized local models. The fully Bayesian training scheme allows for automatic complexity control and provides posterior distributions on the hyperparameters describing the correlation length scales in the problem. This makes feasible the modeling of epistemic uncertainty in the responses due to limited training data.

Steven Atkinson, Nicholas Zabaras  
University of Notre Dame  
satkinso@nd.edu, nzabaras@gmail.com

### PP3

#### Polynomial Filtering for Large, Sparse SVD Computations

We present a new technique for accelerating the computation of a few singular values and vectors of a large, sparse matrix using polynomial filters. Numerical experiments show that our method is competitive with existing algorithms and can offer significant advantages for computing small or interior singular values, especially when the matrix is large enough that methods based on shift-and-invert

techniques are not practical.

Anthony P. Austin  
Argonne National Laboratory  
austina@anl.gov

Jared Aurentz  
Instituto de Ciencias Matemáticas  
jaredaurentz@gmail.com

Vasileios Kalantzis  
University of Minnesota  
kalan019@umn.edu

### PP3

#### The "Porous Media Initiative" (OPM) Software

In this poster we will give an overview of the "Open Porous Media Initiative" (OPM, <http://opm-project.org>) software and open data modules. It is a collaborative effort of both academia and industry to foster collaboration between both and thus facilitate innovation using an open source approach. The current development focus is CO<sub>2</sub> sequestration and improved and enhanced oil recovery. Among the contributors are IRIS, SINTEF ICT, Statoil, and Ceetron Solutions. The OPM software, in part based on DUNE (<http://dune-project.org>), contains various solvers for porous media flow (IMPES, two-phase incompressible transport, a reordering-based two-phase polymer solver, and a parallel fully implicit black-oil solver). In addition solvers for flow-based upscaling of both permeability (single-phase upscaling) and relative permeability (two-phase upscaling), including capability for steady-state upscaling are included. For the fast and easy visualization of reservoir simulations it also provides the tool ResInsight. Together with the software OPM provides various SPE test data sets, and the data of the Norne oil field as open data to allow easy experiments and benchmarking. The software licensed under GNU General Public License (GPL) version 3 and the data under the Open Database License.

Markus Blatt  
Dr. Blatt - HPC-Simulation-Software & Services  
markus@dr-blatt.de

Atgeirr Rasmussen  
SINTEF ICT  
atgeirr.rasmussen@sintef.no

### PP3

#### The Reduced Collocation Method for Nonlinear Steady-State PDEs

Reduced basis methods are a class of numerical methods developed for settings which require a large number of solutions to a parameterized problem. In this work, we consider settings which require many queries to a numerical solver for a parameterized steady-state partial differential equation. We adapt an existing method for linear steady-state problems, the reduced collocation method, to the nonlinear case.

Chris Bresten  
University of Massachusetts Dartmouth  
tba

Yanlai Chen  
Department of Mathematics  
University of Massachusetts Dartmouth

yanlai.chen@umassd.edu

eugenio.aulisa@ttu.edu

**PP3****Neutrino Radiation Hydrodynamics in Genesis**

Neutrino radiation hydrodynamics is a key element of the physics governing environments with hot and dense nuclear matter. Such extreme environments include the universe itself, during primordial nucleosynthesis just after the big bang; compact object mergers following the inspiral of neutron star / neutron star or neutron star / black hole binaries; and core-collapse supernovae. Neutrino radiation hydrodynamics is particularly important to the explosion mechanism of core-collapse supernovae. In this poster, we will present our current implementation of a software element to solve radiation hydrodynamics problems on modern supercomputing featuring hybrid architectures. This radiation hydrodynamics functionality is being developed within GenASiS (General Astrophysical Simulation System), a new code being developed to facilitate the simulation of astrophysical problem on supercomputers. We will present the multi-threading infrastructure, kernel offload to the GPU and coprocessors for the implicit integration solver in radiation-matter coupling, draft implementation of the Implicit-Explicit (IMEX) Runge-Kutta scheme, and test problems to verify our implementation. We will also showcase the capabilities and solvers provided by GenASiS Mathematics, a division of GenASiS we make available as open source.

Reuben Budiardja  
University of Tennessee  
reubendb@utk.edu

Christian Cardall, Eirik Endeve  
Oak Ridge National Laboratory  
cardallcy@ornl.gov, endevee@ornl.gov

Daniel Murphy  
University of Tennessee  
rmurph16@vols.utk.edu

**PP3****Numerical Simulations for Fluid-Structure Interaction Modeling of Artery Aneurysms**

We present numerical simulations of fluid-structure interaction (FSI) problems involving artery aneurysms, which are common vascular problems with fatal implications. The physics of the problem is described using a monolithic approach and both the fluid flow and the hyperelastic material are considered to be incompressible. The deformation of the fluid domain is taken into account according to an Arbitrary Lagrangian Eulerian (ALE) scheme. A Newton-Krylov solver has been implemented in the finite element in-house code FEMuS. The solver is combined with a geometric multigrid preconditioner with smoothing based on domain decomposition. The configurations studied consist of realistic aortic and brain aneurysms where, for the latter, the corrective effect of stents has been investigated.

Sara Calandrini  
Texas Tech University  
sara.calandrini@ttu.edu

Eugenio Aulisa  
Department of Mathematics and Statistics.  
Texas Tech University

**PP3****Modeling Multiphase Buoyancy Driven Plume Migration During Geologic CO<sub>2</sub> Injection**

Geologic CO<sub>2</sub> sequestration requires no more than 1% of injected CO<sub>2</sub> escape within 1000 years after injection. To predict long-term retention of CO<sub>2</sub> in a reservoir, the interaction of geochemical and geomechanical effects of injection must be investigated via numerical simulation. Transport of gas phase CO<sub>2</sub> through micrometer scale fractures in porous sandstone and shale caprock is one concern that could lead to unwanted release of injected CO<sub>2</sub> into the atmosphere. We model the buoyancy driven flow of a two-phase system consisting of a CO<sub>2</sub>-H<sub>2</sub>O vapor mixture phase and an aqueous phase composed of formation water, dissolved CO<sub>2</sub>, and charged solutes formed from the dissolution of quartz, feldspars, carbonate, and clay minerals. This two-phase system forms a plume of CO<sub>2</sub> that can migrate upward due to differences in density between CO<sub>2</sub>-rich phases and the surrounding formation fluid. We model the gas phase CO<sub>2</sub>-H<sub>2</sub>O composition using a Redlich and Kwong equation of state (EOS) with mixing rules, and the aqueous phase composition using the revised Helgeson Kirkham Flowers model for approximating thermodynamic properties of aqueous electrolytic solutions at high temperature and pressure. Water density and electrostatic properties are computed using the virial EOS developed by Haar, Gallagher, and Kell outside the critical region and the Levelt and Sengers EOS within the critical region. Pitzer equations are used to compute osmotic and solute activity coefficients.

Kyle Campbell, Christopher Paolini  
San Diego State University  
kylejaycampbell@gmail.com,  
paolini@engineering.sdsu.edu

Jose Castillo  
Computational Science Research Center  
San Diego State University  
jcastillo@mail.sdsu.edu

**PP3****A Parallel Particle Tracking Algorithm for Finite Element Applications**

Numerical simulations of a parallel particle tracking algorithm on unstructured finite element grids are presented. The algorithm is designed to work for both 2D and 3D applications, with linear and quadratic spatial discretization. To determine the position of the particle relative to the mesh, a new inclusion test algorithm has been designed to work with parallel computing and finite element applications. The advection is performed on the physical domain in order to treat completely unstructured grids. As a consequence, the inversion of the isoparametric finite element mapping is requested. We comply with this demand implicitly using Newton-Raphson's iteration for all Lagrangian finite elements and all degrees of freedom. We report tests made to investigate the performance of our algorithm and a fluid flow application of pure advection.

Giacomo Capodaglio  
Texas Tech University  
giacomo.capodaglio@ttu.edu

Eugenio Aulisa

Department of Mathematics and Statistics,  
Texas Tech University  
eugenio.aulisa@ttu.edu

### PP3

#### Why Is the Ensemble Adjoint Approach to Sensitivity Analysis Not Practical in Chaotic Systems?

The ensemble adjoint method proposed years ago by Lea, Allen and Haine still remains in perspective as a useful sensitivity analysis tool in chaotic systems scoring over the only other candidate, the Least Squares Shadowing (LSS) method, in terms of conceptual simplicity and ease of implementation. In this work, we perform a part-computational, part-mathematical analysis of the ensemble sensitivity estimator with an objective of revealing its feasibility for use in chaotic fluid flow simulations. Our analysis is based on relating the rate of convergence of the estimator to the Lyapunov exponents and the spectrum of the Liouville operator of the chaotic system. We apply the analysis to the Lorenz'63 attractor to prove that the ensemble sensitivity approach converges slower than a Monte-Carlo simulation, as confirmed by Eyink et al's numerical simulations. We then extend our analysis technique to more complex systems such the Lorenz'96 and a chaotic flow over a NACA 0012 airfoil.

Nisha Chandramoorthy, Qiqi Wang  
Massachusetts Institute of Technology  
nishac@mit.edu, qiqi@mit.edu

### PP3

#### A Novel Speech-Based Diagnostic Test for Parkinson's Disease Integrating Machine Learning with Application Development for Cloud Deployment

Parkinsons disease remains one of the most poorly diagnosed neurological conditions despite the critical need of early diagnosis for effective management and treatment. This work presents a new method of diagnosing Parkinsons disease and accompanying scalable web and mobile applications towards the goal of employing this diagnostic test on the cloud. This method provides a more simple, inexpensive, and time-effective approach than traditional diagnosis strategies by requiring the patient to only speak into a microphone attached to their computer or mobile device before providing a highly accurate diagnosis within seconds. This work employs speech processing algorithms, an artificial neural network for machine learning, and an application framework with a user-friendly interface that packages the speech test and diagnosis results for easy access by patients and physicians. This test was specifically designed for rapid in-home Parkinsons testing, a capability that is currently not offered by any existing tests. The diagnosis test developed was tested with actual patient data and was shown to be 96.55% accurate. The results of this research thus take a significant step towards building the first globally available speech test for the diagnosis of Parkinsons disease.

Pooja Chandrashekar  
Harvard University  
pchandrashekar@college.harvard.edu

### PP3

#### Protein NMR Reference Correction: A Statistical

#### Approach for an Old Problem

Accurate chemical-shift assignments are a vital requirement for protein structure determination. The accuracy of the chemical-shifts value depends on a correct reference value. Internal reference standard exists but the We are developing a statistical-based algorithm to correct referencing by: 1.calculating composition probabilities of investigating protein alpha and beta carbon (CA/CB) resonance pairs from the NMR data; 2.summing the probabilities across all resonance pairs to give an estimate of amino acid (AA) composition; and 3.employing a grid search method to find a minimum difference (correct referencing value) between predicted and actual protein AA composition. We show that CA/CB resonance covariance (dependence) is a potent statistic and that oxidized/reduced cysteine residues should be treated separately. Our results demonstrate that the overall approach is feasible and will provide the biomolecular NMR field with a unique tool allowing spectral referencing to be corrected and refined at the beginning of protein NMR data analysis without using chemical shift assignments or protein structure as needed by current retrospective referencing correction methods. Thus, our method should improve both the speed and quality of protein resonance assignment and downstream NMR-based analyses including structure determination. Currently we are developing a shiny web app that will further simplify this protein NMR reference correction procedure.

XI Chen  
University of Kentucky  
billchenxi@gmail.com

Andrey Smelter  
Department of Computer Engineering and Computer Science  
University of Louisville  
andrey.smelter@louisville.edu

Hunter Moseley  
Department of Molecular and Cellular Biochemistry  
University of Kentucky  
hunter.moseley@uky.edu

### PP3

#### Numerical Solutions of Basic Boundary-Contact Problems of Tumor-Brain Interface Tissues

This paper is devoted to the development of computational method for approximation solution of basic boundary-contact problems of coupled-elasticity for isotropic inhomogeneous tumor-brain interface tissues. The tools applied in this development are based on generalized potential method and singular integral equations. Numerical solutions are solved by using the generalized Fourier's series method.

Manana Chumburidze  
Akaki Tsereteli State University  
maminachumb02@gmail.com

David Lekveishvili  
Akaki tsereteli State University  
dlekveishvili@gmail.com

### PP3

#### Horsetail Matching: A Flexible Approach to Optimization Under Probabilistic and Interval Uncer-

**tainties**

It is important to design engineering systems to be robust with respect to uncertainties in the design process. Many previously proposed approaches for robust design aim to achieve this by treating uncertainties probabilistically and optimizing statistical moments of the quantity of interest. However, over-reliance on statistical moments in the formulation of a robust optimization can result in designs that are stochastically dominated by other feasible designs. Additionally, it is often difficult to assign probability distributions to uncertainties without making strong assumptions, and so in many problems some uncertainties are more appropriately represented with intervals. We propose a formulation for optimization under any mix of probabilistic and interval uncertainties that minimizes the difference between a design's horsetail plot, which consists of two curves representing the upper and lower bound on the true CDF, and a target. The target offers flexibility to a designer to specify the desired behavior under uncertainty, but the approach can also be used without a target to produce a design that both lies on the classical robust Pareto front and that is not stochastically dominated. A numerical discretization is developed that delivers a single, differentiable metric as the objective function for optimization.

Laurence W. Cook

University of Cambridge  
Massachusetts Institute of Technology  
lwc24@cam.ac.uk

Jerome P. Jarrett  
University of Cambridge  
jpp1001@cam.ac.uk

Karen E. Willcox  
Massachusetts Institute of Technology  
kwillcox@MIT.EDU

**PP3****Accelerating Multiplication of Small Or Skinny Matrices with Intel Math Kernel Library Packed GEMM Routines**

Matrix-matrix multiplication is ubiquitous in scientific computing across many application areas. Due to the high compute-density of the operation, General Matrix-Matrix Multiplication (GEMM) routines in optimized libraries are capable of obtaining very high computational efficiency for large matrix sizes. However, in many application areas GEMM routines are primarily applied to matrices for which performance can be heavily impacted by data-packing routines that are necessary to obtain optimal performance in the GEMM kernels. In particular, machine learning applications often perform GEMM operations on non-square matrices with one or more small dimensions. To increase the performance of applications with these GEMM cases, the Intel Math Kernel Library has introduced a Packed GEMM API. The Packed API takes advantage of the fact that for some applications, e.g. machine learning and LU factorization, the same matrices are re-used in many GEMM calls by separating the data-packing stage from the compute stage of the GEMM kernels. In this poster we present the Packed GEMM API and demonstrate performance improvements for matrix sizes typical of machine learning applications.

Timothy B. Costa, Murat E. Guney, Sarah Knepper,  
Shane Story  
Intel Corporation

timothy.b.costa@intel.com, murat.e.guney@intel.com,  
sarah.knepper@intel.com, shane.story@intel.com

**PP3****A Mathematical Model of Hemostasis**

Hemostasis is the process by which a blood clot forms to prevent bleeding. The formation time, size and structure of a clot depends on the local hemodynamics and the nature of the injury. Our previous computational models were developed to study intravascular clot formation, a process confined to the interior of a single vessel. Here we present the first computational model of extravascular clot formation (hemostasis) in which blood through a single vessel initially escapes through a hole in the vessel wall and out a separate injury channel. The model consists of a system of partial differential equations that describe blood coagulation biochemistry, platelet aggregation, and hemodynamics, solved via the finite element method. In the model, formation of a blood clot occludes the injury channel and stops flow from escaping while blood in the main vessel retains its fluidity. We discuss the different biochemical and hemodynamic effects on clot formation using distinct geometries representing intra- and extravascular injuries.

Nicholas Danes

Applied Mathematics & Statistics  
Colorado School of Mines  
ndanes@mines.edu

Karin Leiderman  
Department of Applied Mathematics & Statistics  
Colorado School of Mines  
kleiderman@mines.edu

**PP3****An IDS Study of Flow over A Leading Edge**

Near the leading edge of a sharp flat plate over which fluid flows at hypersonic speeds, the physical phenomena that occur in the region is marked by regions of near free molecular flow, transitional flow and merged layer flow. It has long been suspected that the continuum-based Navier-Stokes solvers are incapable of capturing the physics within the first two regions. This study investigates the ability of an advanced numerical Navier Stokes solver to capture the flow physics within these regions. The main purpose of this investigation is to demonstrate new features of the Navier-Stokes solver by comparing the new predictions of the solver with previous experimental works. In this study, the flow over a 0.128 meter flat plate is solved using the advanced Navier-Stokes solver based off Integral Differential Scheme (IDS) "[Elamin, G.A., 2008]". The flow boundary conditions are set up to be the same as that of the experimental works of Becker & Boylan "[Becker, M. and D.E. Boylan, 1966]". Modifications necessary to obtain the results that match the experimental results are made in the boundary conditions at the leading edge of the plate and not in the Navier Stokes solver itself. The results validate the robustness of the IDS Navier-Stokes solver in studying the flow phenomena in the hypersonic leading edge problem.

David Dodoo-Amoo, Julio C. Mendez, Frederick Ferguson  
Department of Mechanical Engineering  
North Carolina A&T State University  
dndodooa@aggies.ncat.edu, jcmendez@aggies.ncat.edu,



ferguso@ncat.edu

### PP3

#### Higher Order Asymptotic for Burgers Equation

We derive an approximate solution, suitable for computations, to the viscous Burgers equation via heat equation. Using generalization of the truncated moment problem to a complex measure space, we construct asymptotic N-wave approximate solution to the heat equation subject to the initial data whose moments exist upto the order  $2n + m$  and i-th order moment vanishes, for  $i = 0, 1, 2 \dots m$ .

Satyanarayana Engu  
National Institute of Technology Karnataka  
satyaw1@gmail.com

### PP3

#### A Bi-Fidelity, Low-Rank Approximation Technique for Uncertainty Quantification

The use of model reduction has become widespread as a means to reduce computational cost for uncertainty quantification of PDE systems. In this work we present a model reduction technique that exploits the low-rank structure of the solution of interest, when it exists, for fast propagation of high-dimensional uncertainties. To construct this low-rank approximation, the proposed method utilizes models with lower fidelities (hence cheaper to simulate) than the intended high-fidelity model. After obtaining realizations to the lower fidelity models, a reduced basis and an interpolation rule are identified and applied to a small set of high-fidelity realizations to obtain this low-rank, bi-fidelity approximation. In addition to the construction of this bi-fidelity approximation, we present convergence analysis and numerical results.

Hillary Fairbanks  
Department of Applied Mathematics  
University of Colorado, Boulder  
hillary.fairbanks@colorado.edu

Jerrad Hampton  
University of Colorado, Boulder  
jerrad.hampton@colorado.edu

Alireza Doostan  
Department of Aerospace Engineering Sciences  
University of Colorado, Boulder  
Alireza.Doostan@Colorado.EDU

Akil Narayan  
University of Utah  
akil@sci.utah.edu

### PP3

#### Boundary Integral and Image-Moment Hybrid Method for Simulations of Solvated Proteins

Solving the Poisson or Poisson-Boltzmann equation with varying coefficients has been a bottleneck for simulating proteins within the framework of implicit solvent model due to the dielectric inhomogeneity. We will present our recent work on a numerical Boundary Integral Equation (BIE) method, and a semi-analytical Image-Moment hybrid method for efficient and accurate simulations of electrostatic fields in systems consisting multiple proteins in solvent. The BIE is coupled with treecode for fast

kernel summation, and is applicable to arbitrary shaped dielectric interface; while the hybrid method combines analytical image charge solution of dielectric spheres with the Method of Moments, and is accelerated by the fast multipole method (FMM). Consistent numerical results from both methods will be presented, and we will show how the induced (or polarization) charge on the dielectric interfaces can significantly change the solvation/interaction energy of solvated proteins.

#### References:

Zecheng Gan  
Department of Mathematics, University of Michigan  
zecheng@umich.edu

Weihua Geng  
Department of Mathematics  
South Methodist University  
wgeng@smu.edu

Robert Krasny  
University of Michigan  
Department of Mathematics  
krasny@umich.edu

### PP3

#### Asynchronous Optimized Schwarz Method for Poisson Equation in Rectangular Domains

Optimized Schwarz methods are Domain Decomposition methods in which the transmission conditions between subdomains are chosen in such a way that the convergence properties of the method are enhanced with respect to classical Schwarz Methods. Asynchronous iterative algorithms are parallel iterative algorithms in which communications and iterations are not synchronized among processors. Thus, as soon as a processor finishes an iteration, it starts the next one with the latest data received during a previous iteration, without waiting for any other processor. These algorithms increase the number of iterations in some processors (with respect to the synchronous case) but suppress all the idle times, which can result in a smaller (execution) time for convergence. In this work, we analyze the convergence properties of the Asynchronous Optimized Schwarz method applied as a solver for the solution of the Poisson Equation in a bounded rectangular domain. To our knowledge, this is the first time that a convergence proof of Optimized Schwarz for a problem defined in a bounded domain and for an arbitrary number of subdomains is presented.

Jose C. Garay  
Temple University  
tue62222@temple.edu

Daniel B. Szyld  
Temple University  
Department of Mathematics  
szyld@temple.edu

### PP4

#### Design of a Parallel AMR Infrastructure for Multi-Accelerator Computing

Design of a Parallel AMR Infrastructure for Multi-Accelerator Computing : Adaptive mesh refinement (AMR) is a powerful technique to increase the spatial

resolution where it is needed the most. AMR capabilities are available to researchers through a few open-source packages. However current AMR packages have to be re-designed and implemented to make efficient use of accelerators. We propose the building brick method (BBM) as a new AMR package for clusters of accelerators. Our goal is to develop an open-source user-friendly library so that it can easily be adopted by other researchers. The software infrastructure is specifically developed for multi-level parallelism on clusters of accelerators. BBM is flexible to be used in finite difference, finite volume and finite element methods with little modification. The software infrastructure of BBM makes use of abstraction in terms of objects and classes available in the C++ language. These features are useful in developing and designing user friendly functionalities for engineering applications. Polymorphism and operator overloading features of the language become useful in designing codes that are easy to understand and maintain. The tree structure in BBM is stored implicitly and the 2:1 consistency criterion is enforced. Additionally, the tree structure is kept local to each accelerator to minimize the cost of intra- and inter-node communication.

Jaber J. Hasbestan  
boise state university  
jaberjavanshir@hotmail.com

Inanc Senocak  
Mechanical & Biomedical Engineering  
Boise State University  
senocak@boisestate.edu

#### PP4

##### **A Single-Phase Slightly Compressible Flow and Multicomponent Transport Model in Porous Media at Laboratory Scale**

In this work the development of a single-phase slightly compressible flow and multicomponent transport model in porous media at laboratory scale is presented. The development methodology consisted on first establishing a conceptual model where the hypothesis, goals, objectives and limitations of the model are considered. Subsequently, applying the systematic formulation approach of continuum medium modeling the mathematical model is derived resulting a partial differential equation system with initial and boundary conditions [Herrera, I. and Pinder, G.F., *Mathematical models of science and engineering*, John Wiley & Sons, New York, USA, 2012]. After that, the mathematical model is discretized using a finite element method [Chen, Z., *Finite Element Method and their Applications*, Springer-Verlag Berlin Heidelberg, 2005]. Finally, its computational implementation is performed on the open source software FEniCS project [Anders Logg, et. al., *Automated Solution of Differential Equations by the Finite Element Method*, Springer-Verlag Berlin Heidelberg, 2012]. The numerical solutions were compared with the commercial software COMSOL Multiphysics. A case of study for water injection into a vertical cylinder in two dimensions where water is injected on the bottom and produced from the top is showed.

Eduardo Linares-Pérez  
Posgrado en Ciencias de la Tierra, Instituto de Geofísica  
Universidad Nacional Autónoma de México  
eduardo\_linares@comunidad.unam.mx

Mario Noyola-Rodríguez  
Posgrado en Ciencias de la Tierra, Instituto de Geofísica  
Universidad Nacional Autónoma de México (UNAM)

estaduagro@gmail.com

Martin A. Diaz-Viera  
INSTITUTO MEXICANO DEL PETROLEO  
mdiazv@imp.mx

#### PP4

##### **New Second-Order Time Schemes for Optimal Control of Time-Dependent PDEs**

In this poster, we will present two new second-order finite difference time schemes within the framework of iterative one-shot methods for PDE-Constrained optimization. (1) A new second-order leapfrog scheme in time with a multi-grid solver is proposed for solving the optimality KKT system of the parabolic optimal control problems. (2) A new second-order implicit scheme in time with a preconditioned GMRES solver is developed for solving the optimality KKT system of the hyperbolic wave optimal control problems. Numerical tests are given to illustrate our numerical algorithms in terms of both accuracy (of the scheme) and efficiency (of the solver).

Jun Liu  
Southern Illinois University  
jun.liu@jsums.edu

#### PP4

##### **An Assembly-Free Heterogeneous Computing Method for Simulation of Heat Conduction in Heterogeneous Materials**

The simulation of heat flow through heterogeneous material is important for the design of structural and electronic components. Analytical solutions to the heat equation PDE are not known for such domains, even with relatively simple geometries. The finite element method, however, can provide approximations to a continuum solution, with increasing efficacy as the number of degrees of freedom in the model increases. This comes at a cost of increased memory usage and computation time for even sparse matrix approaches. We present a method that does not require the explicit assembly of any system matrices, which has been adapted from recent approaches in solving problems in structural mechanics. The method is highly parallelizable, and has been implemented for the use of multiple GPU using the OpenCL computing framework. Furthermore, it lends itself to the simulation of heterogeneous material with minimal added complexity. We compare the scalability of single- and double-GPU implementations of our method with a serial sparse matrix approach.

