



*Searchable  
Abstracts  
Document*

Conference on  
**Parallel Processing for  
Scientific Computing**

February 23 – 26, 2022

*This document was current as of February 18, 2022. Abstracts appear as submitted.*



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

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

**IP1****Welcome Remarks and Presentation: Julia: The Power of Language as an Alternative to MPI**

When I was a graduate student, Jack Dongarra presented MPI at a big meeting, I think at Oak Ridge. I remember raising my hand and asked whether MPI would just be for some people, and there would be an easier approach for everyone else. As I recall, everybody nodded sure, yes, that's right. Somehow that did not happen. We very much believe this hard problem has stood in the way of progress in high performance computing. We believe that Julia is part of the solution to this problem.

Alan Edelman

Department of Mathematics  
Massachusetts Institute of Technology and Julia  
Computing  
edelman@mit.edu

**IP2****Life after Fugaku—What Have we Learned, and How do we Proceed with End of Moore's Law Approaching?**

Circa 2021, Fugaku is the first of the many of the so-called 'exascale' machines that had been under planning across the globe since the end of the 2000s, and have not only exhibited supreme performance and scalability in traditional benchmarks e.g., Linpack, driving nearly 8 million high-performance Arm cores, but also, have demonstrated such in other benchmarks as new breed of applications such as those involving big data and AI. However, the next step after exascale still remains unclear, despite that there will be several by the mid 2020s, and R&D for the next generation should already have started. This is due to the end of Moore's law approaching, and traditional method of increasing the number of cores would cease to be effective, due to power and cost limitations. At R-CCS, in collaboration with other partners we are investigating innovative means of increasing performance beyond Fugaku, based on the data obtained from Fugaku as well as other exascale machines, and undertaking a series of projects to push the supercomputing technology forward towards the 2030s.

Satoshi Matsuoka

Tokyo Institute of Technology  
matsu@is.titech.ac.jp

**IP3****Software as an Instrument of Science**

Ever since high-performance computing has enabled exploration of the world hitherto impossible, scientific software has been a de-facto instrument of science, but without gaining the necessary recognition as such. A culture grew where the next scientific deadline drove development, often at the cost of quality and reusability of software. The care that experimentalists expend on their instruments was almost never accorded by computational scientists to their software. This laissez faire approach to scientific software development also extended to computational experiments, where provenance and reproducibility were often an afterthought. Two external factors are causing the community to re-examine this culture. Firstly, codes and platforms are simultaneously growing in complexity. Secondly, some high visibility failures have placed a spotlight on policy and decision-making dependent upon scientific software. We have now entered an era where we are re-

examining our assumptions about software and scrutinizing our sustainability and quality approaches for efficiency and efficacy. This is an opportunity for the high-performance computing community to apply lessons learned from the missteps of the past and fully recognize software as the scientific instrument it is. In this presentation I will describe my experience and wisdom gained as a software architect of a high-performance multiphysics community code. The code, FLASH, began as an astrophysics code, and then went on to serve several additional communities as diverse as solar physics, cosmology, computational fluid dynamics, and laser plasma experiments. A new code Flash-X, derived from FLASH and incorporating our acquired wisdom, is being developed under the Exascale Computing Project to continue to serve most of these communities on heterogeneous platforms. I will describe my take on the blueprint of what our future software endeavors could be if we are intentional about software quality and sustainability, and how we are applying it to Flash-X.

Anshu Dubey

Argonne National Laboratory  
adubey@anl.gov

**IP4****A Quantum Future of Computation**

Still in early development, quantum computing is already overturning our contemporary notions of computational methods and devices. Using new concepts of computing based in quantum physics, quantum computers will be able to solve certain problems that are completely intractable on any imaginable classical computer, such as accurate simulations of molecules and materials, or breaking public key encryption. While this potential is real, quantum computers are best viewed as special purpose accelerators for specific problem classes. In an effort to bring clarity to the fast-growing field of quantum computing, I will describe the hardware and software architecture of quantum computers and discuss how they differ from conventional classical high performance computers. Based on this, I will also attempt to dispel myths and hype surrounding the field and present a realistic assessment of the potential of these devices and the specific application areas on which they are expected to have a large impact. I will end by showing that quantum computing already generates value today, through quantum inspired approaches. These are quantum approaches implemented on classical hardware that outperform the state of the art of classical methods known before, with applications in health care, logistics, chemistry and other areas.

Matthias Troyer

Microsoft, U.S.  
mtroyer@microsoft.com

**IP5****Nonnegativity Constrained Low Rank Approximations for Scalable Data Analytics on Distributed Memory Parallel Environment**

Constrained Low Rank Approximation (CLRA) is a powerful foundation for important data analytic tasks such as topic modeling and community detection. Some advantages of CLRA include scalable algorithms and software based on advances in numerical linear algebra and parallel computing. Nonnegativity constraints allow more judicious formulation, interpretable results, and effective methods whether the input data is in feature-data relationship as in

the Nonnegative Matrix Factorization (NMF) or data-data relationship as in spectral clustering or symmetric NMF (SymNMF). A common foundation of CLRA also allows a hybrid method for information fusion called JointNMF from merging the objective functions of the NMF and SymNMF for multi-view data sets with both content and connection information. In this talk, a distributed memory parallel algorithm and software framework, PLANC (Parallel Low-rank Approximation with Nonnegativity Constraints), is described for nonnegativity constrained matrix and tensor low rank approximations. PLANC uses parallel distributions and algorithms designed to optimize communication and computation, allows extension regarding data characteristics, algorithm, constraints, and architecture beyond the above mentioned NMF variants. Efficiency and scalability of PLANC and some new knowledge discoveries from large real-world data sets are presented.

Haesun Park

Georgia Institute of Technology  
hpark@cc.gatech.edu

## IP6

### Supercomputing, Cloud and Edge Security Models Convergence using APIs and Container Technologies

Scientific applications and workflows today are exploiting a range of distributed resources including supercomputing, cloud, and edge technologies, such as scientific instruments. User access and security models across these resource domains vary considerably resulting in overhead for the development of scientific applications and workflow engines, often leading to productivity losses. Recently, adoption of zero trust security models have accelerated due to prevalent usage of off-premises cloud and edge technologies. The model relies on having multiple layers of trust and verification for applications, users, devices, and networks. HPC applications, networks, storage, and access controls have unique architectural features for delivering close-to-metal performance and to achieve scalability. In this talk, I introduce motivating use cases, discuss challenges in adopting zero trust security models for HPC, and present solutions that we have developed at the Swiss National Supercomputing Centre. These include a RESTful services gateway called FirecREST and a user-friendly, HPC container engine called Sarus. FirecREST serves as a secure, web-accessing interface to HPC resources providing a coupling layer between the identity and access management of cloud and the essential POSIX interfaces used in HPC. Sarus fulfills unique HPC applications portability requirements offering native hardware performance while improving isolation and security in multi-tenant environments.

Sadaf Alam

CSCS  
alam@cscs.ch

## IP7

### Graph Neural Network Research at AWS AI

In the course of just a few years, Graph Neural Networks (GNNs) have emerged as the prominent supervised learning approach that brings the power of deep representation learning to graph and relational data. An ever-growing body of research has shown that GNNs achieve state-of-the-art performance for problems such as link prediction, fraud detection, target-ligand binding activity prediction, knowledge-graph completion, and product recommenda-

tions. As a result, GNNs are quickly moving from the realm of academic research involving small graphs to powering commercial applications and very large graphs. This talk will provide an overview of some of the research that AWS AI has been doing to facilitate this transition, which includes developing the Deep Graph Library (DGL) an open source framework for writing and training GNN-based models, improving the computational efficiency and scaling of GNN model training for extremely large graphs, developing novel GNN-based solutions for different applications, and making it easy for developers to train and use GNN models by integrating graph-based ML techniques in graph databases.

George Karypis

Amazon  
University of Minnesota Twin Cities  
gkarypis@amazon.com

## SP1

### SIAG Best Paper Prize: Optimizing Communication-Avoiding Sparse Lu Factorization on Multi-Gpu Clusters

We present a highly optimized implementation of the communication-avoiding sparse LU factorization algorithm, specifically targeting pre-exascale multi-GPU clusters such as Summit Supercomputer at Oak Ridge National Laboratory. Prior to this work, distributed memory sparse LU factorization used GPUs mostly as a co-processor because of the relatively smaller DRAM capacity available and limited hardware support for GPU-aware message passing on older GPUs. The current pre-exascale multi-GPU clusters have relatively higher DRAM capacity and hardware support for GPU-aware MPI that allows performing the entire sparse LU factorization on GPU. The challenge is, sparse LU factorization consists of many operations on small and irregular size operands, which makes it difficult to effectively use GPU during all phases on sparse LU factorization. To overcome such challenges, we (a) redesigned the data structure to reduce the cost of index-algebra on GPUs; b) combined streams with the so-called tree parallelism to schedule multiple operations, and; c) exploit high bandwidth GPU-2-GPU and GPU-aware MPI with the *look-ahead* factorization techniques to effectively overlap communication with computation. Our proposed optimizations improve the performance of communication-avoiding sparse LU factorization by up to 3× over offload based GPU acceleration of the same algorithm on single and multiple node configurations on the Summit supercomputer.

Piyush Sao

Oak Ridge National Laboratory  
saopk@ornl.gov

XiaoYe Li

Lawrence Berkeley National Laboratory  
xsli@lbl.gov

Richard Vuduc

Georgia Institute of Technology  
richie@cc.gatech.edu

## SP2

### SIAG/Supercomputing Career Prize: Scaling in Space and Time

My career has centered on developing algorithms and

software that scale with the next generation of high-performance computers and high-fidelity simulations. As such, multigrid methods have been a major research focus since they have the potential to exhibit optimal-order parallel and computational complexity. In this talk, I will discuss the role of multigrid methods in scientific computing, which have served mainly as fast solvers for ill-conditioned spatial systems and are now being developed in other areas such as parallel time integration, an approaching paradigm shift for computing.

Robert D. Falgout  
Lawrence Livermore National Laboratory  
falgot2@llnl.gov

### CP1

#### Implementation and Scalability of a High-Order Immersed Interface Method on HPC Architectures

Immersed interface methods locally modify a finite difference scheme to account for discontinuities in the solution of a PDE while preserving the accuracy of the initial scheme. They are particularly useful in simulating interface-coupled multiphysics problems such as fluid-structure interaction and multiphase flow. However, the local modifications to the discretization widen the stencil of the finite difference scheme, introducing additional communication requirements in a distributed-memory environment. In this work we identify the specific communication requirements of high-order immersed interface methods and propose a scalable approach for managing this communication on block-structured Cartesian grids. We then demonstrate its effectiveness through fourth-order accurate simulations of PDE-based interface problems on adaptive multiresolution grids.

James Gabbard, Thomas Gillis, Wim van Rees  
MIT  
jgabbard@mit.edu, tgillis@mit.edu, wvanrees@mit.edu

### CP1

#### Openmp Target Task: Tasking and Target Offloading for Heterogeneous Linear Algebra Implementations

In this talk, it is evaluated the use of OpenMP tasking with target GPU offloading as a potential solution for the implementation of heterogeneous linear algebra algorithms. Also, it is proposed a new OpenMP specification to make the implementation of heterogeneous codes simpler by using OpenMP target task, which integrates both OpenMP tasking and target GPU offloading in a single OpenMP pragma. As a test case, we use one of the most popular and widely used Basic Linear Algebra Subprogram Level-3 routines: triangular solver (TRSM). To benefit from the heterogeneity of the current high-performance computing systems, we propose a different parallelization of the algorithm by using a nonuniform decomposition of the matrices to be computed. We evaluate the use of target GPU offloading inside OpenMP tasks to address the heterogeneity found in the hardware. This new approach can outperform the state-of-the-art algorithms, which use a uniform decomposition of the data, on both the CPU-only and hybrid CPU-GPU systems, reaching speedups of up to one order of magnitude. The performance that this approach achieves is faster than the IBM ESSL math library on CPU and competitive relative to a highly optimized heterogeneous CUDA version. One node of Oak Ridge National Laboratory's supercomputer, Summit, was used for performance

analysis.

Pedro Valero-Lara  
Oak Ridge National Laboratory  
valerolarap@ornl.gov

### CP1

#### Exascale-Enabled Physical Modeling for Post-Moore Devices

We have developed scalable simulation tools to enable leadership computing systems to model emerging post-CMOS microelectronic devices (electronic, spintronic, nanomagnetic and nanomechanical). Our open-source code, ARTEMIS (Adaptive mesh Refinement Time-domain ElectrodynaMics Solver), contains support for dispersive material properties, user-defined excitations and boundary conditions, and heterogeneous physical coupling that can be present in next-generation devices. The basic module of ARTEMIS solves Maxwells equations for the time-evolving electric and magnetic fields using a finite-difference time-domain approach on structured Cartesian grids. ARTEMIS leverages the developments of two Exascale Computing Projects: (1) the AMReX software libraries for scalable block-structured adaptive mesh refinement applications and (2) the WarpX electromagnetic Particle-In-Cell code. Thus, it is portable and scales well on many-core/GPU-based supercomputers that are far beyond the reach of commercial tools, allowing us to capture larger-scale spatial disparities inherent to realistic circuits. We have demonstrated algorithmic flexibility by developing a micromagnetics module, making it suitable for nonlinear spintronic applications. Our current efforts include upgrading the functionality of ARTEMIS to accurately describe the working mechanism of new devices such as a magnetoelectric spin-orbit logic, enabling ARTEMIS to serve as the device-level design and optimization tool.

Zhi Jackie Yao, Revathi Jambunathan, Andrew Nonaka  
Lawrence Berkeley National Laboratory  
jackie\_zhiyao@lbl.gov, rjambunathan@lbl.gov, ajnonaka@lbl.gov

### CP2

#### Predictions of Compressible Steady and Unsteady Fluid Flow using PointNet Deep Learning Backbone

Computational fluid dynamics (CFD) is known as the analytic process of predicting fluid flow, mass transfer, and other related phenomena during the design or manufacturing process. Aggressive use of CFD is credited with drastic reductions in wind tunnel time and lower numbers of experimental rig tests. CFD has saved millions of dollars for industries, governments, and national laboratories. CFD offers the potential of delivering an insight into the critical physical phenomena limiting component performance, thus opening new frontiers in multiple fields, especially vehicle design. One of the key strengths of CFD is its ability to produce simulations (forward steps) that can be used in an inverse design. However, each forward step can be very time-consuming to converge. To enable more efficient and scalable CFD forward steps, we leverage the universal approximation property of deep neural networks (DNNs) to estimate a surrogate to the CFD simulation. We pass the point cloud of a mesh through several layers of shared multi-layer perceptrons. The network learns an end-to-end mapping between spatial positions and CFD quantities. After training, predictions are instantaneous on

new objects with arbitrary geometries and a selected number of time steps. We train the neural network on compressible laminar steady and unsteady flows past various shapes. The network predicts the flow fields hundreds of times faster than our conventional CFD solver while maintaining good accuracy.

Shyma Y. Alhuwaidar

KAUST

Masters Student

shyma.alhuwaidar@kaust.edu.sa

Mohammed Sayyari, Matteo Parsani

King Abdullah University of Science and Technology (KAUST)

mohammed.alsayyari@kaust.edu.sa,

mat-

teo.parsani@kaust.edu.sa

## CP2

### Parallel Domain Decomposition Techniques Applied to Multivariate Functional Approximation of Discrete Data

Compactly expressing large-scale datasets through Multivariate Functional Approximations (MFA) can be critically important for analysis and visualization to drive scientific discovery. We present scalable domain partitioning approaches to compute MFA representations, by reducing the total work per task through domain-decomposition strategies utilizing Schwarz-type iterative schemes for converging the interface data. For the underlying MFA, we utilize a tensor expansion of non-uniform B-spline (NURBS) basis in multiple dimensions to reduce the functional approximation error in the input data. Previous work on adaptive NURBS-based MFA rely primarily on post-processing techniques to blend discontinuities across subdomains boundaries after computing the MFA. In contrast, we explore a robust constrained parallel solution infrastructure to impose higher-order continuity directly on the MFA representation computed. We demonstrate effectiveness of the approach with an overlapping Restricted-type Additive Schwarz method (RASM), with a l-BFGS subdomain solver to minimize the local error residuals, and more specifically to recover continuity across non-matching boundaries. The analysis of the presented scheme for analytical and scientific datasets in 1-D and 2-D are also presented. Additionally, scalability studies are also shown for several analytical and realistic, scientific 1d and 2d datasets to evaluate the parallel speedup of the algorithm on large computing clusters.

Vijay S. Mahadevan, David Lenz, Iulian Grindeanu,

Thomas Peterka

Argonne National Laboratory

mahadevan@anl.gov, dlennz@anl.gov, iulian@anl.gov,

tpeterka@mcs.anl.gov

## CP2

### Computational Investigation of the Effect of Chemistry on Mars Retropropulsion Environments Using a Massively Parallel GPU Approach

In this work, we investigate the effects of chemistry on a human-scale Mars lander concept using scale-resolving computational fluid dynamics (CFD) with finite-rate chemistry and a graphics processing unit (GPU)-enabled implementation of the NASA FUN3D flow solver, enabling run-times of a few days for the simulations presented. Simulations are carried out on Summit at Oak Ridge Leader-

ship Computing Facility using thousands of GPUs. Retropropulsion ground tests require significant compromises on physical scale, instrumentation, configuration, and environments. Ground tests of retropropulsion configurations thus far have neglected effects of chemistry due to physical constraints of wind tunnel models and facilities; most experiments use inert simulant gases at low temperatures. As such, a strong reliance on high-fidelity computational analyses such as those presented in this work is required to expand the knowledge of retropropulsion aerodynamics. An overview of the GPU approach will be presented. Results are compared to a previous scaled perfect gas (air) campaign.

Gabriel Nastac

NASA

gabriel.c.nastac@nasa.gov

Aaron Walden, Eric Nielsen, Ashley Korzun

NASA Langley

aaron.c.walden@nasa.gov, eric.j.nielsen@nasa.gov,

ashley.m.korzun@nasa.gov

Christopher Stone

HPCMP PETTT ACE On-site (Contractor)

christopher.stone@nianet.org

Patrick Moran

NASA Ames

patrick.moran@nasa.gov

## CP2

### Performance Optimization Methods for a Memory-Bound, Unstructured CFD Application on Massively Parallel GPU Platforms

Computational performance of the FUN3D unstructured-grid computational fluid dynamics (CFD) application on massively parallel GPU environments is memory-bound and highly dependent upon efficient reads from and atomic updates to the irregular cell-, edge-, and node-based data structures. In this talk, we present recent efforts into optimizing select performance-critical kernels on NVIDIA Tesla V100 and A100 GPUs and AMD CDNA MI100 GPUs. A novel use of L2 cache residency controls and asynchronous loads into on-chip shared memory are explored on the A100 GPU for the sparse iterative solver, which is dominated by mixed-precision, sparse matrix vector multiplication. Demonstrations show that these methods improve global memory bandwidth utilization by 13.5% on the A100 GPU. Several techniques are also presented that use registers and/or shared memory to facilitate array transposition and aggregation which combine to reduce the frequency and increase the cache efficiency of floating-point atomic updates to the irregular data structures. These methods are demonstrated to improve the kernel throughput by nearly 500% on select kernels on the AMD MI100 over atomic updates directly to global memory. Overall, both V100 and A100 GPUs outperformed the MI100 GPU on kernels dominated by double-precision atomic updates; however, the techniques demonstrated here reduced the performance gap and improved the MI100 performance.

Christopher Stone

HPCMP PETTT ACE On-site (Contractor)

christopher.stone@nianet.org

Aaron Walden

NASA Langley

aaron.c.walden@nasa.gov

Mohammad Zubair  
Old Dominion University  
zubair@cs.odu.edu

Eric Nielson  
NASA Langley Research Center  
eric.j.nielson@nasa.gov

Gabriel Nastac  
NASA  
gabriel.c.nastac@nasa.gov

Christopher Stone  
HPCMP PETTT ACE On-site (Contractor)  
christopher.stone@nianet.org

### CP3

#### Prediction of Optimal Solvers for Sparse Linear Systems using Deep Learning

Yannick Funk  
Karlsruhe Institut für Technologie, Germany  
yannickfunk@yahoo.de

### CP3

#### Deep Learning and Spectral Embedding for Graph Partitioning

We present a graph bisection and partitioning algorithm based on graph neural networks. For each node in the graph, the network outputs probabilities for each of the partitions. The graph neural network consists of two modules: an embedding phase and a partitioning phase. The embedding phase is trained first by minimizing a loss function inspired by spectral graph theory. The partitioning module is trained through a loss function that corresponds to the expected value of the normalized cut. Both parts of the neural network rely on SAGE convolutional layers and graph coarsening using heavy edge matching. The multilevel structure of the neural network is inspired by the multigrid algorithm. Our approach generalizes very well to bigger graphs and has partition quality comparable to METIS, Scotch and spectral partitioning, with shorter run-time compared to METIS and spectral partitioning.

Alice Gatti  
Lawrence Berkeley National Laboratory  
agatti@lbl.gov

Zhixiong Hu  
University of California, Santa Cruz, U.S.  
zhu95@ucsc.edu

Tess Smidt  
Massachusetts Institute of Technology  
tsmidt@mit.edu

Esmond G. Ng  
Lawrence Berkeley National Laboratory  
engng@lbl.gov

Pieter Ghysels  
Lawrence Berkeley National Laboratory  
Computational Research Division

pghysels@lbl.gov

### CP3

#### Tuned Preconditioners for the Spectral Element Method on GPUs

The Poisson pressure solve resulting from the spectral element discretization of the incompressible Navier-Stokes equation requires fast, robust, and scalable preconditioning. In the current work, a parallel scaling study of Chebyshev-accelerated Schwarz and Jacobi preconditioning schemes is presented, with special focus on GPU architectures, such as OLCF's Summit. Convergence properties of the Chebyshev-accelerated schemes are compared with alternative methods, such as low-order preconditioners combined with algebraic multigrid. Performance and scalability results are presented for a variety of preconditioner and solver settings. The authors demonstrate that Chebyshev-accelerated-Schwarz methods provide a robust and effective smoothing strategy when using  $p$ -multigrid as a preconditioner in a Krylov-subspace projector. The variety of cases to be addressed, on a wide range of processor counts, suggests that performance can be enhanced by automated run-time selection of the preconditioner and associated parameters.

Malachi Phillips  
University of Illinois, Urbana-Champaign  
malachi2@illinois.edu

Stefan Kerkemeier  
Argonne National Laboratory  
kerkemeier@anl.gov

Paul Fischer  
University of Illinois at Urbana-Champaign  
fischerp@illinois.edu

### CP3

#### Gptuneband: Multi-Task and Multi-Fidelity Autotuning for Large-Scale High Performance Computing Applications

Autotuning aims at automatically finding code parameters that optimize the runtime, memory, communication cost, or accuracy for complex, black-box functions. We propose a novel multi-task and multi-fidelity autotuning framework, GPTuneBand, for tuning large-scale expensive high performance computing (HPC) applications. GPTuneBand combines a multi-task Bayesian optimization algorithm with a multi-armed bandit strategy, well-suited for tuning expensive HPC applications such as numerical libraries, scientific simulation codes and machine learning (ML) models, particularly with a very limited tuning budget. Our numerical results show that compared to other state-of-the-art autotuners, which only allows single-task or single-fidelity tuning, GPTuneBand obtains significantly better performance for numerical libraries and simulation codes, and competitive validation accuracy for training ML models. When tuning the Hypre library with 12 parameters, GPTuneBand wins over its single-fidelity predecessor GPTune on 62.5% tasks, with a maximum speedup of 1.2x, and wins over a single-task, multi-fidelity tuner BOHB on 72.5% tasks. When tuning the MFEM library on large numbers of CPU cores, GPTuneBand obtains a 1.7x speedup when compared with the default code parameters.