Andrew Loeb, Chris Earls  
Cornell University  
ael89@cornell.edu, cje23@cornell.edu

#### PP4

##### **Coins Classification Using Image Processing and Linear Vector Quantization**

Coins detection and classification methods have been significantly improved lately. While traditional approaches rely on precise devices for measuring the coin weight or diameter, alternative methods include even processing the sound echoed by a coin in a controlled collision. The present paper discusses an approach for recognizing and classifying Brazilian coins by using a Learning Vector Quantization (LVQ) algorithm, which receives, as input,

features from their images, in particular (i) the components Cr and Cb of the average YCrCb color, (ii) the normalized diameter and (iii) the color variation in terms of Cr and Cb measured from the circle with  $r \leq 0.4 \cdot r_C$  to the annulus  $r | 0.85 \cdot r_C \leq r \leq 0.75 \cdot r_C$ , where  $r_C$  is the coins most external radius. The method has rotational invariance and is agnostic to which of the coins faces is exposed to the camera. The solutions performance is evaluated with static images captured from 100 coins, as well as in real time, in which case the coins are exposed to the camera, one by one, while the application shows the total amount of money and outputs it as synthetic voice.

Leandro R. Mattioli

Universidade Federal de Uberlândia  
leandro.mattioli@gmail.com

#### PP4

##### Low Dispersion Mimetic Discretizations of Maxwell's Equations in Media with Linear Constitutive Laws

We consider the numerical simulation of electromagnetic wave propagation in a material whose constitutive laws are modeled by a system of linear ordinary differential equations for example a Debye or Lorentz media. These materials are dispersive, i.e. wave speed depends on frequency rather than being constant. In wave propagation problems, long time errors are dominated by dispersion error—non-physical dispersion which is a result of discretizing the equation. In a material which exhibits physical dispersion these errors especially important to quantify and minimize. We consider a discretization using Mimetic Finite Differences (MFD) in space and Exponential Time Differences in time, we call this method this ET-MFD. The MFD allows for the construction of a parameterized family of methods with formally equivalent error properties, i.e. second order accurate on regular meshes. This family includes many well known discretizations including the Yee-Scheme and Nedélec Finite elements. Exploiting the non-uniqueness of the MFD we are able to construct a method which minimizes the dispersion error of the ET-MFD further the error minimizing parameters are oblivious to the underlying constitutive law. We present numerical results which support our theoretical results.

Duncan A. McGregor

Sandia National Laboratories  
Computational Multiphysics  
damcgre@sandia.gov

#### PP4

##### Markov Chain Monte Carlo Optimization for Fitting Excitable Cells Current-Voltage Relations to Voltage Clamp Data

Models of voltage-sensitive ionic currents in electrically excitable cells, such as neurons, contain parameters that depend on properties of the cell. Voltage clamp, a commonly used protocol in which the experimenter cancels out deviations in a cell's membrane potential from a desired potential, can be used to extract current-voltage relations for excitable cell ion channels. We demonstrate a Markov Chain Monte Carlo method to fit parameters in models of voltage-sensitive ionic currents to current-voltage relations obtained from voltage clamp experiments. The method that we present is an example of a stochastic nonlinear optimization method that has the advantage avoiding local minimizers to find the global minimum of a highly non-

linear cost function. 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.

Joseph Mckenna

Department of Mathematics  
Florida State University  
jmckenna@math.fsu.edu

#### PP4

##### Adaptive Sparse Grid Stochastic Collocation for Random Ordinary Differential Equations

Random effects in applications get more and more important and can be modeled in multiple ways. In the context of differential equations, random ordinary differential equations (RODEs) represent an alternative approach to classes of stochastic ODEs by solving massive amounts of deterministic paths. Here, stochastic effects are usually encoded by random processes. To actually compute the RODEs, an approach via a KL expansion approximation and classical Uncertainty Quantification is possible. Since the number of terms in the expansion directly refers to the stochastic dimension we, however, quickly run into the curse of dimensionality. One could use Monte Carlo methods but they converge rather slowly. Furthermore, the KL expansion shows exponential decrease in its eigenvalues which we want to leverage. We, therefore, use adaptive sparse grid surrogates. Building up on the sparse grid stochastic collocation previously presented at SIAM CSE15 [Menhorn, Neckel, UQ in Incompressible Flow using Sparse Grids, 2015], we present an adaptive sparse grid method based on [Gerstner, Griebel, Dimension Adaptive Tensor Product Quadrature, 2003] applied on RODE examples. Due to the decrease in the eigenvalues, the adaptive algorithm implicitly constructs an anisotropic grid with higher refinement in the direction of the largest eigenvalues. Hence, we reach a lower approximation error for the same number of collocation points compared to regular sparse grid or Monte Carlo methods.

Friedrich Menhorn

Department of Informatics, Technische Universität München  
menhorn@in.tum.de

Tobias Neckel

Technische Universität München  
neckel@in.tum.de

#### PP4

##### $N$ Bugs on a Surface

The bugs on a square problem considers  $N$  ordered bugs in a plane that continuously move in the direction of their neighbor. We extend this classic problem to  $N$  bugs on closed curves and surfaces. This extension makes it possible for bugs to not collapse to a point, but instead they can form a closed rotating loop. We show how the dynamics depend on both the geometry, the number of bugs, and the bugs initial condition.

David J. Miller

Florida State University  
dm11ad@my.fsu.edu

Bryan D. Quaife

Scientific Computing

Florida State University  
bquaife@fsu.edu

#### PP4

### Towards the Ultimate Finite Element Method for the Stokes Equations

We discuss an algorithm for solving the Stokes system based on the Scott-Vogelius iterated penalty method. We show that by using a continuous function space to approximate the pressure, we can recover an accurate solutions while retaining a divergence-free velocity. Thus, this algorithm could be of particular interest when modeling incompressible fluids that are thought to have a continuous pressure. Full details of our implementation in the automated finite element software FEninCS and numerical results are given.

Hannah M. Morgan

University of Chicago, department of Computer Science  
hmmorgan@uchicago.edu

Ridgway Scott  
Department of Computer Science  
University of Chicago  
ridg@cs.uchicago.edu

#### PP4

### Efficient and General Parallel Solver for Boundary Integral Equations

We present a high-order, parallel boundary integral equation solver, based on an efficient scheme to evaluate integral operators associated with constant-coefficient elliptic PDE boundary value problems in three dimensions. Our method can compute singular and nearly-singular integrals with high-order accuracy and can handle complex geometric domains. Our solver uses a parallel kernel-independent Fast Multipole Method implementation (PvFMM) to achieve high scalability. We present results on both the effectiveness and efficiency of our solver, including performance benchmarks, numerical solution accuracy, and parallel scaling results for a range of equation and geometric domain types.

Matthew J. Morse

Courant Institute - New York University  
mmorse@cs.nyu.edu

Abtin Rahimian  
Courant Institute of Mathematical Sciences  
New York University  
arahimian@acm.org

Denis Zorin  
Computer Science Department  
Courant Institute, New York University  
dzorin@cs.nyu.edu

#### PP4

### Multiscale Energy-Conserving Finite Elements for Atmospheric Flows

The issue of the numerical treatment of unresolved scales is of fundamental importance in geophysical fluid dynamics, where the multiscale nature of the problem and the limits in computational power lead to under-resolved simulations. At low resolutions, numerical simulations of atmospheric

flows should conserve energy almost exactly, whereas the unresolved scales should ideally trigger a non-local energy transfer from the smallest resolved scales to the most energetic large ones. We introduce an upwind-stabilised finite element discretisation for the incompressible Euler equations [A. Natale and C.J. Cotter, A variational H(div) finite element discretisation for perfect incompressible fluids, 2016, arXiv:1606.06199], which enforces this behaviour. In particular, we derive the scheme by defining a Lagrangian functional, representing the kinetic energy of the fluid, and applying Hamilton's principle of stationary action. Then, the resulting discretisation preserves the Lagrangian by construction and it can be easily generalised to more complex Hamiltonian fluid models. Moreover, we show that the variational derivation defines a form of upwind stabilisation that implicitly induces an energy transfer from small to large scales, reproducing in this way the characteristic behaviour of two-dimensional turbulent flows.

Andrea Natale

Imperial College London  
a.natale14@imperial.ac.uk

Colin J. Cotter  
Imperial College London  
Department of Mathematics  
colin.cotter@imperial.ac.uk

#### PP4

### Numerical Approximations to Katz Centrality

Graphs and networks are prevalent in modeling relationships across many fields. Using iterative solvers to approximate graph centrality measures (specifically Katz Centrality) allows us to obtain a ranking vector on the nodes of the graph, consisting of a number for each vertex in the graph identifying its relative importance. Given an approximate solution to the vector, we can use the residual to accurately certify how much of the ranking matches the ranking given by the exact solution. In many applications, it is primarily the highest ranked nodes (vertices with the highest centrality scores) that are of importance. We obtain bounds on the accuracy of the approximation to the exact solution with respect to the highly ranked nodes and apply numerical analysis to the computation of Katz Centrality with iterative methods. This relates the numerical accuracy of the linear solver to the data analysis accuracy of finding the correct ranking given by the exact solution. In particular, we answer the question of which pairwise rankings are reliable, given an approximate solution to the linear system. Experiments on several real-world networks validate our theory and show that it is able to accurately certify large portions of the approximation. By analyzing convergence error and using probabilistic matrix norms, we can understand the effect that numerical analysis has in developing confidence in the ranking schemes of data mining.

Eisha Nathan

Georgia Institute of Technology  
enathan3@gatech.edu

Geoff Sanders, Van Henson  
Lawrence Livermore National Laboratory  
sanders29@llnl.gov, henson5@llnl.gov

#### PP4

### An Adjoint Capable Solver for the Stefan Problem

The dynamics of ice ocean interaction (including melting at

the interface and implied geometry changes) can be modeled as an implicit surface moving in accordance with conservation laws; this is known as the Stefan problem. In practice there are numerous sources of uncertainty in the problem (initial conditions, boundary conditions and material properties etc.); a Bayesian framework provides a way to estimate these parameters and indeed quantify their uncertainty. In order to solve the Bayesian inverse problem described here an adjoint capable formulation must be constructed. While many robust methods are available for solving the forward Stefan problem (finite difference based level set methods, fast marching methods etc.), these methods are not well equipped to handle the inverse problem. In this work a PDE constrained bi-level variational optimization formulation is proposed. This formulation is well suited to a Bayesian inverse problem framework, as well as extension to more general conservation laws such as those formulated in ocean circulation models etc. Preliminary results will be shown for a known analytic solution of the Stefan problem as well as for test cases with simple geometries.

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

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

Umberto Villa  
University of Texas at Austin  
uvilla@ices.utexas.edu

Patrick Heimbach  
University of Texas - Austin  
heimbach@ices.utexas.edu

#### PP4

##### On Applying Effective Parallelization of Nonlinear Programming to Topology Optimization

straints on application physics and additional manufacturing constraints (e.g., mass), with applications in design for additive manufacturing (e.g., 3D printing). Topology optimization formulations are typically characterized by large numbers of decision variables (equal to a small constant times the number of PDE degrees of freedom) and computationally complicating integral constraints that act on the entire material domain. Recent results demonstrate that solving formulations with high resolution PDE meshes yields designs that outperform those obtained via existing approaches using IPOPT on coarse meshes. However, the methods used to solve these large-scale formulations are typically limited to simple physics. We propose a domain decomposition approach to topology optimization for both structural physics and electromagnetic physics applications that uses specialized parallel linear algebra kernels within an interior-point method, and will generalize to topology optimization problems with more complex physics. We focus on these applications, and present designs obtained using this approach for both of these applications. We compare our approach to IPOPT, which is commonly used in topology optimization. Our overarching goal is to develop a distributed memory optimization solver that matches or surpasses the parallelism present in the physics PDE solver.

Geoffrey M. Oxberry, Mark L. Stowell, Ryan Fellini,

Daniel White  
Lawrence Livermore National Laboratory  
oxberry1@llnl.gov, stowell1@llnl.gov, fellini1@llnl.gov,  
white37@llnl.gov

Cosmin G. Petra  
Lawrence Livermore  
Mathematics and Computer Science Division  
petra1@llnl.gov

#### PP4

##### Dynamical Polynomial Chaos Expansions with Applications to Long Time Evolution of SDEs and SPDEs

Polynomial chaos expansions (PCE) allow us to propagate uncertainties in the coefficients of differential equations to the statistics of their solutions. While their main advantage is to replace stochastic equations by systems of deterministic equations, their main challenges are the computational cost incurred by high dimensional stochastic variables and the efficiency in long-time computations. In the setting of Markovian random forcing, we propose a new numerical algorithm based on PCE to address these challenges. The method uses a restart procedure and constructs evolving chaos expansions based on polynomials of projections of the time dependent solution and the random forcing with the following consequences: (i) our dynamical expansions retain their optimality for long times; and (ii) the curse of dimensionality is mitigated. To illustrate the efficiency of our algorithm, we apply it to compute solutions of nonlinear stochastic differential equations and stochastic partial differential equations driven by Brownian motion. Long-time simulations include numerical verification of the existence of invariant measures in some cases.

H. Cagan Ozen  
Columbia University  
hco2104@columbia.edu

Guillaume Bal  
Columbia University  
Department of Applied Physics and Applied Mathematics  
gb2030@columbia.edu

#### PP4

##### On the Singular Values Decay of Solutions to a Class of Generalized Sylvester Equations and Efficient Krylov Methods

We are interested in the study of generalized Sylvester matrix equations, that is, the considered linear map is the sum of a Sylvester operator  $\mathcal{L}(X) := AX + XB$  and a positive operator  $\Pi(X) := \sum_{i=1}^m D_i X E_i$ . This kind of equations classically arises in the analysis of bilinear systems and stochastic control but it was recently shown that the discretization of certain elliptic PDEs yields generalized Sylvester equations. Assuming  $\mathcal{L} + \Pi$  to be a convergent splitting of the overall operator, that is  $\rho(\mathcal{L}^{-1}\Pi) < 1$ , a closed-form of the solution  $X$  can be derived and, in this poster, we show how to exploit it to estimate the singular values decay of  $X$ , often observed experimentally. If a low numerical rank is predicted, Krylov solvers can be employed and, under further assumptions on  $\Pi$ , new approximation spaces are introduced. Numerical experiments illustrate the effectiveness of the proposed strategies.

Davide Palitta  
Universita' di Bologna

davide.palitta3@unibo.it

Elias Jarlebring  
KTH Stockholm  
eliasj@kth.se

Giampaolo Mele  
KTH Royal Institute of Technology  
Department of Mathematics, Numerical analysis group  
gmele@kth.se

Emil Ringh  
KTH, Royal Institute of Technology  
erinh@kth.se

#### PP4

##### **A Scalable Linear Solver Based on the FSAI Preconditioner**

A hybrid MPI/OpenMP implementation of a set of preconditioners based on the Factorized Sparse Approximate Inverse (FSAI) technique is developed in this work. These preconditioners are coupled to a conjugate gradient method with the aim to build a scalable linear solver for symmetric positive definite problems. The proposed solver exploits an original algorithm for the optimal dynamic construction of the sparsity pattern used by the approximate inverse and is combined with a two-level domain decomposition approach, thus ensuring a nearly optimal scalability in both strong and weak sense for problems having up to a few millions of unknowns. A detailed description of the algorithms and configuration parameters is provided. Fluid-dynamic and geomechanical real-world numerical models are considered as test cases for the evaluation of the proposed linear solver. Fine-grained numerical tests solved with the one-level FSAI preconditioner demonstrate a nearly optimal strong scalability while coarse-grained tests solved by the two-level FSAI preconditioners show that also the weak scalability is obtained.

Victor A. Paludetto Magri, Andrea Franceschini  
University of Padova  
victor.magri@dicea.unipd.it,  
andrea.franceschini@dicea.unipd.it

Carlo Janna, Massimiliano Ferronato  
Dept. ICEA - University of Padova  
carlo.janna@unipd.it, massimiliano.ferronato@unipd.it

#### PP4

##### **A Two Level Preconditioner for Helmholtz Equation in High Frequency Regime**

The Helmholtz equation is frequently used to model propagation of waves in applications such as acoustic, seismic and electromagnetic systems. Solution to this equation in high frequency regime requires inverting a linear system with a highly indefinite and non-symmetric  $A$  matrix, where standard Krylov methods like GMRES have poor convergence rates without an appropriate preconditioner. A breakthrough in preconditioning this system occurred with the invention of Shifted Laplacian Preconditioner which is based on a Helmholtz-type differential operator with a complex term. This preconditioner can be effectively inverted using a multigrid algorithm and it is known to work well for low to medium range frequencies. However, as the frequency increases the spectrum of the preconditioned system contains many eigenvalues close to

zero which hamper the convergence of the Krylov method. In this work we use a deflation based preconditioner on the Shifted Laplacian preconditioned system to move these eigenvalues away from zero. Numerical tests indicate that such a two level preconditioning technique can work well for high frequencies and large problem sizes. We also discuss the efficient ways of constructing the deflation preconditioner.

Premkumar Panneerchelvan  
The University of Texas at Austin  
premkumarpc@utexas.edu

Laxminarayan Raja  
Dept. of Aerospace Engineering and Engineering  
Mechanics  
The University of Texas at Austin  
lraja@mail.utexas.edu

#### PP4

##### **Generating Long-Term Wind Scenarios Conditioned on Sequential Forecasts**

The large uncertainty in wind power forecasts makes it difficult to incorporate this energy source into the power grid. As a result, control problems have been formulated to operate a controllable energy source in conjunction with wind power to help reduce the uncertainty in future total power generation. The efficacy of these models should be evaluated by rigorously estimating their long-term performance. Existing methods can generate synthetic wind scenarios conditioned on a wind forecast, which is typically between one and three days. By calculating a models optimal controls conditioned on a wind forecast the synthetic scenarios can be used to rigorously estimate the short-term performance. Conversely, by calculating a models optimal controls in a rolling-horizon fashion using historical sequential wind forecasts, the models performance can be estimated using the actual historical wind realization. This approach yields a long-term performance estimate, but only using the single wind realization to evaluate the performance can make it difficult to draw statistically significant results. We develop methods to generate long-term wind scenarios conditioned on a set of sequential short-term forecasts in order to evaluate long-term performance with statistical significance. These methods are tested to ensure that they have statistically similar properties as compared to the historical data. We also investigate methods to generate synthetic sequential wind forecasts.

Kyle Perline  
Cornell University  
krp73@cornell.edu

Christine Shoemaker  
National University of Singapore  
shoemaker@nus.edu.sg

#### PP4

##### **Advancements in Simulation-Based Optimal Experimental Design**

We present recent advancements on a measure-theoretic perspective for Simulation-Based Optimal Experimental Design (SB-OED). This perspective focuses on optimizing quantifiable geometric properties of Quantity of Interest (QoI) maps from model inputs to observable model output data. The two geometric properties considered in the formulation of the SB-OED are related to the accuracy and

precision in solving stochastic inverse problems based on modeling uncertain QoI data using probability measures. We define these geometric properties precisely, provide numerical results illustrating various concepts, and discuss the current and future work of SB-OED.

Michael Pilosov

University of Colorado: Denver  
michael.pilosov@ucdenver.edu

#### PP4

##### Computing of the Effective Coefficients via Multi-scale Discontinuous Galerkin Method

Determination of the macroscopic overall characteristics of heterogeneous media is an essential problem in many engineering applications. The overall behavior of the media depends strongly on the geometrical and physical properties of the microstructural inclusions. To predict the macroscopic behavior of heterogeneous materials various homogenization techniques are typically used but most of them have limitations and are not suitable for complex media. The aim of the work is to develop a computational homogenization technique for the multiscale modeling of heat-transfer problem. The computation scheme for solving the problem is based on multiscale discontinuous Galerkin method (MDG). The solution is obtained on the direct sum of continuous (coarse) and discontinuous (fine) spaces. One of the advantages of the method is that the method allows naturally use the idea of multilevel methods for solving the system of linear equations to significantly reduce the time of the calculation. In the work numerical experiments are carried out for different cases in comparison with the classical finite element method and with analytical methods. The numerical results show that MDG can be successfully used for solving various problems of thermal physics.

Stanislav Polishchuk

Novosibirsk State Technical University  
exsidia@gmail.com

#### PP4

##### Decay Bounds for Functions of Structured Non-Hermitian Matrices

The derivation of decay bounds for the entries of functions of matrices has attracted the interest of many researchers working in a wide range of applications, such as numerical analysis, harmonic analysis, quantum chemistry, signal processing, quantum information theory, multivariate statistics, queuing models, control of large-scale dynamical systems, quantum dynamics, random matrix theory, and others; see [Benzi and Simoncini, Decay bounds for functions of hermitian matrices with banded or kronecker structure, SIAM J. Matrix Anal. Appl, 2015] and references therein. While decay bounds for functions of Hermitian banded matrices have been known for some time, only recently more structured sparsity patterns have been analyzed. Moreover, the non-Hermitian case is an especially challenging setting. By using Faber polynomial series we first extend results for Hermitian matrices to banded non-Hermitian matrices. Several special cases are treated, together with an application to the inexact Krylov approximation of matrix function evaluations. The poster is based on [Pozza, Simoncini, Decay bounds for functions of banded non-Hermitian matrices, Submitted, Aug 2016,

arXiv:1605.01595v3 [math.NA]].

Stefano Pozza

Dipartimento di Matematica  
Università di Bologna  
stefano.pozza@unibo.it

Valeria Simoncini

Università di Bologna  
valeria.simoncini@unibo.it

#### PP4

##### Balanced Iterative Solvers for Linear Systems Arising from Fem Approximation of Pdes with Random Data

This poster discusses the design and implementation of efficient solution algorithms for nonsymmetric linear systems associated with finite element approximation of partial differential equations. The novel feature of our preconditioned preconditioned GMRES and BICGSTAB( $\ell$ ) solver is the incorporation of error control in the natural norm (associated with the specific approximation space) in combination with a reliable and efficient a posteriori estimator for the PDE approximation error. This leads to a robust and balanced inbuilt stopping criterion: the iteration is terminated as soon as the algebraic error is insignificant compared to the approximation error.

Pranjal Pranjal

School of Mathematics  
University of Manchester  
pranjal.pranjal@manchester.ac.uk

David Silvester

University of Manchester  
d.silvester@manchester.ac.uk

#### PP4

##### Modeling the Mitigation of Zika and Chikungunya by Infecting Mosquitoes with Wolbachia Bacteria

We develop and analyze a partial differential equation (PDE) model to investigate the transmission dynamics of releasing Wolbachia-infected mosquitoes to establish an endemic infection in a population of wild uninfected mosquitoes. Wolbachia is a genus of endosymbiotic bacteria that can infect mosquitoes and reduce their ability to transmit Zika or Chikungunya virus. A key dimensionless parameter that determines if a Wolbachia infection can be sustained in the wild is the ratio of infected to uninfected adult mosquitoes. Therefore, accounting for the spatial heterogeneity of the infected and uninfected mosquito population is critical to accurately predict if a release can be sustained. We create a PDE model for the spread of a Wolbachia epidemic within the mosquito population that captures both the complex vertical transmission cycle (both adult and aquatic mosquito life-stages) and the horizontal transmission (spatial diffusion) of Wolbachia. To study the proposed PDE model and analyze how the magnitude of the diffusion coefficient can determine if a locally highly infection region will grow, or shrink, we study the one-dimensional version of the equations with cylindrical symmetry to investigate if a circular patch shrinks or grows. We then look into the full two-dimensional simulations and analysis of stability of traveling fronts in both transverse and normal directions.

Zhuolin Qu

Tulane University  
zqu1@tulane.edu

Ling Xue  
Harbin Engineering University  
lxue@hrbeu.edu.cn

James Hyman  
Tulane University  
mhyman@tulane.edu

#### PP4

##### **New Rigorous Error Bound For Reduced Basis Approximations**

The reduced basis method is a model order reduction technique that is suitable and efficient for parameter-dependent systems in the many-query or real-time contexts. In most cases, for the construction of the reduced basis the so called standard error estimator is used. It can be shown, that this error estimator is rigorous, but as a disadvantage the inf-sup constant has to be estimated, which can be done with the so called successive constraint method (SCM). To avoid the computation of the inf-sup constant or the SCM, we present a new rigorous error bound, which is independent of the inf-sup constant. This error bound is applied on the stationary helmholtz equation and time-dependent wave equation.

Mladjan Radic, Stefan Hain  
PhD  
mladjan.radic@uni-ulm.de, stefan.hain@uni-ulm.de

#### PP4

##### **Decentralized Framework for Sensor-Driven Optimization in Power Systems**

Optimization problems in modern Cyber-Physical Systems is increasingly challenged by two key aspects, large number of physical subsystems, and huge amount of real-time sensor data. As a result, recent decentralized approaches that capture real-time dynamics of the constituent system components are becoming mainstream. This is particularly true in power generation networks, where the integration of sensor information from highly distributed assets is key to reliability and profitability. In this paper, we develop an asynchronous decentralized optimization framework for the problem of generator maintenance in power networks. Sensor data analytics are used to derive an optimal maintenance schedule for fleets of generators distributed across geographically remote regions. Asynchronous communication mechanisms are used to enable scalability of decentralized methods for solving large-scale optimization problems that also leverage the dynamic nature of the sensor data. We demonstrate the computational performance and scalability of our framework relative to the conventional techniques.

Paritosh P. Ramanan  
School of Computational Science and Engineering  
Georgia Institute of Technology  
paritoshpr@gatech.edu

Murat Yildirim  
School of Industrial and Systems Engineering  
Georgia Institute of Technology  
murat@gatech.edu

Edmond Chow  
School of Computational Science and Engineering  
Georgia Institute of Technology  
echow@cc.gatech.edu

Nagi Gebraeel  
School of Industrial and Systems Engineering  
Georgia Institute of Technology  
nagi@isye.gatech.edu

#### PP4

##### **Stochastic Analysis of Turbulent Mix**

For multi-fluid mixing problems, the interface between the two fluids or the front serves as the separation between the fluids. This interface is tracked dynamically using Stony Brook's Front-tracking algorithm. The evolving front between the two fluids can be used to indicate the level and the type of mix. We explore the spatial distribution of the mix by looking at the front's statistical properties. A random point along the front is chosen and using the front normal as the search direction we measure the distance until the first intersection of the front occurs. This process can then be repeated to generate the distance until the second crossing and so on. Using this distance (crossing distance), we produce a dataset of crossing distances. We investigate the relationship between first and second crossings and analyze whether the overall process is memoryless or not, i.e., if the governing process is Poisson or not. We also repeat this analysis across a range of grid resolutions and Reynolds numbers and study the dependence of the crossing length scales with regard to these parameters. Our results show that both the first and second crossings follow exponential distribution individually, but are not independent of each other. We observe a sub-linear increase in the rate parameter for exponential distribution as the mesh is refined, however the rate parameter is almost unchanged as the Reynolds number is increased.

Pooja Rao  
Stony Brook University  
prao@ams.sunysb.edu

Jeremy A. Melvin  
University of Texas - Austin  
jmelvin@ices.utexas.edu

Hyunkyung Lim, James Glimm  
Stony Brook University  
hyulim@ams.sunysb.edu, glimm@ams.sunysb.edu

David Sharp  
Los Alamos National Laboratory  
dco@lanl.gov

#### PP4

##### **Estimating the Uncertainty of Imprecise Computer Models Using Optimization Methods**

Predictions made in the engineering and scientific community rely extensively on computer simulations of complex systems. However, every computer-based solution is influenced by numerical, parameter, and structural uncertainty. Having a measure of credibility for a given simulation allows scientists to generate results upon which to make sound predictions. A typical approach to address this issue is to use Bayesian Inference, which is considered to be computationally demanding in practice. An alternative



approach is to view the task as an optimization problem and to estimate the structural uncertainty using a polynomial; in literature this is referred to as discrepancy. Such an approach has been implemented using Particle Swarm Optimization(PSO); these results merit further research. Global optimization methods for non-convex problems will be implemented to increase the rate of convergence to a viable solution while mitigating computational cost relative to the implementation of PSO. Research for a more agreeable form of the discrepancy polynomial will be undertaken; possible function forms include but are not limited to Legendre and Bessel functions.

Lee R. Redfean

Clemson University  
Department of Mathematical Sciences  
lredfea@g.clemson.edu

#### PP4

##### Containers for Scientific Computing: From Laptop to HPC

Complex software stacks are difficult to maintain, even for the developers. Libraries may have many dependencies, and the availability and reliability of these dependencies can vary widely between systems. Containerization makes it easy to encapsulate and distribute a complete software environment, containing all required dependencies for a particular piece of software, and run it on many diverse platforms. Docker ([www.docker.com](http://www.docker.com)) has become a very popular container system recently, as it allows containers in the docker format to be run on any suitable Linux host (or Linux VM in Windows or macOS), as well as providing a central registry where container images can be shared. Shifter, a NERSC project, allows docker images to run on HPC systems. We demonstrate the use of docker images with the FEniCS Finite Element package, and show that it is possible to use the same container to run on a laptop as on a HPC system, with no appreciable loss of performance compared to running natively. HPC systems can be difficult to configure: using containers brings a clear benefit in accessibility, and can also boost performance.

C N. Richardson

University of Cambridge  
chris@bpi.cam.ac.uk

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

Jack Hale  
University of Luxemburg  
jack.hale@uni.lu

#### PP4

##### A Finite Element Discrete Fracture Model to Simulate Fluid Flow Through Fractured Porous Media

In this work a methodology is presented for modeling a fractured porous media and to simulate fluid flow through multiple fractures and rock matrix, with drastically different properties, at a scale such that the fractures could be modeled explicitly. To represent the porous matrix a continuum approach is employed whereas for the fractures a discrete one. The necessity of representing the thickness of each of the fractures in the computational model is eliminated, given that it is explicitly considered in the

mathematical model. Thus, the fractured porous media is modeled with mixed dimensions elements, representing the fractures as elements of n-1 dimensions immersed in a porous matrix of n dimensions by isolated internal boundaries, where the equations that govern the flow of fluids in the matrix and the fractures are coupled by means of jump and average equations, taking into account interactions between the fractures and the surrounding porous media. The derived flow model in fractured porous media is single phase, based on the fluid pressure and considering the Darcys law for fluid velocity in the porous media and fractures, it is assumed that the fractures are filled with a porous material. For the numerical solution is applied a Finite Element Method and its computational implementation is carried out in Python using FeniCS project. Finally, the discrete fracture model is numerically validated in a case study in two dimensions.

Carlos Romano-Pérez

Posgrado de Ingeniería, UNAM  
Exploración y Explotación de Recursos Naturales  
carlos\_alberto@comunidad.unam.mx

Martin Díaz-Viera  
Instituto Mexicano del Petróleo  
mdiazv@imp.mx

#### PP4

##### Neural Networks As Reduced Models for Physical Systems and Inverse Problems

Artificial neural networks have exploded out of the world of machine learning into a wide array of industrial and academic applications, but their efficacy in modeling physical systems as represented by parameterized partial differential equations has not been definitely established. In particular, if a neural network can reliably model a nonlinear parameter-to-output mapping of a system governed by PDEs, then using the network as a reduced model for the PDEs can provide a significant computational advantage when exploring parameter space. The latter is important for inverse problems, optimal control and design, and uncertainty quantification. We demonstrate that a shallow, feedforward neural network can model this nonlinear relationship, as well as a methodology for training these networks to represent the inverse relationship, which is often ill-posed. The application is forward and inverse scattering from layered media, as modeled by the Helmholtz equation.

Bassel Saleh, Umberto Villa  
University of Texas at Austin  
bassel.s.saleh@gmail.com, uvilla@ices.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

#### PP4

##### Adaptive Sparse Grids Interpolation Techniques for Multilevel Stochastic Collocation in Fluid-Structure Interaction Problems

This contribution discusses the application of sparse grid interpolation in the quantification of uncertainty of fluid-structure interaction problems. We model the uncertainty in five input parameters as independent normal random variables and use the established multilevel polynomial chaos-based stochastic collocation method, which minimizes the overall cost via a hierarchical decomposition in

both the spatial and parameter domains. We employ hierarchical finite element meshes for the spatial discretization and we interpolate the underlying solver in the parameter space with either standard sparse grids or the combination technique, the latter being constructed using Leja points. While previous works focused on global Lagrangian-type polynomials, we employ local hierarchical basis functions. In the spatial domain we increase the resolution by a constant factor, whereas in the parameter space we use adaptive refinement, keeping the number of collocation points small. Using the tensor structure of sparse grids and the stochastic independence of the inputs, the problems corresponding multivariate integral breaks into a sum of products of univariate integrals that can be evaluated exactly. Applying the method to the interaction of a fluid with an elastic beam, we obtain the expectation, variance, and Sobol' indices for global sensitivity analysis for the velocity in the entire domain, as well as for the x axis displacement and force at user-defined structure coordinates.

Paul Cristian Sarbu  
 Technical University Of Munich  
 Department of Informatics  
 sarbu@in.tum.de

Ionut-Gabriel Farcas  
 Technical University of Munich  
 Department of Scientific Computing in Computer Science  
 farcasi@in.tum.de

Benjamin Uekermann, Tobias Neckel  
 Technische Universität München  
 uekerman@in.tum.de, 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

#### PP4

##### **Implicit Mesh Discontinuous Galerkin Methods and Interfacial Gauge Methods for High-Order Accurate Interface Dynamics, with Applications to Surface Tension Dynamics, Rigid Body Fluid-Structure Interaction, and Free Surface Flow**

A high-order accurate implicit mesh discontinuous Galerkin framework is presented for fluid interface dynamics, facilitating precise computation of interfacial fluid flow in evolving geometries. The framework uses implicitly defined meshes, wherein a reference quadtree or octree grid is combined with an implicit representation of evolving interfaces and moving domain boundaries, and allows physically-prescribed interfacial jump conditions to be imposed or captured with high-order accuracy. Several topics are highlighted in the design of the framework, including high-order accurate quadrature on implicitly defined mesh elements, optimal-order accurate and symmetric positive definite discretisations of scalar- and vector-valued elliptic interface PDEs, multigrid algorithms, and projection operators for incompressible fluid flow. A variety of 2D and 3D applications of the framework are also presented, including surface tension-driven two phase flow with phase-dependent fluid density and viscosity, rigid body fluid-structure interaction, and free surface flow. By utilising a class of techniques called interfacial gauge methods, high-order accuracy in the maximum norm is achieved. Particular examples include: (i) high Reynolds number soap

bubble oscillation dynamics and (ii) vortex shedding in a water ripple free surface flow problem; these last two examples compare numerical results with that of experimental data and also reveal physical phenomena not seen in the experiments.

Robert Saye  
 Dept. of Mathematics  
 University of California, Berkeley  
 rsaye@lbl.gov

Robert Saye  
 Lawrence Berkeley National Laboratory  
 rsaye@lbl.gov

#### PP4

##### **Moving Objects Through Phase Change Material - a Hybrid Modeling Approach**

Phase change processes are key to a variety of present-day industrial applications. One example are thermal energy storages, in which it is exploited that the so-called phase change material (PCM) absorbs and releases heat at a certain temperature. Another interesting research field is subglacial exploration with thermal melting probes. Here, heat transfer and phase change processes are coupled to fluid flow of the liquid PCM and rigid body dynamics. This multi-physics situation is often referred to as close-contact melting (CCM). In our contribution we present a novel modelling technique to solve the CCM of objects (e.g. a melting probe) that migrate through a PCM. Challenging to CCM is the existence of two spatial scales, given by the objects size and the much smaller width of the melt film. A stand-alone method that resolves both scales would be computationally very expensive. We therefore propose an efficient hybrid approach consisting of two complementary modules. In order to obtain the melting velocity, we use a simplified melt-film model that solves the CCM problem restricted to the liquid domain. The melting velocity, on the other hand, is the prerequisite to a two-phase model that enables us to study the macro-scale thermal state in both the liquid, and the solid PCM. We will present simulation results and compare these to real data. Based on the model, we are able to draw conclusions towards the efficiency of the overall CCM process.

Kai Schüller  
 AICES Graduate School  
 RWTH Aachen University  
 schueller@aices.rwth-aachen.de

Julia Kowalski  
 AICES  
 RWTH Aachen University  
 kowalski@aices.rwth-aachen.de

#### PP4

##### **A Stable Added-Mass Partitioned (AMP) Algorithm for Elastic Solids and Incompressible Flows**

A stable added-mass partitioned (AMP) algorithm is developed for fluid-structure interaction (FSI) problems involving viscous incompressible fluids and compressible elastic solids. The traditional coupling algorithm suffers from the added-mass instability, no matter how heavy the solid, when the grid is sufficiently fine. However, the AMP scheme is stable even for light solids when added-mass effects are large. The main ingredient of the AMP algorithm is a Robin (mixed) interface condition on the fluid pressure.

This condition is derived by combining the outgoing characteristic structure in the solid with the velocity and traction matching conditions. Normal-mode analysis is used to analyze the stability of the traditional, anti-traditional, and AMP schemes with second-order accurate discretizations and a linearized interface. The AMP algorithm is shown to be stable for any ratio of solid and fluid densities, including when added-mass effects are large. The traditional algorithm is shown to be unconditionally unstable as added-mass effects become large with grid refinement. Exact solutions are obtained for the finite amplitude case and are used to verify the stability and accuracy for light, medium, and heavy solids. Numerical examples are considered to showcase the robustness of the AMP algorithm and to demonstrate interesting applications.

Daniel A. Serino

Department of Mathematical Sciences  
Rensselaer Polytechnic Institute  
serind@rpi.edu

Jeffrey W. Banks, William Henshaw  
Rensselaer Polytechnic Institute  
banksj3@rpi.edu, henshw@rpi.edu

Donald W. Schwendeman  
Rensselaer Polytechnic Institute  
Department of Mathematical Sciences  
schwed@rpi.edu

**PP4**

**A Structure Preserving Lanczos Algorithm for Computing the Optical Absorption Spectrum**

We present a new structure preserving Lanczos algorithm for approximating the optical absorption spectrum in the context of solving full Bethe–Salpeter equation without Tamm–Dancoff approximation. The new algorithm is based on a structure preserving Lanczos procedure, which exploits the special block structure of Bethe–Salpeter Hamiltonian matrices. A recently developed technique of generalized averaged Gauss quadrature is also incorporated to accelerate the convergence. We also establish the connection between our structure preserving Lanczos procedure with several existing Lanczos procedures developed in different contexts. Numerical examples are also presented to demonstrate the effectiveness of our Lanczos algorithm.

Meiyue Shao

Lawrence Berkeley National Laboratory  
myshao@lbl.gov

Felipe H. Da Jornada  
UC Berkeley  
jornada@berkeley.edu

Lin Lin  
University of California, Berkeley  
Lawrence Berkeley National Laboratory  
linlin@math.berkeley.edu

Chao Yang  
Lawrence Berkeley National Lab  
cyang@lbl.gov

Jack Deslippe  
National Energy Research Scientific Computing Center  
jdeslippe@lbl.gov

Steven Louie  
University of California, Berkeley  
sglouie@berkeley.edu

**PP4**

**A Weak Galerkin (WG) Method for Maxwell's Equations in the Time Domain**

The WG method was first applied to Maxwell's equations in the frequency domain by Mu, Wang, Ye and Zhang in 2013. Due to the lack of applications and analysis of the WG method with time domain models, we have developed a WG method for the time domain Maxwell's equations. Stability and error estimates for this scheme are provided as well.

Sidney R. Shields, Jichun Li

University of Nevada Las Vegas  
shields3@unlv.nevada.edu, jichun.li@unlv.edu

Eric Machorro  
National Security Technologies  
machorea@nv.doe.gov

**PP4**

**Corrugated Coaxial Cable Modeling with a Nodal Discontinuous Galerkin (NDG) Method**

Due to the lack of well developed time domain models for a corrugated coaxial cable, the cable was modeled by numerically solving axially symmetric Maxwell's equations in the dielectric. Although there has already been many implementations of a nodal Discontinuous Galerkin method for solving the axially symmetric Maxwell's equations, there has yet to be any stability or convergence results proven for this method.

Sidney R. Shields, Jichun Li

University of Nevada Las Vegas  
shields3@unlv.nevada.edu, jichun.li@unlv.edu

Eric Machorro, Jerome Blair  
National Security Technologies  
machorea@nv.doe.gov, jerry@nf6z.net

**PP4**

**Numerical Simulation of Poroelastic Wave Equation with Discontinuous Galerkin Using Upwind and Modified Penalty Flux**

Poroelasticity deals with the propagation of elastic waves in porous media where both the mineral grains (solids) and interstitial fluids are excited as the wave energy dissipates through. The poroelastic wave equation couples elastic wave equation with Darcys law (in the form of a dissipation potential) which results in system of eight coupled hyperbolic PDEs in 2D and thirteen in 3D. The poroelastic system is, however, difficult to implement numerically due to the presence of dissipation potential, which poses a stability problem on numerical scheme. In this article we show how wave propagation in an anisotropic porous system can be expressed in conservation form by combining the Biot's and Hamiltonian mechanics. We also show that high-order nodal discontinuous Galerkin (dG) method yields a stable and accurate solution for the poroelastic system of equations. We also compare and contrast two methods for implementation of natural boundary conditions. First we attempt to solve the Riemann problem (upwind flux) using

spectral decomposition of the Jacobian matrices. In the Poroelastic system, we find that the Jacobian is unstructured and hence, the computation of its eigensystems is quite challenging; it is equivalent to solving a polynomial of the order that is equal to the rank of Jacobian matrix. Second, we circumvent the spectral decomposition by using the standard central flux appended with the penalty terms based on the interior boundary continuity.

Khemraj Shukla  
OKLAHOMA STATE UNIVERSITY  
khemraj@okstate.edu

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

Priyank Jaiswal  
OKLAHOMA STATE UNIVERSITY  
priyank.jaiswal@okstate.edu

#### PP4

##### Fast Dynamic Load Balancing Tools for Extreme Scale Systems

Massively parallel computing combined with scalable simulation workflows that can reliably model systems of interest are central to the continued quest of scientists, engineers, and other practitioners to address advances in scientific discovery, engineering design, and medical treatment. However, to meet their potential, these methods must be able to operate efficiently and scale on massively parallel computers executing millions of processes. Reaching the goal of millions of parallel processes requires new methods in which the computational workload is extremely well balanced and interprocessor communications overheads are minimized. Building on initial efforts to improve dynamic load balancing methods for adaptive unstructured mesh applications, the goal of our research is to develop fast multicriteria dynamic load balancing methods that are capable of quickly producing well balanced computations, with well controlled communications, for a wide variety of applications. To meet this goal we apply specific combinations of methods that consider the existing distributions, concurrency level (i.e. number of processes), memory limitations, and application workflow requirements.

Cameron Smith  
Scientific Computation Research Center  
Rensselaer Polytechnic Institute  
smithc11@rpi.edu

Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

#### PP4

##### Apollo: An Unstructured Framework for Multi-Fluid Plasma Modeling

In certain applications a plasma model more general than the magnetohydrodynamic (MHD) model is needed. A choice that is not too computationally expensive is the multi-fluid plasma (MFP) model. The MFP model equations are derived by taking velocity moments of the Boltzmann equation for each of the components in a plasma, and each species mass density, momentum density and total energy are evolved in time. An unstructured discontin-

uous Galerkin (DG) framework for the MFP model is implemented using the Portable, Extensible Toolkit for Scientific Computation (PETSc). The framework takes advantage of PETScs unstructured data management subclass DMPlex as well as its time stepping library TS to solve ordinary differential equations (ODE) and differential algebraic equations (DAE). APOLLO is designed as a science/engineering framework in and attempt to bridge the gap between the complex geometries of experiments, magnetic coil structures, and basic principles physical models. In addition, the framework is to be coupled with radiation, ionization/recombination and collisional-radiative (CR) models currently being developed at the Air Force Research Laboratory. Distribution A: Approved for public release; distribution unlimited AFTC/PA clearance No. 16062.

Eder M. Sousa  
University of Washington  
eder.sousa.ctr@us.af.mil

#### PP4

##### Adaptive Smoothed Aggregation Episode II: Return of the Bad Guys

Algebraic multigrid (AMG) is a popular solver for large sparse linear systems, often resulting from the discretization of partial differential equations (PDEs). The setup and solve phase are linear in complexity, and AMG has shown to scale to hundreds of thousands of processors for certain problem types. However, AMG relies on a priori knowledge of algebraically smooth modes of the system, i.e. modes associated with small eigenvalues, colloquially known as ‘bad guys,’ for effective convergence. This has motivated the development of adaptive methods, including adaptive smoothed aggregation ( $\alpha$ SA), to generate bad guys when a sufficient set is not known a priori. In this work we extend on the original  $\alpha$ SA framework, developing new methods to construct a set of bad guys, as well as incorporate them into interpolation operators. Numerical results show the potential of  $\alpha$ SA II to improve the robustness of  $\alpha$ SA, while also reducing the setup and solve complexity.

Ben Southworth  
University of Colorado, Boulder  
ben.southworth@colorado.edu

Steve McCormick, John Ruge, Tristan Konoldige  
University of Colorado at Boulder  
stephen.mccormick@colorado.edu,  
john.ruge@colorado.edu, tristan.konoldige@colorado.edu

#### PP4

##### Graph Sparsification Approach for Updating Dynamic Networks on Shared Memory Systems

Network analysis is an important tool for studying large-scale systems of interacting entities that arise in diverse domains such as bioinformatics, and sociology. Properties of networks, such as connectivity, shortest paths between vertices can provide insights into the characteristics of the underlying systems. Since the networks are extremely large, parallel algorithms are essential for timely analysis. However, developing scalable parallel algorithms for networks is very challenging. This is because graph traversal is the primary component of many network algorithms. Traversal over unstructured data, such as networks, lead to irregular memory accesses resulting in low scalability and high com-

putation costs. The problem is even more difficult when the networks are dynamic. In this poster, we present a framework for creating fast and scalable parallel algorithms for updating properties of dynamic networks. Our framework is created using an elegant technique known as graph sparsification. Graph sparsification uses a divide and conquer approach for updating the network properties. We specifically demonstrate how scalable algorithms for minimum spanning tree and single source shortest paths can be developed using graph sparsification.

Sriram Srinivasan

Department of Computer Science  
University of Nebraska, Omaha  
sriram882004@gmail.com

#### PP4

##### **A Plasma-Vacuum Interface Tracking Algorithm for Fully-Implicit Magnetohydrodynamic Simulations**

A level-set based plasma-vacuum interface tracking algorithm is described in the context of fully implicit resistive Magnetohydrodynamic (MHD) simulations. The conventional approach of modelling the expansion of a high density thermal plasma transient into a vacuum background involves approximating the vacuum using a low background density gas and solving the continuum based MHD formulation throughout the domain. We demonstrate that this approach leads to a poor conditioning of the non-linear and associated linear systems formed as a result of the implicit temporal discretization. This issue is overcome by tracking the plasma-vacuum interface and formulating the implicit system exclusively in the cells which are tagged as containing plasma. Such a formulation eliminates the influence of the background density on the conditioning of the linear and non-linear systems, while simultaneously reducing the problem size. As a result, significant improvements are observed in convergence rates as well as net wall-clock time for the simulations. The plasma-vacuum interface is evolved in a physically consistent manner by imposing a face flux at the interface that is derived by approximately solving a Riemann problem with vacuum as one of the initial states.

Vivek Subramaniam

The University of Texas at Austin  
vivek91@utexas.edu

Laxminarayan Raja

Dept. of Aerospace Engineering and Engineering  
Mechanics

The University of Texas at Austin  
lraja@mail.utexas.edu

#### PP4

##### **High-Dimensional Function Interpolation with Gradient-Enhanced Weighted $\ell^1$ Minimization**

Many physical problems involve in approximating high-dimensional functions with a limited number of sampling points. It is seen that the high-dimensional function interpolation problem has various applications such as uncertainty quantification. For example, in order to compute a quantity of interest for the parametric PDE, high-dimensional function approximation is often required. In this poster, we present the work of interpolating a high-dimensional function using the weighted  $\ell^1$  minimization technique when both points for the original function and

its derivatives are sampled. With additional derivative information, we see a numerical improvement over the case only samples points from the original function. A theoretical analysis for this interpolation problem with derivative information incorporated is also presented.

Yi Sui, Ben Adcock

Simon Fraser University  
ysui@sfu.ca, ben\_adcock@sfu.ca

#### PP4

##### **Advances in Nonlinear Solvers for Coupled Systems**

Unstructured polytopal meshes capture critical variability of computational domains, e.g. topography and stratigraphy in terrestrial applications. Advanced spatial discretizations, such as the Mimetic Finite Difference and Nonlinear Finite Volume methods, are required to maintain accuracy, but analytic or automatic differentiation of the discrete system may lead to an unstable Jacobian. We control preconditioner properties for Jacobian-Free nonlinear solvers by discretizing the analytic Jacobian. The efficacy of this approach is demonstrated for coupled systems of energy and flow in watershed modeling.

Daniil Svyatskiy

Los Alamos National Laboratory  
dasvyat@lanl.gov

#### PP4

##### **A Stochastic Permeability Model for Reduced Order Simulation and Representative Volume Element Prediction**

There is an increasing interest in the application of methods of machine learning to reduced order modeling, simulation and uncertainty quantification for geophysical systems. In this work, data from LBM simulation of fluid flow through porous media in the non-Darcy regime is used to construct a reduced order stochastic formulation of the dependence of the momentum resistance tensor on the pore geometry, imposed pressure gradient, and flow rate using the Karhunen-Loeve expansion and its nonlinear generalization. Analysis of the explained variance determined by both linear and kernel PCA reveals that this dependence may be represented by a small number of stochastic modes. Direct comparison of the reduced stochastic model with the prevailing nonlinear corrections to Darcy's law reveal agreement with standard permeability models at the scale of a Representative Volume Element (RVE) for monodisperse isotropic packings of spheres, and enhanced accuracy for polydisperse anisotropic packings of ellipsoids. We demonstrate how the decay in variance described by the stochastic model with the size of the medium furnishes a means of correctly predicting the size of a RVE using a small number of simulations below the scale of the known RVE for monodisperse isotropic media. Finally, we show how the stochastic permeability model may be combined with stochastic collocation methods to facilitate a method for reduced simulation of flow through porous media, and present related results.

Charles Talbot

Lawrence Livermore National Laboratory  
University of North Carolina at Chapel Hill

ctalbot@live.unc.edu

#### PP4

##### Simulation of Coastal Flows and Waves by Integration of Geophysical Fluid Dynamics and Fully 3D Fluid Dynamics Models

In order to simulate multiscale coastal ocean flows, we proposed a modeling system that is a hybrid of a geophysical fluid dynamics model and a fully 3D fluid dynamics model. The former captures background estuary flows and the latter resolves local flows of interest. Such system is the first of its kind and able to simulate in high-fidelity various flow phenomena at spatial scales  $O(1)$  m –  $O(10,000)$  km, especially those complicated, fully 3D flows at small scales, and such simulation is essentially beyond the reach of any other existing models. On the basis of our previous efforts with this regard, the system is further extended to include free surfaces at local flows. In this presentation, the methodology will be outlined, and a few validation cases will be presented, together with an example modeling of propagation of tsunami-like waves all the way from a deep ocean to seashore and then impinging a coastal structure. Difficulties and further development will be also discussed.

Hansong Tang

Dept. of Civil Eng., City College, City Univ. of New York  
htang@ccny.cuny.edu

#### PP4

##### Mesosopic and High Performance Modeling of Biomimetic Polymers

In this poster I will present our latest development on the mesoscopic modeling of polypeptoids, a class of biomimetic polymers, using the Dissipative Particle Dynamics method. A bottom-up parameterization approach based on the Mori-Zwanzig projection will be used to construct the DPD model from low-level molecular dynamics simulations. The model is then used to study the self-assembly behavior of polypeptoids in aqueous solution whose monomer sequence can be precisely controlled using solid phase synthesis. We then compare the bottom-up DPD model with the more common empirical models in terms of accuracy, computational efficiency and parameterization complexity.