Xinran Zhu  
Cornell University

xz584@cornell.edu

#### CP4

##### Scalable GPU H2-Matrix Solvers

Hierarchical matrices allow for memory efficient representation of the data sparse matrices that often appear in scientific applications. The open source H2Opus library provides distributed CPU and GPU implementations of several key operations using the H2-variant of hierarchical matrices, where nested row and column bases allow for asymptotically optimal memory storage requirements. In this talk, we discuss the details of the Hierarchical Adaptive Randomized Approximation (HARA) algorithm in H2Opus, where an H2-matrix is constructed using only fast matrix-vector products in a top-down algorithm. We show how one of the most time consuming phases of the algorithm, the application of low rank updates followed by multiple compressions, can be improved using randomized sketching. We also discuss some initial details of a fully algebraic variant of recursive skeletonization factorization based on sketching for matrices constructed using strong admissibility.

Wajih Boukaram

Lawrence Berkeley National Lab  
wajih.boukaram@lbl.gov

#### CP4

##### Multifrontal Solver with Low-Rank Compression

Some of the major challenges to scale parallel applications are to keep the asymptotic complexity of computation and the memory consumption low. This talk presents a multifrontal solver that leverages hierarchical and non-hierarchical matrices, using low-rank compression methods, to reduce the asymptotic complexity of computation and memory usage. The multifrontal algorithm takes advantage of the low-rank property of the blocks appearing during the factorization of sparse linear systems with the utilization of a compression algorithm. Previously, the hierarchically off-diagonal butterfly (HODBF) format has been leveraged to accelerate compression and factorization in the multifrontal method. HODBF turns out to be very efficient for large-sized fronts. In addition, experiments have shown a reduced flop count for the non-hierarchical Block Low-Rank (BLR) compression of medium size fronts in comparison with hierarchical rank-structured techniques. In this talk, we discuss a hybrid version using the HODBF format in combination with the BLR format. In particular, we apply the HODBF compression to large sized fronts of the multifrontal method and apply the BLR compression to medium sized fronts. This leads to a reduction of the factorization time and memory consumption and allows us to solve larger problem sizes. We apply the factorization as a preconditioner for GMRES and illustrate the performance of this new preconditioner.

Lisa Claus

Lawrence Berkeley National Laboratory  
lclaus@lbl.gov

Pieter Ghysels

Lawrence Berkeley National Laboratory  
Computational Research Division  
pghysels@lbl.gov

Yang Liu, Xiaoye Sherry Li

Lawrence Berkeley National Laboratory  
liuyangzhuan@lbl.gov, xsli@lbl.gov

#### CP4

##### Comparison of a Randomized and a Deterministic Fixed-Precision Low-Rank Approximation Method for Sparse Matrices

Low-rank approximations of large sparse matrices are important in many scientific applications. Fixed-precision methods compute such approximations when a priori, the approximate rank  $k$  is unknown. A relevant question is the cost required for achieving a certain approximation quality for a large sparse matrix. We design and evaluate parallel implementations of a variant of the Randomized QB Factorization (RandQB\_EI) and the Truncated LU Factorization with Column and Row Tournament Pivoting (LU\_CRTP). We experimentally evaluate for different problems the computational cost required for achieving a certain approximation quality. Our results show that LU\_CRTP tends to be particularly competitive for low approximation quality. However, the runtime of LU\_CRTP highly depends on the amount of fill-in produced when updating the input matrix after each iteration, and is therefore outperformed by RandQB\_EI especially for higher approximation quality. Motivated by this observation, we developed a fill-in reducing parallel incomplete LU\_CRTP algorithm where small enough elements of the updated matrix are dropped after each iteration. We illustrate that this approach leads to speedups higher than 20 for some matrices.

Robert Ernstbrunner, Viktoria Mayer

University of Vienna  
robert.ernstbrunner@univie.ac.at,  
viktoria.mayer@univie.ac.at

Wilfried N. Gansterer

Department of Computer Science  
University of Vienna  
wilfried.gansterer@univie.ac.at

#### CP4

##### Block Low Rank and Hierarchical Low Rank Solvers for Fractional Diffusion Equations

Significant computational challenges are encountered when solving multidimensional space-fractional diffusion equations with spatially varying diffusivity or fractional order. The difficulties are due to both the kernel singularity in the fractional integral operator and the resulting dense discretized operators, which quickly become prohibitively expensive to handle because of their memory and arithmetic complexities. In this talk, we present two alternative low rank representation techniques to handle the dense matrix discretization. One is a block low rank representation that exploits the ability to approximate uniformly-sized blocks of the matrix by low rank factorizations. A Cholesky factorization solver operates directly on this representation using the low rank blocks as its atomic computational tiles, and achieves high performance on GPU hardware with  $O(N^{1.5})$  storage and  $O(N^2)$  arithmetic cost. In order to reach optimal complexities at scale, we present another multi-GPU distributed hierarchical matrix representation of the discretized operator, and an AMG-preconditioned Krylov solver that operates on it. We illustrate the scalability of the representation and of the solver on examples

up to 16M degree-of-freedom 2D problems on 64 GPUs.

George M. Turkiyyah  
KAUST  
George.Turkiyyah@kaust.edu.sa

Stefano Zampini  
King Abdullah University of Science and Technology  
stefano.zampini@kaust.edu.sa

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

## CP5

### Accelerating Coupled Climate Simulations using Efficient Remapping Techniques

Predictable simulation of coupled climate models require accurate and conservative transfer of scalar and flux fields defined on unstructured meshes that are often tuned to resolve individual component physics. In the Energy Exascale Earth System Model (E3SM), the Model Coupling Toolkit (MCT) coupler is used for transferring data between components, with the remapping weights computed through an offline pre-processing step. Recent efforts using the Mesh Oriented datABase (MOAB) to compute the remap operator during the simulation has removed several of the bottlenecks in such offline-online workflows. The use of the MOAB coupler has been proven to be efficient at scale (<https://doi.org/10.5194/gmd-13-2355-2020>) without sacrificing discretization accuracy during transfers. In this talk, we describe latest improvements in MOAB-based framework that has scalable intersection methods, use Zoltan mesh partitioners, fast searching methods from ArborX, accurate projections with TempestRemap library and that can be used on both multi-node manycore and multi-GPU hybrid architectures. We also present several performance and accuracy results to emphasize the feasibility of using this infrastructure for E3SM production climate simulations.

Iulian Grindeanu, Vijay S. Mahadevan  
Argonne National Laboratory  
iulian@anl.gov, mahadevan@anl.gov

Karen D. Devine  
Sandia National Laboratories  
kddevin@sandia.gov

## CP5

### GPU Accelerated Sparse Floyd-Warshall Algorithm for Graph All-Pair Shortest Path (APSP) Computation

The Floyd-Warshall (FW) algorithm is an attractive choice for computing Graph All-pair Shortest Path (APSP) on high-performing systems as it can be expressed as dense matrix-matrix multiplication on Min-plus semiring. However, the higher asymptotic complexity of FW for sparse graphs makes it less attractive to work optimal Johnsons algorithm. Recently, we proposed a sparse variant of FW (SpFW) algorithm that improves the asymptotic complexity for several classes of sparse graphs. In this work, we present a GPU accelerated variant of SpFW algorithm based on an efficient min-plus semi-ring matrix multiplication (SrGEMM). Our GPU accelerated SpFW achieves 80% of theoretical performance on Nvidia V100

GPUs. On multi-GPU configuration, our implementation achieves 32TFLOP/sec single precision performance on a single node of Summit supercomputer. Such performance improvement considerably increases the number of cases where SpFW outperforms an efficient multithreaded Johnson's algorithm.

Hao Lu, Piyush Sao  
Oak Ridge National Laboratory  
luh1@ornl.gov, saopk@ornl.gov

Ramakrishnan Kannan, Thomas Potok  
Oak Ridge National Laboratories  
kannanr@ornl.gov, potokte@ornl.gov

Vijay Thakkar  
Georgia Tech  
thakkarv@gatech.edu

Richard Vuduc  
Georgia Institute of Technology  
richie@cc.gatech.edu

## CP5

### Nonblocking Execution in Graphblas

GraphBLAS is a recent standard that allows expression of graph algorithms in the language of linear algebra and enables easy code parallelization and optimization. GraphBLAS operations are executed either in blocking or in non-blocking mode. The blocking mode guarantees that the output of the operation is computed when the function returns. On the other hand, the nonblocking mode aims at performing code optimizations that improve the overall performance by delaying the execution of the operations. Thus, a function is allowed to return although the computation may not be completed. Subsequent GraphBLAS operations can be safely invoked, but they may need to delay their computations as well, until the output of all earlier operations is computed and stored in memory. There exist multiple implementations of GraphBLAS for both shared and distributed memory systems, but none of them supports full nonblocking execution. We present the design and implementation of the first full nonblocking execution in GraphBLAS for shared memory systems. Nonblocking execution is achieved by exploiting lazy evaluation to enable runtime optimizations that improve data locality. Dynamic parallelism is used to group together operations that reuse data in cache and execute them in parallel. Since GraphBLAS operations are memory bound, a preliminary evaluation confirms the importance of nonblocking execution for various datasets of two algorithms, by showing up to 2.5x speedups over blocking execution.

Aristeidis Mastoras  
Zurich Research Center, Huawei Technologies Switzerland  
Computing Systems Laboratory  
aristeidis.mastoras@huawei.com

Albert-Jan N. Yzelman  
Huawei Technologies Switzerland AG  
albertjan.yzelman@huawei.com

## CP5

### Parrsb: Spectral Element Mesh Partitioning at Exascale

Most of scientific simulations start with reading a mesh

which is a representation of the computational domain using smaller units called elements. Next step is distributing these elements among the processors. This element distribution, is known as partitioning and critical for overall efficiency of the simulation. Large simulations over 100 million of elements are already running on pre-exascale systems like Summit at OLCF which is a sign that we will see meshes with about billions of elements on exascale computers that are planned for future. In this talk, we introduce parRSB, a partitioner designed to run on exascale computers. parRSB is a parallel mesh partitioner which uses Recursive Spectral Bisection (RSB) method to partition the mesh. RSB makes the partitioning decision based on a specific Eigenvector of the Laplacian matrix associated with dual graph of the mesh and is known to produce better partitions than methods like Recursive Coordinate Bisection. We talk about how we efficiently evaluate the action of Laplacian in a parallel setting and the preconditioners and other pre-partition optimizations we did to speed up the partitioning process. We also talk about some of the challenges we faced with disconnected components and parallel sorting at scale. We present performance and partition quality results we got with parRSB running at Scale on Summit supercomputer with meshes as large as 100 million elements.

Thilina B. Ratnayaka  
University of Illinois  
rbr2@illinois.edu

Paul Fischer  
University of Illinois at Urbana-Champaign  
fischerp@illinois.edu

### CP6 Asynchronous Global-Local Non-Invasive Coupling

Our aim is to derive an asynchronous version of the non intrusive Global-Local coupling. This method consists in using a coarse description of a structure and enriching the modeling of some zones of interest with enhanced geometry, refined mesh, complex material laws. Mathematically, the method is a stationary iteration which proved to be particularly robust to solve nonlinear problems. An asynchronous version would bring several advantages, the method would be more resilient and less sensitive to load balance, it would make better use of the computational resources whether high performance cluster or network of workstations. Paracontraction technics have been used to prove the convergence of the asynchronous model. From an implementation point of view based, our approach makes use of MPI one-sided communication, each CPU can directly access the memory of the other CPUs to get or to put pieces of information. Using a python code driving the get-fem finite element software and the mpi4py library. Comparison with a synchronous version of the code will be provided. The results are carried out with Ruche supercomputer commonly used by cole CentraleSupélec and ENS Paris Saclay. A weak scalability study will be performed on an academic configuration, and a turbine blade from the high-pressure compressor of a turbomachine will be studied where patches are used to introduce micro-perforations in the geometry of the leading and trailing edges.

Ahmed El Kerim  
ENS Paris Saclay  
ahmed.elkerim@ens-paris-saclay.fr

Gosselet Pierre  
Université de Lille , CNRS, Centrale Lille / LaMcube

pierre.gosselet@univ-lille.fr

Magoules Frédéric  
Université Paris-Saclay, CentraleSupélec, MICS,  
frederic.magoules@centralesupelec.fr

### CP6 Tile Low-Rank Matrix-Vector Multiplication with Architecture-Aware Dynamic Load Balancing

We present Tile Low-Rank Matrix-Vector Multiplication (TLR-MVM) with an architecture-aware dynamic load-balancing. Our TLR-MVM implementation leverages high bandwidth memory technologies and exploits data sparsity and an efficient data layout. Load balancing being a cornerstone of performance, we design a portable MPI+OpenMP framework and balance workloads using the OpenMP dynamic scheduler for node performance. We further accelerate TLR-MVM by explicitly mapping the processes/threads onto the underlying hardware, increasing occupancy. We achieve decent scalability on several x86, ARM, and GPU-based hardware platforms. We highlight TLR-MVM performance impact on computational astronomy and seismic imaging applications.

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

Jesse Cranney, Damien Gratadour  
The Australian National University  
jesse.cranney@anu.edu.au, damien.gratadour@obspm.fr

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

Matteo Ravasi  
King Abdullah University of Science and Technology  
matteo.ravasi@kaust.edu.sa

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

### CP6 A Swept Framework for Performing Memory Efficient Stencil Calculations in Parallel on Unstructured Meshes

Stencil calculations lie at the heart of many scientific computing applications, such as the solution of hyperbolic partial differential equations (PDEs) via explicit time-stepping methods, and the solution of linear systems of equations via stationary iterative methods (e.g. Jacobi iteration, Gauss-Seidel). Although straightforward to implement, stencil updates are difficult to perform efficiently on high performance computing architectures as they require much more data access relative to computation (making them memory bound). Many studies for improving the memory efficiency of stencil calculations have focused primarily on structured domains. In this talk, we present an algorithm for performing stencil calculations on a general unstructured mesh efficiently in parallel (i.e. on parallel CPUs or a GPU). On a GPU, this efficiency is achieved through the use of cached on-chip shared memory for all stencil updates. At each stage of the algorithm (four stages in total),

an appropriate partition of the mesh is determined to maximize the number of stencil updates possible within each subdomain. Subdomain data is allocated to the shared memory of GPU blocks where stencil updates can be performed in parallel. We illustrate the algorithm behavior on two examples: a structured square domain and an unstructured mesh around an airfoil. The algorithm only requires knowledge of the mesh connectivity, making it applicable to general unstructured problems.

Mohammad S. Islam, Qiqi Wang  
Massachusetts Institute of Technology  
moislam@mit.edu, qiqi@mit.edu

## CP6

### Parallel Asynchronous Reinforcement Learning Algorithm for Distributed Learning

Reinforcement Learning (RL) has not been heavily applied to scientific simulations, partly due to computational constraints. Here, we present a distributed and accelerated RL algorithm for tackling a materials science challenge simulated via a kinetic Monte-Carlo model. Our parallel RL algorithm works as follows. An RL computation is composed of RL agents, simulations, and models. In the case of a single process, we place RL agents on GPUs and simulations run on the CPU, with each agent updating their model based on completed simulation runs from every other agent. In the case of distributed memory parallelism, we partition RL agents among different processes to reduce inter and intra-node communication. When trying to scale to many processes, time imbalance between different agents becomes a major bottleneck. We propose an asynchronous update algorithm, where all agents need not synchronize in every iteration. The asynchronous update algorithm shows much better load balancing, which effectively minimizes idling time as agents wait for simulations to complete before they can proceed. We find this load balancing can result in at least 20% efficiency gains for highly stochastic simulations, such as those common to kinetic modeling in materials science. We demonstrate the efficacy of our approach for the Stein variational policy gradient algorithm, although our strategy can be readily adapted to other RL algorithms and readily modified based on convergence needs.

Rama K. Vasudevan  
Center for Nanophase Materials Sciences  
Oak Ridge National Laboratory  
vasudevanrk@ornl.gov

Piyush Sao  
Oak Ridge National Laboratory  
saopk@ornl.gov

Ramakrishnan Kannan  
Oak Ridge National Laboratories  
kannanr@ornl.gov

## CP7

### Chronos: a GPU-Accelerated Adaptive AMG for Extreme Size Simulations

The solution of linear systems of equations is a central task in several scientific and engineering applications. For extreme-size simulations, nowadays accounting for hundreds of millions or even billions of unknowns, the algebraic multigrid (AMG) methods are the most common choice as linear solvers because of their optimal or nearly opti-

mal convergence. In terms of HPC hardware, most of the top-ranked modern supercomputers leverage on GPU accelerators, mounting GPU cards on each compute node. So far, there are few AMG solvers that fully exploit GPUs and most are basic, not suited to solve ill-conditioned problems. The porting of complex algorithms to GPUs is not easy and needs continuous updating to preserve performance because it is a dynamically evolving hardware. In this context, Chronos is emerging as a sophisticated AMG framework that relies on several peculiarities: an adaptive test space generation, a sophisticated smoother that is effective even for severely ill-conditioned problems, an adaptive prolongation, and a filtering procedure able to preserve both interpolation and operator main properties. This communication focuses on the enhancement of Chronos by using GPU accelerators on distributed memory systems. A set of real-world problems, including structural mechanics, CFD and underground applications, are solved to assess the efficiency and robustness of Chronos.

Giovanni Isotton, Matteo Frigo, Carlo Janna  
M3E S.r.l.  
g.isotton@m3eweb.it, m.frigo@m3eweb.it,  
c.janna@m3eweb.it

## CP7

### Scalable Resilience Against Node Failures for the Communication-Hiding Conjugate Gradient Method with Deep Pipelines

The growth in the number of nodes in large-scale parallel computers gives rise to two major challenges: global communication operations like global reductions are becoming major bottlenecks due to their limited scalability, and the likelihood of node failures is increasing. We study an approach for addressing these challenges in the context of solving large sparse linear systems. The communication-hiding conjugate gradient method with deep pipelines overlaps communication of global reductions with the computation of matrix-vector products over multiple iterations. We present extensions to this solver which makes it resilient against node failures while fully preserving their communication-hiding properties and thus their scalability. A few redundant copies of local vector elements are efficiently communicated to neighboring nodes with very little overhead. In case of a node failure, these redundant copies are gathered at a replacement node, which reconstructs the lost parts of the solvers state. After that, the parallel solver can continue as in the failure-free scenario. The cost for the reconstruction depends on the length of the pipeline. Experimental evaluations of our approach illustrate on average low runtime overheads compared to the standard non-resilient algorithm if the pipeline length is not too large. This illustrates that scalable algorithmic resilience can be achieved at rather low extra cost.

Viktoria Mayer, Markus Levonyak  
University of Vienna  
viktoria.mayer@univie.ac.at,  
markus.levonyak@univie.ac.at

Wilfried N. Gansterer  
Department of Computer Science  
University of Vienna  
wilfried.gansterer@univie.ac.at

## CP7

### PDDSpase: A New Approach to Probabilistic Do-

## main Decomposition

Probabilistic Domain Decomposition (PDD) is a non-standard, highly-scalable algorithm for large-scale boundary-value problems (BVPs), in which the subdomains are decoupled by first calculating the solution on the fictitious interfaces thanks to probabilistic representations of BVPs. When the BVP coefficients are represented by look-up tables, however, it may happen that they exceed the memory of a single processor, thus preventing PDD from full scalability. A new version of PDD -PDDsparse- that has been tailored to this scenario will be presented. The idea is to express the solution on an interfacial node as a linear combination of the nodal solutions around them, via formal interpolation. Diffusions from a node are confined to very small regions, ensuring they fit in the processor's memory. This leads to a highly sparse system of equations for the interfacial values only (instead of all the values as it is the case with standard domain decomposition). We will formulate the new algorithm, discuss some implementation options, and present proof-of-concept results and scalability tests performed at the CINECA supercomputer facility, which illustrate the promise of this formulation.

Jorge Morón, Francisco Bernal  
 Universidad Carlos III de Madrid  
 jmoron@math.uc3m.es, fcoberna@math.uc3m.es

Juan Acebrón  
 Instituto Universitário de Lisboa  
 juan.acebron@iscte-iul.pt

## CP7

### Cache-Efficient Hybrid-Parallel Algorithm for Sparse Deep Neural Network Inference

We investigate parallel algorithms for sparse neural network inference, that is, the task of classifying a number of input data items using a sparse neural network. While sparsity is typically exploited to reduce the memory footprint by compressing the matrix storage, we also use it to reduce the communication volume at parallel inference. We propose the staircase model of a neural network, based on rearranging the layers into a diagonal-like form, that allows finding optimal partitionings of the nonzeros of a neural network using hypergraph partitioning software. We propose the hybrid-parallel algorithm which (1) implements model-parallel inference based on partitioning the network into  $p_0$  parts via the staircase model and (2) replicates such model-parallel inference  $p_1$  times, with each group of  $p_0$  processors computing inference on a different batch of input feature matrix, akin to data-parallelism. The proposed hybrid-parallel algorithm on generated deep neural networks consisting of up to 1920 layers on MNIST dataset is up to 23% faster on a 44-core 2-socket Intel machine and up to 73% faster on a 96-core 2-socket ARM versus a baseline data-parallel algorithm, confirming that hybrid-parallelism is necessary to obtain scalable performance. To improve performance, we use tiling techniques that cache the intermediate results and/or the neural network layers and estimate  $p_0$  and  $p_1$  using a computational model.

Filip I. Pawłowski  
 Huawei Technologies Switzerland  
 filip.igor.pawlowski@huawei.com

Rob H. Bisseling  
 Utrecht University  
 Mathematical Institute

r.h.bisseling@uu.nl

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

Albert-Jan N. Yzelman  
 Huawei Technologies Switzerland AG  
 albertjan.yzelman@huawei.com

## CP8

### A Scalable Computational Framework for Wavelet-Based 3D Multiresolution Grids

We present a wavelet-based block-structured 3D multiresolution framework for finite-difference applications, targeted to run on distributed-memory parallel architectures. The wavelet-based multiresolution theory enables a coherent approach to identify and compress spatial scales. Of particular relevance when applied to the simulation of a large range of physical phenomena, it naturally adapts the resolution throughout the domain to reduce the overall computational cost. We have implemented a family of second-generation interpolating wavelet filters, characterized by the interpolation order  $N$  and the conservation of  $\tilde{N}$  moments. We use the wavelet transform both as a criterion for refinement/compression, as well as a consistent way to interpolate across resolution levels. Our specific focus is to provide a scalable implementation of the wavelet transform with second-generation filters in 3D, on rectangular domains with arbitrary aspect ratios. Our implementation aims to reduce the communication overhead of ghost reconstruction, while providing flexibility in the choice of the wavelet order. We demonstrate the scalability of our solver on several thousands of cores, and show the dynamic mesh adaptivity using test cases involving scalar transport and diffusion.

Thomas Gillis, Wim Van Rees  
 MIT  
 tgillis@mit.edu, wregis@mit.edu

## CP8

### Automating Code Generation for Point Evaluations with Finite Elements

Point data turn up all the time in scientific computing problems - perhaps you want to know the value of a field at a particular point or set of points, couple the inputs and outputs of models via a grid of points, or you have point data measurements that you want to assimilate into a model of some phenomenon. Firedrake is a code generation system which enables straightforward and scalable model development for systems of partial differential equations using the finite element method (FEM). In Firedrake, mathematical expressions of the problem to solve are written in Domain Specific Language (DSL) abstractions from which optimised C code is generated. Hitherto it has been difficult to generate code from expressions involving point data: either one would extrapolate the point data to match some field defined over a larger domain or be forced to write bespoke code. We present additions to Firedrake which allow code generation from expressions which include point data. Our framework is compatible with the automated adjoint system dolfin-adjoint/pyadjoint and we demonstrate automated code generation for inverse problems which include point data via a variational data assimilation example. This work is part of a broader project aiming to use

DSLs and code generation to automate diagnostics of geo-scientific models with very large output data sets.