Yu-Hang Tang

Brown University  
yuhang\_tang@brown.edu

Zhen Li

Division of Applied Mathematics  
Brown University  
Zhen\_Li@brown.edu

George E. Karniadakis

Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

#### PP4

##### Quantized Tensor Train and Uniformization Approach for Stochastic Chemical Kinetics

The functioning of a biological cell is determined by the complex interactions of chemical components, whose discrete populations evolve as a continuous-time, discrete-

state Markov process. Finding the probability distribution requires solving the chemical master equation (CME), whose state space grows exponentially with the number of species. We employ the quantized tensor train (QTT) format to represent both the transition rate matrix and the initial condition of the CME. We then seek the solution at every time step within this format, using a combination of uniformization and the alternating minimal energy method.

Huy D. Vo

University of Alabama  
hvo@crimson.ua.edu

Roger Sidje

Department of Mathematics  
University of Alabama  
roger.b.sidje@ua.edu

#### PP4

##### Fast Algorithms and Computational Software for Obtaining Low Rank Matrix Decompositions

Low rank matrix decompositions are useful in a variety of applications, including large scale data processing and classification, inverse problems in geophysics, image enhancement and compression, and video processing. We present several newly developed software packages, which enable the user to compute various popular low rank matrix decompositions, including the low rank SVD, ID, and CUR factorizations of matrices of various sizes. The software utilizes new developments in randomized sampling techniques to achieve highly competitive runtimes; in many cases, far out-accelerating existing software packages. We present several newly developed shared memory codes targeting multi-core and GPU architectures and popular software packages such as R. We also present distributed memory codes suitable for very large matrices implementing a novel block matrix algorithm. All the developed software we present is open source and freely available for use.

Sergey Voronin

Tufts University  
Sergey.Voronin@tufts.edu

N. Benjamin Erichson

University of St Andrews, United Kingdom  
nbe@st-andrews.ac.uk

#### PP4

##### Improving RANS Predictive Capability Based on Machine Learning

Although increased availability of computational resources has enabled high-fidelity simulations, the Reynolds-Averaged Navier-Stokes (RANS) models are still the workhorse tools in the turbulence modeling of industrial flows. However, the model discrepancy due to the inadequacy of modeled Reynolds stresses largely diminishes the reliability of simulation results. In this work, we use a physics-informed machine learning approach to improve the RANS modeled Reynolds stresses and propagate them to obtain the mean velocity field. Specifically, the functional forms of Reynolds stress discrepancies with respect to mean flow features are trained based on an offline database of flows with similar characteristics. The random forest model is used to predict Reynolds stress discrepancies in new flows. Then the improved Reynolds stresses are propagated to the velocity field through RANS equa-

tions. The effects of expanding the feature space using an integrity basis of Galilean tensor invariants are also discussed. The flow in a square duct, which is challenging for standard RANS models, is investigated to demonstrate the merit of the proposed approach. The results show that both the Reynolds stresses and the propagated velocity field is improved over the baseline predictions.

Jianxun Wang, Jinlong Wu, Heng Xiao  
Dept. of Aerospace and Ocean Engineering, Virginia Tech  
vtwjx@vt.edu, jinlong@vt.edu, hengxiao@vt.edu

#### PP4

##### Numerical Computing with Functions in Polar and Spherical Geometries

A synthesis of the classic Double Fourier sphere method and new ideas in low rank approximation is used to construct approximations to functions in polar and spherical geometries. This approach preserves bi-periodicity for functions on the sphere, overcomes issues associated with artificial singularities, enables the use of fast, FFT-based algorithms, and is near-optimal in its underlying interpolation scheme. It has been used to develop a suite of fast, scalable algorithms for numerical computing with functions on the surface of the sphere and on the disk, and these have been implemented in the Chebfun computing system.

Heather D. Wilber  
Boise State University  
hdw27@cornell.edu

Grady B. Wright  
Department of Mathematics  
Boise State University, Boise ID  
gradywright@boisestate.edu

Alex Townsend  
Cornell University  
townsend@cornell.edu

#### PP4

##### Reducing Communication in Distributed Asynchronous Iterative Methods

For computing at the exascale, reducing internode communication may become a more important factor than reducing overall run time. This poster compares two new algorithms for reducing communication in fixed-point iterative methods for solving sparse linear systems. These two algorithms are based off the sequential Southwell method, where instead of relaxing rows in order, i.e., Gauss-Seidel, the row with the maximum residual norm is relaxed. Southwell often requires fewer relaxations to converge than Gauss-Seidel, but, Southwell is not parallel and requires global communication on each step. Our first algorithm, the parallel Southwell method, addresses these two problems. Instead of relaxing a single row per parallel step, a row relaxes if it has the maximum residual norm within its neighborhood, which allows for simultaneous relaxation of rows, and no global communication. In practice, rows communicate residual information to neighbors by appending residual norms to outgoing messages. However, when implemented asynchronously, rows may hold inaccurate copies of neighboring residuals. This motivates our second algorithm, the distributed Southwell method, which requires that rows hold estimates of neighboring residuals, which are updated when the row relaxes, and reset after receiving an update from a neighbor. This poster com-

pares these two algorithms with both asynchronous and synchronous Jacobi, and explores the potential for multi-grid smoothing.

Jordi Wolfson-Pou  
School of Computational Science and Engineering  
Georgia Institute of Technology  
jwp3@gatech.edu

#### PP4

##### Representation of Discrepancies Between Stress Tensors and Its Application in Data-Driven Turbulence Modeling

Reynolds-Averaged Navier–Stokes (RANS) equations are still an important tool for engineering design and analysis applications involving wall bounded turbulent flows. However, the modeled Reynolds stress tensor is known to be a main source of uncertainty, diminishing the reliability of the results of RANS simulations. Recently, there is growth interest in estimating the discrepancies in the RANS modeled Reynolds stress tensor with machine learning techniques. In this work, we discuss the machine learning performances with regard to different approaches in representing the discrepancies between stress tensors. The results show that the learning performances are influenced by both the representation of discrepancies between stress tensors and the machine learning algorithms. Based on the results, a general procedure of choosing proper output quantities of machine learning techniques will be discussed in the context of computational mechanics.

Jinlong Wu, Jianxun Wang, Heng Xiao  
Dept. of Aerospace and Ocean Engineering, Virginia Tech  
jinlong@vt.edu, vtwjx@vt.edu, hengxiao@vt.edu

#### PP4

##### Model Discrepancy for Simulating Complex Physical Systems: From Data Assimilation to Machine Learning

Many complex systems are characterized by physics at a wide range of scales, for which first-principle-based high-fidelity models resolving all the scales are prohibitively expensive to run. Consequently, practical simulations have primarily relied on low fidelity models with approximate closure models, which introduced model-form uncertainties and diminish their predictive capabilities. In this work, we present a comprehensive framework for using data to reduce model uncertainties with emphasis on turbulent flow applications. With online, continuously streamed monitoring data, we use data assimilation and Bayesian inference techniques; With offline data from related flows, we use a physics-informed machine learning technique. The framework is general enough to be extended to other complex physical systems, e.g., climate, composite materials, granular and multiphase flows.

Heng Xiao, Jinlong Wu, Jianxun Wang  
Dept. of Aerospace and Ocean Engineering, Virginia Tech  
hengxiao@vt.edu, jinlong@vt.edu, vtwjx@vt.edu

#### PP4

##### LOBPCG Method in Exact Diagonalization for Hubbard Model: Performance of Communication Avoiding Neumann Expansion Preconditioner

The Hubbard model has attracted a tremendous member of physicists, since the model exhibits a lot of interesting

phenomena such as high-temperature superconductivity. One of the computational approaches to solve the model is the exact diagonalization method which solve the smallest eigenvalue and the corresponding eigenvector of the Hamiltonian derived from the model. Since the Hamiltonian is a huge sparse matrix, iteration methods are usually utilized. The LOBPCG method is one of the most powerful methods, and we have succeeded in solving huge dimensional Hamiltonians. The convergence property of the LOBPCG method strongly depends on preconditioners. We have already proposed a diagonal scaling preconditioner shifted by an approximate eigenvalue and confirmed its validity for some problems. However, the preconditioner does not have much effect on the Hamiltonian whose diagonal elements magnitude are small. In this presentation, we apply Neumann expansion preconditioner and examine its validity. Moreover, we propose the communication avoiding strategy for the preconditioner in consideration of physical properties of the model and evaluate its performance.

Susumu Yamada

Japan Atomic Energy Agency  
yamada.susumu@jaea.go.jp

Toshiyuki Imamura

RIKEN Advance Institute for Computational Science  
imamura.toshiyuki@riken.jp

Masahiko Machida

Japan Atomic Energy Agency  
machida.masahiko@jaea.go.jp

#### PP4

##### Interior Penalties for Summation-by-Parts Discretizations of Linear Second-Order Differential Equations

This work focuses on multidimensional summation-by-parts (SBP) discretizations of linear elliptic operators with variable coefficients. We consider a general SBP discretization with dense simultaneous approximation terms (SATs), which serve as interior penalties to enforce boundary conditions and inter-element coupling in a weak sense. Through the analysis of adjoint consistency and stability, we present several conditions on the SAT penalties. Based on these conditions, we generalize the modified scheme of Bassi and Rebay (BR2) and the symmetric interior penalty Galerkin (SIPG) method to SBP-SAT discretizations. The SIPG penalty has been related to BR2 using straightforward matrix analysis. Numerical experiments are carried out on unstructured grids with triangular elements to verify the theoretical results.

Jianfeng Yan, Jared Crean

Rensselaer Polytechnic Institute  
yanj4@rpi.edu, creanj@rpi.edu

Jason E. Hicken

Rensselaer Polytechnic Institute  
Assistant Professor  
hickej2@rpi.edu

#### PP4

##### A New 4-Point $C^2$ Non-Stationary Subdivision Scheme for Computer Geometric Design

The importance of computer geometric designing is rapidly increasing, as they play vital role in almost every field of life like computer applications, medical image processing, sci-

entific visualization, reverse engineering and robotics etc. For the goal of creating smooth curves or surfaces, subdivision techniques are the only pleasing methods these days and are easy to operate. So in this article, a new ternary 4-point subdivision scheme has been proposed using the space  $\tau = \{1, \sin(x), \cos(x), \sin(\alpha x), \cos(\alpha x)\}$  for some  $0 < \alpha < \frac{\pi}{3}$ . To find its convergence the theory of asymptotic equivalence has been used. It has been observed that the proposed scheme can generate the smooth curves of  $C^2$  continuity. The comparison of the proposed scheme with the existing schemes have been demonstrated using different examples, the graphical results show that the limit curves of the proposed scheme behave more pleasantly and can also generate the conic sections.

Muhammad Younis

University of the Punjab, Lahore Pakistan  
younis.pu@gmail.com

#### PP4

##### An Ale-Level-Set Method for Moving Boundary Problems

Many flow problems involve different types of moving boundaries. For example, as a particle floats at a fluid interface, we need to deal with both the rigid particle surface and the deformable fluid interface. In this work, we develop a hybrid method that tracks the rigid moving boundaries by an arbitrary Lagrangian-Eulerian (ALE) method and the deformable moving boundaries by a level set method. The particle-fluid equations are solved by a continuous Galerkin finite element method on a moving mesh. The level set equation is solved by a discontinuous Galerkin method. To keep the level set function a signed distance function, we solve a Hamilton-Jacobi equation as a conservation law by the discontinuous Galerkin method. We adopt Gmsh to generate the computational mesh and implement the computational code using the finite element library deal.II. Numerical results on floating particles will be presented. This research is supported by NSF DMS-1522604

Jiaqi Zhang

Department of Mathematics, Virginia Tech,  
Blacksburg, VA 24061  
zjiaqi@vt.edu

Pengtao Yue

Virginia Polytechnic Institute and State University  
ptyue@math.vt.edu

#### PP4

##### 1D Diffusion with Fractional Laplacian: Regularity and Numerical Methods

We consider a one dimensional nonlinear diffusion-reaction equation with fractional Laplacian. The regularity of the solution to this equation is characterized in weighted Sobolev spaces rather than in standard Sobolev spaces. The weighted space allows us to accommodate the boundary singularity of the solution and precisely describe the interior regularity. Moreover, we develop spectral Petrov-Galerkin methods for the equation and obtained optimal error estimates of numerical methods.

Zhongqiang Zhang

Worcester Polytechnic Institute



zzhang7@wpi.edu

#### PP4

##### Computational Methods to Study the Pattern Formation in Tissues

We present a computational model to study the pattern formations in tissues with stem cells and differentiated cells. Our model is derived based on the Generalized Onsager Principle, combining one energy dissipative system and several active factors, such as spontaneous polarity states, birth and death of cells, self-propelled motion of differentiated cells and ATP hydrolysis. Firstly, linear stability analysis is conducted to reveal the long-wave instability inherent in the neighborhood of the constant steady states. Secondly, 2D and 3D experiments simulate the process of various pattern formations in the tissues. To solve this complex model, we develop an efficient energy stable numerical scheme and implement it on GPU clusters for high-performance computing. Our model introduces a novel way to investigate the interplay of stem cell division, differentiated cell migration and other active factors from environments. Using our methodology we observe new spatial patterns with spontaneous polarity and study the relations between those new patterns with various active terms in the system.

Xueping Zhao, Qi Wang  
University of South Carolina  
zhaoxueping0802@gmail.com, wangq@mailbox.sc.edu

#### PP4

##### A Semismooth Newton Method for the Solution of a Thermomechanically Coupled Stokes Ice Sheet Model

Modeling the dynamics of polar ice sheets is critical for projection of future sea level rise. Yet, large uncertainties remain in the boundary conditions at the base of the ice sheet due to the lack of direct observations. In particular, the geothermal heat flux has a strong influence on the thermal state of the ice and hence plays a critical role in understanding the dynamics of the ice sheet through its effect on basal and internal ice temperatures. Targeting at inversion of the geothermal heat flux from surface velocity observations, we employ a semismooth Newton method for the solution of a three-dimensional steady-state thermomechanically coupled Stokes ice sheet model with variational inequality boundary conditions at the base, which account for ice either below or at the melting point.

Hongyu Zhu  
Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
zhuhongyu@ices.utexas.edu

Tobin Isaac  
The University of Chicago  
tisaac@ices.utexas.edu

Georg Stadler  
Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

Noemi Petra  
University of California, Merced  
npetra@ucmerced.edu

Thomas Hughes  
Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
hughes@ices.utexas.edu

Omar Ghattas  
The University of Texas at Austin  
omar@ices.utexas.edu

#### PP4

##### The Effect of Memory Layout on Batched BLAS Routines

It is becoming increasingly common in modern high-performance computing (HPC) to solve a large linear algebra problem by subdividing it into thousands of smaller problems which can be solved independently. To address the need for efficient libraries which allow users to perform batches of small BLAS operations in parallel, and to make efficient use of their hardware, many APIs have been investigated. While each of these new APIs has been designed to speed up specific applications on a target computing platform, the linear algebra community has recently begun standardization efforts to settle on a common batched BLAS interface. As part of the batched BLAS standardization working group, we analyze the most promising batched standard proposals and present some recommendations. The main difference between the batched BLAS proposals are how they store the thousands of small matrices. This work focuses its attention on three relevant memory layouts and demonstrates their impact on hardware, from traditional many-cores to accelerators such as GPUs and the new self-hosted Intel Xeon Phi processors, code named Knights Landing (KNL).

Mawussi Zounon  
University of Manchester  
mawussi.zounon@manchester.ac.uk

Jack J. Dongarra  
University of Tennessee, Oak Ridge National Laboratory,  
USA  
dongarra@icl.utk.edu

Nicholas J. Higham  
School of Mathematics  
The University of Manchester  
nick.higham@manchester.ac.uk

Sven Hammarling  
The University of Manchester  
sven.hammarling@btinternet.com

Samuel Relton  
University of Manchester, UK  
samuel.relton@manchester.ac.uk

Pedro Valero-Lara  
University of Manchester  
pedro.valero-lara@manchester.ac.uk

#### PP5

##### Improving Ab Initio Gene Prediction in Prokaryotic Genomes

Although computational prediction of prokaryotic genes is sometimes considered a solved problem, the rate of prediction errors of even state-of-the-art tools is not negli-

ble. In particular, short genes and gene initiation sites are often cited as difficult to predict, and atypical genes, i.e. those whose GC composition differs significantly from the remaining genes in the genome, are prone to significant errors. Here, we describe a new gene-finding algorithm called GeneMarkS-2 (the successor to GeneMarkS) whereby, given an unannotated genome, it attempts to predict the start and end coordinates of individual genes. It starts off by employing local heuristic models with parameters adjusted to local GC content. This sets of an iterative procedure, where each iteration re-estimates parameters by determining the features of transcription and translation mechanisms using a Gibbs Sampling-based probabilistic sequence alignment, as well as parameters estimations for genetic information encoded in the coding and non-coding regions. The algorithm also controls the balance between sensitivity and specificity by adjusting hyperparameters and thresholds of the underlying generalized Hidden Markov Model. With genes of atypical codon usage composition being a particular focus of the algorithm, our tests show that the accuracy of GeneMarkS-2 is favorably compared with other state-of-the-art gene prediction tools.

Karl Gemayel  
Computational Science and Engineering  
Georgia Institute of Technology  
karl@gatech.edu

Shiyuyun Tang  
School of Biology  
Georgia Institute of Technology  
tangshi06@gmail.com

Alexandre Lomsadze  
Department of Biomedical Engineering  
Georgia Institute of Technology  
alexandre.lomsadze@bme.gatech.edu

Mark Borodovsky  
Georgia Institute of Technology  
Wallace H. Coulter Department of Biomedical  
Engineering  
borodovsky@gatech.edu

## PP5

### Ridge Approximation Using Variable Projection

A ridge approximation seeks to mimic a function of many variables by another function that only depends on a few linear combinations of the original variables; i.e.,  $f(\mathbf{x}) \approx g(\mathbf{U}^T \mathbf{x})$  where  $\mathbf{U}$  is a rectangular matrix with orthonormal columns. One application of ridge approximations are used as a response surface model, providing an inexpensive surrogate for a complex model with many parameters. Here, we present a new algorithm for constructing ridge approximations given input-output pairs  $\{\mathbf{x}_i, f(\mathbf{x}_i)\}_i$  where  $g$  is a polynomial. Using Variable Projection [Golub and Pereyra; 1973] the polynomial coefficients are implicitly removed leaving an optimization over  $\mathbf{U}$  in the Grassmann manifold. The optimization of  $\mathbf{U}$  is accomplished following [Edelman, Arias, and Smith; 1998] with a Gauss-Newton approximation of the Hessian. This new algorithm offers substantially improved performance over existing approaches that alternate between minimizing over  $\mathbf{U}$  and fitting the polynomial coefficients of  $g$  [Constantine, Eftekhari, Ward; to appear]. For example, wall clock time to fit a polynomial of degree five on a five dimensional subspace has decreased from over an hour to approximately

one second. This permits higher order ridge approximations and repeated optimizations to avoid local minima, increasing the value of a ridge approximation as a response surface model.

Jeffrey M. Hokanson, Paul Constantine  
Colorado School of Mines  
hokanson@mines.edu, paul.constantine@mines.edu

## PP5

### Simulation of Capacitively Coupled Plasmas using a High Performance Parallelized Particle-in-Cell Simulation

Nowadays, the research interests in plasma-assisted materials processing attract lots of attention on plasma simulations to reduce time and cost for the development of technology. Also, the research trend to apply plasmas to bio-medicine also needs simulation research because the experimental diagnostics are difficult for the small size and high gas pressure devices. Recently, the simulation speed was boosted up by parallel computing skill, and thus it is possible to utilize time-consuming particle-in-cell (PIC) simulations for the investigation of low-temperature plasmas for materials processing instead of conventionally utilized fluid simulations. Especially, nonlinear, stochastic, and transient effects can be accurately analyzed with a PIC simulation instead of fluid simulations. In this study, the method of high-performance PIC simulation of low-temperature plasmas are presented for plasma-aided microelectronics processing like etching and deposition. Also, the non-linear heating mode transition was simulated for the variation of driving frequency in a capacitively coupled plasma. Moreover, the techniques for the parallelization of a PIC code with graphics processing units (GPUs) and many integrated cores (MICs) are also explained.

Young Hyun Jo, Chang-Ho Kim, Yoon Ho Lee, Geon Woo Park, Jung Yeol Lee, Hae June Lee  
Electric and Electronic Engineering  
Pusan National University  
bcjo17@gmail.com, kimch1060@naver.com,  
ioer14@naver.com, gunto01@naver.com, lee-  
jungyel@pusan.ac.kr, haejune@pusan.ac.kr

## PP5

### A Tunably-Accurate Spectral Method with Linear Complexity for Multi-Term Fractional Differential Equations on the Half Line

We present a new tunably-accurate Laguerre Petrov-Galerkin spectral method for solving linear multi-term fractional initial value problems with derivative orders at most one and constant coefficients on the half line. Our method results in a matrix equation of special structure which can be solved in  $\mathcal{O}(N \log N)$  operations. We also take advantage of recurrence relations for the generalized associated Laguerre functions (GALFs) in order to derive explicit expressions for the entries of the stiffness and mass matrices, which can be factored into the product of a diagonal matrix and a lower-triangular Toeplitz matrix. The resulting spectral method is efficient for solving multi-term fractional differential equations with arbitrarily many terms. We apply this method to a distributed order differential equation, which is approximated by linear multi-term equations through the Gauss-Legendre quadrature rule. We provide numerical examples demonstrating the spectral con-

vergence and linear complexity of the method.

Anna Lischke  
Brown University  
anna\_lischke@brown.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

Mohsen Zayernouri  
Michigan State University  
zayern@egr.msu.edu

## PP5

### Likelihood Approximation with Hierarchical Matrices for Large Spatial Datasets

In this work we use available measurements to estimate unknown hyper-parameters (variance, smoothness parameter and covariance length) of the covariance function. We do it by maximizing the joint log-likelihood function. To overcome cubic complexity in the linear algebra, we approximate the discretized covariance function in the hierarchical (H-) matrix format. The H-matrix format has a log linear computational cost and storage  $O(kn \log n)$ , where the rank  $k$  is a small integer and  $n$  is the number of locations. Previous results show that the H-matrix technique is very robust for approximating the matrix itself, its inverse, its Cholesky decomposition and the Schur complement (conditional covariance). Other motivating factors for applying H-matrix techniques are: [i] The class of H-matrices is more general than other classes of matrices; [ii] The H-matrix technique allows us to compute not only matrix-vector products, but also quite general classes of functions, such as inverse, LU decomposition, determinant, resolvents, and many others; [iii] H-matrices are relatively new, but already a well studied technique; [iv] Approximation accuracy is fully controlled by the rank. Full rank ( $k=n$ ) gives an exact representation; [v] Keep advantages after matrix operations, such as computing conditional covariance matrix (Schur complement can be again approximated in the H-matrix format).

Alexander Litvinenko  
SRI-UQ and ECRC Centers, KAUST  
alexander.litvinenko@kaust.edu.sa

Ying Sun, Marc Genton  
KAUST  
Saudi Arabia  
ying.sun@kaust.edu.sa, marc.genton@kaust.edu.sa

David E. Keyes  
KAUST  
david.keyes@kaust.edu.sa

## PP5

### Evaluation of An Improved Numerical Technique for Solving the Hypersonic Boundary Layer/Shockwave Interaction Problem

This work demonstrates the robustness and effectiveness

of the scheme introduced by [Elamin Gafar, The Integral-differential Scheme (IDS): A New CFD Solver for the System of the Navier-Stokes Equations with Applications, 2008] referred to as the *Integro-Differential Scheme* (IDS). IDS merges the traditional differential and integral representation of the conservation equations. This scheme is built on the premise that the numerical control volume, cell and nodes are hardwired and, the temporal fluxes are obtained through a consistent averaging procedure. The problem of interest to this study is the hypersonic boundary layer/shockwave interaction problem. The Mach number was 5.0, the Reynolds number based on the length of the plate ( $Re_{\infty L}$ ) was in the order of  $10^6$ . The numerical solution is obtained by solving the two-dimensional compressible Navier-Stokes equations (NSE) using an explicit formulation. Parametric studies indicate that the scheme reproduce accurately the flow features for a wide range of problems and conditions. Unlike other schemes, this method is independent of the mathematical classification of the flow, allowing us to implement a wide range of boundary and initial conditions without tweaking the scheme. This feature highlight the potential of this scheme to solve a variety of fluid dynamics problems with the same numerical method.

Julio C. Mendez, David Doodoo-Amoo, Frederick Ferguson  
Department of Mechanical Engineering  
North Carolina A&T State University  
jcmendez@aggies.ncat.edu, dndoodoo@aggies.ncat.edu, fferguso@ncat.edu

## PP5

### Experiences, Optimizations, and Future Directions with Petsc on the 2nd Generation ("Knights Landing") Intel Xeon Phi Processor

As the high-performance computing community pushes towards the exascale horizon, power and heat considerations have driven the increasing importance and prevalence of fine-grained parallelism in new computer architectures. The Intel Many Integrated Core (MIC) architecture utilizes many small, low power x86 cores within a single Intel Xeon Phi processor to achieve power-efficient high performance. The second generation ("Knights Landing") Intel Xeon Phi Processor offers a very high degree of parallelism in a single, bootable CPU, with up to 72 cores (288 hardware threads), each of which has two vector processing units supporting 512-bit SIMD operations. It also incorporates very high-bandwidth on-package memory, which facilitates good utilization of those cores when working with assembled sparse matrices. The high degree of fine-grained parallelism and more complicated memory hierarchy considerations of such "manycore" processors present some challenges to existing scientific software. Here, we consider how the widely-used Portable, Extensible Toolkit for Scientific Computation (PETSc) can best take advantage of such architectures. We will discuss some key architectural features of the Intel Xeon Phi, relate experiences drawn from a variety of application areas using PETSc on it, discuss ongoing code optimization and algorithmic development work, and outline possible ways that PETSc should evolve to best utilize this and future manycore architectures.

Richard T. Mills

Intel Corporation  
richardtmills@gmail.com

Mark Adams  
Lawrence Berkeley Laboratory  
mfadams@lbl.gov

Jed Brown  
Mathematics and Computer Science Division  
Argonne National Laboratory and CU Boulder  
jed@jedbrown.org

Maurice Fabien  
Rice University  
fabien@rice.edu

Tobin Isaac  
University of Chicago  
tisaac@uchicago.edu

Matthew Knepley  
Rice University  
knepley@rice.edu

Karl Rupp  
Institute for Microelectronics  
Vienna University of Technology  
me@karlrupp.net

Barry F. Smith  
Argonne National Lab  
MCS Division  
bsmith@mcs.anl.gov

Hong Zhang  
Argonne National Laboratory  
hongzhang@anl.gov

## PP5

### On An Inverse Diffusion Coefficient Problem Arising in Geochronology

We consider a problem arising in geochronology, a branch of geology which deals with the dating of rock formations and geological events. In particular, we investigate the reconstruction of temperature histories of rocks by solving a time-dependent inverse diffusion coefficient problem for parabolic partial differential equations with an integral overspecification. We show the existence and uniqueness of classical solutions using fixed point theory. Our numerical algorithm employs the implicit Euler method with variable time step and a finite element discretization in space. We present some numerical results including the errors and convergence rates to illustrate the accuracy of the proposed algorithm.

Sedar Ngoma, Dmitry Glotov  
Auburn University  
nzb0015@auburn.edu, dvg0001@auburn.edu

A. J. Meir  
Department of Mathematics, Southern Methodist University

ajmeir@mail.smu.edu

WILLIS E. HAMES  
Department of Geosciences, Auburn University  
hameswe@auburn.edu

## PP5

### An Effect of Turbulence on Zonal Jet Flows in Forced 2D and Quasi-Geostrophic Shallow Water Models on a Beta Plane

In randomly forced two-dimensional Navier-Stokes turbulence on a rotating sphere, it is known that a multiple zonal-band structure, i.e. a structure with alternating eastward and westward jets, emerges in the course of time development. The multiple zonal-band structure then experiences intermittent mergers and disappearances of zonal jets, and a structure with only a few large-scale zonal jets is realised as an asymptotic state (Obuse et al., 2010). With the view of understanding the long-time behaviour of the zonal jets, especially the merging and disappearing processes of the zonal jets, Obuse et al. (2011) considered large-scale zonal flows superposed upon a homogeneous zonal flow and a small-scale sinusoidal transversal flow on a beta plane, which was originally introduced by Manfroi and Young (1999), then discussed the merging and disappearing processes of zonal jets by investigating the linear stability of analytical steady isolated zonal jet solutions. In this talk, we extend the Manfroi-Young model by taking account of the spatial variation of the disturbance in the zonal direction, and the surface variation of fluid layer, in order to make the model a little more realistic. The linear stability analysis of analytical steady isolated zonal jet solutions suggests the instability of zonal jets due to back-ground turbulence effect is widely common on beta plane.

#### Kiori Obuse

Graduate school of environmental and life science  
Okayama university  
obuse@okayama-u.ac.jp

Shin-Ichi Takehiro, Michio Yamada  
Research Institute for Mathematical Sciences  
Kyoto University  
takepiro@kurims.kyoto-u.ac.jp, yamada@kurims.kyoto-u.ac.jp

## PP5

### A Computational Model for Sound Source Recognition

Hearing is an important part of normal human interaction, yet we understand surprisingly little about how our brains make sense of sound. The ability of a normal human listener to recognize objects in the environment from only the sounds they produce is extraordinarily robust with regard to characteristics of the acoustic environment and of other competing sound sources. Robust listening requires extensive contextual knowledge, but the potential contribution of sound-source recognition process has largely been neglected by researchers. As a stepping stone for an artificial listener, a computational model was developed to recognize

isolated sound sources. The recognition process includes statistical inference and used a fingerprinting method in the sound source identification. This system is computer-based prototype that can lead to surveillance systems, music composing and ultimately to an artificial listener. This computational model was built on MATLAB platform and Graphical user interface was developed by using GUIDE tool in MATLAB. Since this system was born in a PC it can be easily run in any common computer which has got a regular audio input facility. This project contains a brief outline about the problems faced and how the solutions were achieved.

Chathurangi H. Pathiravasan

Department of Mathematics

Southern Illinois University Carbondale Illinois  
chathurangi@siu.edu

Dinuka Gallaba

Department of Physics

Southern Illinois University Carbondale Illinois  
dinuka4@siu.edu

## PP5

### Applicability Analysis of Validation Evidence for Biomedical Computational Models

Computational modelling has the potential to revolutionise medicine the way that it transformed physics and engineering, but despite decades of work there has been limited progress to successfully translate research to patient care. One major difficulty with biomedical models is often an inability to perform validation in a setting that closely resembles how the model will be used. For example, for a model that makes in vivo clinical predictions, ‘direct validation’ of predictions may be impossible for ethical, technological or financial reasons. The unavoidable difference between how a biomedical model may be validated versus how it is used, can lead to difficulty in rigorously assessing validation evidence, and lack of trust in biomedical models. While the engineering literature provides some guidance regarding ‘applicability’ - whether validation evidence supports the model for a specific context of use (COU) - current methods are motivated by engineering problems and not very relevant to biomedical problems. We propose a novel framework for performing applicability analysis, that is, the systematic assessment of the applicability of a computational model to a COU given the validation evidence. The framework provides a step-by-step method for breaking down the broad question of applicability into a series of tractable questions. The proposed framework is relevant to a wide range of biomedical models and models from other disciplines, and a wide range of underlying physics.

Pras Pathmanathan, Richard Gray, Leonardo Angelone,  
Tina Morrison

U.S. Food and Drug Administration

pras.pathmanathan@fda.hhs.gov,

richard.gray@fda.hhs.gov, leonardo.angelone@fda.hhs.gov,

tina.morrison@fda.hhs.gov

## PP5

### Kronecker Product Preconditioners for Very High

### Order Discontinuous Galerkin Methods

Although the DG method generalizes to arbitrary orders, there are several challenges preventing the use of very high degree polynomial bases. For explicit methods, the CFL condition requires that the time step satisfy approximately  $\Delta t \leq Ch/p^2$ , where  $h$  is the element size, and  $p$  the polynomial degree. For implicit methods, the number of DOFs per element scales as  $p^d$  in  $d$  dimensions. The resulting linear system can be considered as a sparse matrix with dense  $p^d \times p^d$  blocks, requiring  $p^{3d}$  operations per linear solve using dense linear algebra, rendering the problem intractable for very large  $p$ . If we can approximate these  $p^d \times p^d$  blocks with sums of Kronecker products of smaller matrices, it would allow for asymptotically more computationally efficient linear algebra. We describe an implicit DG method with a tensor-product structure whose computational cost per DOF scales linearly with the degree  $p$ . This method requires a tensor-product basis on quadrilateral or hexahedral meshes. The matrix corresponding to the linear system is not explicitly constructed. Fast matrix-vector products are performed as the kernel of the GMRES solver. Such systems are often preconditioned using the block Jacobi preconditioner. To avoid inverting the diagonal blocks and thus incurring the above-mentioned  $\mathcal{O}(p^{3d})$  operations, we make use the Kronecker product SVD to approximate this block by a sum of lower-dimensional tensor products.

Will Pazner

Division of Applied Mathematics

Brown University

will\_pazner@brown.edu

Per-Olof Persson

Dept. of Mathematics

University of California, Berkeley

persson@berkeley.edu

## PP5

### Uncertainties in An Inverse Source Problem for Pde-Constrained Optimization with Inequality Constraints

The research focuses on solving an inverse source problem by applying PDE-Constrained Optimization (PCOpt) with inequality constraints. Solving an optimization problem with inequality constraints in optimization is mostly straight forward but not in PCOpt due to dependency of the problems on discretization and uncertainty issues. Our preliminary research found uncertainties as follows: 1) the solver may or may not need regularization, 2) the globalization may be needed, 3) preconditioners are needed, and 4) further investigation in parallelization is needed, which involves the effect of the discretization in the PDEs to the PDE-constrained optimization. This is due to the fact that different grids may result different solutions.

Widodo Samyono

Jarvis Christian College

Hawkins, Texas

wsamyono@jarvis.edu

## PP5

### Optimizing and Vectorizing Rank Reduction Techniques in Quantum Chemistry

Quantum Mechanics (QM) provides a fundamental means for calculating properties of atoms and molecules. These calculations yield the most accurate chemical predictions of societal interest, but QM methods are computationally expensive. Therefore, an urgent CSE application involves improving the scaling of QM codes. This study examines the computational cost when using essential 4-center electron repulsion integrals, which are a major bottleneck for QM codes due to storage costs. By utilizing Density Fitting, a rank reduction technique, these 4-center integrals can be reduced to a 3-center form. Two steps must take place in constructing these 3-center tensors, in interchangeable order. One step is the tensor's contraction with the inverse Coulomb metric, taking the form:  $b_{\sigma\nu}^P, C^{PQ} \rightarrow b_{\sigma\nu}^Q$  and costing  $O(N^4)$ . The other step requires the tensor to be transformed into a new basis:  $c_{p\sigma}, b_{\sigma\nu}^Q, c_{\nu q} \rightarrow b_{pq}^Q$ , costing  $O(N^4)$ . This study reveals speedups via optimized workflows involving these two operations. We found the optimized order in which they occur is context dependent; different workflows optimize Iterative and Perturbative quantum methods. Moreover, an enhanced Schwarz screening was developed which provides superior vectorization in the transformation phase. We implemented and tested these optimizations using the open-source electronic structure package, Psi4 ([github.com/psi4/psi4](https://github.com/psi4/psi4)).

Matthew C. Schieber

Georgia Institute of Technology  
School of Computational Science and Engineering  
mschieber3@gatech.edu

Daniel Smith, David Sherrill  
Georgia Institute of Technology  
dgasmith@gatech.edu, sherrill@gatech.edu

## PP5

### Nodal Integral Method for Complex Geometries Using Higher Order Elements

Fluid flow over complex geometries in computational fluid dynamics (CFD) can be predicted with very fine grids while nodal integral method (NIM) which is developed for solving neutron-transport equation, yields better approximation in reasonable coarse meshes. In this approach one-onto mapping is done using Lagrange interpolation functions to transform higher order quadrilateral elements to square elements. This transformation is utilized in NIM for complex geometries. Higher order shape functions conserve the nature of underlying geometry of the problem to greater extent. This approximation leads to transform the complete partial differential equation (PDE) from global cartesian domain  $(x, y)$  to local domain  $(\xi, \eta)$  called cells or nodes. Then applying transverse integration procedure (TIP) over the transformed equation to generate the cell analytical solution over the cell domain. These cell analytical solutions are used to develop further numerical scheme. The proposed scheme predict results with coarser grids.

Diffusion and Convection-diffusion equations in polar cavity with Neumann as well as Dirichlet boundary conditions are chosen for verification and comparison of results. Results are in good match with the analytical solutions even for quite coarse grids with minimal error compared to other Numerical schemes like FDM or FVM.

Rishabh P. Sharma

Thapar University, Patiala, Punjab  
Mechanical Engineering Department  
rsharma1\_me15@thapar.edu

Neeraj Kumar  
Mechanical Engineering Department  
Thapar University, Patiala, Punjab, India  
neerajkumar@thapar.edu

## PP5

### A Particle-Based Numerical Method for Solving Vlasov Models in Plasma Simulations

Plasmas are ionized gases that appear in a wide range of applications including astrophysics and space physics, as well as in laboratory settings such as in magnetically confined fusion. Depending on the application of interest, plasmas can be modeled through a variety of models ranging from the Liouville's equation to kinetic models such as the Vlasov equation, and down to fluid models such as the magnetohydrodynamics(MHD) equations. In this work we are interested in developing high-order methods for the kinetic Vlasov equation using so-called particle methods that represent the distribution function as a sum of delta function (aka, the particles). The dominant approach in plasma physics is to evolve these particles using a particle-in-cell(PIC) method. In this work we are interested in solving Vlasov-Poisson and Vlasov-Darwin models by computing through particle interactions appropriate integrals of the particle distribution against the Greens function of the Poisson equation. The numerical method relies on treecode algorithm to solve the Poisson equation for interacting charged particles. The treecode replaces particle-particle interactions by particle-cluster interactions which are evaluated by Taylor expansions. Our goal is to assess the particle-based numerical method in plasma dynamics in comparison with PIC algorithm. We apply the resulting method to several standard plasma test cases, including the virtual cathode, plasma sheath and ion-acoustic wave problems.

Boqian Shen

Iowa State University  
boshen@iastate.edu

James A. Rossmannith  
Iowa State University  
Department of Mathematics  
rossmani@iastate.edu

## PP5

### Positivity-Preserving Limiters for the Piecewise- $P_N$ Equations

Transport problems are difficult to numerically simulate

because they need rich phase space to resolve the kinetic distribution. Since the  $P_N$  equations are rotationally invariant and converge in  $L_2$  sense to the solution of transport equation as  $N \rightarrow \infty$ , the  $P_N$  equations are used to solve the transport problems. However, in the multidimensional setting, the  $P_N$  equations often produce negative particle concentrations, which is physically incorrect. In the literature there exist many approaches to preserve positivity of the particle concentration. In our approach we introduce the piecewise- $P_N$  equations to solve the kinetic transport equation and develop for these equations a modified version of the Zhang-Shu positivity-preserving limiters to achieve physically meaningful positive particle concentrations.

Minwoo Shin

Iowa State University  
mws@iastate.edu

James A. Rossmanith  
Iowa State University  
Department of Mathematics  
rossmani@iastate.edu

## PP5

### Conservative Exponential Integrators for Nonlinear Conservation Laws

Current Climate models use explicit schemes in time over implicit schemes. This provides a reduction in computational cost but severely restricts the size of the time step that can be taken in a simulation. Recently, exponential integrators have been developed that provide explicit-type schemes, while relaxing time step restrictions. This is an attractive property for climate modeling simulations that must simulate to a time horizon of the order of millennia, while being restricted to time steps of the order of minutes to hours. This problem is further complicated by the fact that the time step size restriction is dependent on the spatial resolution as well. As the simulation grid is made finer, to resolve smaller effects, the allowable time step size becomes even smaller. Exponential integrators provide an explicit-type scheme while avoiding the time step size restriction. Although it seems clear that exponential integrators provide an advantage over explicit methods, the conservation properties of the exponential integrators are still not clear. In this work, we aim to show through analysis and computations that a family of exponential integrators possess conservation properties when applied to non-linear conservation laws and the shallow water equations.

Chad Sockwell

Florida State University  
Department of Scientific Computing  
kcs12j@my.fsu.edu

## PP5

### Computation of Electromagnetic Fields Due to Dipoles in Two and Three Layered Media in $R^3$

Ever wondered how to determine the magnetic and electric fields due to dipoles placed in a layered material like air or sand or a combination of the aforementioned? This re-

search derives a concise and explicit formula for the dyadic Greens function, which represents the electric and magnetic fields due to a dipole source placed in a layered medium(s). The resulting formula is compact, more efficient and can be extended to the calculation of these effects for multiple dipole sources placed in numerous mediums. To achieve this feature, the electric and magnetic fields in the spectral domain in each layer using Fresnel reflection and transmission coefficients. Each component of electric field in the spectral domain constitutes the spectral Green's function in layered media. The Green's function in the spatial domain is recovered by taking Sommerfeld integrals on each component in the spectral domain. Using Bessel identities, the number of Sommerfeld integrals is reduced, resulting in a much simpler and efficient formula for numerical implementation. This poster presentation will reiterate the preexisting formula for achieving this feat, showcase the process for deriving the simplified model and include numerical results to compare and contrast both approaches.

Kennedy Udechukwu

Francis College of Engineering  
University of Massachusetts Lowell  
kennedy\_udechukwu@student.uml.edu

Min Hyung Cho  
Department of Mathematical Sciences  
University of Massachusetts Lowell  
minhyung\_cho@uml.edu

## PP5

### Evaluating Production Engineering Application Performance on the Intel Knights Landing Many Core Processor

The deployment of Intel Knights Landing Many-Core Processors in Trinity – the NNSA/ASC's first Advanced Technology System (ATS) – represented a significant change in production application requirements. Instead of being able to assume an MPI-only application development model would provide continued performance on the ASC's leadership computing platforms, developers have begun the task of investing in on-node parallelism with an increased focus on vectorization. While a great deal of work has already been done investing in application development for the Trinity platform, the task of modifying the entire spectrum of applications for future platforms, of which the Knights Landing processor is one, is a significant multi-year effort. In this poster we will present our initial baselines of production engineering application performance with a comparison to MPI-only execution on the more familiar Xeon Haswell processor. The technologies introduced in Knights Landing include high-bandwidth memory, dual wide-vector units and considerably more processor cores. Our results will include an analysis of performance on a collection of these and show which combinations provide for the highest performance. Finally, we will summarize any lessons learned for the broader application development community.

Simon D. Hammond  
Scalable Computer Architectures  
Sandia National Laboratories

sdhammo@sandia.gov

Courtenay T. Vaughan, Dennis Dinge, Paul Lin, Robert Benner, Christian Trott, Doug Pase, Jeanine Cook, Robert J. Hoekstra  
Sandia National Laboratories  
ctvaugh@sandia.gov, dcdinge@sandia.gov,  
ptlin@sandia.gov, rebenne@sandia.gov,  
crtrott@sandia.gov, dmpase@sandia.gov, jea-  
cook@sandia.gov, rjhoeks@sandia.gov

## PP5

### Sharp Convergence Rates of Numerical Solutions of Stochastic Fractional Equations Driven by White Noise

The stability and convergence of numerical approximation of solutions of the stochastic time-fractional PDE

$$\partial_t^\alpha u - \Delta u = f + \partial_t^{\alpha-1} \dot{W}$$

is considered, where  $\partial_t^\alpha u$  denotes the Caputo fractional derivative and  $\dot{W}$  space-time white noise. Such problems arise naturally by considering the heat equation in a material with thermal memory, subject to white noise. For the above model, both the time-fractional derivative and the stochastic process result in low regularity of the solution. Hence, the numerical approximation of such problems and the corresponding numerical analysis are very challenging. In this work, the stochastic time-fractional equation is discretized by a backward Euler convolution quadrature in time for which the sharp error estimate

$$\mathbb{E}\|u(t_n) - u_n\|_{L^2}^2 = O(\tau^{1-\alpha d/2})$$

is established for  $\alpha \in (0, 2/d)$ , where  $d$  denotes the spatial dimension,  $u_n$  the approximate solution at the  $n$ th time step, and  $\mathbb{E}$  the expectation operator. The results indicate optimal convergence rates of numerical solutions for both stochastic subdiffusion and diffusion-wave problems in one spatial dimension. Numerical results are presented to illustrate the theoretical analysis.

Max Gunzburger  
Department of Scientific Computing  
Florida State University  
gunzburg@fsu.edu

Buyang Li  
Department of Applied Mathematics  
The Hong Kong Polytechnic University  
buyang.li@polyu.edu.hk

Jilu Wang  
Department of Scientific Computing  
Florida State University  
jwang13@fsu.edu

## PP5

### Numerical Interface Treatment for Wave Propagation Problems

In numerical simulations of wave phenomena, there are many cases when non-uniform meshes are desirable. For

example, a finer mesh is required in the computational domain where the solution varies rapidly or the geometry is complex. With a finite difference method non-uniform meshes could lead to improved computational efficiency, but often bring the difficulty of numerical non-conforming interface treatment with hanging nodes. In this presentation, we show our effort in improving the accuracy property of such schemes. We in particular consider a finite difference method satisfying a summation-by-parts property. We are also interested in comparing the finite difference method with the discontinuous method, which is known to be flexible with resolving boundary features.

Siyang Wang  
Uppsala University  
siyang.wang@it.uu.se

## PP5

### Analysis of a Preconditioner for Matrices with Block Low-Rank Structure

We present a framework to analyze a preconditioner for symmetric positive definite (s.p.d.) matrices with low-rank off-diagonal blocks. For a general s.p.d. matrix and a partitioning of it, we replace its off-diagonal blocks by low-rank approximations and then apply diagonal compensation to obtain an s.p.d. structured preconditioner. We show how the approximation error of the off-diagonal blocks and the size of the diagonal compensation affect the spectrum of the preconditioned matrix. The result provides a guideline for approximating the off-diagonal blocks and choosing the size of the diagonal compensation for the preconditioner.

Xin Xing  
School of Mathematics  
Georgia Institute of Technology  
xxing33@gatech.edu

Edmond Chow  
School of Computational Science and Engineering  
Georgia Institute of Technology  
echow@cc.gatech.edu

## PP5

### Application of Optimal Transport and the Quadratic Wasserstein Metric to Full-Waveform Inversion

Full-waveform inversion (FWI) is a large-scale data-fitting procedure to achieve high-resolution velocity image. Recent developments in high-performance computing make FWI feasible today. Conventional FWI using the least-squares norm ( $L_2$ ) as a misfit function is known to suffer from cycle skipping. We proposed the quadratic Wasserstein metric ( $W_2$ ) as a new misfit function for FWI. It has been proved to have many ideal properties with regards to convexity and insensitivity to noise. We match the observed and predicted data through an optimal map, and then we design a misfit based on the measured map. Unlike the  $L_2$  norm,  $W_2$  measures not only amplitude differences, but also global phase shifts, which helps to avoid cycle skipping issues. We propose two ways of using the  $W_2$  metric in FWI: trace-by-trace comparison and global



comparison. In this poster, we will demonstrate the applications of this new misfit function on three large-scale 2D models: Camembert, Marmousi, and the 2004 BP model with parallel computation. The inversion follows the usual approach to FWI as a PDE-constrained optimization with l-BFGS method. Numerical results show the effectiveness of  $W_2$  for eliminating cycle skipping issue that are prevalent with the traditional  $L_2$  norm. Both mathematical theory and numerical examples demonstrate that the quadratic Wasserstein metric is a good candidate for a misfit function in seismic imaging and inversion.

Yunan Yang

Department of Mathematics  
The University of Texas at Austin  
yunanyang@math.utexas.edu

Bjorn Engquist

Department of Mathematics and ICES, UT Austin  
engquist@ices.utexas.edu

Junzhe Sun

Formerly at BEG, UT-Austin  
Currently at ExxonMobil Upstream Research Company  
sunjzhe@gmail.com

Brittany Froese

New Jersey Institute of Technology  
Department of Mathematical Sciences  
bdfroese@njit.edu

## PP5

### Positivity-Preserving High Order Discontinuous Galerkin Schemes for Compressible Navier-Stokes

For gas dynamics equations such as compressible Euler and Navier-Stokes equations, preserving the positivity of density and pressure without losing conservation is crucial to stabilize the numerical computation. The L1-stability of mass and energy can be achieved by enforcing the positivity of density and pressure during the time evolution. However, high order schemes such as DG methods do not preserve the positivity. It is difficult to enforce the positivity without destroying the high order accuracy and the local conservation in an efficient manner for time-dependent gas dynamics equations. For compressible Euler equations, a weak positivity property holds for any high order finite volume type schemes including DG methods, which was used to design a simple positivity-preserving limiter for high order DG schemes in Zhang and Shu, JCP 2010. Generalizations to compressible Navier-Stokes equations are however nontrivial. We show that the weak positivity property still holds for DG method solving compressible Navier-Stokes equations if a proper penalty term is used in the scheme. This allows us to obtain the first high order positivity-preserving schemes for compressible Navier-Stokes equations.

Xiangxiong Zhang  
Purdue University

zhan1966@purdue.ed

## PP101

### Minisymposium: Semi-Lagrangian Solution of the Gyrokinetic Vlasov Equation: Field-Aligned Interpolation And Splitting in Complex Geometry

State-of-the-art global turbulence simulations of magnetic fusion devices, based on the solution of the 5D gyrokinetic Vlasov equation, are computationally intensive because the thermal ion Larmor radius must be resolved. The computational burden may be reduced by aligning one coordinate with the local magnetic field line, along which the gradients are known to be small. Unfortunately this methodology poses restrictions on the poloidal mesh, and cannot easily handle complex magnetic field configurations found in diverted Tokamaks and Stellarators. An alternative and more flexible approach was developed in [Hariri and Ottaviani, 'A flux-coordinate independent field-aligned approach to plasma turbulence simulations', CPC 184(11):2419, 2013], where local field-aligned differentiation/interpolation was performed between adjacent poloidal planes. Such a method was adapted to the semi-Lagrangian context and combined with dimensional splitting in [Latu, Mehrenberger, Güçlü, Ottaviani, and Sonnendrücker, 'Field-aligned interpolation for semi-Lagrangian gyrokinetic simulations', hal.inria.fr/hal-01315889], and we now extend it to general flux-conformal curvilinear coordinates. We describe here the general mathematical formulation, the parallel numerical framework in Selalib, and the details of our field-aligned interpolation algorithm. Our code is verified with a linear dispersion analysis for the ITG instability in screw-pinch configuration.