Reuben W. Nixon-Hill  
Department of Mathematics  
Imperial College London  
reuben.nixon-hill10@imperial.ac.uk

Daniel Shapero  
Polar Science Center  
Applied Physics Lab  
shapero@uw.edu

Colin J. Cotter  
Imperial College London  
Department of Mathematics  
colin.cotter@imperial.ac.uk

David Ham  
Department of Mathematics  
Imperial College London  
david.ham@imperial.ac.uk

### CP8

#### An Accurate Method to Select Floating-Point Precision and Control Round-off Error on Upwind Scheme

HPC and exaflop machines offer computational capabilities that allow simulations to be performed on large meshes. With a large number of meshes and a large number of iterations, the weight of rounding errors on the global error increases and must be controlled so that the results do not deviate too much from the solution. Moreover, current machine architectures perform better with single precision than with double precision. If the code implementation can take advantage of this performance, then the simulations will be able to lower their memory and energy consumption and the restitution time. The problem of measuring round-off errors is therefore critical for verifying simulation codes. In this study, we precisely measured the effect of rounding errors on the upwind scheme on different test cases of the transport equation and derived a method to ensure that round-off errors remain small compared to discretization errors. This method allows, for example, to select a precision adapted to the calculation (single or double) and to take advantage of the HPC performances while guaranteeing the results.

Ben Salem-Knapp Louise  
CEA  
louise.bensalemknapp@protonmail.com

Guillaume Perrin  
COSYS, Université Gustave Eiffel  
guillaume.perrin@univ-eiffel.fr

William Weens  
CEA  
william.weens@protonmail.com

### CP8

#### Support for Fully Parameterized Construction of $\mathcal{P}_r^-$ Finite Elements in PETSc

Efficient use of the Finite Element method requires that the discrete spaces for the computation be carefully selected to match properties of the continuous problem. From a de-

veloper's standpoint this means that we should attempt to provide as many tools for constructing a wide variety of spaces as possible. We present our recent work in the PETSc code base which allows users to construct function spaces by applying operations to other function spaces which have been previously implemented. This work focuses on the implementation of the Koszul differential operator which makes possible full parametric construction of the  $\mathcal{P}_r^-$  family of finite elements — including elements such as the Raviart-Thomas and Nédélec elements as well as some more recent constructions — which are useful for constructing  $\mathcal{H}(\text{div})$  or  $\mathcal{H}(\text{curl})$  conforming discretizations. In this work, we describe the computational implementation of the Koszul differential operator. Since the null space of the Koszul differential is not necessarily empty we have developed a procedure for identifying the null space via random sampling of the dual basis. This new procedure is also described and we provide estimates of time complexity and proofs of correctness and reliability of the algorithm.

Brandon M. Whitchurch

Georgia Tech  
Department of Computational Science and Engineering  
bwhitchurch3@gatech.edu

### CP9

#### Parallel Minimum Spanning Forest Computation using Sparse Matrix Kernels

Formulations of graph algorithms using sparse linear algebra have yielded highly scalable distributed algorithms for problems such as connectivity and shortest path computation. We develop the first formulation of the Awerbuch-Shiloach parallel minimum spanning forest (MSF) algorithm using linear algebra primitives. We introduce a multilinear kernel that operates on an adjacency matrix and two vectors. This kernel updates graph vertices by simultaneously using information from both adjacent edges and vertices. In addition, we explore optimizations to accelerate the shortcutting step in the Awerbuch-Shiloach algorithm. We implement this MSF algorithm with Cyclops, a distributed-memory library for generalized sparse tensor algebra. We analyze the parallel scalability of our implementation on the Stampede2 supercomputer.

Tim Baer, Raghavendra Kanakagiri, Edgar Solomonik  
University of Illinois at Urbana-Champaign  
tbaer2@illinois.edu, raghaven@illinois.edu,  
solomon2@illinois.edu

### CP9

#### Performance of Low Synchronization Orthogonalization Methods in Anderson Accelerated Fixed Point Solvers

Anderson Acceleration (AA) is a method to accelerate the convergence of fixed point iterations for nonlinear, algebraic systems of equations. Due to the requirement of solving a least squares problem at each iteration and a reliance on modified Gram-Schmidt for updating the iteration space, AA requires extra costly synchronization steps for global reductions. Moreover, the number of reductions in each iteration depends on the size of the iteration space. In this work, we introduce three low synchronization orthogonalization algorithms into AA within SUNDIALS that reduce the total number of global reductions per iteration to a constant of 2 or 3, independent of the size of the iteration space. A performance study demonstrates the reduced time required by the new algorithms

at large processor counts with CPUs and demonstrates the predicted performance on multi-GPU architectures. Most importantly, we provide convergence and timing data for multiple numerical experiments to demonstrate reliability of the algorithms within AA and improved performance at parallel strong-scaling limits.

Shelby Lockhart

University of Illinois at Urbana-Champaign  
sll2@illinois.edu

David J. Gardner, Carol S. Woodward  
Lawrence Livermore National Laboratory  
gardner48@llnl.gov, woodward6@llnl.gov

Stephen Thomas  
National Renewable Energy Laboratory  
stephethomas@gmail.com

Luke Olson  
University of Illinois Urbana-Champaign  
lukeo@illinois.edu

### CP9

#### Batched Second-Order Adjoint Sensitivity for Reduced Space Methods

This talk presents an efficient method for extracting the second-order sensitivities from a system of implicit nonlinear equations on graphical processing units (GPU). We design a custom automatic differentiation (AutoDiff) backend that targets highly parallel architectures by extracting the second-order information in batch. When the nonlinear equations are associated to a reduced space optimization problem, we leverage the parallel reverse-mode accumulation in a batched adjoint-adjoint algorithm to compute efficiently the reduced Hessian of the problem. We apply the method to extract the reduced Hessian associated to the balance equations of a power network, and present an application to the resolution of the optimal power flow problem.

Francois Pacaud

Argonne National Laboratory, U.S.  
fpacaud@anl.gov

### CP9

#### QR Decomposition on GPUs in Multiple Double Precision

The aim is to compensate the cost overhead of multiple double arithmetic using Graphics Processing Units (GPUs) capable of teraflop performance. Following the GPU acceleration of the blocked QR decomposition of [Kerr, Campbell, and Richards, GPGPU'09], the multiple double arithmetic from the QD library [Hida, Li, Bailey, Arith-15 2001], extended with code generated by the CAMPARY software [Joldes, Muller, Popescu, and Tucker, ICMS 2016] is applied to accelerate the QR decomposition on the NVIDIA P100 and V100 GPUs. Because the problems become compute bound, teraflop performance is already observed in double double precision for matrices of size 1,024.

Jan Verschelde

University of Illinois at Chicago  
Chicago, Illinois USA

janv@uic.edu

### CP10

#### Parallel High-Resolution Compact Partial FFT-Type Direct Algorithms for Subsurface Scattering Problems

This research presents a direct parallel partial FFT-type algorithm for the numerical solutions of the two and three-dimensional Helmholtz equations. The governing equations are discretized by high-order compact finite-difference or finite-element methods. The resulting discretized system is indefinite, making the convergence of the most iterative methods deteriorate with the increase of the frequency. In this situation, the parallel direct approaches become a better alternative, especially for the systems with discontinuous and singular right-hand sides. The focus is the efficient parallel implementation of the proposed algorithm in the shared (OpenMP) and distributed (MPI) memory environment. The complexity and scalability of the direct parallel method are investigated on scattering problems with realistic ranges of parameters in soil and mine-like targets. The results are also compared with the Krylov-type preconditioned methods developed in our previous publications.

Ron Gonzales

Idaho State University  
gonzrona@isu.edu

### CP10

#### Parallel FFT-Krylov Subspace Method for Helmholtz Equation

In this talk, we will present an efficient parallel implementation of the iterative compact high-order approximation numerical solver for 3D Helmholtz equation on multicore computers. The developed algorithms will have many applications such as imaging mine-like targets using ground penetrating radar. Our high-order parallel iterative algorithm is built upon a combination of Krylov subspace-type method with a direct parallel Fast Fourier transform (FFT) type preconditioner from our previous work. The resulting numerical method enables a natural and fast implementation on multicore computers. The numerical results and scalability of our implementation in the MPI programming environment will be presented.

Yun Teck Lee, Yury Gryazin, Ronald Gonzales

Idaho State University  
leeyunt@isu.edu, yurygryazin@isu.edu, ronaldgonza-  
les@isu.edu

### CP10

#### Improving the Performance of Reverse Time Migration Based on the First-Order Formulation of the Wave Equation using MWD-TB

Widely used in the oil and gas exploration industry, Reverse Time Migration (RTM) is a powerful technique for imaging complex geologic structures. It is conventionally based on directly solving the wave equation in the time domain using seismic datasets recorded at various locations. The computation of RTM concentrates on the evaluation of finite-difference stencils. Spatial blocking is a standard method for accelerating stencil computation that reduces memory bandwidth pressure by enabling data reuse in spatial directions. On the other hand, Multicore Wavefront

Diamond (MWD) is a new temporal blocking technique that enables further data reuse among multiple successive time steps. Previous work shows that MWD tiling can be integrated into RTM based on the second-order acoustic wave equation. This integration has a snowball effect on the absorbing boundary conditions, the I/O operations, and the imaging condition in order to maintain the performance superiority of temporal blocking and the correctness of the migrated image. In this talk, we present its integration into the RTM using the first-order formulation and discuss its performance on various contemporary architectures.

Long Qu  
King Abdullah University of Science and Technology (KAUST)  
long.qu@kaust.edu.sa

Rached Abdelkhalek  
Nvidia  
rabdelkhalek@nvidia.com

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

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

## CP10

### Scalability Studies of Albany Land Ice: a Performance Portable, Ice Sheet Solver

High resolution simulations of polar ice sheets play a crucial role in the ongoing effort to develop more accurate and reliable Earth-system models for probabilistic sea-level projections. These simulations often require a massive amount of memory and computation from large, heterogeneous supercomputing clusters to provide sufficient accuracy and resolution. In an effort to avoid architecture specific programming, the open-source, C++, finite element code base called Albany uses high level abstractions to integrate Trilinos libraries and the Kokkos programming model for performance portable code across a variety of different architectures. Albany Land Ice utilizes this framework to provide a scalable solver for advanced analysis of the first-order Stokes equations on Greenland and Antarctica. The primary challenge in simulating ice sheets at scale is solving the linear system associated with a thin, high-aspect ratio mesh. To address this problem, a preconditioner is constructed with matrix-dependent multigrid to semicoarsen to a single layer and algebraic multigrid to coarsen the layer further. This presentation focuses on the extensions to Albany Land Ice and matrix-dependent multigrid to provide a performance portable solver. A scalability study on NERSC Cori and OLCF Summit is used to provide a benchmark for future performance enhancements.

Jerry Watkins, Max Carlson  
Sandia National Laboratories  
jwatkin@sandia.gov, maxcarl@sandia.gov

Carolyn Kao  
Stanford University  
chkao831@alumni.stanford.edu

Irina K. Tezaur

Sandia National Laboratories  
ikalash@sandia.gov

Ray S. Tuminaro  
Sandia National Laboratories  
Computational Mathematics and Algorithms  
rstumin@sandia.gov

## CP11

### Subtractive Meshing

This paper introduces *subtractive meshing*, a novel approach to compute structured and predominantly hexahedral meshes well suited for finite element methods. The main application of this method is mesh generation for structural simulations in mechanical engineering. Geometries in mechanical engineering are to a large extent driven by manufacturability. Many parts in industrial use are manufactured with subtractive processes such as milling. Thus, the property of a part being manufactured subtractive can be exploited in meshing it. Given that a good mesh for the initial geometry before machining can be trivially computed as this geometry is usually sweepable this approach is promising. To leverage this potential an algorithm to perform subtractive operations on meshes was developed. The described algorithm consists of five stages. In mesh cutting all nodes which are cut away are identified. The element segmentation stage bisects all edges of the mesh and creates sub elements for all elements, which guarantees exclusion of all cut nodes from the mesh and absence of hanging nodes. The third stage assigns mesh entities to all geometric entities using Dijkstra's algorithm and answer set programming. Then the matched mesh entities are forced onto the respective geometry in order to improve geometric accuracy. At last, the mesh is optimized by moving nodes.

Henrik J. Stromberg  
Institute of Mechanical Engineering  
TU Clausthal  
stromberg@imw.tu-clausthal.de

Valentin Mayer-Eichberger  
Technische Universität Berlin  
valentin@mayer-eichberger.de

Armin Lohrengel  
Institute of Mechanical Engineering  
TU Clausthal  
lohrengel@imw.tu-clausthal.de

## CP11

### Signal Processing for a Reverse-Gps Wildlife Tracking System: CPU and GPU Implementation Experiences

We present high-performance implementations of signal-processing tasks performed by a wildlife tracking system called ATLAS. The system tracks radio transmitters attached to wild animals by estimating the time of arrival of radio packets to multiple receivers (base stations). Time-of-arrival estimation of wideband radio signals is computationally expensive, especially in acquisition mode. We developed a sequential high-performance CPU implementation of the computations a few years back, and more recently a parallel GPU implementation. Both strive to balance performance with simplicity, maintainability, and development effort, as most real-world codes do. The talk

will describe the two implementations and carefully evaluate their performance. The GPU implementation dramatically improves performance and power-performance relative to the sequential CPU implementation running on a desktop CPU typical of the computers in current base stations. Performance improves by more than 50X on a high-end GPU and more than 4X with a GPU platform that consumes almost 5 times less power than the CPU platform. Performance-per-Watt ratios also improve (by more than 16X), and so do the price-performance ratios. The results have already been published in a journal (Concurrency and Computation Practice and Experience (2021) but the presentation is highly appropriate for SIAM PP 2022.

Sivan A. Toledo  
Tel Aviv University  
stoledo@tau.ac.il

Yaniv Rubinpur  
Blavatnik School of Computer Science  
Tel Aviv University  
yaniv.ru@gmail.com

## MS1

### Independent Parallel Particle Layer (IPPL) Implementation and Observations

IPPL is a generic library written in C/C++ and MPI for high performance Eulerian, Lagrangian and mixed Eulerian-Lagrangian simulations. Since its development around the turn of the millennium, the variety of hardware architectures has grown and software developers are challenged in the design and maintenance of their codes. In this talk, I discuss our efforts to rewrite IPPL in the latest C++ standard and to enhance the library with the performance portable capabilities of Kokkos. I will also show first benchmarks performed on CPUs and GPUs.

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

Matthias Frey  
Paul Scherrer Institut (PSI), Switzerland  
matthias.frey@psi.ch

Muralikrishnan Sriramkrishnan  
Paul Scherrer Institut, Switzerland  
sramkrishnan.muralikrishnan@psi.ch

## MS1

### Alpine: A Portable Plasma and Particle-in-Cell Mini-App for Exascale

The goal of the ALPINE project is to create an electrostatic particle-in-cell mini-app that simulates a set of benchmark problems in plasma physics targeting exascale architectures. The mini-app serves as a proxy for real plasma physics applications. Hence, it can be used as a sandbox for implementing new algorithms and measuring performance on modern computing architectures. It is written using the library independent parallel particle layer (IPPL) which has been recently enhanced with performance portable capabilities. We show the performance for an idealized uniform plasma test case, Penning trap and Landau damping benchmark plasma physics problems in 3D on modern CPU

and GPU-based machines.

Sriramkrishnan Muralikrishnan  
Laboratory for scientific computing and modeling  
Paul Scherrer Institute  
sramkrishnan.muralikrishnan@psi.ch

Matthias Frey  
University of St Andrews  
mf248@st-andrews.ac.uk

Alessandro Vinciguerra  
ETH Zurich  
vincigua@student.ethz.ch

Antoine Cerfon  
Courant Institute NYU  
cerfon@cims.nyu.edu

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

## MS1

### Load-Balancing in Cabana and MPM/PIC Specifically

For a more efficient usage of computation resources the load balancing library ALL is integrated into Cabana/Cajita. The load balancer returns a new domain decomposition with a better balanced workload. Only full cells can be moved across processes. Although common in PIC based codes, MD codes do not have an inherent cell structure. They should use a very fine internal grid so the boundaries can be moved with high resolution allowing a finer load balancing. The performance gain of two setups is shown as well as the necessity of being cautious with memory allocations on the GPU. The proxy apps ExaMPM and CabanaMD are used to demonstrate the performance gains on summit. Cabana is backed by kokkos and hence these performance improvements are also portable to other architectures.

Stephan Schulz  
Ruhr-Universität Bochum  
stephan.schulz@phetarer.de

## MS1

### Pumipic: Infrastructure for Unstructured Mesh Pic on GPUs

Developing particle-in-cell (PIC) based fusion particle physics applications that effectively utilize thousands of GPUs in leadership class systems requires a flexible and extensible infrastructure that gives applications control over inter- and intra-node parallelism. PUMIPic, the parallel unstructured mesh infrastructure for PIC is a C++ Kokkos based library that provides control of inter node parallelism through fast particle load balancing methods from EnGPar, and mesh partitioning methods that support buffer layers to satisfy data dependencies of particle operations using Omega.h. Intra-node parallelism is provided in PUMIPic through an abstraction layer that allows application developers to quickly implement GPU accelerated kernels that operate on mesh and particle data without needing to specialize the code for the underlying data structure. The PUMIPic particle data structures group particles by the mesh element they exist within to improve

locality to mesh data during particle operations and minimize data movement. Performance results for these inter- and intra-node methods using four different particle data structures (Sell-C-Sigma, CSR, ECP-COPA Cabana, and DPS) will be presented on leadership class systems.

Cameron W. Smith  
Scientific Computation Research Center  
Rensselaer Polytechnic Institute  
smithc11@rpi.edu

Gerrett Diamond  
Rensselaer Polytechnic Institute  
diamog2@rpi.edu

Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

## MS2

### **A Software Ecosystem Approach to Increasing Progress and Addressing Challenges in Extreme Scale Computing and Big Data**

As computational science become more multi-disciplinary, multi-scale and multi-domain, and as scientific applications must execute on increasingly diverse computational platforms, the challenges for any particular team increases dramatically. Teams must have a variety of skills beyond a specific scientific domain: computer science, software engineering, data science and, for larger teams, social and cognitive sciences. In this presentation we describe efforts to create ecosystems and sustainable scientific software communities and products to address these challenges.

Michael A. Heroux  
Sandia National Laboratories  
St. John's University  
maherou@sandia.gov

## MS2

### **Algebraic Auxiliary Matrix Construction for Subspace Correction and Deflation based on Error Vector Sampling for a Sequence of Linear Systems**

In this talk, 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 dependence of the right-hand side vector on the solution vector of the prior linear system. For the problem, we investigate the subspace correction and deflation methods to accelerate the convergence of the Krylov subspace method. Practically, these acceleration methods work well when the auxiliary matrix consists of eigenvectors corresponding to small eigenvalues of the coefficient matrix. We have developed a new auxiliary matrix construction method to identify the approximation of the eigenvectors with small eigenvalues using error vector sampling in the prior solution step. Numerical tests confirm that both subspace correction and deflation methods with the generated auxiliary matrix accelerate the convergence of the iterative solver.

Takeshi Iwashita, Kota Ikehara, Takeshi Fukaya  
Hokkaido University  
iwashita@iic.hokudai.ac.jp, k\_ikehara@eis.hokudai.ac.jp,

fukaya@iic.hokudai.ac.jp

## MS2

### **Wisteria/BDEC-01 and h3-Open-BDEC: Innovative Scientific Computing by Integration of (Simulation+Data+Learning)**

We propose an innovative method for computational science for sustainable promotion of scientific discovery by supercomputers in the Exascale Era by combining (Simulation + Data + Learning (S+D+L)). In May 2021, we start operation of the Wisteria/BDEC-01 system with 33+PF at University of Tokyo. It is a Hierarchical, Hybrid, Heterogeneous (h3) system, which consists of computing nodes for CSE with A64FX and those for Data Analytics/AI with NVIDIA A100 GPUs. We develop a software platform h3-Open-BDEC for integration of (S+D+L) and evaluate the effects of integration of (S+D+L) on the Wisteria system. The h3-Open-BDEC is designed for extracting the maximum performance of the supercomputers with minimum energy consumption focusing on (1) innovative method for numerical analysis with high-performance/high-reliability/power-saving based on the new principle of computing by adaptive precision, accuracy verification and automatic tuning, and (2) Hierarchical Data Driven Approach (hDDA) based on machine learning. Integration of (S+D+L) by h3-Open-BDEC enables significant reduction of computations and power consumption, compared to those by conventional simulations.

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

## MS2

### **Sparse Linear Algebra for High Performance Machine Learning**

Sparse linear algebra is a crucial part of several applications since decades, and is now becoming an important topic for machine learning and Big Data. Classical example are the PageRank methods and Graph Neural Network, and, to avoid over-smoothing, the Neural Networks are also generating such problems. In this talk we propose, in particular, an efficient algorithm to compute sequences of sparse matrix vector products for stochastic matrices, such as matrices involved in the PageRank methods. We discuss experimental results obtain in multicore processors and in a large parallel supercomputer. Then, we discuss some sparse linear algebra methods in the same environment, such as, for example, for Graph Convolutional Network.

Serge G. Petiton  
University Lille  
serge.petiton@univ-lille.fr

## MS3

### **A Matrix-Free Approach for Algebraic Multigrid**

We present a matrix-free approach for algebraic multigrid for finite element methods on advanced architectures. We implement the algorithm on simple physics problems with an emphasis on the concept and extensions. We show numerical results on our advanced architecture testbeds to support the concepts and discuss room for future work.

Graham Harper

Computational Mathematics Dept.  
Sandia National Laboratories  
gbharpe@sandia.gov

### MS3

#### Novel Solver Algorithms for Nearly Singular Linear Systems Arising in Combustion Modelling

Direct Numerical Simulations of realistic combustion devices are extremely challenging due to the wide separation of scales in the simulation, for example an internal combustion (IC) engine chamber, and the flame thickness of a high-pressure flame. The PeleLMex solver uses adaptive mesh refinement (AMR) to evolve multi-species reacting flows in the low Mach number limit at the Exascale and relies on an embedded boundary (EB) approach to represent complex geometries. In that framework, the EB geometries often give rise to very small cut-cells along the boundary, which translate into extreme ill-conditioning of the pressure-projection, with eigenvalues that span 15-16 orders of magnitude. In this talk, we focus on the case of a typical IC piston bowl geometry for which we present on a novel approach towards solving these nearly singular linear systems with ILU-based, C-AMG smoothers on massively parallel architectures. In particular, we use scaling and equilibration algorithms to handle the non-normality of the upper triangular factors. This enables us to approximate the highly sequential triangular solve algorithm, embedded in the AMG smoothing-solve phase, with Jacobi iterations. This approximation can be written as a convergent Neumann series whose terms are composed of highly parallel sparse matrix vector multiplications. The result is an algorithm that substantially decreases setup and solve time, compared to state-of-the-art, for these challenging linear systems.

Paul Mullaney, Stephen Thomas  
National Renewable Energy Laboratory  
paul.mullaney@nrel.gov, stephethomas@gmail.com

Arielle K. Carr  
Lehigh University, U.S.  
arg318@lehigh.edu

Katarzyna Swirydowicz  
Pacific Northwest National Laboratory  
kasia.swirydowicz@pnnl.gov

Marcus Day  
National Renewable Energy Laboratory  
marcus.day@nrel.gov

Lucas Esclapez  
National Renewable Energy Lab  
lucas.esclapez@nrel.gov

### MS3

#### A New Semi-Structured Algebraic Multigrid Method in HyPre

Multigrid methods are well suited for solving linear systems on large massively parallel computer architectures since they are mathematically optimal and display excellent parallelization properties. In modern HPC architectures, taking advantage of problem structure has become crucial for performance since most of the computational power originates from GPUs. The hyPre software library provides high-performance multigrid preconditioners and

solvers through conceptual interfaces, including a semi-structured interface that describes matrices primarily in terms of stencils and logically structured grids. In this talk, we present a new semi-structured algebraic multigrid (SSAMG) method that leverages problem structure and which is built on top of the semi-structured interface. We evaluate this method's numerical convergence and performance for a set of semi-structured problems such as structured adaptive mesh refinement test cases. SSAMG achieves comparable results to other solver options in hyPre, but has better potential for GPU architectures than hyPre's unstructured AMG solver (BoomerAMG). In addition, the new method can solve more complex problems than hyPre's structured solvers such as PFMG and SMG.

Victor A. Paludetto Magri, Ulrike Meier Yang, Robert D. Falgout  
Lawrence Livermore National Laboratory  
paludettomag1@llnl.gov, yang11@llnl.gov, falgout2@llnl.gov

### MS3

#### Solving Sparse Linear Systems on FPGAs using OneAPI

Intel's OneAPI and Data-Parallel C++ (DPCPP) offer the potential to program Intel FPGA devices with the ease of high-level C++ programming, rather than having to write Verilog or other low-level languages. We consider the solution of linear systems derived from partial differential equations using modified versions of HPCG and MiniFE on Intel Stratix10 FPGA devices. We discuss implementation considerations of various kernels (e.g. matrix vector multiplication, symmetric Gauss-Seidel smoothing and sparse direct solves) that would be necessary in a multigrid-preconditioned conjugate gradient solve and how implementation choices relate to FPGA performance.

Christopher Siefert, Clayton Hughes, Nicholas Miller, Stephen Olivier, Gwendolyn Voskuilen  
Sandia National Laboratories  
csiefer@sandia.gov, chughes@sandia.gov,  
nicmill@sandia.gov, slolivi@sandia.gov, grvosku@sandia.gov

### MS4

#### Multiphysics and Multiscale Statistical Emulation of Complex Computer Models

Investigating uncertainties in computer simulations can be prohibitive in terms of computational costs, whenever dealing with multi-scale and multi-physics simulations. First, so-called linked Gaussian processes offer a way to build analytical emulators for systems of computer models such as multi-physics models. We present new exact formula linked Gaussian process under the squared exponential kernel to a class of Matérn kernels, essential in advanced applications, with orders of magnitude gains. An iterative procedure to construct surrogate models for any feed-forward systems of computer models is also introduced and illustrated, with again large gains. We also introduce an adaptive design algorithm that increases the approximation accuracy of linked Gaussian process surrogates with further reduced computational costs. Second, we present a novel approach to building an emulator using multiple levels of resolution (or fidelity). We employ an sequential design of experiments at each level, and combine training data of different degrees of sophistication in a sequential mul-

tilevel approach, using Gaussian process emulators. This dual strategy allows us to allocate efficiently limited computational resources over simulations of different levels of fidelity and build the GP emulator efficiently. We theoretically prove the validity of our approach. Gains of orders of magnitudes in accuracy for medium-size computing budgets are demonstrated in numerical examples.

Deyu Ming  
University College London, UK  
deyu.ming.16@ucl.ac.uk

Ayao Ehara, Serge Guillas  
University College London  
ayao.ehara.18@ucl.ac.uk, s.guillas@ucl.ac.uk

#### MS4

##### Information-Preserving Bayesian Models for Efficient and Robust Learning

Modern deep neural networks are in general quite brittle, and hence less robust to noise in the input or adversarial perturbation applied to them, as well as out-of-distribution data. From an information-theoretic point of view, the ability of the model to achieve better generalization and robustness will depend on whether the model learns more semantically meaningful information and compresses the nuisance information. Typically, the DNNs are trained by reducing an empirical loss function which can lead to nuisance or irrelevant information being memorized, that can have a detrimental impact on the model's robustness. To that end, we adopt an information-theoretic Bayesian approach that poses this intuition as a constrained optimization problem and hence seeks to learn a compressed (latent) encoding of input that is maximally informative about our target as measured through their mutual information. An important question is the choice of the latent dimension of the encoding so that the nuisance information is reduced. This question also appears in the context of popular generative models such as the variational autoencoders. To address this, we developed a Bayesian approach to model the joint distribution of the latent encoding dimension and the latent variable distribution through a spike and slab distribution and its variational formulation, which has the ability to systematically quantify the uncertainty across dimensions.

Sandeep Madireddy  
Argonne National Laboratory  
smadireddy@anl.gov

#### MS4

##### Parallelized Multilevel Markov Chain Monte Carlo

The ExaHyPE Engine is a hyperbolic PDE engine capable of solving systems of first order hyperbolic PDEs. The engine provides a space-tree discretization of the computational domain, higher-order DG schemes and a-posteriori subcell limiters. ExaHyPE was initially designed with forward modeling in mind. In order to utilize expertise from the uncertainty quantification (UQ) community we have extended it with an interface to the MIT UQ library (MUQ). I will present a parallelization strategy for multilevel Markov chain Monte Carlo, a state-of-the-art, algorithmically scalable UQ algorithm for Bayesian inverse problems. The integration between MUQ and ExaHyPE allows for large-scale parallelism across forward model evaluations and the UQ algorithms themselves. The main scalability challenge presents itself in the form of strong data

dependencies introduced by the MLMCMC method, prohibiting trivial parallelization. I will demonstrate the effectiveness of the method by inferring the most probable locations for the initialization of a tsunami.

Anne Reinarz  
Durham University  
anne.k.reinarz@durham.ac.uk

Linus Seelinger  
Heidelberg University  
Interdisciplinary Center for Scientific Computing (IWR)  
mail@linusseelinger.de

#### MS4

##### Deep Gaussian Processes for Uncertainty Quantification in Computer Experiments

Deep Gaussian processes (DGPs) are increasingly popular as predictive models in machine learning for their non-stationary flexibility and ability to cope with abrupt regime changes in training data. The layered structure of the DGP likelihood makes direct inference impossible. Existing variational inference methods offer thrifty predictions but oversimplify uncertainties, particularly in the low data settings common in computer experiments. To achieve full uncertainty quantification, we present a novel elliptical slice sampling Bayesian posterior inferential scheme for DGP surrogates. Elliptical slice sampling is particularly suited for sampling latent Gaussian layers as it is free of tuning parameters and is able to bounce readily between multiple modes. Efficient computation relies on parsimonious layouts of latent layers, effective mixing of MCMC chains, and careful utilization of parallel processing. Our methods are illustrated on simulation data and real computer experiments of varying input dimensionality. We provide an open source implementation in the "deepgp package on CRAN.

Annie Sauer  
Virginia Tech  
anniees@vt.edu

#### MS5

##### Lessons Learned During My 10 Year Quest to Enable Collaborative, Sustainable and Reproducible ML Systems R&D

In this talk, I will describe multiple problems that I have faced while trying to enable collaborative, sustainable and reproducible ML Systems research since 2010. I will also present the Collective Knowledge framework (CK) that I have developed to address some of these problems with the help of the community. The CK concept is to organize all research projects as a database of reusable components, automation recipes and portable workflows with common APIs and extensible meta descriptions based on DevOps and FAIR principles. Such approach helped researchers and practitioners to decompose complex ML Systems into reusable automation tasks and meta packages for ML models, data sets, AI frameworks, libraries and compilers. We could then assemble portable ML workflows to automatically build, test, benchmark and auto-tune diverse ML Systems across continuously changing software and hardware. Furthermore, the community could collaboratively fix reproducibility issues in common CK workflows and components encountered during crowd-benchmarking of ML Systems thus enabling their sustainable and evolutionary development. I will conclude with lessons learned

when using the CK technology to automate and reproduce the MLPerf(tm) benchmark - a recent community effort to develop a common ML benchmark that provides consistent, reproducible, and fair measurements of accuracy, speed, and efficiency across diverse ML models, data sets, hardware, and software.

Grigori Fursin  
OctoML.ai  
grigori.fursin@ctuning.org

#### MS5

##### **Trust Tools: Experiences in Reproducibility in Undergraduate Research**

Computational results must be trustworthy to hold merit. As hardware and software evolve, results that were considered reproducible can no longer be obtained. With some extra effort, this problem can be avoided. After integrating reproducibility tools into a variety of computational workflows, our experience shows that these tools move the burden of executing common ad-hoc tasks from the human to the machine, therefore minimizing human error that could degrade the trustworthiness of computational results.

Patrick J. Hesse  
St. John's University  
phesse001@csbsju.edu

#### MS5

##### **Reproducibility and Transparency Practices in HPC**

Several systems research conferences now incorporate an artifact description and artifact evaluation (AD/AE) process as part of the paper submission. Authors of accepted papers optionally submit a plethora of artifacts: documentation, links, tools, code, data, and scripts for independent validation of the claims in their paper. An artifact evaluation committee (AEC) evaluates the artifacts and stamps papers with accepted artifacts, which then receive publisher badges. Does this AD/AE process serve authors and reviewers? Is it scalable for large conferences such as SCxy? Using the last three SCxy Reproducibility Initiatives as the basis, this talk will analyze the AD/AE process using a data-driven approach. We will distinguish studies that benefit from AD, i.e., increased transparency versus areas that benefit from AE. We will present a vision for the resulting curated, reusable research objects—how such research objects are a treasure in themselves for advancing computational reproducibility and making reproducible evaluation practical in the coming years.

Tanu Malik  
none  
tanu.malik@depaul.edu

#### MS5

##### **Toward a Better Ecosystem for Understanding and Reusing Scientific Workflows**

Workflow environments have re-emerged as a key topic in scientific computing with the advent of machine learning, multi-disciplinary code coupling, multi-site compute infrastructures, and Big Data for science. A wide range of tools and approaches have been developed and tested within organizations and communities that can help with reuse and refactorization of complex workflow components. However, we don't have a shared ecosystem for scientific workflows

that allows end users to easily reason about how to connect, adapt, and refactor existing workflow components and fragments from other communities. In this talk, we outline some approaches for building a new systems abstraction for workflow management that leverages insights from metadata semantics, data management systems and middleware, and model-driven code generation techniques. Building off insights from our previous and current work with computational biology, materials, and other sciences, we point to new opportunities for improving the software process for individual workflows while also supporting a broader push to develop a more open and FAIR community for scientific workflows.

Matthew Wolf, Jeremy Logan  
Oak Ridge National Laboratory  
wolfmd@ornl.gov, loganjs@ornl.gov

#### MS6

##### **Load Balancing Parallel Fock Matrix Computation in Nwchemex using Submodular B-Matching**

Computational load balancing in a Parallel environment is challenging. In this talk, we will show how submodular optimization can implicitly provide load balancing. We will formulate a load balancing problem, that arises in parallel Fock matrix computation in NWChemEx, as a submodular b-matching, and will provide an efficient approximation algorithm for it. Using this assignment, we then demonstrate the empirical scalability of Fock matrix computation in 14k processors of Summit supercomputer at ORNL.

S M Ferdous  
Purdue University  
sferdou@purdue.edu

#### MS6

##### **Graph Algorithms for the Exascale: A Case Study with Influence Maximization**

Availability of unprecedented amounts of data has ushered in a new era of data-driven sciences, and graph analytics has emerged as one of the leading tools for analysis of unstructured data. Availability of exascale computing platforms will provide an unprecedented access to computing power that will lead to novel discoveries in science. Generalization of graph algorithms in the form of combinatorial optimization has numerous applications in scientific computing and data-driven discovery. Despite widespread use, efficient parallel tools for graph analytics are hard to come by, especially when targeting the hybrid CPU-Graphics Processing Unit (GPU) architectures at extreme scales. In this talk, we will present our ongoing work on distributed multi-GPU systems for influence maximization, a prototypical graph problem. We will demonstrate substantial gains in performance on the current #2 supercomputer, Summit.

Marco Minutoli  
Pacific Northwest National Lab  
marco.minutoli@pnnl.gov

Mahantesh Halappanavar  
Pacific Northwest National Laboratory  
hala@pnnl.gov

Ananth Kalyanaraman  
Boeing Centennial Chair, Associate Professor  
School of EECS, Washington State University

ananth@wsu.edu

## MS6

### Scalable Algorithms for Controlling Epidemics on Networks using Multiplicative Weights Update (MWU) Method

As the COVID-19 outbreak has shown, effective interventions strategies such as vaccinations and social distancing are crucial to controlling epidemic outbreaks. We study the Epidemic Control problem, which involves designing interventions, to minimize the expected number of infections in an epidemic outbreak. This is a challenging stochastic optimization problem in the context of the SIR epidemic model on a network. We developed efficient approximation algorithms for this problem using the sample average approximation (SAA) technique from stochastic optimization and linear programming based rounding. We show that this algorithm scales up to networks with about a million nodes corresponding to county-level populations. This algorithm needs to solve the Linear Program (LP) relaxation of the problem using an LP solver, which restricts the scalability of this approach. Therefore, instead of using an LP solver, we design an algorithm that adapts the multiplicative weights update (MWU) framework to approximately solve the LP relaxation for the problem instance. We also present a memory-efficient version of this algorithm which allows it to scale well to networks corresponding to state- and even country-level populations. Also, we show that this approach provides good performance guarantees in practice.

Prathysh Sambaturu  
University of Virginia  
pks6mk@virginia.edu

## MS6

### Scalable Preconditioner for Graph Laplacian

The problem of finding a solver and preconditioner for graph Laplacian linear systems is well studied. However, not much attention has been devoted to the implementation of such algorithms, especially on HPC systems. In this talk we propose a viable and practical approach for the construction of a graph Laplacian preconditioner. The proposed algorithm is effective in reducing the iteration count and utilizing parallel resources. In addition, we demonstrate how the proposed preconditioner is applied inside LOBPCG in order to find eigenvalues of the graph Laplacian.

Katarzyna Swirydowicz, Mahantesh Halappanavar  
Pacific Northwest National Laboratory  
kasia.swirydowicz@pnnl.gov, hala@pnnl.gov

Stephen Thomas  
National Renewable Energy Laboratory  
stephethomas@gmail.com

## MS7

### Performance-Driven Memory Management for Disaggregated Memory Architectures

Disaggregated memory architectures promise scalable, efficient, and flexible integration of various processing elements in memory-centric architectures. Today's high-performance computing (HPC) systems suffer from severe memory underutilization, which is expected to further grow

with the use of dense high-capacity emerging non-volatile memories (NVMs). Moreover, the nature of emerging workloads where large graphs or data structures are operated on collaboratively would need a scalable memory-centric architecture, which can be realized efficiently by disaggregating memory and using shared memory pools. Finally, with emerging NVMs possessing the features of storage (data persistence) and memory (byte-addressability and speed), they can be leveraged to host large files which can be accessed by many different compute nodes. Thus, hosting such files in a location that is directly and efficiently accessible by many compute nodes is advantageous. In this talk, I will summarize our efforts in investigating memory management and access control mechanisms for such disaggregated memory architectures. I will also describe our ongoing efforts for further improving the performance, isolation, and security of disaggregated memory architectures. In the talk, I will describe various performance bottlenecks and challenges when deploying such architectures and the various ways to overcome them.

Amro Awad

North Carolina State University  
ajawad@ncsu.edu

## MS7

### Workflow Critical Path, a Runtime Metric for Dynamic Resource Management

Optimizing scientific application performance in HPC environments is a complicated task which has motivated development of many performance analysis tools over the past decades. Many tools were designed to analyze the performance of a single parallel code using common approaches such as message passing (MPI), multithreading (OpenMP), acceleration (CUDA), or a hybrid approach. However, current trends in HPC such as the push to exascale, convergence with Big Data, and growing complexity of HPC applications, have created gaps that these performance tools do not cover, particularly involving HPC workflows comprising multiple codes, paradigms, or platforms. To address this performance monitoring gap, we define a new metric called Workflow Critical Path (WCP), a data-oriented, runtime critical path metric for Holistic HPC Workflows. WCP constructs graphs using data states as vertices and data mutations as edges. I will present a fully functional prototype called Crux, a distributed analysis tool that uses cloud-based technologies for calculating and visualizing WCP. Crux data store can be queried throughout the workflow runtime, thus providing the current critical path as needed for dynamic resource management. Our experiments with a workflow simulator on AWS show Crux is scalable and capable of correctly calculating WCP for common Holistic HPC workflow patterns. I will explore the use of WCP and discuss how Crux could be used with production HPC platforms and resource managers.

Karen L. Karavanic  
Portland State University  
karavan@pdx.edu

## MS7

### Power-Aware Resource Management

Power is a critical problem as the supercomputing community ventures toward exascale. Several system, hardware and application-level power requirements make it challenging to schedule jobs effectively with high throughput and

utilization on large-scale supercomputers. In this talk, I will present two power-aware scheduling strategies. The first is a dynamic power-aware backfilling algorithm implemented as a simulation study within the popular resource manager, SLURM, and the second is a power-aware manager being developed as part of LLNL's novel resource management framework for El Capitan, called Flux. Additionally, I will also briefly discuss a variation-aware scheduling algorithm implemented within Flux that addresses manufacturing variability under power constraints. Using real-world datasets from large supercomputers at LLNL and University of Tokyo, I will demonstrate the effectiveness of these power-management scheduling algorithms.

Tapasya Patki  
Lawrence Livermore National Laboratory  
patki1@llnl.gov

### MS7

#### Reliable Performance Testing for Multi-Tenant Clouds

To effectively utilize cloud computing, cloud resource management requires accurate knowledge of the performance of cloud applications. However, due to the random performance fluctuations, obtaining accurate performance results in the cloud is extremely difficult. To handle this random fluctuation, prior research on cloud performance testing usually relied on advanced statistical tools. However, the effectiveness of these advanced statistical tools depends heavily on the statistical characteristics of cloud performance testing data. In this talk, we present two performance testing methodologies. The first methodology is designed to obtain the performance distribution of a cloud application. It relies on statistical stability and multinomial likelihood theory to handle the non-typical performance distributions of cloud applications. The second methodology is designed to obtain single-point estimates of performance and relies on block bootstrapping to handle the internal dependency of cloud performance data. Evaluation results show that both methodologies can provide testing results with more than 90% accuracy on average.

Wei Wang  
The University of Texas at San Antonio  
wei.wang@utsa.edu

### MS8

#### Experience-Driven Performance Tuning of GPU Kernels for Dense Linear Algebra

High performance GPU kernels are often highly parameterized in order to achieve performance portability across different architectures. However, this comes at the cost of a parameter space that is often too large to search entirely for every possible use-case of the considered kernel. We present our work on the MAGMA library, where instead of pruning the search space, we perform exhaustive searches on specific use-cases that are known to be very important in the originating applications. As an example, rank-k updates in matrix multiplication (GEMM) are crucial for the performance of dense matrix factorizations, and so are considered during the tuning process of the GEMM kernel. Examples for other kernels will be presented as well.

Ahmad Abdelfattah  
University of Tennessee, Knoxville

ahmad@icl.utk.edu

### MS8

#### Analytical, Statistical, and Machine Learning Models for Compiler Optimization of Convolutional Neural Networks

We discuss the role of performance modeling for predicting the most effective code transformation sequence for the common problem of optimizing convolutional neural networks (CNNs) when dimensions of the network are known. Convolution dominates performance of CNNs, and the performance of convolution improves significantly as a result of compiler optimizations including loop permutation, tiling, unroll-and-jam, parallelization and vectorization. Deriving the optimal code transformation sequence for a specific size of network is challenging. Deep loop nests and large numbers of dimensions of inputs and output mean that there are many different choices, and the best choice is heavily influenced by network size and relative size of different dimensions. This talk will describe a number of approaches we have explored for predicting the best optimization sequence: an analytical model utilizing decades of prior compiler research and architecture-specific heuristics, statistical modeling combined with autotuning, and recent work on models derived from machine learning. We present a machine learning model that has demonstrated success in selecting loop order, and how to extend that model to support additional transformations. We show that all approaches are effective for optimizing convolution as compared to manually-tuned libraries, and the tradeoff between approaches involves the amount of developer and system time to construct the model.

Mary Hall  
School of Computing  
University of Utah  
mhall@cs.utah.edu

### MS8

#### Search Space Reduction Through Analytical Modeling

Increasing complexity in computer architecture means that the space of high performance implementation is also increasing rapidly. This makes it impractical to search the space exhaustively for good implementations. In this talk, we discuss using analytical models to reduce the search space of possible implementations. Our analytical models capture knowledge made by experts in optimizing for different hardware features. This reduces the search space to parameters where models are too complicated, inaccurate, or not available; and auto-tuning or ML-based approaches may be applied. This can result in decreased search-time and potentially reduced size of the ML models.

Tze Meng Low  
Carnegie Mellon University  
lowt@andrew.cmu.edu

### MS8

#### Deep Gaussian Processes with Multi-Task and Transfer Learning for Performance Optimization

Bayesian learning is a commonly used framework for the performance modeling and optimization of HPC applications. Deep Gaussian processes combine the uncertainty quantification advantage of Gaussian processes with the

predictive power of deep learning. Moreover, multitask and transfer learning allow for improved prediction of machine learning models when several similar tasks are to be learned simultaneously or when previous learning are sought to help in the learning of new tasks. In this paper, we combine deep Gaussian processes with multitask and transfer learning to allow for both an improved tuning of an application parameters on problems of interest but also the prediction of parameters on any potential problem the application might encounter.

Wissam Sid-Lakhdar  
University of Tennessee  
Innovative Computing Lab  
wissam@icl.utk.edu

## MS9

### Domain Decomposition Solvers for Nonlocal Equations

We present a domain decomposition method for the efficient simulation of nonlocal problems. Our approach is based on a multi-domain formulation of a nonlocal diffusion problem where the subdomains share “nonlocal” interfaces of the size of the nonlocal horizon. This system of nonlocal equations is first rewritten in terms of minimization of a nonlocal energy, then discretized with a meshfree approximation and finally solved via a Lagrange multiplier approach in a way that resembles the finite element tearing and interconnect method. Specifically, we propose a distributed projected gradient algorithm for the solution of the Lagrange multiplier system, whose unknowns determine the nonlocal interface conditions between subdomains. Several two-dimensional numerical tests on problems as large as 191 million unknowns illustrate the strong and weak scalability of our algorithm, which outperforms the standard approach to the distributed numerical solution of the problem. This work is the first rigorous numerical study in a two-dimensional multi-domain setting for nonlocal operators with finite horizon and, as such, it is a fundamental step towards increasing the use of nonlocal models in large scale simulations.

Xiao Xu  
The University of Texas at Austin  
xiaoxu42@utexas.edu

Christian Glusa, Marta D’Elia  
Sandia National Laboratories  
caglusa@sandia.gov, mdelia@sandia.gov

John Foster  
The University of Texas at Austin  
jfoster@austin.utexas.edu

## MS9

### Overlapping Domain Decomposition Methods: New Algebraic Coarse Spaces

The design of robust multilevel preconditioners often relies on the choice of appropriate transfer operators. When it comes to overlapping domain decomposition methods, there is a wide range of choices, but the construction of these operators is often based on extra information not always available algebraically such as rigid body modes or per-element matrix assembly. In this talk, I will present the construction of new algebraic coarse spaces for different linear systems, such as sparse least-squares problems, symmetric positive definite systems, or nonsymmetric prob-

lems. Furthermore, the integration of this work in PETSc will be presented and showcased for different applications, using the proposed methodology as a black-box solver, with comparisons against other state-of-the-art solvers such as GAMG and BoomerAMG.

Pierre Jolivet  
CNRS  
pierre.jolivet@enseeiht.fr

Hussam Al Daas  
INRIA  
hussam.al-daas@stfc.ac.uk

## MS9

### Adaptive Nonlinear Preconditioning for PDEs with Error Bounds on Output Functionals

Nonlinear preconditioning refers to transforming a nonlinear algebraic system to a form for which Newton-type algorithms have improved success through quicker advance to the domain of quadratic convergence. We place these methods, which go back at least as far as the Additive Schwarz Preconditioned Inexact Newton (ASPIN, 2002) in the context of a proliferation distinguished by being left- or right-sided, multiplicative or additive, and partitioned by subdomain, field type, or other criteria. We present the Nonlinear Elimination Preconditioned Inexact Newton (NEPIN, 2021), which is based on a heuristic bad/good heuristic splitting of equations and corresponding degrees of freedom. We augment basic forms of nonlinear preconditioning with two features of practical interest: an adaptive switchover to ordinary Newton as the domain of convergence is approached and error bounds on output functionals of the solution. Various nonlinearly stiff algebraic and model PDE problems are considered for insight and we illustrate performance advantage and scaling potential on challenging two-phase flows in porous media.