Edoardo Zoni

Max-Planck-Institut für Plasmaphysik, NMPP division,  
and  
Technische Universität München, Department of  
Mathematics  
edoardo.zoni@ipp.mpg.de

Yaman Güçlü

Department of Mathematics  
Michigan State University  
yaman.guclu@ipp.mpg.de

Michel Mehrenberger

Institut de Recherche Mathématique Avancée (IRMA),  
University of Strasbourg  
mehrenbe@math.unistra.fr

Eric Sonnendrücker

Max-Planck-Institut für Plasmaphysik, NMPP division,  
and  
Technische Universität München, Department of  
Mathematics  
eric.sonnendruecker@ipp.mpg.de

## PP101

### Minisymposium: Eulerian Algorithms for the

### Discretization of Plasma Kinetic Equations

While fluid models are common tools in the study of plasmas, many of these systems, whether in astrophysics or the lab, are only weakly collisional and far from equilibrium, making them more accurately described by kinetic equations. Kinetic equations can be computationally demanding due to the need to solve for the distribution function of the particles in a higher dimensional phase space, with position and velocity coordinates. Despite this challenge, the motivation for solving the plasma kinetic equation is large as there remains a vast array of questions concerning collisionless dynamics in real plasma systems. Here we present algorithms in an Eulerian framework for the discretization of the plasma kinetic equation, using a discontinuous Galerkin finite element method due to its arithmetic intensity and parallelizability. Scaling and performance of the algorithm are discussed, especially in comparison to the traditional particle-in-cell method. We stress in particular the use of novel basis sets in the finite element expansion, which when combined with sparse grids for the quadrature routines, heavily reduce the computational cost of solving a five or six dimensional system on structured grids. Other features of the algorithm, including a locally implicit time-stepping scheme which allows us to avoid stringent velocity space time step constraints, and hyperdiffusion for avoiding numerical recurrence, are also highlighted. Benchmarks of the algorithm are presented as well.

James Juno, Jason TenBarge  
IREAP

University of Maryland  
jjuno@umd.edu, jtenbarg@umd.edu

Ammar Hakim  
Princeton Plasma Physics Laboratory  
ahakim@pppl.gov

William Dorland  
University of Maryland  
bdorland@umd.edu

### PP101

#### Minisymposium: New Hybrid Method for the Vlasov-Maxwell System: Coupling of Spectral and Pic Methods

The Vlasov-Maxwell (VM) system is a system of nonlinear partial differential equations which describes the time evolution of the plasma distribution function in a six dimensional phase space driven by self-consistent electromagnetic fields. A new hybrid method based on coupling spectral (SM) and particle-in-cell (PIC) methods is proposed for the solution of the VM system. The SM part uses a moment-based expansion in Hermitian basis while the PIC part employs Lagrangian particles to approximate the distribution function. The main idea is to improve the convergence properties of SM by treating the complex/highly disrupted part of the distribution function with particles. The merits and properties of the new numerical method are compared to spectral method properties on the example of weak electron-beam plasma interaction problem.

Oleksandr Koshkarov

University of Saskatchewan, Canada  
University of Saskatchewan  
olk548@mail.usask.ca

Gian Luca Delzanno, Gianmarco Manzini  
Los Alamos National Laboratory  
delzanno@lanl.gov, gmanzini@lanl.gov

Vadim Roytershteyn  
Space Science Institute  
vroyersh@gmail.com

### PP101

#### Minisymposium: Spectral Based-Discontinuous Galerkin Discretizations of the Vlasov-Poisson System

We present a new approach for solving the 1D-1V Vlasov-Poisson system describing a collisionless plasma of singly charged ions and electrons. For the numerical approximation of the Vlasov equation we combine a discontinuous Galerkin (DG) discretization of the physical space with a spectral representation of the velocity space based on Legendre polynomials or Hermite functions. To control the filamentation effect we introduce an artificial dissipation term that acts mainly on the higher modes of the spectral expansion. The Poisson equation is integrated directly by a numerical quadrature rule. The time-dependent set of ordinary differential equations for the coefficients of the spectral-DG expansion is discretized by the Crank-Nicolson method. The non-linear problem resulting from coupling the Vlasov equation with the Poisson equation is solved iteratively at any time cycle by a non-linear Newton-Krylov Jacobian-free method. We prove theoretically that the  $L^2$  stability of the Legendre-DG method can be enforced through a penalty term. We also analyze theoretically the structure of the main conservation laws, e.g., mass, momentum, and energy, for the Hermite and Legendre-based methods and investigate the conservation property numerically. Numerical results for standard test problems confirm the effectiveness of this approach.

Gianmarco Manzini, Gian Luca Delzanno  
Los Alamos National Laboratory  
gmanzini@lanl.gov, delzanno@lanl.gov

### PP101

#### Minisymposium: Large-Scale Implicit Particle-in-Cell Simulations of Magnetospheres with iPIC3D

A magnetosphere is a region of space filled with plasma around a planet with a dipolar magnetic field. Magnetospheres are shaped by the microscopic interaction phenomena between the solar wind and the dipolar magnetic field of the planet. To describe these interactions correctly, we need to model phenomena occurring over a large range of time and spatial scales. In fact, magnetospheres comprise regions with different particle densities, temperatures and magnetic field, where the characteristic time scales (plasma period, electron and ion gyro period) and spatial scales (Debye length, ion and electron skin depths) vary considerably. In this poster, we show how to address the problem

of modeling multi-scale phenomena during the formation of a magnetosphere by using an implicit Particle-in-Cell code, called "iPIC3D". We describe the implicit numerical method, the boundary conditions and the results of fully kinetic simulations of magnetosphere formation.

Stefano Markidis

KTH, Sweden  
markidis@kth.se

Ivy Bo Peng, Erwin Laure

KTH  
Sweden  
bopeng@kth.se, erwinl@kth.se

Elin Eriksson, Andreas Johlander, Yuri Khotyaintsev,

Andris Vaivads  
IRF - Uppsala  
eline@irfu.se, andreasj@irfu.se, yuri@irfu.se,  
andris@irfu.se

Gian Luca Delzanno

Los Alamos National Laboratory  
delzanno@lanl.gov

#### PP101

##### **Minisymposium: Computational Ideal Magneto-hydrodynamics with FV-FEEC Discretizations**

We present novel numerical methods for the ideal magneto-hydrodynamics (MHD) problem based on the coupling of Galerkin schemes for the electromagnetic fields via finite element exterior calculus (FEEC), with finite volume methods for the conservation laws of fluid mechanics. For the generalized transient advection problem, we rely on an Eulerian method of lines with explicit timestepping and spatial upwind discretization based on the duality between the contraction of differential forms and the extrusion of chains. The balance laws for the fluid form a system of conservation laws with the magnetic induction field as a variable coefficient. We design finite volume schemes based on approximate Riemann solvers and adapted to accommodate the electromagnetic contributions to the momentum and energy conservation. On a set of tests for 2D planar ideal MHD, the resulting lowest order fully coupled scheme proves first order accurate for smooth solutions, conservative and stable.

Cecilia Pagliantini

none  
cecilia.pagliantini@sam.math.ethz.ch

Ralf Hiptmair

ETH Zurich  
ralf.hiptmair@sam.math.ethz.ch

Siddhartha Mishra

ETH Zurich  
Zurich, Switzerland  
siddhartha.mishra@sam.math.ethz.ch

#### PP101

##### **Minisymposium: Kinetic Simulations of Astro-**

##### **physical Plasma Turbulence**

We briefly describe recent progress in fully kinetic simulations of astrophysical plasma turbulence. Specifically, we focus on problem of turbulent energy dissipation in weakly collisional plasmas. While fluid effects dominate the large-scale dynamics of plasma turbulence, the dissipation of turbulence in weakly collisional plasmas is inherently linked to kinetic effects. Since the local conditions are set by the large-scale dynamics, a correct description of dissipation requires bridging a very large gap in scales. We discuss petscale hybrid and fully kinetic particle-in-cell simulations of this problem and the insights gained.

Vadim Roytershteyn

Space Science Institute  
vroyersh@gmail.com

#### PP102

##### **Minisymposium: Two Families of $H(\text{div})$ Mixed Finite Elements of Minimal Dimension on Quadrilaterals**

We present two new families of mixed finite elements on quadrilaterals. The new families are inf-sup stable, and they approximate optimally the velocity, pressure, and divergence of the velocity. The spaces are of minimal dimension subject to the approximation properties and finite element conformity (i.e., they lie in  $H(\text{div})$  and are constructed locally). The two families give full and reduced  $H(\text{div})$  approximation, like Raviart-Thomas and BDM spaces. The two families are identical except for inclusion of a minimal set of vector and scalar polynomials needed for higher order approximation of the divergence of the velocity and pressure, and thereby we clarify and unify the treatment of finite element approximation between these two classes on quadrilaterals.

Todd Arbogast

University of Texas at Austin  
arbogast@ices.utexas.edu

Maicon R. Correa

University of Campinas  
UNICAMP, Brazil  
maicon@ime.unicamp.br

Zhen Tao

Institute for Computational Engineering and Sciences  
The University of Texas at Austin  
taozhen@ices.utexas.edu

#### PP102

##### **Minisymposium: Approximation on Quadrilateral and Hexahedral Meshes**

Classical finite element shape functions are defined by a transformation from a reference element associated to a mapping of it onto the actual element, e.g., a Piola transformation or a simple composition, depending on the context. When the mapping is required to be affine, the resulting approximation theory is simple. Affine mappings are sufficient for simplicial finite elements, but for quadrilateral and hexahedral finite elements multilinear maps are

required, and the approximation theory is much more complex. We survey joint work with Boffi and Bonizzoni using the framework of finite element exterior calculus to investigate the approximation properties of finite element spaces constructed in this way. The theory indicates a severe decrease in the rate of convergence obtained by shape functions transformed via multilinear mappings (as used for general quadrilateral and hexahedral meshes) compared to that obtained with affine mappings. The loss occurs even for simple finite elements, such as the 8-node serendipity quad, but the severity of the loss increases with the differential form degree, so, for example, it is much more extreme for  $H(\text{div})$  finite elements in three dimensions. We illustrate the consequences of this analysis for various finite elements, including the classical cubic edge and face elements and the serendipity family of finite element differential forms introduced by Awanou and the author.

Douglas N. Arnold  
School of Mathematics  
University of Minnesota  
arnold@umn.edu

#### PP102

##### **Minisymposium: Bernstein-Bézier Basis for $H(\text{div})$ and $H(\text{curl})$ Finite Elements on Hypercubes**

We present a set of Bernstein-Bézier basis for the tensor-product  $H(\text{div})$ - and  $H(\text{curl})$ - finite elements on square and cubes. The basis functions have clear geometric interpretations, and separates the gradient fields (curl-free) and non-gradients for the  $H(\text{curl})$  finite elements, and separates the curl fields (divergence-free) and non-curves for the  $H(\text{div})$  finite elements.

Guosheng Fu  
School of Mathematics  
University of Minnesota  
guosheng\_fu@brown.edu

Mark Ainsworth  
Division of Applied Mathematics  
Brown University  
Mark\_Ainsworth@brown.edu

#### PP102

##### **Minisymposium: Trimmed Serendipity Finite Elements**

We introduce the family of trimmed serendipity finite element differential form spaces which are defined on cubical meshes in any dimension, for any polynomial degree, and for any form order. The (non-trimmed) serendipity family was first introduced by Arnold and Awanou [*Math. Comp.* 83(288) 2014] and the relation between the trimmed and (non-trimmed) serendipity families is analogous to the relation between the trimmed and (non-trimmed) polynomial finite element differential form families on simplicial meshes. For these new spaces, we provide degrees of freedom in the general setting and prove that they are unisolvent for the trimmed serendipity spaces in all cases of immediate relevance to application: spatial dimension  $n$  up to 4, any differential form order  $k$ , and polynomial order  $r$

up to 10. In these cases, the sequence of trimmed serendipity spaces with a fixed polynomial order  $r$  provides an explicit example of a system described by Christiansen and Gillette [*ESAIM:M2AN* 50(3) 2016], namely, a minimal compatible finite element system on  $n$ -dimensional cubes containing order  $r - 1$  polynomial differential forms.

Tyler Kloefkorn, Andrew Gillitte  
University of Arizona  
Department of Mathematics  
tkloefkorn@math.arizona.edu, agillette@math.arizona.edu

#### PP102

##### **Minisymposium: A Progress Report on Construction of Smooth Generalized Barycentric Coordinates**

Many kinds of generalized barycentric coordinates (GBC) have been constructed in the literature for various applications in geometric design and numerical solution of partial differential equations. However, none of GBC can be pieced together to form a smooth  $C^1$  surface. It is a long standing problem how to find a smooth GBC. We present a report on our progresses toward constructing smooth GBC. We first explain how to find smooth GBC which can reproduce constants. Then we explain how to construct smooth GBC to reproduce linear functions. Construction of smooth GBC which can reproduce quadratic and cubic polynomials are difficult and we make some progress on this direction. This is a joint work with James Lanterman.

Ming-Jun Lai  
University of Georgia  
Department of Mathematics  
mjlai@math.uga.edu

#### PP102

##### **Minisymposium: Divergence-Free Stokes Elements on Quadrilateral Meshes**

Using a smooth de Rham complex as a guiding tool, we construct stable, low-order, and conforming finite element spaces on general convex quadrilateral partitions that yield divergence-free approximations. The proposed spaces are related to both the de Veubeke spline space and the quadratic serendipity finite element space. The resulting methods produce exactly incompressible flow, and the velocity errors are decoupled from both the pressure approximations and the viscosity. Extensions to three dimensions will also be discussed.

Michael J. Neilan  
University of Pittsburgh  
Department of Mathematics  
neilan@pitt.edu

Duygu Sap  
University of Pittsburgh  
duygusap@gmail.com

#### PP103

##### **Minisymposium: ForestClaw : Parallel Library for Solving Pdes on Mapped Multiblock Quadtree**

## Grids

We show recent results for solving PDEs using finite volume methods on mapped, adapt quad tree grids. In particular, we show examples from the GeoClaw ([www.geoclaw.org](http://www.geoclaw.org)) extension of ForestClaw.

Donna Calhoun

Boise State University  
donnacalhoun@boisestate.edu

Carsten Burstedde  
Universität Bonn  
burstedde@ins.uni-bonn.de

## PP103

### Minisymposium: Simulations of Asteroid-Generated Tsunamis Using GeoClaw

Following the Chelyabinsk meteor event of 2013, there has been renewed interest in the question of the dangers posed by relatively small asteroids that explode in the atmosphere, creating a strong pressure wave that can do damage on the ground, or perhaps create a tsunami. We will present simulations of tsunami propagation from asteroid-generated air bursts, using the open source GeoClaw software. The simulations span a range of conditions, from 5 to 250 MT of kinetic energy, in water of varying depths, and using bathymetry from the real coastlines. We have developed a radially symmetric one-dimensional test problem to better explore the nature and decay rate of waves generated by air burst pressure disturbances traveling at the speed of sound in air, which is much greater than the gravity wave speed of the resulting tsunami, and compared numerical simulations with analytical results. One-dimensional simulations along a transect of bathymetry are also used to explore the resolution needed for the full two-dimensional simulations, which are much more expensive (even with the use of adaptive mesh refinement) due to the short wave lengths of these tsunamis. For this same reason, the shallow water equations often used for modeling earthquake-generated tsunamis are generally found to be inadequate and we also discuss dispersive effects.

Randall LeVeque

University of Washington  
Applied Math  
rjl@uw.edu

Marsha Berger  
Courant Institute of Mathematical Sciences  
New York University  
berger@cims.nyu.edu

## PP103

### Minisymposium: Embedding Protective Mechanisms in Coastal Flooding Simulations

Flooding in many coastal communities has become a central concern due to the growing threat of climate change, in particular sea level rise. To combat this effective mitigation strategies are needed that are optimized for flood risk reduction but the question remains as to what strategies are the most effective. In this poster we present a method-

ology for simulating sea walls, dunes, and other protective strategies as an aspect of this optimization process with an eye towards computational efficiency due to the large number of runs (upwards of a 1000) that are required to run modern optimization approaches that will take into account climate change and the substantial uncertainty in the problem.

Kyle T. Mandli, Jiao Li

Columbia University  
Applied Physics and Applied Mathematics  
kyle.mandli@columbia.edu, jl4170@columbia.edu

## PP103

### Minisymposium: Clawpack and GeoClaw - Software Developments and Applications - Numerical Prediction of Water Level and Hydrodynamic Loads in Coastal Communities During a 500-Year Csz Tsunami

Prediction of water level and hydrodynamic loads on coastal structures during a tsunami inundation is critical for natural tsunami hazard mitigation. The numerical models must incorporate the constructed environment to model the interaction between the structures and the fluids, which is challenging for many numerical modelers because of the scale and complexity of the physical problem. A two-dimensional (2D) depth-averaged model is typically used to model waves offshore but is seldom applied to predict the inundation on land because of the complex flow with transient variance in the vertical direction around constructed environment. On the other hand, a more complex three-dimensional (3D) model, which is often used to solve the complex flow around a single coastal structure like a pier, is often too computationally expensive, with hundreds of structures being modeled at once in this case. In this study, the ability of a 2D model and a 3D model are investigated and discussed by comparing the numerical results with measurements from an experiment on a 1:50 model-scale physical model of part of Seaside, OR, USA. Both models predict the flow parameters accurately, in general, except for those at time near the initial impact, which is challenging due to the complexities in the flow, especially for the 2D model. The comparison of predicted forces indicates that force prediction from the momentum flux in a 2D model is not always reliable in such a complicated case.

Xinsheng Qin

University of Washington, Seattle  
xsqin@uw.edu

Michael Motley  
University of Washington  
mrmotley@uw.edu

Randall LeVeque  
University of Washington  
Applied Math  
rjl@uw.edu

Frank Gonzalez  
University of Washington

figonzal@uw.edu

### PP103

#### **Minisymposium: A Combined Model for Sediment Transport In Coastal Hazard Events (GeoClaw-Striche): Theoretical Formulation and Validation**

GeoClaw-STRICHE is designed for simulating the physical impacts of tsunamis as it relates to erosion, transport and deposition. GeoClaw-STRICHE comprises GeoClaw for the hydrodynamics and the sediment transport model we refer to as STRICHE, which includes an advection diffusion equation as well as bed-updating. Multiple grain sizes and sediment layers are added into GeoClaw-STRICHE to simulate grain-size distribution and add the capability to develop grain-size trends from bottom to the top of a simulated deposit as well as along the inundation. Unlike previous models based on empirical equations or sediment concentration gradient, the standard Van Leer method is applied to calculate sediment flux. We tested and verified GeoClaw-STRICHE with flume experiment by Johnson et al. (2016) and data from the 2004 Indian Ocean tsunami in Kuala Meurisi as published in Apotsos et al. (2011). The comparison with experimental data shows GeoClaw-STRICHE's capability to simulate sediment thickness and grain-size distribution in experimental conditions, which builds confidence that sediment transport is correctly predicted by this model. The comparison with the data from the 2004 Indian Ocean tsunami reveals that the pattern of sediment thickness is well predicted and is of similar quality, if not better than the established computational models such as Delft3D.

Hui Tang  
Virginia Tech  
tanghui@vt.edu

Robert Weiss  
Virginia Tech  
Department of Geosciences  
weisrz@vt.edu

### PP103

#### **Minisymposium: Seismic Modeling to Improve Tsunami Prediction in Geoclaw**

Earthquake-generated tsunamis can carry with them a powerful, destructive force. One of the most well-known and recent examples is the tsunami generated by the Tohoku earthquake, which was responsible for the nuclear disaster in Fukushima. The Geoclaw software package was created to assist in tsunami simulation and forecasting, a necessary element of emergency procedure planning and execution. It currently uses the Okada solution for a homogeneous, elastic half-space as an initial condition for the sea-surface deformation. While this solution has been shown to work well in many cases, the Cascadia Subduction Zone has largely varying sea-floor topology and plate densities. Given that these violate the assumptions of the Okada solution, this work replaces that solution with a time-dependent, sea-floor deformation generated by simulating the original seismic event directly. This allows

for the interaction of the water column with both time-dependent seismic waves and varying sea-floor topology. The results of Geoclaw coupled with these seismic simulations are compared against those using the Okada solution. Simulations specific to the Cascadia Subduction Zone are to be used in a early earthquake warning initiative for the Pacific Northwest.

Chris Vogl, Randy LeVeque  
University of Washington  
cvogl@uw.edu, rjl@uw.edu

### PP104

#### **Minisymposium: Efficiently Exploring the Conformational Space of Proteins Using the Concurrent Adaptive Sampling Algorithm**

Molecular dynamics (MD) simulations are useful in sampling thermodynamic and kinetic properties of bio-molecules. However, our sampling of the bio-molecules properties is severely limited by the timescale barrier, and MD simulations routinely get stuck in metastable free energy minima. While there are several existing methods to overcome these issues, problems remain in regard to being able to sample unknown systems, deal with high-dimensional space, and focus effort on slow timescales. Hence, a new sampling method, called the Concurrent Adaptive Sampling algorithm (CAS), has been developed to tackle these three main problems and efficiently obtain conformations and pathways. The beauty of CAS is that it adaptively constructs macrostates, requires little a priori knowledge about the system, and considers an arbitrary number of general collective variables. In addition, CAS uses the second eigenvector of the transition matrix to maintain a fine discretization along pathways while using importance sampling in orthogonal directions to control computational cost. In this talk, we introduce the new algorithm and show new results about the triazine polymers [Grate et al, Triazine-Based Sequence-Defined Polymers with Side-Chain Diversity and Backbone-Backbone Interaction Motifs] that were found using CAS.

Surl-Hee Ahn  
Stanford University  
sahn1@stanford.edu

Jay Grate  
Pacific Northwest National Laboratory  
jwgrate@pnnl.gov

Eric F. Darve  
Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

### PP104

#### **Minisymposium: A Coupling Strategy for Non-local and Local Models**

We develop an optimization-based method for the coupling of nonlocal and local problems, with applications to static peridynamics. The approach formulates the coupling as a control problem where the states are the solutions of the

nonlocal and local equations, the objective is to minimize their mismatch on the overlap of the nonlocal and local domains, and the controls are virtual volume constraints and boundary conditions. We implement our method using Sandia's agile software components toolkit that provides the groundwork for the development of engineering analysis tools. Also, we present numerical results for non-local material models in three-dimensions that illustrate key properties of the coupling method.

Marta D'Elia  
Sandia National Laboratories  
mdelia@sandia.gov

Pavel Bochev  
Sandia National Laboratories  
Computational Math and Algorithms  
pbboche@sandia.gov

Mauro Perego  
CSRI Sandia National Laboratories  
mperego@sandia.gov

David Littlewood  
Sandia National Laboratories  
djlittl@sandia.gov

#### PP104

##### **Minisymposium: Numerical Exterior Calculus Methods for Fluctuating Hydrodynamics Within Curved Fluid Interfaces**

We use exterior calculus of differential geometry to formulate fluctuating hydrodynamic equations to account phenomena within curved fluid interfaces. For manifolds of spherical topology, we present spectral methods that provide a discrete approximation of exterior calculus operators such as the exterior derivative, Hodge star, co-differential, or Hodge Laplacian. Our numerical approximation of the exterior calculus for surfaces of spherical topology is based on hyperinterpolation and Lebedev quadratures. We show how our methods can be used to formulate fluctuating hydrodynamic descriptions to investigate phenomena within curved fluid interfaces.

Paul J. Atzberger  
University of California-Santa Barbara  
atzberg@gmail.com

Ben J. Gross  
UC Santa Barbara  
bgross09@yahoo.com

#### PP104

##### **Minisymposium: Computing the Non-Markovian Coarse-Grained Interactions Derived from the Mori-Zwanzig Formalism in Molecular Systems: Application to Polymer Melts**

We consider the construction of coarse-grained models in the form of generalized Langevin equations (GLE) and non-Markovian dissipative particle dynamics (DPD), using the Mori-Zwanzig formalism. In this work, we compare GLE

and non-Markovian DPD models for the dynamics of a single coarse particle and the long-time hydrodynamic effects. Computationally, a direct evaluation of the non-Markovian terms requires storing historical information, which significantly increases the computational complexity. For DPD, this could be troublesome, since many pairwise interactions are considered. To address this, we add a few auxiliary variables per pairwise interaction; this allows replacing the non-Markovian dynamics with a Markovian dynamics in a higher dimensional space, leading to a much reduced memory footprint and computational cost. In our numerical benchmarks, the GLE and non-Markovian DPD models are constructed from molecular dynamics (MD) simulations of star-polymer melts. Results show that a Markovian dynamics with auxiliary variables successfully generates equivalent non-Markovian dynamics consistent with the reference MD system, while maintaining a tractable computational cost. Transient subdiffusion of the star-polymers can be reproduced by the CG models. However, to correctly reproduce the long-time hydrodynamics with an algebraic decay of the velocity autocorrelation function (VACF), the single-particle GLE model is insufficient, and a DPD model is required.

Hee Sun Lee  
Stanford University  
hslee88@stanford.edu

Zhen Li  
Division of Applied Mathematics  
Brown University  
Zhen\_Li@brown.edu

Eric F. Darve  
Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

#### PP104

##### **Minisymposium: Quantifying Quasi-Equilibrium and Non-Equilibrium Properties of Biomolecule System**

Biomolecules exhibit conformation fluctuations near equilibrium states, inducing uncertainty in various biological properties near metastable states as well as transition between the states. We have developed a general method to quantify the uncertainty of target properties induced by conformation fluctuations. For local properties, to alleviate the high dimensionality of the conformation space, we propose a method to increase the sparsity by defining a set of collective variables within active subspace, which increases the accuracy of the surrogate model. For dynamic properties, we develop a data-driven method to evaluate the memory kernel of the energy-dissipation process based generalized Langevin Equation. The method is demonstrated on solvation properties and is generalizable to investigate

uncertainty in numerous biomolecular properties.

Huan Lei, Xiu Yang, Nathan Baker  
Pacific Northwest National Laboratory  
huan.lei@pnnl.gov, xiu.yang@pnnl.gov,  
nathan.baker@pnnl.gov

Xiantao Li  
Department of Mathematics  
Pennsylvania State University  
xli@math.psu.edu

#### PP104

##### **Minisymposium: Mesh Refinement and Model Reduction: Two Sides of the Same Problem**

Despite the remarkable increase in computational power, most real-world systems are still too complex to simulate in full detail. In such cases, the hope is to construct numerical algorithms which can retain the salient features of a system while reducing the complexity. Two different ways of dealing with complexity, namely mesh refinement and model reduction are based on the common concept of transfer of activity from larger scales to smaller ones. This allows the unified development of reduced models and mesh refinement schemes. Examples involving singularity detection and tracking as well as uncertainty quantification for systems exhibiting bifurcations will be presented for illustration purposes.

Jing Li, Panos Stinis  
Pacific Northwest National Lab  
jing.li@pnnl.gov, stinis.panos@pnnl.gov

#### PP104

##### **Minisymposium: Mesoscopic Modeling, Concurrent Coupling and Multiscale Framework**

We will present the developments on the mesoscopic modeling, concurrent coupling and multiscale framework taking place in the Collaboratory on Mathematics for Mesoscopic Modeling of Materials, or CM4, with both fundamental contributions to mathematical models and practical algorithms for multiscale simulations. The new Dissipative Particle Dynamics methods enable accurate modeling of transport and fluctuations at the mesoscale, including reactive transport and charged polymers. We further can scale up or down the DPD method by domain decomposition that allows concurrent coupling. A novelty of our approach is that we go beyond mean field theory and obtain seamless transition of fluctuations across heterogeneous domains. In addition to developing the proper interface conditions, we have also developed a multiscale universal interface (MUI) that provides an efficient framework to connect heterogeneous solvers for modeling multiscale phenomena. The integration of these models and numerical algorithms using the MUI framework paves the way for investigation of the important materials modeling problems related to mesoscopic transport processes in complex materials and at materials interfaces.

Zhen Li  
Division of Applied Mathematics  
Brown University

Zhen\_Li@brown.edu

Xin Bian, Yu-Hang Tang  
Brown University  
xin.bian@tum.de, yuhang\_tang@brown.edu

George E. Karniadakis  
Brown University  
Division of Applied Mathematics  
george\_karniadakis@brown.edu

#### PP104

##### **Minisymposium: A Multifidelity Approach to Effectively Compute Steady-State Flow Of Ionic Solutions**

Steady-state flow and ion spatial distributions are typically calculated using the Poisson-Nernst-Planck (PNP) model, where the ions are assumed to be point charges. A more accurate, but much more expensive model, is the classical density functional theory (cDFT) model which accounts for correlations between ions due to their finite size and to electrostatic correlations. Since the effects described by cDFT are primarily important near charged surfaces, one can use the PNP model in large part of the computational domain and resort to cDFT only near the surfaces. To this end, we propose an alternating Schwarz procedure to couple the nonlocal cDFT equations with the PNP equations. We present numerical results showing that the proposed approach leads to significant memory and computational savings, while retaining the descriptive capabilities of a full cDFT approach

James Cheung  
Florida State University  
jcheung@sandia.gov

Amalie Frischknecht  
Sandia National Laboratories  
alfrisc@sandia.gov

Mauro Perego  
CSRI Sandia National Laboratories  
mperego@sandia.gov

Pavel Bochev  
Sandia National Laboratories  
Computational Math and Algorithms  
pbboche@sandia.gov

#### PP104

##### **Minisymposium: Fast Hierarchical Solvers for Dense Linear Systems**

Solving a large linear system of equations is often the most expensive part of a scientific calculation. Depending on the application, the linear system can either be sparse (e.g., partial differential equations) or dense (e.g., integral equations and fractional partial differential equations). Conventional direct solver methods have quadratic/cubic computational complexity, which is prohibitive for large problems. On the other hand, iterative methods demonstrate better complexity and scalability, however, suffer from highly



problem dependent convergence. We present new linear solve methods that can be deployed for both sparse and dense matrices. We compute a hierarchical representation of the LU factorization of the given matrix, using sequence of low-rank approximations. The solvers are approximate, and the error can be controlled and made as small as needed. Therefore, the proposed method can be used both as a stand-alone solver, and a pre-conditioner in conjunction to other iterative methods. We present variety of applications in the realm of scientific computing, for which our solver exhibit linear complexity and outperforms other solvers/preconditioners.

Hadi Pouransari

Stanford  
hadip@stanford.edu

Pieter Coulier  
Stanford University  
pieter.coulier@gmail.com

Eric F. Darve  
Stanford University  
Mechanical Engineering Department  
darve@stanford.edu

#### PP104

##### **Minisymposium: Particle Methods for the Mesoscale**

We present recent work from the Collaboratory on Mathematics for Mesoscopic Modeling of Materials (CM<sup>4</sup>) using particle methods to study material science problems occurring at mesoscopic lengthscales. Particle methods provide an ideal framework for many of these problems, where a careful treatment of thermal fluctuations, complex geometry, and coarse-graining is necessary. We will show recent results developing new methods, establishing rigorous mathematical foundations for existing methods, and using particle approaches to study multiphysics problems at the mesoscale.

Nathaniel Trask

Brown University  
Division of Applied Mathematics  
nat.trask@gmail.com

Martin Maxey  
Division of Applied Mathematics,  
Brown University  
martin\_maxey@brown.edu

Alexandre M. Tartakovsky  
PNNL  
alexandre.tartakovsky@pnnl.gov

Wenxiao Pan  
Pacific Northwest National Laboratory  
wenxiao.pan@pnnl.gov

George E. Karniadakis  
Brown University  
Division of Applied Mathematics

george\_karniadakis@brown.edu

#### PP104

##### **Minisymposium: Effect of Uncertainties in Radii and Atomic Charges on the Solvation Calculations**

Atomic radii and charges are the two main parameters used in implicit solvent electrostatics and energy calculations. The optimization problem for finding good leads to uncertainty in the values of these parameters. As such, the results of electrostatics and solvation energy calculations can vary due to the uncertainty in the charge and radius parameters. We presents a method for quantifying the uncertainty in solvation energies using surrogate models based on generalized polynomial chaos (gPC) expansions. There are relatively few types of radii parameters used in implicit solvation calculations so surrogate models for these low-dimensional spaces could be constructed using least-squares fitting. However, there are many more types of atomic charges; construction of surrogate models for the parameter space required a compressed sensing approach combined with an iterative rotation method to enhance problem sparsity. We present results for the uncertainty in small molecule solvation energies based on these approaches.

Xiu Yang, Huan Lei, Peiyuan Gao, Dennis Thomas  
Pacific Northwest National Laboratory  
xiu.yang@pnnl.gov, huan.lei@pnnl.gov,  
peiyuan.gao@pnnl.gov, dennis.thomas@pnnl.gov

David Mobley  
UC Irvine  
dmobley@mobleylab.org

Nathan Baker  
Pacific Northwest National Laboratory  
nathan.baker@pnnl.gov

#### PP105

##### **Minisymposium: Performance Comparisons of Application Codes on Modern Computer Architectures**

State-of-the-art distributed-memory clusters today contain multi-core CPUs with 8 to 16 cores and co-processors such as massively parallel GPUs with thousands of computational cores or many-core Intel Xeon Phi with 60 to 70 cores. The second-generation Phi, code named Knights Landing, can have over 70 cores, connected by a 2D mesh network, and can also be used as stand-alone processor. We compare the performance of an application code for a system of time-dependent partial differential equations on these hardware choices and discuss their reduced energy consumption, which is a key advantage to using accelerators. The application code models calcium induced calcium release (CICR) in a heart cell, and its matrix-free implementation is the key feature that allows it to run on the co-processors with their limited memory.

Ishmail Jabbie  
University of Maryland, Baltimore County

Department of Mathematics and Statistics  
 jjabbiel@umbc.edu

George Owen  
 Louisiana State University  
 Department of Mathematics  
 georgeowen2@gmail.com

Benjamin Whiteley  
 University of Maryland, Eastern Shore  
 Department of Engineering and Aviation Sciences  
 bwhiteley2012@gmail.com

Jonathan Graf, Xuan Huang  
 Department of Mathematics and Statistics  
 University of Maryland, Baltimore County  
 jongraf1@umbc.edu, hu6@umbc.edu

Matthias K. Gobbert, Samuel Khuvis  
 University of Maryland, Baltimore County  
 Department of Mathematics and Statistics  
 gobbert@umbc.edu, khsa1@umbc.edu

#### PP105

##### **Minisymposium: Performance Analysis and Numerical Method Tuning for a System of Non-Linear Time-Dependent Advection-Diffusion-Reaction Equations**

Our problem of interest is a system of non-linear, coupled, time-dependent PDEs modeling calcium induced calcium release (CICR) in a heart cell. To study the behaviors requires large numbers of parameter studies run to long final times. This demands highly efficient numerical methods, tuning of the numerical method parameters, and the use of modern parallel architectures. This is particularly challenging as the model includes thousands of point sources at which calcium can be released over the three dimensional domain. We examine the interplay between the numerical methods within a method of lines approach with FEM and FVM through an examination of the time stepping behavior. The parallel implementation of this method in C with MPI and OpenMP is matrix-free and has demonstrated speedup on modern CPUs.

Jonathan Graf  
 Department of Mathematics and Statistics  
 University of Maryland, Baltimore County  
 jongraf1@umbc.edu

Matthias K. Gobbert  
 University of Maryland, Baltimore County  
 Department of Mathematics and Statistics  
 gobbert@umbc.edu

#### PP105

##### **Minisymposium: Coupling the Electrical Excitation and Calcium Signaling in a Heart Cell**

Calcium dysregulation is a significant cause of fatal cardiac arrhythmias, but it is an incompletely understood phenomenon and difficult to predict. Cardiac calcium levels can be modeled as a system of partial differential equations

linking the electrical excitation, calcium signaling, and mechanical contraction effects on the calcium dynamics of a heart cell. Our full model with 8 PDEs is an extension of the previous model to introduce links between the calcium and mechanical systems as well as the link from calcium to electrical systems. We present simulations with only the link from electrical to calcium system enabled and a parameter study on the strength of the feedback connection with both links between calcium signaling and electrical excitation enabled. Over strengthening the feedback connection from the current generated by calcium efflux results in physiologically unrealistic voltage behavior in the system.

Uchenna Osia  
 UMBC  
 osiaul@umbc.edu

Kallista Angeloff  
 University of Washington  
 Department of Atmospheric Sciences  
 angeloff@uw.edu

Carlos Barajas  
 Olivet College  
 Department of Mathematics and Computer Science,  
 cbarajas@olivetcollege.edu

Alexander Middleton  
 Winthrop University  
 Department of Mathematics,  
 middletona6@mailbox.winthrop.edu

Jonathan Graf  
 Department of Mathematics and Statistics  
 University of Maryland, Baltimore County  
 jongraf1@umbc.edu

Matthias K. Gobbert  
 University of Maryland, Baltimore County  
 Department of Mathematics and Statistics  
 gobbert@umbc.edu

Zana Coulibaly  
 University of California, Davis  
 Department of Pharmacology  
 zcoulibaly@ucdavis.edu

#### PP105

##### **Minisymposium: Performance Comparison of Intel Xeon Phi Knights Landing**

The Intel Xeon Phi is a many-core processor family with over 60 computational cores and a theoretical peak performance of approximately 1 TFLOP/s in double precision. This project compares the performance of code on the Intel Xeon Phi. The benchmark code solves a classical elliptic test problem, the two-dimensional Poisson equation with homogeneous Dirichlet boundary conditions. Specifically, we benchmark a reference code on the first and second generation of the Intel Xeon Phi, code-named Knights Corner (KNC) and Knights Landing (KNL), respectively. Key advantages of the second-generation KNL are a 2D network

connecting the 68 cores and an on-chip MCDRAM memory. The KNL using MCDRAM is dramatically faster in all cases. Despite DDR4 being a slower form of memory, KNL using DDR4 is comparable in most cases to KNC using GDDR5. For both MCDRAM and DDR4 on the KNL, using more threads than MPI processes is significantly faster than the inverse for this code.

George Owen  
Louisiana State University  
Department of Mathematics  
georgeowen2@gmail.com

Ishmail Jabbie  
University of Maryland, Baltimore County  
Department of Mathematics and Statistics  
ijabbie1@umbc.edu

Benjamin Whiteley  
University of Maryland, Eastern Shore  
Department of Engineering and Aviation Sciences  
bwhiteley2012@gmail.com

Jonathan Graf  
Department of Mathematics and Statistics  
University of Maryland, Baltimore County  
jongraf1@umbc.edu

Matthias K. Gobbert, Samuel Khuviv  
University of Maryland, Baltimore County  
Department of Mathematics and Statistics  
gobbert@umbc.edu, khsa1@umbc.edu

#### PP106

##### **Minisymposium: Eigenvalue Computation for 4th Order Sturm Liouville Equations: Magnus Methods**

Greenberg and Marletta produced a fortran code [see Greenberg and Marletta, Algorithm 775: The Code SLEUTH, ACM Trans Math Software 23 (4), 1997, 453-493] for computing the eigenvalues of 4th order Self-Adjoint differential operators, with regular endpoints. The basic shooting algorithm employed was based on using a piecewise constant approximation to the coefficient functions on each mesh interval, similarly to the Pruess method for 2nd order equations. Here we consider the 4th order equation

$$y^{(4)} - (S(x)y')' + Q(x)y = \lambda y$$

on regular intervals with  $S, Q$  continuous, convert to a first order system,  $Y = AY$  in the usual way, and employ the Magnus Method MG4 of Iserles, et. al., to write the matrix exponential for the solution on each mesh interval in terms of Lie brackets. The resulting shooting algorithm is then of order 4. We discuss the performance of this approach on several problems, including squares of 2nd order Sturm-Liouville equations, and compare with the SLEUTH code.

Ahmad A. Alalyani  
Florida Institute of Technology, Melbourne, Florida

aalalyani2013@my.fit.edu

#### PP106

##### **Minisymposium: A Sturm Liouville Problem for Relativistic Electrons Inside Thunderstorms**

Under the influence of strong electric fields inside thunderclouds, electrons can be accelerated to multi MeV energies by overcoming the drag force in air. In 2003, a Monte Carlo computer code was developed to model all the collisional energy losses that these electrons experience along their path of travel [Dwyer (2003), A fundamental limit on electric fields in air, Geophys. Res. Lett., 30, p.2055]. To understand this simulation, a transport equation was developed to describe the steady state energy spectrum of relativistic electrons as they move through Earth's atmosphere [Cramer-Dwyer-Arabshahi-Vodopyanov-Liu-Rassoul (2014), An analytical approach for calculating energy spectra of relativistic runaway electron avalanches in air, J. Geophys. Res. Space Physics, 119, 77947823]. This equation is in the form,

$$-(p(\epsilon)f'_{re}(\epsilon))' + q(\epsilon)f_{re}(\epsilon) = \lambda r(\epsilon)f_{re}(\epsilon)$$

where  $f_{re}(\epsilon)$  is the differential electron energy spectrum (number of electrons per unit energy). Here we show attempts to find an analytical solution to this SL problem by making assumptions about the total force and the seeding of the relativistic electrons inside the acceleration region.

Eric S. Cramer  
University of Alabama, Huntsville  
eric.cramer@uah.edu

#### PP107

##### **BE Posters TBA**

Poster titles and presenters to be announced by 12/2/2016.

Mary Ann E. Leung  
Sustainable Horizons Institute  
mleung@shinstitute.org

#### PP108

##### **Minisymposium: Sustainable Productivity in the Fenics Development Team**

The FEniCS project aims to provide a high productivity environment for development of finite element based simulation software. Techniques applied to achieve this goal include mixed language programming and code generation, which enables writing high performance programs in a high level language. End-user productivity is a high priority goal in our software designs. To sustain the productivity of the multinational team of part-time developers (mainly researchers and students) is paramount to the long term survival of the project. To minimize the developer workload while making the process open and accessible to new contributors and users, we regularly question which tools are the best available for our needs. On this poster we will present our current tool choices and work flows for developers and the wider FEniCS community. This list includes version control, build systems, testing, release management, team communication, documentation, and end

user support. The most recent addition to our toolbox are developer curated Docker images. We are investigating their usefulness in testing infrastructure, end user deployment, HPC cluster deployment, and as reproducible software environments to accompany journal publications. We welcome discussion on alternatives that can simplify our lives.

Martin Alnæs

Simula Research Lab  
martinal@simula.no

Jan Blechta

Charles University in Prague  
blechta@karlin.mff.cuni.cz

Jack Hale

University of Luxemburg  
jack.hale@uni.lu

Anders Logg

Chalmers University of Technology  
logg@chalmers.se

C N. Richardson

University of Cambridge  
chris@bpi.cam.ac.uk

Johannes Ring

Simula Research Laboratory  
johannr@simula.no

Garth Wells

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

#### PP108

##### **Minisymposium: A Sustainable Software Architecture for Scalable Nonlinear Boundary Element Method Simulations**

Boundary element code is most often bespoke, using few underlying libraries. Consequently, BEM research progress often requires "ground up" rewrites. Our flexible software, modular architecture enables N-body algorithms, advanced discretizations (higher order methods), coupling to other models (e.g. FEM) or applications (PDE-constrained optimization), straightforward efficiency analysis via work-precision diagrams, and extensions to new classes of theories, e.g. nonlinear BIEs.

Matthew G. Knepley

Rice University  
knepley@gmail.com \*Preferred

Jaydeep P. Bardhan

Northeastern University  
j.bardhan@neu.edu

#### PP108

##### **Minisymposium: The Application of Tribits to the Development and Integration Processes of**

##### **Larger Componentized Multi-Organization Scientific and Engineering Software Projects**

The Tribal Build, Integrate, and Test System (TriBITS) is a Framework built using and on top of the open-source CMake tools. TriBITS is designed to manage large, mixed-language (C, C++ & Fortran) distributed projects composed of many different (git) repositories created and maintained by separate development teams. The TriBITS framework accomplishes by hierarchically decomposing CMake-configured codes into different software repositories, packages, and sub-packages to create composable CMake projects. For example, the independent CASL DOE Innovation Hub related codes MPACT, COBRA-TF, SCALE, and Trilinos all use TriBITS as their native CMake build systems which also allows them to be seamlessly integrated with each other and to form the full CASL VERA CMake project. In addition, the TriBITS system allows the easy set up and maintenance of multi-repository integration processes with frequent (almost daily) repository integrations that provide the stability that developers and users need to be productive. TriBITS and the development and integration processes that it enables greatly improves software productivity and sustainability for larger ecosystems of co-developed CSE software compared to more traditional "everyone for themselves" approaches.

Roscoe A. Bartlett

Oak Ridge National Laboratory  
rabartl@sandia.gov

#### PP108

##### **Minisymposium: Doecode: The New Department of Energy Software Center**

The U.S. Department of Energy's Office of Scientific and Technical Information (OSTI) is charged with cataloging and preserving the software artifacts, among other things, produced by the Department's researchers. The fast pace at which the software world changes - including the rise and dominance of Open Source Software - has presented OSTI with the opportunity to update and expand its Software Center. These updates will provide a platform for the research community that combines social coding with OSTI's vast database of papers, authors, data, and legacy software to provide an integrated, dynamic environment for scientific software developers. This poster will describe the on-going development effort around that platform, as well as the requirements that define it, and how OSTI is engaging the open research community at large in its development. This poster will also present a description of scientific software in a broader context to clarify further why such a platform is needed and the extremely great benefits expected from its deployment.

Jay Billings

Oak Ridge National Laboratory  
billingsjj@ornl.gov

Mark Martin

DOE Office of Scientific and Technical Information

martinm@osti.gov

**PP108****Minisymposium: PEACHPY.IO, a Web App for Performance Tuning**

As the progress in transistor scaling slows down, low-level performance tuning becomes an important tool to accelerate compute-intensive codes. However, the community of experts who could optimize the codes on low-level is tiny, and their ability to contribute optimizations to open-source software is often restricted by lack of access to relevant hardware. We present PeachPy.io, an IDE for low-level performance tuning that works inside a Web browser. On PeachPy.io developers can modify assembly kernels, run them on different hardware platforms, analyze performance bottlenecks through hardware performance counters, and compare against alternatives without any time-consuming client-side configuration. PeachPy.io crowd-sources performance tuning, and lets anyone interested try their skills in low-level optimization without spending time on local software configuration or investing financial resources to acquire hardware.

Marat DukhanGA Tech  
mdukhan3@gatech.eduRichard Vuduc  
Georgia Institute of Technology  
richie@cc.gatech.edu**PP108****Minisymposium: Managing the Scientific Software Ecosystem with Spack**

Modern scientific software is far from monolithic – the largest packages comprise 70 or more dependency libraries. Building and deploying such packages for a variety of architectures, compilers, and dependency libraries is tedious and error-prone, and the complexity of the task impedes reuse and reproducibility. In other fields, package managers are routinely used to manage large collections of software, but HPC software has unique requirements, such as highly specialized compilers and finely tuned, ABI-incompatible interfaces, that make the use of a traditional package manager difficult. Spack is an open source package manager for HPC. Originally developed at Lawrence Livermore National Laboratory, Spack now spans over 30 organizations and over 100 contributors. Spack can install multiple versions, configurations, and finely tuned builds, allowing scientists to rebuild and reuse even the most complex packages across different supercomputers. Spack's dependency model and its support for combinatorial package versioning simplify the use of HPC software for developers, end users, and facilities staff. This poster describes Spack, its adoption in the scientific community this far, and the development and testing methods we have used to sustain Spack as a large open source project.

Todd Gamblin

Lawrence Livermore National Laboratory

tgamblin@llnl.gov

**PP108****Minisymposium: External Code Review of PFLOTRAN: Lessons Learned**

A code review can be a valuable exercise for maintaining consistency across a diverse team of developers (e.g. readability, coding style, exception handling) and validating feature design and implementation. These reviews can be internal to the project with team members probing each others code for design flaws or inspecting low-level implementation (peer review), or external where the focus is more on software quality and the promotion of best software practices. The overarching goal is to improve software quality while accelerating the training of new developers. In the context of open source scientific computing where software is being written and maintained by a heterogeneous mix of computer scientists, applied mathematicians and domain scientists who may be remotely located around the world and funded on a diverse number of projects, code review can be challenging. Recently, a code review was conducted on PFLOTRAN ([www.pflotran.org](http://www.pflotran.org)), an open source, massively-parallel code employed to simulate geoscientific processes in the Earth's subsurface as an exercise within the IDEAS project ([ideas-productivity.org](http://ideas-productivity.org)). PFLOTRAN currently consists of 213K lines of source code and is developed and maintained by numerous scientists funded on various projects at research institutions around the world. This presentation reports on lessons learned during this review exercise with the intent of improving the code review process for scientific applications in the future.

Glenn HammondSandia National Laboratories  
gehammo@sandia.gov**PP108****Minisymposium: Regression and Performance Testing and Continuous Integration for Scientific Codes**

This poster will discuss motivations, tools, best practices, and past experiences with automated testing for scientific software. Because information about testing software in general is easy to be found, focus will be on issues specific to scientific codes. I will present different unit testing frameworks and other tools (like ctest and numdiff) and discuss automated performance testing. I will explain how testing is done in the various projects I am involved in: the c++ finite element library deal.II, the mantle convection code ASPECT for the geoscience community, and a smaller mineral physics code BurnMan written in Python, among others.

Timo HeisterClemson University  
Mathematical Sciences  
heister@clemson.edu**PP108****Minisymposium: CSE Complete: R & D for**

## Productivity Improvement

Computational Science and Engineering (CSE) is about delivering valuable scientific results. CSE activities are meaningful only to the extent they deliver these results. At the same time, certain practices, tools and processes, while not directly about producing CSE results, actually improve our ability to deliver results over a sufficiently long span of time. In this poster we discuss efforts in the IDEAS Scientific Software Project that are intended to improve the overall effectiveness of CSE efforts. IDEAS is focused on R & D in new productivity improvements for CSE software, including establishing a scientific software ecosystem (xSDK) and creating, providing and demonstrating the value of content that can be used by CSE software teams to improve their software efforts.

### Michael Heroux

Sandia National Laboratories  
maherou@sandia.gov

Lois Curfman McInnes  
Mathematics and Computer Science Division  
Argonne National Laboratory  
curfman@mcs.anl.gov

David E. Bernholdt  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
bernholdtde@ornl.gov

David Moulton  
Los Alamos National Laboratory  
Applied Mathematics and Plasma Physics  
moulton@lanl.gov

## PP108

### Minisymposium: HPC Software Productivity Based on High-Level Abstraction in FEniCS

We are developers of the general adaptive stabilized G2 General Galerkin / Direct FEM methodology and of the FEniCS framework, specifically focusing on HPC and the Unicorn solver targeting turbulent continuum models. We are able to demonstrate a breadth of advanced HPC applications, such as adaptive turbulent flow, variable-density water-air interaction and aeroacoustics, with good scaling on supercomputers and with minimal manual implementation where we exploit UFL, the high-level abstract form language in FEniCS and the generality of G2/Direct FEM. Unicorn/FEniCS is available as open source, and gives the CSE community the ability to easily build and extend general computational applications on leading supercomputer hardware in the world, via e.g. the EU PRACE program.

Johan Jansson  
School of Computer Science and Communication  
Royal Institute of Technology KTH  
jjan@kth.se

Johan Hoffman  
Royal Institute of Technology KTH  
jhoffman@kth.se

Niclas Jansson  
School of Computer Science and Communication  
Royal Institute of Technology KTH  
njansson@kth.se

## PP108

### Minisymposium: Improving Software Productivity of Uintah Through Task-Based Architectures, Performance Portability Libraries and Modern C++ Features

The use of a carefully-layered software architecture in the Uintah framework is seen to make it possible to provide software sustainability through a clear separation of concerns between physics discretization tasks and a runtime system. At the same time productivity is enhanced if portability is straightforward. This challenge can be met through the use of portability layers such as Kokkos. Re-configuring a large framework to use Kokkos is not straightforward and requires changes both to the physics code loops and to the runtime system. Finally the use of modern C++ features makes it possible to dramatically improve legacy older C++ codes through removing redundant code.