Lulu Liu  
King Abdullah University of Science and Technology  
lulu.liu@kaust.edu.sa

Li Luo  
Shenzhen Institutes of Advanced Technology  
Chinese Academy of Sciences  
li.luo@siat.ac.cn

Feng-Nan Hwang  
National Central University  
Department of Mathematics  
hwangf@math.ncu.edu.tw

Xiao-Chuan Cai  
University of Colorado, Boulder  
Dept. of Computer Science  
cai@cs.colorado.edu

Rolf Krause  
Università della Svizzera italiana, Switzerland  
rolf.krause@usi.ch

David E. Keyes  
King Abdullah University of Science and Technology  
(KAUST)

david.keyes@kaust.edu.sa

**MS9****Adaptive Nonlinear Domain Decomposition Methods**

In this talk, different nonlinear domain decomposition methods are applied to nonlinear problems with highly-heterogeneous coefficient functions with jumps. In order to obtain a robust nonlinear domain decomposition method with respect to nonlinear as well as linear convergence, adaptive coarse spaces are employed. First, as an example for a nonlinearly left-preconditioned domain decomposition method, the two-level restricted nonlinear Schwarz method H1-RASPEN (Hybrid Restricted Additive Schwarz Preconditioned Exact Newton) is combined with an adaptive generalized Dryja-Smith-Widlund (GDSW) coarse space. Second, as an example for a nonlinearly right-preconditioned domain decomposition method, a nonlinear FETI-DP (Finite Element Tearing and Interconnecting - Dual Primal) method is equipped with an edge-based adaptive coarse space. Both approaches are compared with the nonlinear domain decomposition methods with classical coarse spaces as well as with the Newton-Krylov methods with adaptive coarse spaces. For two-dimensional problems with different spatial coefficient distributions, it can be observed that the best linear and nonlinear convergence can only be obtained when combining the nonlinear domain decomposition methods with adaptive coarse spaces.

Martin Lanser

Universitaet zu Koeln  
Mathematisches Institut  
martin.lanser@uni-koeln.de

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

Axel Klawonn  
Universität zu Köln  
axel.klawonn@uni-koeln.de

**MS10****Large-Scale Optimization with Structured Hessian Estimates**

For certain large-scale optimization problems, in which the coupling of variables results in a sparse block arrowhead Hessian matrix, exploiting structure is important. Even with sparse blocks using back-substitution for the arrowhead pattern yields a dense intermediate matrix. In order to overcome forming a dense system we instead approximate it using the compact representation of a structured BFGS (S-BFGS) quasi-Newton matrix. Because the S-BFGS matrix is of advantageous structure (i.e., multiple of identity plus low-rank), the Sherman-Morrison-Woodbury inverse is used to solve with it. The proposed methods are described for stochastic optimization when decomposed into scenarios, yet structured approximations may be more widely applicable.

Johannes J. Brust

Argonne National Laboratory  
Mathematics and Computer Science Division

jbrust@anl.gov

**MS10****Efficient Preconditioners for Interior Point Methods via a New Schur Complement Strategy**

For the past few decades, interior point methods (IPM) have gained wide appreciation due to their remarkable success in solving linear and non-linear optimization problems. These iterative methods are dominated in cost by the solution of the Karush-Kuhn-Tucker (KKT) equations at every iteration. We propose a new interior point algorithm for convex quadratic programming problems which aims to reduce the overall cost by re-using the factorization of a fixed KKT subsystem across interior point iterations. Our method uses preconditioned iterative methods to solve the linear systems that arise at every Newton step. These preconditioned conjugate gradient methods operate on an implicit Schur complement of the KKT system at each iteration. In contrast to standard approaches, the Schur complement we consider enables the reuse of the factorization of the Hessian of the equality-constraint Lagrangian across all interior point iterations. Further, the resulting reduced system admits preconditioners that directly alleviate the ill-conditioning associated with the strict complementarity condition in interior point methods. The two preconditioners we propose also provably reduce the number of unique eigenvalues for the coefficient matrix (CG iteration count). One is efficient when the number of equality constraints is small, while the other is efficient when the number of remaining degrees of freedom is small.

Samah Karim

University of Illinois at Urbana-Champaign  
swkarim2@illinois.edu

**MS10****Randomized Sketching for Low-Memory Dynamic Optimization**

In this talk, we develop a novel limited-memory method to solve dynamic optimization problems. The memory requirements for such problems often present a major obstacle, particularly for problems with PDE constraints such as optimal flow control, full waveform inversion, and optical tomography. In these problems, PDE constraints uniquely determine the state of a physical system for a given control. While the control is often low dimensional, the state is typically more expensive to store. To reduce the memory requirements, we employ randomized matrix approximation to compress the state as it is generated. The compressed state is then used to compute approximate gradients and to apply the Hessian to vectors. The approximation error in these quantities is controlled by the target rank of the compressed state. This approximate first- and second-order information can readily be used in any optimization algorithm. As an example, we develop a sketched trust-region method that adaptively learns the target rank using a posteriori error information and provably converges to a stationary point of the original problem. To conclude, we apply our randomized compression to the optimal control of a linear elliptic PDE and the optimal control of fluid flow past a cylinder.

Drew P. Kouri

Optimization and Uncertainty Quantification  
Sandia National Laboratories

dpkouri@sandia.gov

### MS10

#### Sketching with Efficient Tensor Network Embeddings

Tensor network structured data is commonly seen in multiple applications, including quantum simulation, statistical learning and data mining. We introduce tensor network embeddings, which are efficient for the dimensionality reduction of tensor network structured data. We show that the binary tree structured embedding is efficient in terms of both the output sketch size and the leading order computational cost to multiply the embedding with the data. Moreover, we provide algorithms to construct computationally efficient and accurate binary tree embeddings for given tensor network structured data.

Linjian Ma

University of Illinois at Urbana Champaign  
lma16@illinois.edu

### MS11

#### Landau Collision Operator in the Cuda Programming Model Applied to Thermal Quench Plasmas

The Landau form of the Fokker-Planck equation is the gold standard for plasmas dominated by small angle collisions, but its  $O(N^2)$  work complexity has limited its use. This talk extends previous work on a conservative finite element discretization of the Landau operator with new implementations in the CUDA programming model, in both CUDA and Kokkos. Experiments on NVIDIA and AMD GPUs, and preliminary results on a Fujitsu vector processor, are presented. This talk presents new optimizations of the Landau kernel and implementations of the sparse matrix assembly and algebraic solver on the GPU, and shows that a well resolved Landau collision time advance is practical for kinetic plasma applications. We apply this solver to a reduced Vlasov-Maxwell-Landau model of a plasma thermal quench for runaway electron studies. Both the fully implicit Landau time integrator and the Vlasov-Maxwell-Landau model are available in the PETSc (Portable, Extensible, Toolkit for Scientific computing) numerical library.

Mark Adams

Lawrence Berkeley National Laboratory  
mfadams@lbl.gov

### MS11

#### Leveraging Mixed Precision to Accelerate High-Order Finite Element Methods on Gpus

For high-order finite element methods, matrix-free approaches to applying the resulting operators achieve more efficient performance than building a sparse matrix explicitly. libCEED, a product of the Center for Efficient Exascale Discretizations (CEED), offers a flexible, low-level specification for high-order matrix-free finite element methods. libCEED has several GPU backends targeting different architectures. In the case of tensor-product basis functions, its kernels are memory bound, preventing full use of the hardware's computational peak. When lower precision can be tolerated by the application (e.g. for preconditioning or when the computation is known to be stable and full double-precision accuracy in the final result is not necessary), a mixed-precision framework offers a way

to break through the double-precision memory bandwidth barrier. This can entail using lower precision in memory movement between kernels, and potentially also in computation. Non-tensor bases involve more computationally-heavy matrix-matrix multiplications, which can also be accelerated through lower computational precision, sometimes in a hardware-specific way. This talk will discuss the efforts to provide low- and mixed-precision functionality to libCEED's GPU backends, along with performance implications for CEED benchmarks and applications.

Natalie N. Beams

University of Tennessee  
nbeams@icl.utk.edu

Jed Brown, Jeremy Thompson

University of Colorado Boulder

jed@jedbrown.org, jeremy.thompson@colorado.edu

Yohann Dudouit

LLNL

dudouit1@llnl.gov

Will Pazner

Center for Applied Scientific Computing

Lawrence Livermore National Laboratory

pazner1@llnl.gov

### MS11

#### On Portability and Performance Versatility in Non-linear Solid and Fluid Mechanics Using libCEED and PETSc

libCEED is an open-source mathematical software library that is layered in abstractions that handle complexities like multigrid solvers and nonlinearity in real-world simulations without losing performance on different computational device types, including CPUs and GPUs. In this talk, we explore and compare the performance spectrum of some physical scenarios on CPU and GPU machines, as problem size, polynomial degree, machine scale, and physical complexity vary. In doing so, we have used libCEED and PETSc to develop examples in computational fluid dynamics (CFD) and solid mechanics, with which we perform trade-off studies between accuracy, cost, and execution time.

Leila Ghaffari, Jeremy Thompson, Valeria Barra, Rezar

Shakeri, Karen Stengel, Jed Brown

University of Colorado Boulder

leila.ghaffari@colorado.edu,

jeremy.thompson@colorado.edu,

valeria@caltech.edu,

rezgar.shakeri@colorado.edu, karen.stengel@colorado.edu,

jed@jedbrown.org

### MS11

#### Reverse-Mode Automatic Differentiation and Optimization of GPU and Heterogeneous Parallel Programs via Enzyme

Computing derivatives is key to algorithms in scientific computing and machine learning such as optimization, uncertainty quantification, and stability analysis. Enzyme is a LLVM compiler plugin that performs reverse-mode automatic differentiation (AD) and thus generates high performance gradients of programs in languages including C/C++, Fortran, Julia, and Rust. Enzyme is the first fully automatic reverse-mode AD tool to generate gradients of GPU kernels. Moreover, since unlike other tools En-

zyme performs automatic differentiation within a general-purpose compiler, we are able to introduce several novel GPU and AD-specific optimizations. To show the generality and efficiency of our approach, we compute gradients of five GPU-based HPC applications, executed on NVIDIA and AMD GPUs. All benchmarks run within an order of magnitude of the original program's execution time. Without GPU and AD-specific optimizations, gradients of GPU kernels either fail to run from a lack of resources or have infeasible overhead. We demonstrate that increasing the problem size by either increasing the number of threads or increasing the work per thread, does not substantially impact the overhead from differentiation. We will also showcase ongoing work to allow Enzyme to be integrated into any fork-join parallel framework, including OpenMP, MPI, and RAJA, as well as any combination of parallel frameworks (e.g. GPU+MPI).

William S. Moses

Massachusetts Institute of Technology  
wmoses@mit.edu

## MS12

### Exascale Particle-in-Cell Simulations with PIconGPU

Exascale Particle-in-Cell simulations are around the corner. We present the computing challenges one faces when dealing with next generation architectures and our solutions to these challenges. We focus on our own 3D3V electromagnetic particle-in-cell code PIconGPU to highlight the major steps necessary to scale to Exascale. We then present the opportunities that open up when going full Exascale, showcase central applications from laser plasma accelerators to high energy density physics that require Exascale capabilities and finally address the big elephant in the room: What if you can make optimum use of Exascale systems and produce high quality data with it?

Michael H. Bussmann, Alexander Debus, Thomas Kluge, Rene Widera, Sergei Bastrakov, Klaus Steiniger, Richard Pausch  
Helmholtz-Zentrum Dresden-Rossendorf  
m.bussmann@hzdr.de, a.debus@hzdr.de,  
t.kluge@hzdr.de, r.widera@hzdr.de, s.bastrakov@hzdr.de,  
k.steiniger@hzdr.de, r.pausch@hzdr.de

## MS12

### openPMD - Open and F.A.I.R. I/O for Particle-Mesh Data at the Exascale

This talk presents openPMD, and open and F.A.I.R. standard for particle-mesh data, and its impact in Exascale scientific workflows. The openPMD standard is made accessible to scientific software via the openPMD-api, a library for the description of scientific data. It approaches recent challenges posed by hardware heterogeneity by decoupling the data description in domain sciences, such as plasma physics simulations, from concrete implementations in hardware and IO. This concept helps us build a transition path from file-based IO to streaming-based workflows of scientific applications in an HPC environment. The streaming backend is provided by the ADIOS2 framework, developed at Oak Ridge National Laboratory. This talk discusses two openPMD-based loosely coupled setups to demonstrate flexible applicability and to evaluate performance. In loose coupling, as opposed to tight coupling, two (or more) applications are executed separately, e.g. in individual MPI contexts, yet cooperate by exchanging

data. This way, a streaming-based workflow allows for standalone codes instead of tightly-coupled plugins, using a unified streaming-aware API and leveraging high-speed communication infrastructure available in modern compute clusters for massive data exchange. The presented setups show the potential for a more flexible use of compute resources brought by streaming IO as well as the ability to increase throughput by avoiding filesystem bottlenecks.

Franz Poeschel

Helmholtz-Zentrum Dresden-Rossendorf  
Technical University Dresden  
f.poeschel@hzdr.de

## MS12

### alpaka, LLAMA and More Solutions for Exascale Performance Portability

Heterogeneous hardware landscapes will define the Exascale era. At the same time, keeping scientific libraries and applications portable across different hardware setups while maintaining high performance is no trivial matter. Vendor-provided programming platforms often cannot target accelerators from other vendors, and different hardware types like CPUs, GPUs and FPGAs require differently tuned algorithms for optimal performance. A solution to these issues can be found in abstraction layers that provide the user with a single programming interface while still maintaining portability and performance. In this talk, we introduce the Caravan HPC ecosystem. With the *alpaka* abstraction library for accelerator programming at its core and many sibling libraries for related use cases — such as the memory access abstraction layer *LLAMA* or the C++ primitives library *vikunja* — the Caravan ecosystem is an ideal choice for scientists and programmers setting out to tackle the challenges of the Exascale era.

Jan Stephan

TBD  
j.stephan@hzdr.de

## MS13

### Sparse Neighborhood Collectives on Heterogeneous Architectures

Sparse solvers and simulations are often bottlenecked by inter-process boundary exchanges. Typically, these boundary exchanges are implemented with point-to-point communication, leaving optimizations up to the application programmer. Sparse neighborhood collectives, such as the `MPI_Neighbor_alltoallv`, can be utilized to add necessary optimizations within MPI. Optimization challenges arise on heterogeneous architectures, with large numbers of paths between sets of GPUs. During this talk, I discuss optimizations for sparse neighborhood collectives on heterogeneous architectures, and how these optimizations can affect applications.

Amanda Bienz

University of New Mexico  
bienz@unm.edu

## MS13

### A New Performance-Complexity Landscape: JIT-based Fusion for Matrix-Free Multigrid Solvers, with a Case Study in Bonded Granular Media

Matrix-free multigrid solvers are an increasingly necessary to efficiently use modern hardware for science and engi-

neering. Just-in-Time (JIT) compilation with kernel fusion makes it possible to separate concerns while making debuggable high-performance code, but this fusion may only be beneficial when the fused kernels can be scheduled with acceptable occupancy. We'll investigate abstraction leaks in convergence and performance tuning of matrix-free multigrid solvers as the problem specification and complexity of material models changes, all via a case study in bonded granular media.

Jed Brown, Valeria Barra  
University of Colorado Boulder  
jed@jedbrown.org, valeria@caltech.edu

Natalie N. Beams  
University of Tennessee  
nbeams@icl.utk.edu

Yohann Dudouit  
LLNL  
dudouit1@llnl.gov

Rezgar Shakeri, Karen Stengel, Jeremy Thompson  
University of Colorado Boulder  
rezgar.shakeri@colorado.edu, karen.stengel@colorado.edu,  
jeremy.thompson@colorado.edu

### MS13

#### Portability of Solvers for Future Architectures

The last decade and a half have seen the growth of architectures that favor single instruction, multiple threads (SIMT) execution models. The current HPC eco-system of mostly GPU-based architectures has demanded that solvers be built to favor shared-address threaded compute models. Some solvers however have struggled to perform under this paradigm due to data dependencies. Performance is hampered without the ability to fill GPUs with execution threads. This talk will explore what challenges will be faced by developers as hardware moves beyond von Neuman architectures.

Fernanda Foertter  
NextSilicon  
fernanda.foertter@nextsilicon.com

### MS13

#### Composable Solvers in the Ginkgo Library

The Ginkgo Numerical Linear Algebra Library provides a modern C++ interface and kernels for high performance sparse matrix-vector products, iterative solvers and preconditioners for modern GPUs. Inside Ginkgo, all of these operations share the same abstract representation as the application of a linear operator (LinOp)

$$y = Ax$$

with  $A$  being a system matrix, a matrix-free operator, its inverse or an approximation thereof. Just as a simple Krylov solver is composed of a system matrix, a factorization-based preconditioner is composed of the solvers for the individual factors, and a multigrid solver is composed of smoothers, restriction/prolongation operators and a coarse-grid solver. We will study the use of operator composition in the design of novel solvers and preconditioners for massively parallel architecture that trade in synchronization overheads, which can be very costly on these architectures, for higher overall computational effort

in iterative solvers. Finally, we show how these approaches can be extended to distributed and batched solvers as well as mixed-precision algorithms.

Tobias Ribizel  
Karlsruhe Institute of Technology  
tobias.ribizel@kit.edu

### MS14

#### Doubt and Redundancy Kill Soft Errors - Replication-Based Resilience in Exahype

Replication is an established method to improve resiliency, but typically unwelcome in HPC, due to a substantial increase in compute resources. We will present an approach for task-based replication, implemented in the ExaHyPE engine to solve hyperbolic PDE systems, that mitigates replication overhead by sharing task outcomes between teams of MPI ranks. Thus, ideally each team handles only half of the workload. In this setup, we use physics-based error criteria to identify soft errors. We only share task outcomes that yield small error criteria values - otherwise we consider the outcome dubious and compare it with that of the replica task. Disagreeing criteria thus indicate soft errors, but also which replica is likely correct. By switching to the likely correct task outcomes, we obtain a self-healing, resilient algorithm that promises to compensate silent floating-point errors with low overhead regarding performance, I/O and memory footprint.

Philipp Samfass  
Technical University of Munich  
samfass@in.tum.de

Anne Reinartz, Tobias Weinzierl  
Durham University  
anne.k.reinartz@durham.ac.uk,  
tobias.weinzierl@durham.ac.uk

Michael Bader  
Technical University of Munich  
bader@in.tum.de

### MS14

#### Ft-Gcr: a Fault-Tolerant Generalized Conjugate Residual Elliptic Solver

With the steady advance of high performance computing systems featuring smaller and smaller hardware components, the systems and algorithms used for numerical simulations increasingly contend with disruptions caused by hardware failures and bit-levels misrepresentations of computing data. In numerical frameworks exploiting massive processing power, the solution of linear systems often represents the most computationally intensive component. Given the large amount of repeated operations involved, iterative solvers are particularly vulnerable to bit-flips. A new method named FT-GCR is proposed here that supplies the preconditioned Generalized Conjugate Residual Krylov solver with detection of, and recovery from, soft faults. The algorithm tests on the monotonic decrease of the residual norm and, upon failure, restarts the iteration within the local Krylov space. Numerical experiments on the solution of an elliptic problem arising from a stationary flow over an isolated hill on the sphere show the skill of the method in addressing bit-flips on a range of grid sizes and data loss scenarios, with best returns and detection rates obtained for larger corruption events. The simplicity of the method makes it easily extendable to other solvers and

an ideal candidate for algorithmic fault tolerance within integrated model resilience strategies.

Tommaso Benacchio

MOX- Dipartimento di Matematica  
Politecnico di Milano  
Tommaso.benacchio.ext@leonardocompany.com

Mike Gillard  
Loughborough University  
m.gillard@lboro.ac.uk

#### MS14

##### Online Fault Tolerance with Kokkos Resilient Execution Spaces

Hardware heterogeneity in HPC is increasing as we approach the dawn of the exascale era, with performance portable APIs such as Kokkos and RAJA rising to meet the challenge. Accompanying the increase in heterogeneity is an increase in transient faults and soft errors, and industry leaders have identified software resiliency as a critical research need. We propose a framework to tolerate soft errors during execution based on a primary Kokkos abstraction: the resilient execution space. The resilient execution space allows applications already using Kokkos to take advantage of our near-seamless API integration to obtain a user-specified level of resiliency in parallel execution. Soft error detection and mitigation is handled by majority vote in resilient versions of Kokkos parallel execution abstractions. We show performance results on various benchmarks written for Kokkos and modified to use the resilient execution space.

Elisabeth Giem

Sandia National Laboratories  
eagiem@sandia.gov

Nicolas M. Morales  
Sandia National Labs  
nmmoral@sandia.gov

Keita Teranishi  
Sandia National Laboratories  
Livermore, California  
knteran@sandia.gov

#### MS14

##### Asynchrony and Failure Masking via Pseudo-Local Process Recovery in MPI Stencil Applications

Previous work has demonstrated improved scalability of resilience when only failed processes roll back to a checkpoint while other processes continue execution. This technique enables stencil solvers (with only neighbor communication) to exhibit “failure masking” due to limited propagation of local recovery delays. However, localizing the impact of process recovery after a hard failure is challenging because such failures are not transparent to the MPI runtime. In this talk, we describe techniques to perform checkpointing and recovery “pseudo-locally” (along with requisite message logging and replay) and maintain asynchronous progress in MPI applications by adapting the Fenix library, built on User Level Failure Mitigation (ULFM). For a prototype stencil solver with emulated hard failures, results show clear evidence of failure masking in this new setting, which should enable applications to withstand much higher failure rates than possible with global recovery. Our work shows how existing fault-tolerance infrastructure designed

for global checkpoint/restart can be repurposed to enable greater efficiency in a resilience-aware application. [SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.]

Hemanth Kolla, Jackson Mayo  
Sandia National Laboratories  
hnkolla@sandia.gov, jmayo@sandia.gov

Matthew Whitlock  
College of Computing  
Georgia Institute of Technology  
mwhitlo@sandia.gov

Keita Teranishi  
Sandia National Laboratories  
Livermore, California  
knteran@sandia.gov

Rob Armstrong  
Sandia National Laboratories  
rob@sandia.gov

#### MS15

##### Designing Efficient Graph Algorithms on GPUs Through Proxy-Driven Codesign

Developing scalable graph algorithms on GPUs is challenging, due to the inherent irregularities in the graph structure and memory-access intensive computational pattern. Proxy application driven software-hardware codesign plays a vital role in driving innovation among the developments of applications, software infrastructure and hardware architecture. Proxy applications are self-contained and simplified codes that are intended to model the performance-critical computations within applications. In this talk, we discuss facilitating software-hardware codesign through proxy applications with the goal of improving the performance of graph analytics workflows on heterogeneous systems. However, even representative proxy applications may be insufficient to diagnose performance bottlenecks of common graph computational patterns at scale. Therefore, we also discuss the role of derivative benchmarks in enhancing graph applications on GPUs. We will drive the discussion using two applications—Graph matching and clustering, which have applications in the domains of proteomics, computational biology, cybersecurity, numerical analysis and other data science scenarios.

Sayan Ghosh  
Pacific Northwest National Laboratory  
sayan.ghosh@pnnl.gov

#### MS15

##### Sparse Factorization-Based Solvers and Preconditioners on Modern GPUs

We present recent efforts porting the sparse direct solver and preconditioners from the STRUMPACK solver library to GPU architectures. The preconditioners are based on approximate multifrontal factorization using rank-structured matrix compression, which leads to robust and scalable methods applicable to a wide range of problems. We discuss the algorithmic changes and the implementation details that were required to make these algorithms run efficiently on modern NVIDIA and AMD GPUs, using CUDA and HIP/ROCm. We also present initial performance results obtained on Intel accelerators using SYCL and oneAPI. Our implementations rely on emerging stan-

dards for variable sized batched dense linear algebra routines. Our codes run on distributed memory and multi-GPU architectures.