Martin Berzins, Alan Humphrey  
Scientific Computing and Imaging Institute  
University of Utah  
mb@sci.utah.edu, ahumphrey@sci.utah.edu

Dan Sunderland  
Sandia Corporation  
dansunderland@gmail.com

Brad Peterson  
University of Utah  
bradpeterson@gmail.com

## PP108

### Minisymposium: CSE Software Maintenance and Automation: Beyond Testing

We identify key maintenance operations for CSE software that affect productivity and sustainability, including the work to merge contributions while maintaining correctness, performance, and a coherent style. Issues important to CSE are highlighted, such as the storage of data files for regression testing. We present advanced techniques to automate certain maintenance tasks including version control operations, source formatting, and data file generation. Distributed version control workflows are compared in terms of both maintainer and contributor productivity, as well as their ability to prevent regressions.

Dan A. Ibanez  
Rensselaer Polytechnic Institute  
SCOREC  
ibaned@rpi.edu

## PP108

### Minisymposium: The Impact of Structured AMR Representation on Software Design

We discuss the integration of p4est into the PETSc DM

interface, and the extension of the Plex model to handle non-conforming meshes. This presents the extension of an existing software framework with large user base, and its enhancement with significant new functionality. The close integration with discretization and scalable solvers will increase the productivity of a modeler using structured AMR, while integration into the existing PETSc library framework improves maintainability and usability of the interface.

Tobin Issac  
University of Chicago  
tisaac@uchicago.edu

Matthew Knepley  
Rice University  
knepley@rice.edu

### PP108

#### Minisymposium: SlabGenerator: Improving Productivity for Subduction Modeling

With the advent of efficient and scalable mantle convection codes, a significant of user time in subduction modeling is spent constructing a problem instance through data assimilation, most often manually. Many different data modalities need to be combined with analytical modeling. In addition, any software must also interoperate with complex, existing packages for mantle convection. The SlabGenerator software generates thermal and viscosity structures for large scale mantle convection simulations, processing many different data sources, greatly increasing user productivity.

Margarete A. Jadamec  
Brown University  
Department of Geological Sciences  
mjadamec@central.uh.edu

Matthew Knepley  
Rice University  
knepley@rice.edu

### PP108

#### Minisymposium: System Testing for PDE Frameworks - Tools and Experiences

Testing is acknowledged as indispensable support for scientific software development and assurance of software quality to produce trustworthy simulation results. Most of the time, testing in software frameworks developed at research facilities is restricted to either unit testing or simple benchmark programs. However, in a modern numerical software framework, such as deal.II, FEniCS, or Dune, the number of possible feature combinations constituting a program is vast. Only system testing, meaning testing within a possible end user environment also emulating variability, can assess software quality and reproducibility of numerical results. We discuss tools to define system tests including both runtime and compile time variation. We furthermore discuss implementation of quality measures tailored to numerical frameworks for the solution of PDEs. We will also share experiences on using continuous integration systems

(GitLab CI) for numerical software frameworks.

Dominic Kempf  
IWR, Heidelberg University  
dominic.kempf@iwr.uni-heidelberg.de

Timo Koch  
IWS, University of Stuttgart  
timo.koch@iws.uni-stuttgart.de

Bernd Flemisch  
University of Stuttgart, Germany  
bernd.flemisch@iws.uni-stuttgart.de

Peter Bastian  
Interdisciplinary Center for Scientific Computing  
University of Heidelberg  
peter.bastian@iwr.uni-heidelberg.de

### PP108

#### Minisymposium: Devito: Towards An Efficient and Sustainable Finite Difference DSL

The use of domain specific languages (DSL) constitutes an appealing strategy to enable high-productivity environments for computational scientists. The main idea of this approach is to decouple the problem specification from its low level implementation to create a separation of concerns between domain scientists and HPC specialists that has a direct payoff in productivity and sustainability. Inspired by the complexity of real-world seismic imaging problems and by popular frameworks such as FEniCS and Firedrake, we introduce Devito, a finite difference DSL in which high level equations are expressed using symbolic Python (SymPy) expressions. Complex equations are automatically manipulated, optimized, and translated into highly optimized C code that aims to perform comparably or better than hand-tuned code. All this is transparent to users, who only see concise symbolic mathematical notation. Since Devito is primarily targeted at generating fast wave propagation codes it offers an abstraction hierarchy that enables not only the generation of fast stencil kernels, but also a range of domain-specific features, such as sparse grid point interpolation. To overcome the limitation in scope many DSLs suffer from, Devito provides a second-level API that allows additional C-like expressions to be inserted into the generated code, allowing users to enrich their applications with custom operations without leaving the comforts of the Python software ecosystem.

Michael Lange  
Department of Earth Science and Engineering  
Imperial College London  
michael.lange@imperial.ac.uk

Navjot Kukreja  
SENAI CIMATEC  
Salvador, Brazil  
navjot.kukreja@fbter.org.br

Mathias Louboutin  
PhD student at the University of British Columbia  
mloubout@eos.ubc.ca

Fabio Luporini  
 Department of Computing  
 Imperial College London  
 f.luporini12@imperial.ac.uk

Felippe Vieira  
 SENAI CIMATEC  
 Salvador, Brazil  
 felippe.vieira@fieb.org.br

Vincenzo Pandolfo  
 Department of Computing  
 Imperial College London, UK  
 vincenzo.pandolfo15@imperial.ac.uk

Paulius Velesko  
 College of Electrical and Computer Engineering  
 University of Oklahoma, USA  
 pvelesko@ou.edu

Paulius Kazakas  
 Department of Computer Science  
 University of York, UK  
 pk627@york.ac.uk

Gerard J. Gorman  
 Imperial College London  
 g.gorman@imperial.ac.uk

#### PP108

##### **Minisymposium: Lessons Learned from Integrating Scientific Libraries Within a Plugin-Based Architecture**

Designing interoperable software is critical to boosting scientific productivity in many research applications, especially when exposed through flexible interfaces such as a plugin or component based framework. Such scalable applications with complex dependency chains often require rigorous source configuration, continuous testing and flexible deployment processes that cover a wide range of platforms (Linux/OSX/Windows) and environment variations. We will discuss hurdles in achieving interoperability using lessons learned from developing a serial/parallel MOAB database plugin for VisIt. The development of this plugin has led us to question some library design choices, and emphasized the need for creating better processes (versioning and build configuration) that are resilient to software interface and lifecycle changes. Taking the lessons learned from this new VisIt plugin, we present best practices and guidelines that are more broadly applicable to scientific software development in CSE applications.

Vijay S. Mahadevan  
 Argonne National Laboratory  
 mahadevan@anl.gov

Mark Miller  
 Lawrence Livermore National Laboratory  
 miller86@llnl.gov

Iulian Grindeanu  
 Argonne National Laboratory

iulian@mcs.anl.gov

#### PP108

##### **Minisymposium: xSDK: Working toward a Community CSE Software Ecosystem**

As CSE increasingly incorporates multiscale and multi-physics modeling, simulation, and analysis, the combined use of software developed by independent groups has become imperative. However, sharing software is difficult due to inconsistencies in configuration, installation, and third-party packages, as well as deeper challenges when interoperability requires control inversion and controlling data across packages. This poster explains how the Extreme-scale Scientific Software Development Kit (xSDK, <https://xsdk.info>) addresses these difficulties and provides the foundation of a CSE software ecosystem, as we work toward a collection of complementary software elements developed by diverse, independent teams. We demonstrate how the xSDK facilitates investigating climate impacts on the Upper Colorado River System through coupled models for surface-subsurface hydrology and reactive transport. Alquimia, an application-specific xSDK package, provides a common interface library for codes like Amanzi/ATS and ParFlow to access biogeochemistry capabilities from codes such as PFLOTRAN and CrunchFlow. In turn, these applications require the combined use of xSDK numerical libraries, including hypre, PETSc, SuperLU, and Trilinos. A key aspect of work is a set of draft xSDK compliance standards, which improve code quality, sustainability, usability, and interoperability. We invite the CSE community to provide feedback on compliance standards and contribute to the xSDK.

Lois Curfman McInnes  
 Mathematics and Computer Science Division  
 Argonne National Laboratory  
 curfman@mcs.anl.gov

Michael Heroux  
 Sandia National Laboratories  
 maherou@sandia.gov

Roscoe A. Bartlett, Mark Berrill  
 Oak Ridge National Laboratory  
 rabartl@sandia.gov, berrillma@ornl.gov

Jeffrey N. Johnson  
 Applied Numerical Algorithms Group  
 Lawrence Berkeley National Laboratory  
 jnjohnson@lbl.gov

Alicia Klinvex  
 Sandia National Laboratories  
 amklinv@sandia.gov

X. Sherry Li  
 Lawrence Berkeley National Laboratory  
 xsli@lbl.gov

J. David Moulton  
 Los Alamos National Laboratory  
 moulton@lanl.gov



Daniel Osei-Kuffuor  
Lawrence Livermore National Laboratory  
oseikuffuor1@llnl.gov

Jason Sarich  
Argonne National Laboratory  
sarich@mcs.anl.gov

Barry F. Smith  
Argonne National Lab  
MCS Division  
bsmith@mcs.anl.gov

James Willenbring  
Sandia National Laboratories  
jmwille@sandia.gov

Ulrike Yang  
Lawrence Livermore National Laboratory  
yang11@llnl.gov

## PP108

### Minisymposium: Accelerating Application Software Development Through Software Productivity and Sustainability Plans (PSPs)

Scientific software is playing an increasingly important role in both accelerating scientific discovery, and developing a predictive understanding of complex systems needed to inform policy decisions. To support this role amidst the increasing complexity of the system models, and the disruptive changes in hardware and software, improvements in software development productivity and sustainability are needed. To address this critical need, the Office of Biological and Environmental Research (BER) has begun exploring the development and use of Software Productivity and Sustainability Plans (PSPs) within its Subsurface Biogeochemical Research Program and the Interoperable Development of Extreme-scale Software Applications (IDEAS) project. Guidance was developed for the content of a complete software PSP that identifies the critical aspects of the software development process and lifecycle that must be documented, but not how a PSP should be implemented. The application codes used in the IDEAS project have developed software PSPs and are using them to guide improvements in their software development practices, and to develop metrics to quantify this improvement. In this poster, we outline a typical software PSP; highlight implementation content from specific applications; and identify where these plans have led to improved or enhanced practices, such as formal code reviews, new version control workflows, and continuous integration with automated testing.

#### David Moulton

Los Alamos National Laboratory  
Applied Mathematics and Plasma Physics  
moulton@lanl.gov

Jeff Johnson  
Lawrence Berkeley National Laboratory  
jnjohnson@lbl.gov

Ethan T. Coon  
Los Alamos National Laboratory  
ecoon@lanl.gov

Carl Steefel, Sergi Molins  
Lawrence Berkeley National Laboratory  
cisteefel@lbl.gov, smolins@lbl.gov

Glenn Hammond  
Sandia National Laboratories  
gehammo@sandia.gov

Reed M. Maxwell  
Department of Geology and Geological Engineering and IGWMC  
Colorado School of Mines  
rmaxwell@mines.edu

Carol S. Woodward  
Lawrence Livermore Nat'l Lab  
woodward6@llnl.gov

Laura Condon  
Syracuse University  
lecondon@syr.edu

## PP108

### Minisymposium: Sirepo: An Open Source Platform for Portable Reproducible Simulations

Sirepo is a novel system for managing, sharing, and executing simulations. We use software containers as an efficient mechanism for reproducible execution environments. We have packaged seven open source simulation codes for the particle accelerator and high-energy light source community into open source Docker ([rsl.link/cse17/1](https://rsl.link/cse17/1)) and Vagrant ([rsl.link/cse17/2](https://rsl.link/cse17/2)) images. We have created a modern HTML5 and Javascript user interface to configure, execute, and render simulations for three accelerator codes ([rsl.link/cse17/3](https://rsl.link/cse17/3)). We also provide a command line interface to all seven codes via a JupyterHub cluster ([rsl.link/cse17/4](https://rsl.link/cse17/4)). To facilitate sharing, we are developing a self-extracting simulation archive format that runs in any browser and allows users to rerun simulations on our cloud-based cluster or a locally installed instance of Sirepo.

Robert Nagler  
Radiasoft LLC  
nagler@radiasoft.net

David Bruhwiler  
Radiasoft LLC  
bruhwiler@radiasoft.net

Paul Moeller  
Bivio Software, Inc.  
moeller@bivio.biz

## PP108

### Minisymposium: The Journal of Open Source

**Software**

This poster describes the motivation and progress of the Journal of Open Source Software (JOSS), a free, open-access journal designed to publish brief papers about research software. The primary purpose of JOSS is to enable developers of research software to receive citation credit equivalent to typical archival publications. JOSS papers are deliberately short, and are required to include a short abstract describing the purpose and functionality of the software, authors and their affiliations, and key references, as well as link to an archived version of the software (e.g., DOI obtained from Zenodo). Upon acceptance, papers receive a CrossRef DOI. Rather than a review of a lengthy software paper (including, e.g., methodology, validation, sample results), JOSS submissions undergo rigorous peer review of both the abstract and software itself, including documentation, tests, continuous integration, and licensing. The JOSS review process is modeled on the established approach of the rOpenSci collaboration. The entire submission and review process occurs openly on GitHub; papers not yet accepted remain visible and under review until the authors make appropriate changes for acceptance—unlike other journals, papers requiring major revision are not rejected. Since its public release in May 2016, JOSS has published 26 accepted papers as of September 2016, with an additional 20 submitted and under review.

Arfon Smith  
GitHub, Inc.  
arfon@github.com

Lorena Barba  
George Washington University  
labarba@email.gwu.edu

George Githinji  
KEMRI-Wellcome Trust Research Programme  
githinji@kemri-wellcome.org

Melissa Gymrek  
University of California, San Diego  
gymrek@broadinstitute.org

Kathryn Huff  
University of Illinois at Urbana-Champaign  
kdhuff@illinois.edu

Daniel S. Katz  
University of Illinois at Urbana-Champaign  
Daniel.S.Katz@jpl.nasa.gov

Christopher Madan  
Boston College  
christopher.madan@bc.edu

Abigail Cabunoc Mayes  
Mozilla Science Lab  
abigail@mozillafoundation.org

Kevin Moerman  
MIT Media Lab  
kmoerman@mit.edu

Kyle E. Niemeyer  
Oregon State University  
Kyle.Niemeyer@oregonstate.edu

Pjotr Prins  
University of Tennessee  
pjotr.public01@thebird.nl

Karthik Ram  
University of California Berkeley  
karthik.ram@gmail.com

Tracy K. Teal  
Data Carpentry  
tkteal@datacarpentry.org

Jake Vanderplas  
University of Washington  
jakevdp@cs.washington.edu

**PP108****Minisymposium: Software Citation Principles for Credit and Reuse**

In this poster, we present software citation as a method to ensure credit for and identification, discovery, and reuse of software in scientific and engineering research. This poster documents efforts on developing software citation principles by the FORCE11 Software Citation Working Group and WSSSPE working group on software credit. We discuss the six principles of software citation, related work, and key challenges/research directions, including suggestions for metadata necessary for software citation and some examples.

Kyle E. Niemeyer  
Oregon State University  
Kyle.Niemeyer@oregonstate.edu

Arfon Smith  
GitHub, Inc.  
arfon@github.com

Daniel S. Katz  
University of Illinois at Urbana-Champaign  
Daniel.S.Katz@jpl.nasa.gov

**PP108****Minisymposium: Developing Fast Code Through High-Level Annotations**

As scientific libraries prepare for the next generation of machines and architectures, there is pressure on both application programmers and algorithm developers to express a variety of optimizations in the code. This presents a number of challenges when considering multiple programming paradigms and legacy code reuse, from performance to correctness. In this poster an annotation framework is presented, called the Illinois Code Environment (ICE), that facilitates the use of optimized code segments. Several tools are shown that automate the optimization process when used with the ICE annotation framework. This includes Moya, a just-in-time compilation (JIT) package;

hand-tuned kernels for specific architectures; and loop optimization for vectorized efficiency. The overall structure of ICE provides an entry point for a number of potential automation tools, while allowing the development of a stable, tested track in the code base. The annotation system is detailed for several small examples as well as a full simulation code, called PlasComCM, which is in use in the PSAAP2 Center for the Exascale Simulation of Plasma-Coupled Combustion at the University of Illinois at Urbana-Champaign.

Luke Olson

Department of Computer Science  
University of Illinois at Urbana-Champaign  
lukeo@illinois.edu

Bill Gropp, Dan Bodony, Jon Freund, Thiago Teixeira, Tarun Prabhu, Carl Evans, Sam White, Simon Garcia, Cory Mikida, Michael Anderson, Matthew Smith, John Larson

University of Illinois at Urbana-Champaign  
wgropp@illinois.edu, bodony@illinois.edu,  
jbfreund@illinois.edu, tteixei2@illinois.edu,  
tprabhu2@illinois.edu, gcevans@illinois.edu,  
white67@illinois.edu, gredgnz2@illinois.edu,  
cmikida2@illinois.edu, manders2@illinois.edu, mj-smith6@illinois.edu, jlarson1@illinois.edu

**PP108**

**Minisymposium: Jupyter Notebooks Facilitating Productivity, Sustainability, and Accessibility of Data Science**

Jupyter notebooks provide a document-centric web-based environment for data science and education. The web-based platform allows Jupyter to be used as a local desktop application or portal to large scale computing resources. The document aspect facilitates recording, sharing, and reproducing analyses. The integration of notebooks into existing CSE research and education workflows can improve reproducible workflows and communication and collaboration. The interactive nature of the notebook environment enables productive exploration of libraries and analyses, and being based on web technologies allows the same productivity to extend even to large-scale computing resources, which often pose a challenge to productivity and accessibility. The document nature of Jupyter notebooks facilitates preserving and sharing computations and results. Being a free, public, open source project enables equal access to all students and researchers, and the web environment is more familiar and accessible to a much wider set of current and prospective researchers than traditional terminal environments. Due to Jupyter's language agnostic protocol, with support from dozens of programming languages, notebooks can be used in many scientific and mathematical domains, which can have diverse language preferences.

Min Ragan-Kelley

Simula Research Laboratory  
IPython, Jupyter  
benjaminrk@simula.no

Jupyter Development Team

Project Jupyter  
projectjupyter@gmail.com

**PP108**

**Minisymposium: GraphFlow: Workflow-Based Big Data Processing**

In recent years, analysis of large-scale graphs has become an integral research component in a wide variety of disciplines including bioinformatics, social sciences, and epidemiology. Unfortunately, in practice, there is a significant gap between the services provided by the graph mining toolkits and the actual needs of the domain experts. Most graph mining packages only provide a set of algorithms that can be used as a black box. However, domain experts expect to construct complex workflows that involve executing several algorithms over a large amount of data. Current graph mining frameworks do not provide sufficient support for creating, reusing, and extending such complex workflows. Moreover, mining big graph data requires expertise in the area of high-performance computation and parallel processing, which is not trivial for a significant group of data scientists. We introduce GraphFlow, a big data processing toolkit that encapsulates details in order to encode complex research experiments as a set of high-level diagrams. GraphFlow uses the Spark big data processing platform, which is backed by an active community and supporters, as the backend engine and offers a set of components for big graph processing and machine learning through the Galaxy workflow management system. GraphFlow introduces a novel interaction model among its components that enables the users to construct complex workflows. We demonstrate the expressiveness of GraphFlow through several case studies.

Sara Riazzi

Department of Computer and Information Science  
University of Oregon  
riazi@uoregon.edu

Boyana Norris  
University of Oregon  
norris@cs.uoregon.edu

**PP108**

**Minisymposium: Maximizing Usability and Performance in Numerical Software Packages**

As the number of high-performance numerical software packages grows, larger and more complex computational problems can be solved by scientists and engineers. This rapid growth in capabilities, however, makes the task of applying the best solution method to a given problem increasingly challenging because of the number and complexity of solutions of available. In our work we focus on the usability challenges the current implementations of numerical software pose. First, to select a suitable numerical toolkit, a user must read substantial documentation, attend tutorials, seek online help and typically engage in trial and error learning. Using these packages may also require significant background in software engineering, numerical analysis and other domains which may be outside the users area of expertise. Second, once the toolkit has

been chosen and installation hurdles overcome, the complexity involved in using them effectively can be huge. For example, one may have to choose between hundreds of valid solution methods, each of which must be configured to test with a particular application. To conclude, even complex high-performance numerical software should be simple, user-friendly, efficient, adaptable and reusable. We present example ways in which usability and performance can be improved through organizing software into searchable taxonomies, annotated with performance information with use cases for PETSc and Trilinos.

Kanika Sood

University Of Oregon  
1991  
kanikas@cs.uoregon.edu

Boyana Norris  
University of Oregon  
norris@cs.uoregon.edu

Elizabeth Jessup, Pate Motter  
University Of Colorado  
jessup@colorado.edu, pate.motter@colorado.edu

**PP108**

**Minisymposium: World SpatioTemporal Analytics and Mapping Project (WSTAMP): Cloud Implementation of Open Source Algorithms and Data Stores for Sustainable, Scalable Analysis of Space-Time Data**

Application of spatiotemporal (ST) analytics to integrated ST data from major sources such as the World Bank and United Nations holds tremendous potential for shedding light on the evolution of global socio-economic and geopolitical landscapes. In response to this opportunity, an enterprise data science capability comprised of a visually rich, intuitive web front end powered by a scalable technology stack for exploration, analysis, and mining of ST data known as World STAMP was developed. The challenge for WSTAMP architects was to implement an open-source solution with long term sustainability, portability, and scalability. Adoption of major open resources such R, PostgreSQL, D3.js, Apache Solr, and Docker tie WSTAMP to rapidly emerging capabilities in algorithms, data storage, and visualization. Architecting these capabilities as Analytics as a Service and Data as a Service within a cloud computing framework means that custom tailored solutions can be introduced without major revisions to the workflow. This portable architecture can be migrated to and repurposed in new cloud computing domains where entirely new algorithms and data may be required. The current status capabilities, architecture, interface are presented here with thoughts on future directions.

Robert N. Stewart

Oak Ridge National Laboratory  
Oak Ridge National Laboratory  
stewartrn@ornl.gov

Aaron Myers, David Axley  
Oak Ridge National Laboratory  
myersat@ornl.gov, axleyde@ornl.gov

Alex Sorokine  
Oak Ridge National Laboratory  
sorokina@ornl.gov

Jesse Piburn  
Oak Ridge National Laboratory  
piburnjo@ornl.gov

**PP108**

**Minisymposium: Numfocus Sustainability Project**

Open Source infrastructure is at the foundation of much industry and scientific software and research efforts. This infrastructure includes the tools that help developers and researchers build software and conduct analyses - programming languages, libraries, frameworks and even training materials. These are projects that have great value now; they're being used on a daily basis by thousands of people, but it's not clear how to determine or find the necessary resources to support and sustain Open Source Projects. In this poster we describe the NumFOCUS Sustainability Project. This project aims to connect the NumFOCUS Projects to each other to jointly develop and share information on sustainability strategies, connect Project leads with people with relevant expertise and networks, provide training on business and financial planning, marketing strategies and effective communication, develop and disseminate an Open Source Projects Guide to fundraising and project management and support infrastructure that would help the Projects more effectively manage finances and client and business relationships.

Andy R. Terrel  
Continuum Analytics  
andy.terrel@gmail.com

Tracy K. Teal  
Data Carpentry  
tkteal@datacarpentry.org

**PP108**

**Minisymposium: Practical Approaches to Improve Program Understanding and Software Productivity of Scientific Code**

Software complexity has become a real barrier that impedes scientific software development, such as adding new features and functions, validating domain knowledge incorporated in the software systems, as well as redesigning and refactoring code for emerging computational platforms (i.e., exascale computers). In this poster, we analyze several environmental software systems and demonstrate software design and refactoring challenges encountered by scientific community. We also summarize practices that leverage compiler technologies for better program understanding and enhanced software productivity. These technologies have been used to (1) understand existing scientific models, (2) modularize complex code, (3) provide instrumentation mechanisms to improve code portability, (4) generate functional testing unit for key software modules, and (5) add new features for model validation. These activities are supported by the Accelerated Climate Modeling

for Energy (ACME) project and the Interoperable Design of Extreme-scale Application Software (IDEAS) project. We believe our experience can benefit broader scientific communities that are facing the challenge of complex code.

Hindmarsh  
Lawrence Livermore National Laboratory  
banks12@llnl.gov, gardner48@llnl.gov, peles2@llnl.gov,  
hindmarsh@llnl.gov

Dali Wang  
Oak Ridge National Laboratory  
wangd@ornl.gov

Yu Pei  
University of Tennessee  
peiyo@ornl.gov

Youngsung Kim  
National Center for Atmospheric Research  
youngsung@ncar.edu

Oscar Hernandez  
Oak Ridge National Lab  
oscar@ornl.gov

David E. Bernholdt  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
bernholdtde@ornl.gov

## PP108

### **Minisymposium: Software Productivity Strategies for the Sundials Suite of Time Integrators and Nonlinear Solvers**

The SUNDIALS Suite of Nonlinear Differential-Algebraic Solvers and Integrators is an open source software library including highly robust and adaptive time integration methods for ODEs and DAEs as well as robust nonlinear solvers. SUNDIALS has a long history of user deployment with over 3,500 downloads worldwide from the LLNL site yearly. Supporting the SUNDIALS suite within the DOE laboratory environment poses certain challenges, including limited control of software environment, testing infrastructure, and web presence, restricted releasability of code in development, and limited funding for user interactions. To address these challenges, the SUNDIALS team has adopted a number of sustainability strategies. In this poster, we will discuss these strategies, continuing issues, and future plans for strengthening SUNDIALS sustainability. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC. LLNL-ABS-702579.

Carol S. Woodward  
Lawrence Livermore Nat'l Lab  
woodward6@llnl.gov

Daniel R. Reynolds  
Southern Methodist University  
Mathematics  
reynolds@smu.edu

L. Eddy Banks, David J. Gardner, Slaven Peles, Alan