Pieter Ghysels  
Lawrence Berkeley National Laboratory  
Computational Research Division  
pghysels@lbl.gov

Lisa Claus  
Lawrence Berkeley National Laboratory  
lclaus@lbl.gov

Xiaoye S. Li  
Computational Research Division  
Lawrence Berkeley National Laboratory  
xsli@lbl.gov

### MS15

#### Performance Portable Batched Sparse Linear Solvers in Kokkos Kernels

Kokkos Kernels is a library within the Kokkos ecosystem and focuses on the implementation of portable dense and sparse linear algebra algorithms. The library is typically used by higher level math libraries that provide distributed linear algebra algorithms or directly by applications. The later has revealed the need for batched interfaces of linear algebra algorithms. These are then typically used to perform many local operations on a large number of small linear algebra operators. Such use cases arise in computational mechanics to solve problems at integration points, in particle dynamics to compute collisions or when chemical reactions are occurring in fluid flows. In this presentation we discuss the implementation and the performance gains of batched Krylov-based solvers using hierarchical parallelism in Kokkos-Kernels. We explain how the extra dimension of batched linear systems can be used to improve the throughput of the computation and demonstrate performance on multiple architectures.

Kim Liegeois  
University of Liege  
kim.liegeois@ulg.ac.be

Siva Rajamanickam, Luc Berger-Vergiat  
Sandia National Laboratories  
srajama@sandia.gov, lberge@sandia.gov

### MS15

#### PETSc and Kokkos-Kernels Integration

PETSc is a widely used math library for scientific applications modeled by partial differential equations. Traditionally, PETSc relies on MPI for portable and scalable computation between CPU processes and single-threaded C programming for portable and efficient CPU execution within a process. This strategy has served us very well until recently a variety of computer architectures, represented by graphic processing units (GPUs), became the main compute power provider for scientific computations. In preparing PETSc for exascale computing, we met challenges of how to make PETSc still portable and efficient on new architectures, but also avoid code duplication. In this presentation, we report on the effort we use Kokkos and Kokkos-Kernels as a portability layer within PETSc. We'll talk about our design, lessons learned and will also

show some application performance results.

Junchao Zhang  
Argonne National Laboratory  
jc Zhang@anl.gov

Mark Adams  
Lawrence Berkeley National Laboratory  
mfadams@lbl.gov

### MS16

#### Some Domain Decomposition Algorithms in Machine Learning

Domain decomposition (DD) is a general methodology for solving problems involving a large number of variables that are globally coupled. It decomposes a global problem into a set of overlapping or non-overlapping subproblems that are easier to solve than the original problem, and sometime these subproblems can be solved in parallel since they are often independent of each other. A coupling algorithm, which is often an iterative process, is then employed to form the solution of the global problem from the solutions of the subproblems. The methodology has been applied successfully for many problems described by partial differential equations. In this talk, we discuss some preliminary studies of some domain decomposed principal component analysis (PCA) and linear discriminant analysis (LDA) approaches and their applications in unsupervised and supervised learning. This is a joint work with Jingwei Li.

Xiao-Chuan Cai  
University of Macau  
Department of Mathematics  
xccai@um.edu.mo

### MS16

#### Application of Domain Decomposition Algorithms to Physics-Informed Neural Networks

Physics-informed neural networks are new approaches for solutions of partial differential equations and present promising features for many application problems. On the other hand, there are many issues to address for the new approaches, such as accuracy of the approximate solutions and efficient training of parameters. To deal with such issues, the solution can be formed by using local neural networks and the local neural networks are found as solutions of suitable local problems. The idea is similar to that of domain decomposition algorithms. In this talk, we present some experimental study on application of domain decomposition algorithms to find solutions of physics-informed neural networks in pursuit of enhancement of accuracy and efficiency of existing methods.

Hyea Hyun Kim  
Kyung Hee University  
1 Seocheon-dong, Giheung-gu, Youngin-si, Gyeonggi-do,  
446-70  
hhkim@khu.ac.kr

Hee Jun Yang  
Kyung Hee University  
yhjj109@khu.ac.kr

### MS16

#### Deep Learning in Adaptive Domain Decomposition

## Methods

The convergence rate of domain decomposition methods is generally determined by the eigenvalues of the preconditioned system. For second-order elliptic partial differential equations, coefficient discontinuities with a large contrast can lead to a deterioration of the convergence rate. A remedy can be obtained by enhancing the coarse space with elements, which are often called constraints, that are computed by solving small eigenvalue problems on portions of the interface of the domain decomposition, i.e., edges in two dimensions or faces and edges in three dimensions. In general, it is difficult to predict where these constraints have to be computed, i.e., on which edges or faces. Here, a machine learning based strategy using neural networks is suggested to predict the geometric location of these edges or faces in a preprocessing step. This reduces the number of eigenvalue problems that have to be solved during the iteration. Numerical experiments for model problems and realistic microsections using regular decompositions as well as those from graph partitioners are provided, showing very promising results.

### Axel Klawonn

Universität zu Köln  
axel.klawonn@uni-koeln.de

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

Martin Lanser  
Universität zu Köln  
Mathematisches Institut  
martin.lanser@uni-koeln.de

Janine Weber  
Universität zu Köln  
Department of Mathematics and Computer Science  
janine.weber@uni-koeln.de

## MS16

### Two-Level Group Convolution

Group convolution has been widely used in order to reduce the computation time of convolution, which takes most of the training time of convolutional neural networks. However, it is well known that a large number of groups significantly reduce the performance of group convolution. In this talk, we propose a new convolution methodology called "two-level" group convolution that is robust with respect to the increase of the number of groups and suitable for multi-GPU parallel computation. We first observe that the group convolution can be interpreted as a one-level block Jacobi approximation of the standard convolution, which is a popular notion in the field of numerical analysis. In numerical analysis, there have been numerous studies on the two-level method that introduces an intergroup structure that resolves the performance degradation issue without disturbing parallel computation. Motivated by these, we introduce a coarse-level structure which promotes intergroup communication without being a bottleneck in the group convolution. We show that all the additional work induced by the coarse-level structure can be efficiently processed in a distributed memory system. We compare the proposed method to various approaches for group convolution in order to highlight the superiority of the proposed method in terms of execution time, memory efficiency, and

performance.

Chang-Ock Lee  
Department of Mathematical Sciences  
KAIST  
colee@kaist.edu

Youngkyu Lee, Jongho Park  
KAIST  
lyk92@kaist.ac.kr, jongho.park@kaist.ac.kr

## MS17

### Machine-Learning-Based Search for Autotuning

Autotuning has the potential to significantly improve the performance portability of scientific applications on the diverse exascale hardware. Within the DOE exascale project, we have been developing YTOPT/SuRF, a search software package for autotuning. It comprises machine-learning-based search methods that consist of sampling a small number of input parameter configurations, evaluating them, and progressively fitting a surrogate model over the input-output space until exhausting the user-defined time or maximum number of evaluations. In this talk, we will present recent developments on sampling-based Bayesian optimization under algebraic constraints and Monte Carlo tree search for LLVM/Pollys composable loop optimization transformations.

Prasanna Balaprakash, Jaehoon Koo, Xingfu Wu,  
Michael Kruse  
Argonne National Laboratory  
pbalapra@anl.gov, jkoo@anl.gov, xingfu.wu@anl.gov,  
michael.kruse@anl.gov

Paul D. Hovland  
Argonne National Laboratory  
MCS Division  
hovland@anl.gov

Mary Hall  
School of Computing  
University of Utah  
mhall@cs.utah.edu

## MS17

### UFL to GPU: Generating Near-Roofline Finite Element Action Kernels

Recent GPUs are capable of providing peak performance up to 6 TFLOps/s (DP), making them an attractive computational target for Finite Element Methods. However, effectively mapping an FEM solver to a GPU remains challenging due to the scattered memory access, large amounts of on-chip state space (e.g. registers) required for efficient execution, and the inherently large algorithmic variety encountered in local assembly kernels for variational forms. For a general class of finite element operators expressed in UFL (the Unified Form Language), in 1/2/3D and at various approximation orders, we focus on the generation of efficient GPU code for the application of the operator action (i.e. the matrix-vector product). We describe a parametrized family of transformation strategies targeting these kernels, a heuristic cost model, and an auto-tuning strategy that enables us to achieve near-roofline performance for a wide variety of variational forms across domains such as fluid dynamics, solid mechanics, and wave motion. We have integrated our algorithm into the Fire-drake Project, a UFL-based FEM solver framework for

broad accessibility of our work and make a case for software frameworks based on code-generation that progressively lower high-level mathematical abstractions into high-performance codes.

Andreas Kloeckner  
University of Illinois  
andreask@illinois.edu

Kaushik Kulkarni  
University of Illinois at Urbana-Champaign  
kgk2@illinois.edu

### MS18

#### Temporal Alignment: Your Performance Depends On It!

To gain understanding of how HPC resources are being utilized, identify bottlenecks, stay within power bounds, and improve system design, significant effort has gone into monitoring, analysis, and visualization of system resource state. Application developers have also been instrumenting applications to understand sub-component performance, identify bottlenecks, and develop more efficient algorithms and programming methods. While the latter approach can drive innovations in algorithms and improved parallelism, it won't address system resource availability and health issues which have been shown to cause large variation in HPC application run times. Likewise, while the former approach can be used to identify degraded resource performance and bottlenecks, it cannot provide operations staff with insight into what in the associated workflow mix is causing problems. Only by combining time aligned application performance information and system resource state information can correlations between the two be discovered, root causes identified, and problems mitigated. I present work in low-overhead capture of timestamped system resource and application performance data along with tools for run time analysis and time-aligned visualization of both to enable visual correlation of features. This work is foundational for improving application performance and system throughput through automated correlation, root cause analysis, and mitigating feedback to system resources and applications.

Jim Brandt  
Sandia National Laboratories  
brandt@sandia.gov

### MS18

#### A Runtime System for Orchestrating Data Movement in Heterogeneous Memory and Storage Hierarchies

With the emergence of heterogeneous computing elements, such as CPUs, GPUs, and special-purpose FPGAs, data movement among their memories in high performance computing (HPC) systems is complex. In addition, storage devices that are used for storing data for longer term than an applications runtime are becoming heterogeneous, with new non-volatile and storage-class memories. Moreover, traditional parallel file systems are here to stay as capacity storage. Software that glue these various levels of memory and storage devices are diverse. For instance, various memory devices are managed by processors or operating systems, while storage devices are managed by multiple layers of I/O libraries and parallel file systems. In these complex HPC environments, runtime systems play a critical role in orchestrating data movement with the goal of

achieving superior performance and keeping the data closer to processing. With this objective, we have been developing systems, such as Proactive Data Containers (PDC), and enhancing existing libraries, such as HDF5, with asynchronous data movement, caching, and prefetching techniques. In this talk, we will describe our efforts, success stories, lessons learned, and future plans of a robust data movement runtime system.

Suren Byna  
Lawrence Berkeley National Laboratory  
sbyna@lbl.gov

### MS18

#### Autotuning of Execution Parameters and Automatic Response: the Performance Analysis Tool Perspective

Software users and application developers would like to have faith that all parametric choices in software, both hard wired and user configurable, have optimal defaults. In a perfect world, the defaults would be analytically tested values that are guaranteed to provide optimal performance in response to all reasonable inputs on all supported platforms. Sadly, the world is not perfect. While the default settings are usually reasonable, and may even be set after some initial testing and heuristic guidance, these magic number settings rarely provide optimal performance for all inputs on all systems - especially on architectures that didn't exist when the parameter choices were made. Software autotuning has a long, rich history involving several active areas of research. In this talk, we will discuss some research from our group into the area of both offline and online autotuning of software parameters, code transformations, and system parameters to improve performance and/or efficiency. We will also discuss some of the challenges in the area, and requirements for reconfigurable software that can enable more aggressive optimizations for potential future input and system challenges.

Kevin Huck  
Dept. of Computer and Information Science  
University of Oregon  
khuck@cs.uoregon.edu

### MS18

#### Temporal Data Mining in Dynamic Resource Management

Dynamic resource management in an HPC platform needs a responsive resource manager that tracks the changes in applications over the run-time, as well as over life-time (i.e. versions) of the applications. Application owners (i.e. researchers and developers) freely change their applications, hence such tracking inherently needs unsupervised methods to discover signature patterns of an application for identification, distinguishing patterns to characterize changes, and response synthesis based on the discovered patterns. We consider discovering signature patterns from resource utilization traces and develop algorithms to exploit temporal motifs to accurately categorize/cluster jobs into applications. The algorithm uniquely warps the time scale in order to allow for various delays unrelated to the applications. The algorithm exploits multiple resource types (i.e. memory, processor, network) at a higher sampling rate (i.e. 1Hz) to capture a wide variety of signature patterns. On a test dataset collected at a Sandia National Laboratory, the algorithm achieved 90% accuracy in grouping the jobs in applications with common user and working

directory. We demonstrate use cases of additional temporal data mining algorithms in other domains and envision that these algorithms can be transformative in dynamic resource management.

Abdullah Mueen

University of New Mexico.  
mueen@cs.unm.edu

### MS19

#### Scalable Algorithms for Learning Graph Representations

Graphs arise as a natural representation and modeling tool of complex systems across the sciences, such as social, biological, and technological systems, where vertices represent the components of the system, and edges represent their observed interactions. However, the growing data volumes and rates often require a significant boost in the performance of graph algorithms. In particular, graph representation learning is a very promising emerging research area with many industrial applications and it would benefit significantly from more scalable algorithms. In this talk, we will address the challenges of learning structured representations from massive graphs, and we will discuss scalable algorithms for learning structured features and training graph neural networks.

Nesreen Ahmed

Intel Labs  
nesreen.k.ahmed@intel.com

### MS19

#### Sparse Matrices and High-Performance Computing Meet Biology

Recently, the benefit of high-performance computing (HPC) for science has grown rapidly, beyond traditional simulations to data analysis, for example, in genomics. Given the vast amount of data and computation involved in such applications, they can require the full computational power and memory of institutional or agency-wide HPC systems. One of the most data- and compute-intensive challenges in genomics is de novo genome assembly, i.e., reconstructing an unknown genome from redundant, erroneous genomic sequences. Here we introduce the first distributed memory assembler for long-read sequencing data, called ELBA. ELBA introduces sparse matrices as the main abstraction in this context and makes extensive use of sparse linear algebra computation and probabilistic modeling. ELBA is up to 2× faster on CPU than an algorithm based on distributed hash tables, which are harder to parallelize. ELBA integrates GPU support in the most compute-intensive stages of the pipeline to take advantage of today's HPC heterogeneous hardware.

Giulia Guidi

University of California, Berkeley  
gguidi@berkeley.edu

### MS19

#### Graph Neural Network based Framework for Identifying Candidate nodes of Influence Maximization Problem

A lot of Machine learning algorithms have been developed for graphs in recent times which demonstrate high efficacy in solving various graph theory problems. Especially, the graph combinatorial problems have been targeted because

of the high computational complexity and scalability issues of the existing algorithms. Influence maximization (IM) is one such combinatorial problem which is typically solved with the greedy Hill-climbing algorithm (GHC). GHC and in that matter any IM algorithm iteratively search over a large pool of nodes for identifying influential seed nodes. The search space increases drastically with the size of the network leading to a significant rise in the algorithm running time. In a typical graph, only a certain fraction of nodes are likely to contribute to the IM solution set and, the rest of the nodes are nonessential. In this talk, we will present our ongoing work on identifying such essential nodes known as candidate nodes. Specifically, we have framed this identification problem as a binary node classification problem where node classes are "candidate" and "non-candidate". A Graph neural network model is trained to predict node classes and thereafter, only candidate nodes are searched to identify Influential seed nodes. We will demonstrate the substantial reduction of algorithm running time without compromising the quality of seed nodes.

Sai Munikoti

Kansas State University  
saimunikoti@ksu.edu

Marco Minutoli

Pacific Northwest National Lab  
marco.minutoli@pnnl.gov

Bala Natarajan

Kansas State University  
bala@ksu.edu

Mahantesh Halappanavar

Pacific Northwest National Lab  
mahantesh.halappanavar@pnnl.gov

### MS19

#### Large Scale Structural Graph Representation Learning

Graph representation learning methods have attracted an increasing amount of attention in recent years. These methods focus on learning a numerical representation of the nodes in a graph. Learning these representations is a powerful instrument for scientific computing applications involving core tasks such as graph mining, visualization, hashing, etc. Graph representation learning is of particular interest because they facilitate the direct use of standard machine learning models on graphs. Graph representation learning methods can be divided into two main categories, methods preserving the connectivity information of the nodes and methods preserving nodes' structural information. Connectivity-based methods focus on encoding relationships between nodes, with connected nodes being closer together in the resulting latent space. While methods preserving structure generate a new space where nodes serving a similar structural function in the network are encoded close to each other, independently of them being connected or even close to each other in the graph. While preserving node connectivity is a well-established field, structural graph representation learning still has many open problems. This talk provides an overview of the most recent structural representation approaches, their complexity in relation to their expressive power, and gives parallel computing direction to enhance the expressive power of graph

representation learning.

Edoardo Serra  
Boise State University  
edoardoserra@boisestate.edu

## MS20

### Scaling AMReX Applications using a Bittree Framework

Flash-X is a multiphysics code with octree based adaptive mesh refinement (AMR), and AMReX is a block structured AMR library with no constraint on how refined regions are organized. Integration of AMReX into FLash-X has provided an opportunity to examine its performance in comparison to FLash-Xs native Octree mode with Paramesh. AMReXs native refinement algorithm provides better refinement efficiency, while octree has better scaling behavior. We have integrated grid generation using Bittree, a compact octree bitmap representation, into AMReX as an alternative for applications where flexibility of AMReXs native algorithm is not needed. In this talk we present the details of this integration and its impact on AMReXs performance in Flash-X.

Akash V. Dhruv, Tom Klosterman  
Argonne National Laboratory  
adhruv@anl.gov, tklosterman@anl.gov

Wei-qun Zhang  
Lawrence Berkeley National Laboratory  
Center for Computational Sciences and Engineering  
weiqunzhang@lbl.gov

Klaus Weide  
University of Chicago  
klaus@flash.uchicago.edu

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

Anshu Dubey  
Argonne National Laboratory  
adubey@anl.gov

## MS20

### Leveraging Parallel AMR Capabilities in Overset Grid Simulations of Aerospace and Wind Energy Applications

Many aerodynamics problems require multi-scale simulation capabilities to track turbulent vortical fluid structures such as the wakes generated by wind turbines or helicopter rotors. To resolve these complex multi-scale phenomena, parallel adaptive mesh refinement paired with overset grid technologies provides a computationally efficient and scalable solution. We highlight various AMR approaches, such as octree-based and patch-based systems, in tandem with the TIOGA overset grid assembler for simulating these relative-body motion problems. Moreover, we examine the algorithmic characteristics of each approach for providing the necessary data required by the overset grid assembler. Lastly, we'll discuss possible AMR and overset capabilities needed for leveraging next-generation heterogeneous computing systems.

Andrew Kirby  
University of Wyoming

akirby@uwyo.edu

Michael Brazell  
National Renewable Energy Lab  
michael.brazell@nrel.gov

Dimitri Mavriplis  
Department of Mechanical Engineering  
University of Wyoming  
mavripl@uwyo.edu

## MS20

### High Resolution Two-phase Flow Simulations based on a Weakly Compressible Scheme with an Interface-Adapted AMR Method

A fully explicit weakly compressible scheme for solving low Mach number gas-liquid twophase flows with interface-adapted AMR method is proposed to achieve large-scale and high-resolution simulation by avoiding solving pressure Poisson equation. Fully explicit time integration is developed by solving independent hyperbolic pressure evolution derived from isothermal Navier-Stokes equation instead of pressure Poisson equation which has a severe scalability problem. We have developed a GPU code of the tree-based AMR method, which can significantly reduce the computational cost to assign high-resolution mesh to the region of moving interfaces. Many high-resolution applications such as oil jets and flows with liquid films have been successfully calculated.

Shintaro Matsushita  
Department of Mechanical Engineering  
Tokyo Institute of Technology  
matsushita.s.ad@m.titech.ac.jp

Takayuki Aoki  
Tokyo Institute of Technology  
taoki@gsic.titech.ac.jp

## MS20

### Fusanov - a Task-Based, Fused Rusanov Solver over Dynamically Adaptive Meshes

Dynamic adaptive mesh refinement is not a direct fit to accelerators (GPUs) which are optimised streaming architectures, i.e. perform the better the more uniform and regular all operations. Block-structured AMR accommodates this fact by embedding regular patches into the AMR (meta-)mesh yet forces us to compromise between mesh granularity and pre-patch efficiency. We propose an alternative strategy which combines the advantages of large patches with totally dynamic, fine-granular AMR: Enclave tasking identifies cells (aka tasks) within the compute mesh that can be treated with low priorities. Our proposal is to collect these cells into a batches of (uniform) tasks at runtime, and then to schedule them in one go as one compute kernel on a GPU. Such a batch allows us to write compute kernels which are optimised over multiple patches (which can be spatially disconnected). We use this approach to deliver a fused Rusanov solver over multiple cells: a Fusanov solver.

Holger Schulz  
University of Durham  
holger.schulz@durham.ac.uk

Tobias Weinzierl  
Durham University

tobias.weinzierl@durham.ac.uk

## MS21

### FAIR Research Software as a Step Towards Better Research

In recent years, the scholarly community has examined its culture and practices and found a set of overlapping areas to improve. Those that involve research software include open science (making the outputs and processes of scholarly research available), reproducibility (increasing trust in scholarly results by making them repeatable), FAIR (making software findable, accessible, interoperable, and reusable), citation (making software directly citable and encouraging the practice of citation), sustainability (ensuring that software will continue to be available in the future, on new platforms, meeting new needs), and career paths (recognizing and supporting essential roles such as research software engineers, particularly in academia). While the scholarly community is generally supportive of all of these efforts, the degree of support wanes with both the amount of extra work needed and the lack of clear details on how to achieve them, along with misaligned incentives. This talk will focus on FAIR and how it can be applied to research software. This leads to a set of distinct challenges, including scope (defining research software), principles (defining what findable, accessible, interoperable, and reusable mean for research software), implementation (developing guidelines and instructions for how to make research software FAIR), adoption (getting different stakeholders to recognize and use the principles) and metrics (measuring the adoption of FAIR for research software).

Daniel S. Katz

NCSA

University of Illinois

d.katz@iee.org

Michelle Barker

Research Software Alliance

barkermd@outlook.com

Neil P. Chue Hong

Software Sustainability Institute

University of Edinburgh

n.chuehong@software.ac.uk

Leyla Jael Castro

ZB MED

ljgarcia@zbmed.de

Morane Gruenpeter

Software Heritage

morane@softwareheritage.org

Jennifer Harrow

ELIXIR-Europe

jen.harrow@elixir-europe.org

Carlos Martinez

Netherlands eScience Center

c.martinez@esciencecenter.nl

Paula Andrea Martinez

Australia Research Data Commons

orchid00@gmail.com

Fotis E. Psomopoulos

Centre for Research and Technology Hellas

fpsom@certh.gr

## MS21

### UX-Focused Reproducible Analysis Toolkit for Mars 2020

The Mars 2020 Mast Camera Zoom (Mastcam-Z) is a stereoscopic multispectral camera pair mounted on the Perseverance rover. A Mastcam-Z multispectral observation generally consists of 12 stereo-paired 4-megapixel images taken through standard RGB bayer and narrowband filters optimized for the Martian surface. These data are both dense and urgent. They contain a large amount of high-dimensional information that must be analyzed in hours to days to inform tactical decisions, like where to drive the rover or what to target for sample collection. We developed a suite of tools to improve speed and consistency of last-mile reduction of these data. The tools automate many best practices and provide a minimal interface to enable others. Scientists who are not software engineers can interrogate and modify the tools: their operations are defined in markup files that can be edited without reading the entire source. They are fast, easy to use, and create many distinct, fully-reproducible views on the data in formats suitable for both archival and interchange (including at web endpoints). By providing both variety and consistency, they reduce cognitive overhead and facilitate collaboration. We will outline our approaches to project planning, user research, and implementation, and discuss how they could be used to produce software that is durable, reproducible, and responsive to research needs – even in domains less alien than Martian robotics.

Chase Million

Million Concepts

chase@millionconcepts.com

Michael St. Clair

MillionConcepts

mstclair@millionconcepts.com

Melissa Rice

Western Washington University

ricem5@wwu.edu

## MS21

### Productive Performance Portability: Building in Rust with PETSc and libCEED

C/C++ and Fortran are mainstays of high performance computing, but modern languages such as Julia and Rust are gaining popularity for HPC applications. PETSc, the Portable, Extensible Toolkit for Scientific Computation, is a suite of data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations. libCEED provides fast algebra for element-based discretizations, designed for performance portability, run-time flexibility, and clean embedding in higher level libraries and applications. In this talk, we introduce the PETSc and libCEED Rust interfaces as well as a mini-application based upon these interfaces. Via this application, we will demonstrate the productivity benefits when compared to our C based mini-applications, particularly with respect to debugging, documentation, and testing.

Jeremy Thompson, Jed Brown, Zack Jorquera

University of Colorado Boulder

jeremy.thompson@colorado.edu, jed@jedbrown.org, zakary.jorquera@colorado.edu, nurban@bnl.gov

### MS21

#### Impacting Software Quality and Process Through the Extreme-Scale Scientific Software Stack (E4S) and Software Development Kit (SDK) Projects

The E4S and SDK projects improve quality and efficiency within the scientific software ecosystem through containerization, build caches, continuous integration (CI) testing, improving interoperability between packages, and community policies. The E4S containerization effort delivers solutions that raise productivity by providing a tested environment for development and delivering tutorials. The E4S build cache increases productivity, providing thousands of pre-built binaries for a variety of platforms. The containers and build cache also enhance the reproducibility of configurations and errors. E4S promotes the sustainability of the scientific software ecosystem through a multi-platform CI testing infrastructure. Covering many areas including development tools, math libraries, and data and visualization, individual SDKs enhance the productivity of application developers in a variety of ways, including increasing the interoperability and portability of software within the SDK, and providing an additional, complimentary layer of CI testing. Community policies at the E4S level, and SDK-specific policies promote long-term sustainability through minimal software standards in a variety of areas, including building, testing, documentation, and error handling. This talk will introduce many ways in which E4S and the SDK projects are increasing productivity, and improving sustainability and reproducibility, as well as touch on future efforts in these areas.

James Willenbring  
Sandia National Laboratories  
jmwill@sandia.gov

### MS22

#### Distributed, In-Situ Spatial Inference in Climate Simulations with Sparse Gaussian Processes

As memory and storage cost for extreme-scale physics simulation becomes the greatest performance bottleneck, the ability to access full simulation data for post-hoc statistical analysis is becoming increasingly limited. The capability to perform in-situ statistical inference of state variables is critical for comprehensive utilization of information generated by these simulations. In this presentation, we report results fitting spatial extreme value models underpinned by scalable Gaussian process regression to the state information of an expensive simulation in-situ. The inference was performed using Julia coupled to the E3SM climate model and utilizes the parallel, distributed nature of the climate computation for scaling the inference. The demonstrated capability for in-situ climate analysis of the simulation unlocks new tools for UQ in extreme-scale scientific computing.

Michael Grosskopf, Earl Lawrence, Ayan Biswas, Kelin Rumsey  
Los Alamos National Laboratory  
mikegros@lanl.gov, earl@lanl.gov, ayan@lanl.gov, knrumsey@lanl.gov

Nathan Urban  
Brookhaven National Laboratory

### MS22

#### Korali: Efficient and Scalable Software Framework for Bayesian Uncertainty Quantification and Stochastic Optimization

We present Korali, an open-source framework for large-scale Bayesian uncertainty quantification and stochastic optimization. The framework relies on non-intrusive sampling of computational models and enables their exploitation for optimization and decision-making. In addition, its distributed sampling engine makes efficient use of massively-parallel architectures while introducing novel fault tolerance and load balancing mechanisms. We demonstrate these features by showing experimental results where we interface Korali with existing high-performance CPU and GPU-based simulation libraries. These results showcase a highly efficient scaling on up to 512 nodes of the CSCS Piz Daint supercomputer. Finally, we present benchmarks demonstrating that Korali outperforms related state-of-the-art software frameworks.

Petros Koumoutsakos  
Harvard University  
petros@seas.harvard.edu

Sergio Martin, Georgios Arampatzis  
ETH Zürich  
martiser@ethz.ch, garampat@ethz.ch

### MS22

#### Estimation of Galaxy Redshift with Probabilistic Neural Networks

Astrophysical data analyses like estimations of galactic distances often rely on representative data from telescope observations. Data-driven methods of this type are subject to a number of potential deficiencies, such as sample incompleteness and noise. In addition, redshift estimations often struggle with scalability to millions of galaxies and robust uncertainty quantifications. Motivated by these considerations, we propose an approach for augmenting data-based protocols for training machine learning models by using physically motivated synthetic data. A suitable statistical model can then be trained with a matched distribution and realistically modeled synthetic data in hand. A neural network for conditional density estimation was utilized for this purpose, demonstrating a superior accuracy over traditional redshift estimation techniques and interpretable error estimations. With this example, we will discuss some of the crucial aspects of scientific machine learning, with an emphasis on adaptability to high performance computing systems and importance of uncertainty quantification.

Nesar Ramachandra  
Argonne National Laboratory  
nramachandra@anl.gov

### MS22

#### Calibration of a Computationally Intensive Model with Parallel Computing Aspects

Calibration of a computer model widely requires multiple parallel runs of computationally challenging simulations of the same model with different parameter settings. The coordination of parallel runs is important to accelerate the calibration process and efficiently utilize from the compu-

tational resources. In case of failures, an early detection mechanism is required to leave resources idle due to the imbalance in the run times of the parallel tasks. In other situations, irregular sending/receiving of data between manager and workers is necessary to sequentially update the surrogate and calibration model. In this talk, we provide algorithmic insights on the coordination of parallel runs using surmise, a Python package that is designed to provide a surrogate model interface for calibration, uncertainty quantification, and sensitivity analysis, and illustrate these by calibrating computationally expensive low-energy physics code with surrogate-assisted methods under surmise.

Ozge Surer  
Northwestern University  
ozgesurer2019@u.northwestern.edu

### MS23

#### Performance Optimizations of Low-Rank Matrix Computations for 3D Unstructured Mesh Deformation

We present a framework that couples the PaRSEC runtime and the HiCMA library to solve challenging 3D data-sparse problems. This work highlights how a software bundle powered by a task-based programming model can address the heterogeneous workloads engendered by compressing the dense operator. Using Tile Low-Rank approximation, our approach consists in capturing the most significant information in each tile of the matrix using a threshold. Matrix operations are performed on the compressed data layout, reducing memory footprint and algorithmic complexity. Our proposed software solution accommodates a range of data structures of linear algebra from dense and data-sparse to sparse. We introduce two runtime optimizations to address the challenges encountered when confronted with a large rank disparity: (1) a trimming procedure performed at runtime to cut away data dependencies from the directed acyclic graph discovered to be no longer required after compression and (2) a rank-aware diamond-shaped data distribution to mitigate the load imbalance overheads and reduce data movement. We assess our implementation using 3D unstructured mesh deformation based on Radial Basis Function interpolation. Our implementation shows up to 7-fold performance superiority in situations where the 3D unstructured mesh deformation application renders a matrix operator with low density. Our software framework solves a formally dense 3D problem with 52M mesh points on 65K cores in about half an hour.

Rabab Al-omairy  
King Abdullah University of Science and Technology (KAUST)  
rabab.omairy@kaust.edu.sa

Qinglei Cao  
University of Tennessee Knoxville  
qciao3@vols.utk.edu

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

George Bosilca  
University of Tennessee Knoxville  
bosilca@icl.utk.edu

David E. Keyes  
King Abdullah University of Science and Technology

(KAUST)  
david.keyes@kaust.edu.sa

Jack J. Dongarra  
University of Tennessee, Oak Ridge National Laboratory, USA  
dongarra@icl.utk.edu

Yu Pei  
University of Tennessee  
ypei2@vols.utk.edu

### MS23

#### Compression Techniques for Large Three-Dimensional Boundary Element Methods

Modern algorithms for boundary element problems are frequently based on a tree-structured decomposition of the matrix into submatrices that can be approximated by low-rank factorizations. When parallelizing these algorithms, we have to take the tree structure into account, e.g., by using threads to handle subtrees. More advanced algorithms can be rewritten to reduce the number of synchronization barriers. This talk outlines the techniques used to parallelize fast BEM algorithms and presents experiments indicating that fairly large matrices can be treated efficiently without sacrificing the theoretically optimal convergence rate.

Steffen Börm  
University of Kiel  
Germany  
boerm@math.uni-kiel.de

### MS23

#### Uniform-H: Bridging the Gap Between H and H

While H-matrices enable significant storage improvements compared to H-matrices due to shared and nested bases for the low-rank blocks, matrix arithmetic is much more complicated. This is mainly due to the only implicitly defined data of the bases and the potentially large number of blocks in the matrix affected by a single low-rank update. In uniform H-matrices, only low-rank blocks in a single block row or column share a basis, i.e., without basis nestedness. This brings memory consumption close to H-matrices. Furthermore, matrix arithmetic is much simpler as all data is explicitly stored and low-rank updates also affect only a small number of matrix blocks. We will introduce the Uniform-H-matrix format together with a matrix arithmetic and show numerical examples for several different model problems to demonstrate the various properties of this H-matrix format.

Ronald Kriemann  
Max Planck Institute for Mathematics i.t.S.  
rok@mis.mpg.de

### MS23

#### The Submatrix Method within Electronic-Structure Theory and Approximate Computing

We present the submatrix method and a novel linear-scaling electronic-structure method in conjunction with approximate computing, as well as the implementation of the technique in CP2K. Even though initially proposed for inverse p-th roots, it has recently been recognized that the submatrix method represents a general method to approx-

imate arbitrary matrix functions such as the matrix-sign function of large sparse matrices. The Matrix-sign function is the essential workhorse of linear-scaling electronic-structure theory, and we present an intuitive chemical justification for the accuracy of the submatrix method. We will discuss the efficient implementation of the submatrix method into CP2K with a special focus on limiting communication between compute nodes. The resulting compute kernel is the sign function of a relatively small but dense matrix. Our optimized implementation with a simple diagonalization-based evaluation of the sign function of the submatrices outperforms the Newton-Schulz Sign iteration in initial results, especially for larger cutoffs of matrix elements. This observation shows that the submatrix method will be a valuable tool in the context of approximate computing.

Thomas D. Kühne

Paderborn Center for Parallel Computing  
thomas.kuehne@upb.de

#### MS24

##### Multigrid Methods for Block Operators on Structured Meshes

Simulations involving PDEs in computational science and engineering are often carried out on structured meshes if the application permits to do so. The use of structured meshes is beneficial not only to define the grid hierarchy but also to efficiently implement the solvers on modern architectures. Further, the discretization on structured meshes leads to structured matrices that can be analyzed well. Multigrid methods for structured matrices stemming from the discretization of scalar problems have been analyzed in depth in the literature. Very little work exists for block matrices. Block matrices appear in different cases: In the simplest case they appear when scalar PDEs are discretized using higher order discretizations. Of larger interest is the case of systems of PDEs, where the unknowns represent different physical quantities, e.g., in elasticity, fluid flow simulations, and coupled multiphysics problems. Using the well-known variational theory for multigrid method, e.g. [McCormick and Ruge, SIAM J. Numer. Anal., 1982], as well as recent results from [Notay, Numer. Math., 2016] we are able to analyze multigrid methods for both cases, yielding optimal convergence, similar to the results for scalar problems [Aricò, Donatelli, Numer. Math., 2007]. These theoretical findings allow for efficient multigrid methods for structured meshes on modern parallel computers. In the talk the obtained results are presented and discussed.

Matthias Bolten, Isabella Furci

University of Wuppertal  
bolten@math.uni-wuppertal.de, furci@uni-wuppertal.de

#### MS24

##### Recent GPU Developments of Multigrid Reduction in Hypr for Problems in Porous and Fractured Media

In many fully-implicit multi-physics simulators, including those for porous media applications, the main computational cost typically rests on the linear solve. Because of the wide range of processes and couplings that may be involved e.g., formation and propagation of fractures, deformation of the solid porous medium, viscous flow of one or more fluids in the pores and fractures, complicated well sources and sinks, etc. it is difficult to develop general-

purpose and scalable linear solver. This challenge is further aggravated by the range of different discretization schemes that may be adopted, which has a direct impact on the linear system structure. Recent developments of a multigrid reduction (MGR) framework allow one to explore and experiment with scalable preconditioners for a wide range of these coupled problems. In this talk, I will present new results of porting an implementation of MGR in Hypr to GPU. Several numerical examples of difficult, field-scale problems simulated by GEOSX using MGR will be presented, notably: a hybrid discretization of single-phase flow, compositional multiphase flow with and without complex wells, and hydraulic fracturing.

Quan M. Bui

Lawrence Livermore National Laboratory  
bui9@llnl.gov

Francois P. Hamon

Total Exploration & Production  
francois.hamon@total.com

Nicola Castelletto

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

Daniel Osei-Kuffuor, Joshua A. White

Lawrence Livermore National Laboratory  
oseikuffuor1@llnl.gov, white230@llnl.gov

Randolph Settgast

Lawrence Livermore National Laboratory  
settgast1@llnl.gov

#### MS24

##### Preparing Trilinos Solvers for Exascale Wind Farm Simulations

Scalable solvers such as multigrid are a key enabling technology for applications seeking to exploit exascale supercomputers. In this presentation we will discuss recent advances in multigrid preconditioners from the Trilinos project in support of large scale simulations conducted by the ExaWind project (part of the Exascale Computing Project). We will focus on the adaptation of aggregation-based algebraic multigrid within Trilinos to GPU architectures; discuss progress, ongoing work, and challenges; and present numerical results from the Oak Ridge supercomputer Summit.

Jonathan J. Hu

Sandia National Laboratories  
Livermore, CA 94551  
jhu@sandia.gov

Luc Berger-Vergiat, Ichitaro Yamazaki

Sandia National Laboratories  
lberge@sandia.gov, iyamaza@sandia.gov

#### MS24

##### Multiphysics Components for High Fidelity ICF Hohlräum Simulations

In laser-indirect-drive (LID) ICF, laser energy is converted to x-rays by illuminating the inner surface of a hollow metal cylinder (hohlräum) surrounding a capsule. The x-rays bathe the capsule to compress and heat the fuel by

the rocket force produced by ablation of the outer capsule layer. Then, the fuel participates in nuclear fusion. It is now understood that departures from hydro conditions exhibited by hohlraum plasmas (such as kinetic plasma interpenetration and mix), nonlinear laser-plasma interactions (LPI), and electromagnetic (EM) effects pose major uncertainties for the interpretation of experimental data. Integrated simulations of LID ICF implosions are remarkably difficult as they require capturing faithfully the physical processes underlying laser propagation in the hohlraum, laser energy conversion to x-rays and delivery to the ablator surface, the ablation process, subsequent capsule implosion, and finally ignition and thermonuclear burn. All these processes occur in plasmas in different regimes of coupling and collisionality, and demand tailored simulation strategies. We will report on our ongoing efforts to develop a simulation tool capable of addressing these challenges focusing on some aspects of the full problem such as multi-scale methods for kinetic plasma and radiation transport simulations, implicit EM algorithms, as well as a robust coupling strategy between these components at scale.

Bobby Philip, Luis Chacon, Adam Stanier, Guangye Chen  
Los Alamos National Laboratory  
bphilip@lanl.gov, chacon@lanl.gov, stanier@lanl.gov,  
gchen@lanl.gov

Hans R. Hammer  
Los Alamos National Laboratory  
hrhammer@lanl.gov

## MS25

### High Performant Scalable R-Adaptive Particle In Cell Method with Monotonic Reconstruction Remapping

Particle-in-Cell (PIC) methods for flow are good choices when complex moving interfaces and discontinuities interact with complex geometries due to the low-dissipation nature of Lagrangian advection combined with the ability of particles to interact with various geometry descriptions. In our talk, Polynomial Particle in Cell (PolyPIC) method with additional higher order polynomial description of the particle velocity mode is used to increase the numerical accuracy and stability of compressible flow simulations. Standard r-adaptive mesh redistribution is also demonstrated to improve the solution by improving mesh quality around the steep gradients of the solution through the solution of a mesh motion equation. We improve the mesh motion procedure by introducing a new monotonic reconstruction particle remapping to maintain a good distribution of particle along the computational domain, which guarantees stability and accuracy. The method is demonstrated with canonical benchmark problems including Sod Shock, Woodward-Colella wave and other Riemann problems by demonstrating adaptively populated meshes and corresponding monotonically reconstructed particles.

Kwitae Chong  
Oak Ridge National Laboratory  
chongk@ornl.gov

Stuart Slattery  
Computational Sciences and Engineering Division  
Oak Ridge National Laboratory  
slatterysr@ornl.gov

Luka Malencia  
Oak Ridge National Laboratory

tbd

Massimiliano Lupo Pasini  
Department of Mathematics and Computer Science  
Emory University  
lupopasinim@ornl.gov

## MS25

### HACC+Cabana: Performance Portability for PIC Methods in Cosmological N-Body Simulations

The Hardware/Hybrid Accelerated Cosmology Code (HACC) is a high-performance cosmological N-body structure formation code designed to run extreme-scale simulations on the largest supercomputers. Gravity has infinite extent and is unshielded, so it dominates on the largest scales. HACC uses force splitting techniques to divide the gravitational force into long-range and short-range components. The long-range component is calculated with Particle-In-Cell (PIC) methods, where a Particle-Mesh (PM) interaction deposits density onto a grid, and then particle velocities are updated with forces from a spectral Poisson solver that employs distributed-memory 3D FFTs. Careful spectral shaping is used to limit the force handover to roughly 3 grid cell lengths, limiting the number of particle-particle interactions needed for the short-range force. The short-range force kernel is the most FLOPS-intensive portion of HACC, and specific implementations have been developed and optimized by hand for a variety of architectures. HACC developers are building a new Particle-Particle-Mesh (P3M) mini-app with the ECP/CoPA Cabana Particle Toolkit, using Cabana and Kokkos data structures and algorithms whenever suitable. This code is expected to facilitate prototyping algorithmic designs for HACC, exploring optimal configurations for different architectures (eg. ARM, GPUs, MPI-aware accelerators), and evaluating available backends for Kokkos and Cabana (eg. OpenMP threads, CUDA, HIP, SYCL).

Adrian Pope, Esteban Rangel  
Argonne National Laboratory  
apope@anl.gov, erangel@anl.gov

## MS25

### Performance Portable Additive Manufacturing Simulations with the Material Point Method

PicassoMPM is a new material point method application with a focus on high-fidelity simulation of additive manufacturing (AM). We first describe the software design, where PicassoMPM is built on libraries from three parts of the Exascale Computing Project (ECP): Kokkos for performance portability (software technology), Cabana for particle algorithms and communication (co-design), and Picasso for PIC and MPM motifs (ExaAM application). This strategy enables sharing development across the widest set of applications possible, e.g. other particle methods or MPM for a different target problem. This design is also crucial for simulation across HPC architectures, including ORNL Summit and Frontier from which we show initial performance and discuss potential optimizations. Finally, we discuss the calibration, comparison, and predictive capability of PicassoMPM for laser powder bed fusion against experimental data and other simulation methods as part of ExaAM.

Sam Reeve, Kwitae Chong  
Oak Ridge National Laboratory  
reevest@ornl.gov, chongk@ornl.gov

Stuart Slattery  
 Computational Sciences and Engineering Division  
 Oak Ridge National Laboratory  
 slatterysr@ornl.gov

## MS26

### The Hazards and Challenges of Low and Mixed Precision Computation

Within the next few years, exascale-level machines will become a reality. Making use of this available computing power in applications will require a massive effort towards optimizing parallel large-scale computations. This challenge is complicated by the increasing heterogeneity of machines at the node level. Notably is the emergence of low and mixed precision hardware, the use of which offers potentially huge performance gains. A common technique in scientific computing is to introduce inexactness into the computation to reduce computation time, e.g. in model order reduction, sparsification, low-rank approximation, and randomized and communication-avoiding algorithms. In many cases, roundoff error is assumed to be orders of magnitude smaller than any approximation error, and thus the interaction between these different errors is often ignored. In the world of low precision computation, however, this may no longer be a valid assumption. We therefore ask the question: When does floating-point error matter, and what can go wrong if we use low precision blindly? This talk will give examples, summarize some existing work in this area, and explain the remaining gaps in our knowledge.

Erin Claire C. Carson  
 Charles University MFF  
 carson@karlin.mff.cuni.cz

## MS26

### Mixed Precision Strategies for GMRES

Sparse iterative linear solvers are crucial to many applications. In order to effectively make use of new high performance computers, these solvers must be able to exploit hardware for low precision (32 and 16 bit) computations. We provide experimental studies for low precision GMRES wrapped in an iterative refinement loop (GMRES-IR). We discuss the convergence obtained with GMRES-IR and show how it can give up to 1.35 times speedup over traditional fp64 GMRES. Furthermore, we provide a model for speedup of the sparse matrix-vector product (SpMV) and discuss low-precision preconditioning strategies for GMRES. Finally, we will discuss options for mixed precision GMRES in the software library Trilinos.

Jennifer A. Loe, Christian Glusa, Ichitaro Yamazaki  
 Sandia National Laboratories  
 jloe@sandia.gov, caglusa@sandia.gov,  
 iyamaza@sandia.gov

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

Sivasankaran Rajamanickam  
 Sandia National Laboratories  
 srajama@sandia.gov

## MS26

### Mixed Precision Solver Capabilities for the Hypr

## Linear Solver Library

Mixed precision strategies are a class of approximate computing techniques that are well-suited to iterative algorithms. The increasing cost of data movement, coupled with the significant discrepancy in computational throughput between the different levels of precision in HPC hardware, has led to increased interest in the development of mixed-precision strategies as an important tool for designing efficient and numerically stable linear algebra kernels and algorithms. This talk will address some of our efforts to enable mixed-precision algorithm development within the Hypr linear solver library, to improve performance and make efficient use of existing and emerging HPC hardware.

Daniel Osei-Kuffuor, Ulrike Meier Yang, Robert D. Falgout  
 Lawrence Livermore National Laboratory  
 oseikuffuor1@llnl.gov, yang11@llnl.gov, falgout2@llnl.gov

## MS26

### Mixed Precision Iterative Refinement with Approximate Factorization for the Solution of Large Sparse Systems

The increasing availability of low precision arithmetic in hardware and the growing popularity of numerical approximation methods pave the way to new opportunities for improving speed, memory, and energy consumption for the solution of linear systems. State of the art iterative refinement methods, three precisions LU-based and five precisions GMRES-based iterative refinement, are able to compute an accurate solution from an approximate and/or low precision factorization while preserving performance and robustness. In this talk we will showcase the effectiveness of some of these algorithms for the parallel solution of large scale sparse linear systems through an implementation that relies on the MUMPS sparse direct solver. In particular we will investigate the case of Block Low Rank factorization, static pivoting and a combination of the two on problems from a variety of real life industrial and academic applications. In addition, we provide new iterative refinement convergence conditions for a general approximate factorization.

Patrick R. Amestoy  
 Mumps Technologies  
 patrick.amestoy@mumps-tech.com

Alfredo Buttari  
 CNRS-IRIT-Université de Toulouse, France  
 alfredo.buttari@irit.fr

Nicholas J. Higham  
 School of Mathematics  
 The University of Manchester  
 nick.higham@manchester.ac.uk

Jean-Yves L'Excellent  
 Mumps Technologies  
 jean-yves.l.excellent@mumps-tech.com

Theo Mary  
 Sorbonne Université, CNRS, LIP6  
 theo.mary@lip6.fr

Bastien Vieublé  
 Toulouse INP, IRIT

bastien.vieuble@irit.fr

## MS27

### Accelerating Space and Space-Time Statistical Modeling with Mixed-Precision Arithmetic

Geostatistics predicts desired quantities from geographically distributed data, based on statistical models. A primary computational kernel of spatial statistics is the maximum log-likelihood estimation (MLE) function, whose central data structure is a dense covariance matrix that requires two operations, i.e., inverse and determinant evaluation, with the cost of quadratic storage and cubic arithmetic complexity. To reduce this complexity, we migrate geostatistics to three precisions approximation (double/single/half) based on the distances-based structure of the application. We exploited the advantages of Nvidia Tensor Core technology to accelerate space and space-time modeling on large scale. Primary experiments on 128 Summit nodes show that our method delivers 6.5-8 Pflops/s sustained performance with mixed-precision to obtain the required application accuracy on 1M locations real datasets compared to 3.5 Pflops/s with double-precision execution of a single iteration of the MLE operation.

Sameh Abdulah

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

Qinglei Cao

University of Tennessee Knoxville  
qcao3@vols.utk.edu

Yu Pei

University of Tennessee  
ypei2@vols.utk.edu

Hatem Ltaief

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

George Bosilca

University of Tennessee Knoxville  
bosilca@icl.utk.edu

Ying Sun

KAUST  
Saudi Arabia  
ying.sun@kaust.edu.sa

Marc G. Genton

King Abdullah University of Science and Technology  
Statistics Program  
marc.genton@kaust.edu.sa

David E. Keyes

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

Jack J. Dongarra

University of Tennessee, Oak Ridge National Laboratory,  
USA

dongarra@icl.utk.edu

## MS27

### On Conjugate Bayesian Linear Regression Frameworks for High-Dimensional Spatial Modeling

High-dimensional spatially oriented data sets are prevalent in the physical sciences. We seek to jointly model several variables, each indexed by a spatial location, to capture associations over space, time and among the dependent variables. Multivariate latent spatial process models have proved effective in driving statistical inference and rendering better predictive inference at arbitrary locations for the spatial process. This talk will present extensions of scalable modeling strategies for multivariate processes including for several variables. We pursue Bayesian inference, which is attractive for full uncertainty quantification of the latent spatial process. We discuss computationally efficient parallelizable approaches using conjugate Bayesian matrix-normal distributions in spatial factor models as well as graphical Gaussian processes that deliver inference over a high-dimensional parameter space and for very large numbers of variables. We illustrate the computational and inferential benefits of our algorithms over competing methods using simulation studies and an analysis of some remote sensed data sets.

Sudipto Banerjee

Department of Biostatistics  
UCLA  
sudipto@ucla.edu

## MS27

### High Performance Computing in the Context of General Software for Hierarchical Statistical Modeling

I'll discuss the challenges of doing high performance computing for hierarchical statistical modeling, with a focus on MCMC. I'll start by giving an overview of how different packages for general hierarchical modeling (e.g., JAGS, Stan, PyMC3, NIMBLE) have been designed in terms of achieving both flexibility and speed. I'll specifically discuss the design decisions we've made in developing the NIMBLE package, of which I am a core developer. NIMBLE provides an interface for programming models and algorithms in R from which we programmatically generate C++ via the NIMBLE compiler. I'll close by talking about the challenges involved in making effective use of parallelization in the context of general purpose hierarchical modeling software.

Christopher Paciorek

Adjunct Professor, Department of Statistics,  
University of California Berkeley  
paciorek@stat.berkeley.edu

## MS27

### Scalable Statistical Learning from Dependent Data

Network, spatial, and temporal data are dependent, high-dimensional, and structured data, presenting challenges as well as opportunities for statistical learning. My talk will highlight some of the computational and statistical challenges and opportunities arising from dependent and structured data in high-dimensional settings.

Michael Schweinberger

Department of Statistics, University of Missouri,  
Columbia  
msq5c@missouri.edu

## MS28

### GPU Optimisation of Local Time Stepping on Unstructured Adaptive Meshes

The earthquake simulation software SeisSol adopts a high-order discontinuous Galerkin discretisation to solve the seismic wave equation (in various materials - elastic, acoustic, poroelastic, e.g.) on adaptive tetrahedral meshes. Cluster-oriented local time stepping, i.e., choosing the time step size adaptively depending on element size and material property is a crucial component to bring down the time to solution for realistic models. We will present approaches to optimize local time stepping on modern GPU-based supercomputers and discuss load distribution challenges that stem from non-homogeneous performance on heterogeneous architectures.

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

## MS28

### The p4est Software for Parallel AMR: Algorithms and Interfaces

A forest of octrees is one data structure to represent the recursive adaptive refinement of an initial, conforming coarse mesh of hexahedra. By construction, it generates a space filling curve ordering of the inner and leaf nodes of the refinement forest that can be exploited for fast partitioning and load balancing, adhering to a strictly disjoint parallel storage of leaves. In addition, the structure allows for communication-free and efficient local neighbor finding and general remote object searches through the entire partition. The p4est software is a well-known software library that implements a collection of algorithms for parallel AMR. In this presentation, we will focus on (a) performance improvements by exploiting the AVX instruction set and (b) latest additions to the set of available algorithms, in particular configurably ordered traversals and non-balanced halo gathering and mesh iteration.

Mikhail Kirilin, Gautam Ghosh, Carsten Burstedde  
Universität Bonn  
kirilin@ins.uni-bonn.de, ghosh@ins.uni-bonn.de,  
burstedde@ins.uni-bonn.de

## MS28

### A space-filling curve for pyramids

We present a new space-filling curve for pyramidal adaptive mesh refinement. Developments in the recent years have extended the scalable and efficient tree-based AMR methods from quadrilaterals/hexahedra to various element shapes such as triangles, tetrahedra or prisms. To support the full range of element shapes used in common unstructured meshes, additional support for pyramids is required. To this end, we develop a new Morton type space-filling curve for this kind of element and integrate it into our AMR framework t8code. This development enables us to apply tree-based AMR to fully hybrid input meshes. Two special challenges for the SFC development and AMR library arise from the fact that one pyramid refines into 6 pyramids and 4 tetrahedra. The first one is not using a usual  $2^n$  refine-

ment pattern, making adaptations of the AMR software to allow for flexible new numbers of elements when refining necessary. The second one is that elements of one shape (pyramids) refine into elements of a different shape (tetrahedra). We present general solutions for both challenges that are also applicable to non-pyramid elements and conclude with benchmarks on the JUWELS supercomputer using up to 50 billion elements on 50k MPI ranks.

David Knapp  
Universität Bonn  
david.knapp@dlr.de

Johannes Holke  
German Aerospace Center (DLR)  
johannes.holke@dlr.de

Carsten Burstedde  
Universität Bonn  
burstedde@ins.uni-bonn.de

## MS28

### Developing GPU Based Conforming Curved Mesh Adaptation Procedures in Mfem

Ongoing efforts on the development of conforming mesh adaptation procedures in MFEM will be summarized. MFEM is a high-order finite element library that focuses on performant and scalable execution on the latest GPU accelerated super computing systems. Thus, mesh adaptation must support curved elements of the order required for the order of finite element used. In addition, it is desirable that those procedures execute on the compute nodes GPUs. Building on Omega\_h's GPU based linear simplex mesh adaptation procedures, efforts are ongoing to extend, and add to, its cavity based mesh modification operators to support curved element geometries. To support the interactions needed to properly curve mesh faces and edges to domain boundaries, the EGADS geometry interrogation APIs, a version of which executes on GPUs, will be used. The application of these methods on the latest DOE systems, which include NVIDIA, AMD and Intel GPUs, will be demonstrated for fusion plasma physics simulations.

Morteza Siboni  
Rensselaer Polytechnic Institute  
m.hakimi.siboni@gmail.com

Cameron W. Smith  
Scientific Computation Research Center  
Rensselaer Polytechnic Institute  
smithc11@rpi.edu

Aditya Y. Joshi, Mark S. Shephard  
Scientific Computing Research Center  
Rensselaer Polytechnic Institute  
joshia5@rpi.edu, shephard@rpi.edu

## MS29

### Challenges of Mixed-Precision Fast Fourier Transforms from the Instruction Level to at Scale Computations

The talk attempts to discuss challenges of creating a Fast Fourier transform implementation that fully exploits opportunities offered by state-of-the art GPUs such as reduced and mixed-precision computing using tensor cores. Making proper use of high throughput, reduced precision

compute requires careful reconsideration of tradeoffs taken by FFT implementation. Result accuracy and performance is inherently sensitive to the algorithmic decomposition of the problem and how its parts are mapped into hardware. To achieve performance improvements, new methods require striking a new balance between memory bandwidth, compute bandwidth and latency. All of these new options add to the design space of possible implementations and the enlarged design space allows to push the boundary of the existing limitations. The talk will discuss design choices in the context of real-world examples.

Lukasz Ligowski  
NVIDIA  
lligowski@nvidia.com

## MS29

### Numerical Behavior of GPU Matrix Multiply-Accumulate Hardware

Tensor cores are hardware units on the latest GPUs that perform matrix multiply accumulate (MMA) operation. 116 of the TOP500 supercomputers contain these units and a lot of the numerical libraries begin to utilize them in various algorithms. Tensor cores and similar arithmetic units targeted at machine learning applications are not necessarily IEEE 754-compliant and features such as rounding, normalization, order of operations, subnormal number support and others can differ from a standard software implementation of the matrix multiplication. In this talk I will discuss our recent work on determining various numerical features of MMAs, using tensor cores as an example. We determined various features by testing tensor cores with carefully constructed numerical test cases on the V100, T4 and the A100 NVIDIA GPUs.

Massimiliano Fasi  
Orebro University, Sweden  
massimiliano.fasi@oru.se

Nicholas J. Higham  
School of Mathematics  
The University of Manchester  
nick.higham@manchester.ac.uk

Mantas Mikaitis  
The University of Manchester  
mantas.mikaitis@manchester.ac.uk

Srikara Pranesh  
University of Manchester  
srikara.pranesh@manchester.ac.uk

Florent Lopez  
Ansys  
florent.joan.lopez@gmail.com

Theo Mary  
Sorbonne Universite, CNRS, LIP6  
theo.mary@lip6.fr

## MS29

### Double-Word Arithmetic and Accurate Calculation of Euclidean Norms

We consider the computation of the Euclidean (a.k.a. L2) norm of an  $n$ -dimensional vector in floating-point arithmetic. We review the classical solutions used to avoid

spurious overflow or underflow. We modify a recently-published algorithm to allow for a very accurate solution, free of spurious overflows and underflows. The returned result will be within very slightly more than 0.5 ulps from the exact result, which means that we will almost always provide correct rounding.

Vincent Lefèvre  
INRIA  
vincent.lefevre@inria.fr

Nicolas Louvet  
Université Claude Bernard Lyon 1  
nicolas.louvet@ens-lyon.fr

Jean-Michel Muller  
CNRS  
jean-michel.muller@ens-lyon.fr

Joris Picot  
Ecole Normale Supérieure de Lyon  
joris.picot@ens-lyon.fr

Laurence Rideau  
INRIA Sophia Antipolis  
laurence.rideau@sophia.inria.fr

## MS30

### New Developments of the FROSch Domain Decomposition Solver Package

Schwarz methods are an algorithmic framework for a large class of domain decomposition methods. The software FROSch (Fast and Robust Overlapping Schwarz), which is part of the Trilinos package ShyLU, provides a highly scalable implementation of the Schwarz framework, and the resulting solvers are based on the construction and combination of the relevant Schwarz operators. FROSch currently focusses on Schwarz operators that are algebraic in the sense that they can be constructed from a fully assembled, parallel distributed matrix. This is facilitated by the use of extension-based coarse spaces, such as generalized Dryja-Smith-Wildund (GDSW) type coarse spaces. This talk gives an overview of the FROSch software framework as well as current developments in improving the performance of the solvers, for instance, due to the use of inexact subdomain and coarse solvers. Moreover, recent parallel results for challenging applications, such as coupled multiphysics simulations of land ice in Greenland and Antarctica, are presented.

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

## MS30

### Extending the Parallel Scalability of FROSch Solvers Beyond 100K Cores

Techniques to extend the parallel performance of the three-level Fast and Robust Overlapping Schwarz (FROSch) preconditioners are described. The FROSch framework is part of the Trilinos software library and contains a parallel implementation of different preconditioners with energy minimizing coarse spaces of GDSW (Generalized Dryja-Smith-Wildund) type. A three-level extension for FROSch is constructed by a recursive application of the FROSch preconditioner to the coarse problem. No explicit geomet-

ric information is needed in this recursive application. In particular, for elasticity, the rigid body modes, including the rotations, can be interpolated on the coarse level without additional geometric information. Parallel results for a three-dimensional linear elasticity problem obtained on the Theta supercomputer (ALCF, Argonne, USA) using up to 220000 cores are discussed and compared to results obtained on the SuperMUC-NG supercomputer (LRZ, Garching, Germany).

Oliver Rheinbach

Technische Universitaet Bergakademie Freiberg  
oliver.rheinbach@math.tu-freiberg.de

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

Axel Klawonn  
Universität zu Köln  
axel.klawonn@uni-koeln.de

Friederike Röver  
Technische Universitaet Bergakademie Freiberg  
friederike.roever@math.tu-freiberg.de

**MS30**

**Accelerating FROSch Preconditioners using GPUs**

The Trilinos package FROSch (Fast and Robust Overlapping Schwarz) implements multilevel overlapping Schwarz methods within the Trilinos framework and has been used in various applications. FROSch focusses on the implementation of extension-based coarse spaces, such as GDSW (Generalized Dryja-Smith-Widlund) type coarse spaces, which can be build based on the fully assembled system matrix. Through Tpetra, FROSch can make use of the Kokkos programming model and linear algebra kernels from KokkosKernels, which allow for performance portability on different computer architectures.  $\llcorner$  In this talk, we discuss our experiences and challenges porting FROSch to run on a GPU cluster. We present several techniques used to enhance its performance on GPUs along with performance results on Summit Supercomputer; we consider different types of model problems, ranging from simple Poisson an linear elasticity problems to landice simulations of Greenland and Antarctica.

Ichitaro Yamazaki

Sandia National Laboratories  
iyamaza@sandia.gov

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

Sivasankaran Rajamanickam  
Sandia National Laboratories  
srajama@sandia.gov

**MS30**

**Device Accelerated Solvers with PETSc. Current Status and Future Perspectives**

In this talk, we will review recent developments towards the exploitation of GPUs from the Portable and Extensible Toolkit for Scientific computing (PETSc). Current status and future perspectives will be discussed, with a fo-

cus on Domain Decomposition solvers, Algebraic Multigrid methods, and hierarchical matrices.

Stefano Zampini

King Abdullah University of Science and Technology  
stefano.zampini@kaust.edu.sa

**MS31**

**Leveraging Single Node Linear Algebra GPU Solvers from a Multi-Nodes User Code with Alien**

Alien is a C++ framework that provides high-level APIs to solve linear problems for distributed-memory applications. Alien does not implement a linear solver but provides unified access to several libraries such as Hypre, PETSc, or Trilinos. However, Alien also extends the capabilities of smaller footprint libraries such as Ginkgo or ViennaCL. Alien manages data redistribution to use them from a distributed memory context, allowing a legacy MPI application to leverage modern GPU solvers. We will describe how the different solvers using different low-level technologies are plugged in Alien. We will illustrate how Alien's unified API helps us efficiently exploit different computer hardware by selecting the most adapted solver library.

Cédric Chevalier

CEA/DAM  
cedric.chevalier@cea.fr

Sylvain Desrozières  
IFP Energies Nouvelles  
sylvain.desrozières@ifpen.fr

Pascal Havé  
haveneer  
contact@haveneer.com

Jean Marc Gratien  
IFP  
jean-marc.gratien@ifp.fr

Thomas Guignon  
IFP Energies Nouvelles  
thomas.guignon@ifpen.fr

Nathalie Möller  
CEA/DAM  
nathalie.moller@cea.fr

Xavier Tunc  
IFP Energies Nouvelles  
France  
xavier.tunc@ifpenergiesnouvelles.fr

**MS31**

**Porting Hypre to Heterogeneous Computer Architectures: Strategies and Experiences**

The high performance preconditioner library *hypre* provides a variety of interfaces and linear solvers, including various multigrid methods, that have achieved good scalability on a variety of homogeneous parallel computer architectures. Heterogeneous architectures with nodes that have both CPUs and accelerators provide new challenges, since they require more fine-grained parallelism and reduced data movement between different memories on a single node as well as across nodes. In this talk, we will discuss our experiences and strategies to port *hypre* to heteroge-

neous computers with accelerators, including the design of a new memory model, the use of abstractions, the BoxLoop macros in the structured and semi-structured interfaces, and the restructuring of algebraic multigrid (AMG) into modular components. We present numerical experiments comparing CPU and GPU performance for several test problems.

Robert D. Falgout  
Lawrence Livermore National Laboratory  
falgout2@llnl.gov

Ruipeng Li, Bjorn Sjogreen  
Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory  
li50@llnl.gov, sjogreen2@llnl.gov

Ulrike Meier Yang, Lu Wang  
Lawrence Livermore National Laboratory  
yang11@llnl.gov, wang84@llnl.gov

### MS31

#### **We Needed a Fast Performance-Portable Geometric Search. Here's What We Learned Along the Way**

Fast geometric search is an often overlooked but a critical part in many scientific applications. For example, computational cosmologists learn more about the Universe by examining distribution of clumped areas of dark matter, called halos. We wanted to help, so we developed a geometric search library called ArborX. To avoid doing the hard work of learning to be experts in all modern architectures (from multithreaded CPUs to novel GPUs), we used Kokkos to hide all that complexity. Instead, we focused on developing a robust implementation based on bounding volume hierarchy (BVH). This talk explores many areas we delved in during our quest to performance. We explore different approaches to BVH traversal, low level mechanisms to Kokkos kernel launching, and adapting the library design. Not all approaches work in all cases, so we still have more to learn.

Andrey Prokopenko  
Oak Ridge National Laboratory  
prokopenkoav@ornl.gov

### MS31

#### **Utopia: A C++ Library for Parallel Multilevel Solution Methods and for Constrained Non-Linear Problems**

Utopia is a C++ embedded domain-specific language designed for parallel non-linear solution strategies. Utopia combines the advantages of high-level programming interfaces with the advantages of scripting languages. On the one hand, it allows using high-level abstractions while providing access to the native low-level data structures. On the other hand, it facilitates expressing complex numerical methods by means of few lines of code. This is achieved by separating the model from the computation, thus allowing us to keep the implementation details hidden from the code of applications such as non-linear solution algorithms and finite element assembly. This separation is realized by using C++ meta-programming and particular evaluation strategies which allow mapping an abstract representation of the computation to the actual code computing the result. In this talk, we showcase the library functionalities and their application to numerical geophysics simulations.

In particular, we focus on fractured rock formations in the subsurface which are of crucial importance in a variety of reservoir applications such as geothermal energy extraction, CO<sub>2</sub> sequestration, nuclear waste storage, and unconventional oil and gas recovery.

Patrick Zulian  
ICS, Università della Svizzera italiana, Switzerland  
patrick.zulian@usi.ch

Alena Kopanicakova  
Università della Svizzera italiana  
kopana@usi.ch

Maria Nestola  
Institute of Computational Science, USI, Lugano, Switzerland  
nestom@usi.ch

Nur A. Fadel, Andreas Fink, Daniel Ganellari, Joost VandeVondele  
CSCS - Swiss National Supercomputing Centre  
nur.fadel@cscs.ch, andreas.fink@cscs.ch,  
daniel.ganellari@cscs.ch, joost.vandevondele@cscs.ch

Rolf Krause  
Università della Svizzera italiana, Switzerland  
rolf.krause@usi.ch

### MS32

#### **Convergent Orthogonality Preserving Approximations of the Kohn-Sham Orbitals**

To obtain convergent numerical approximations without orthogonalization operations is of great importance in electronic structure calculations. In this paper, we propose and analyze a class of iteration schemes for the discretized Kohn-Sham model, which preserves the orthogonality of the Kohn-Sham orbitals automatically. With our schemes, the iterative approximations are guaranteed to converge to the Kohn-Sham orbitals without any orthogonalization operations when the initial orbitals are orthogonal. We prove the convergence and get the local convergence rate of the numerical approximations. In addition, we present two approaches to get suitable time step sizes. This is a joint work with Liwei Zhang and Aihui Zhou.

Xiaoying Dai  
Institute of Computational Mathematics  
Chinese Academy of Sciences, Beijing, China  
daixy@lsec.cc.ac.cn

### MS32

#### **A Parallel Two-Stage, Divide and Conquer Algorithm in Slate for the Dense Symmetric Eigenvalue Problem**

The SLATE library provides a distributed parallel and GPU accelerated algorithm to solve the dense symmetric eigenvalue problem. We use a two stage reduction from dense to band, then band to tridiagonal, which casts most operations as Level 3 BLAS. Different algorithms can subsequently be used for the tridiagonal eigensolver. We support both the classical QR iteration and a divide and conquer approach. The two stage reduction to tridiagonal also requires two back transformations of the eigenvectors. SLATE provides portability across GPU architectures, cur-

rently supporting both CUDA and AMD ROCm platforms.

Mark Gates

Computer Science Department  
University of Tennessee  
mgates3@icl.utk.edu

### MS32

#### Revising Pseudospectral Methods for Hybrid Kohn-Sham Density Functional Theory in the Age of Exascale Computing

In large scale hybrid Kohn-Sham density functional theory (KS-DFT) simulations, the primary computational bottleneck in the solution of the self-consistent field (SCF) eigenvalue problem is the assembly of the exact-exchange contribution to the Fock matrix. In atom-centered bases, traditional methods to perform this task involve the evaluation of an enormous number of electron repulsion integrals which introduce a number of challenges when considering modern massively parallel accelerator based architectures. These include difficulties in load balancing in distributed memory contexts and lack of SIMD concurrency for GPU architectures. In this work, we revisit a class of pseudospectral methods which are capable of overcoming the aforementioned hurdles by exploiting spatial locality of the Coulomb interaction and exhibit linear scaling with system size.

David B. Williams-Young

Lawrence Berkeley National Laboratory  
dbwy@lbl.gov

### MS33

#### Non-Blocking Resilient Extensions in MPI

The resilient extensions to the MPI Standard under investigation as part of the User Level Failure Mitigation (ULFM) propose a set of blocking capabilities to detect, notify and recover from permanent process failures. The key principle is that no MPI call (point-to-point, collective, RMA, IO, ...) can block indefinitely after a failure, but must either succeed or raise an MPI error informing the developer of a drastic change in the number of ranks in the parallel application. However, all these extensions are blocking in their nature, resulting in a recovery process that behaves sequentially. A set of new extensions, entirely non-blocking, will allow applications to construct faster-changing worlds, and fulfill the resilient needs of more dynamic applications. We will present these extensions, highlight how they impact applications and provide some early experimental data.

George Bosilca  
University of Tennessee Knoxville  
bosilca@icl.utk.edu

Aurélien Bouteiller

University of Tennessee, Knoxville, US  
bouteill@icl.utk.edu

### MS33

#### Status Report on MANA-2.0: A Future-Proof Design for Long-Running Mpi-Based Simulations for Hpc

Support for long-running computations on supercomputers has long been a pain point. To maintain scheduling flexibility, sysadmins set a maximum resource allocation (e.g., 40

hours) for HPC jobs. Transparent checkpointing at scale on supercomputers has been a long-time dream to please both users and sysadmins. MANA-2.0 (MPI-Agnostic, Network-Agnostic checkpointing) is such an effort. Like the original MANA academic prototype, MANA-2.0 operates over any MPI implementation and network interconnect that supports the MPI API standard. MANA-2.0 is also future-proof, in the sense that it runs independently of the underlying MPI and network libraries. Details of new algorithms required for its robustness will be presented. MANA-2.0 is being tested both on NERSC's Cori supercomputer (proprietary Cray MPI and Cray GNI network) and NERSC's Perlmutter (#5 supercomputer; proprietary Cray MPI and HPE Cray Slingshot network). Like all large projects, this is a years-long collaboration of many participants, to be credited in the talk.

Gene Cooperman

Khoury College of Computer Sciences  
Northeastern University, Boston, MA  
gene@ccs.neu.edu

### MS33

#### Extending Reinit Towards a Formal Specification and Interplay with GPUs

Reinit is a recently proposed extension to MPI to support global-restart recovery through a simple but powerful API that minimizes user involvement and changes in the source code of applications. In this talk, we will present extensions to the Reinit fault-tolerance API towards a formal specification that composes with the rest of the MPI specification and facilitates MPI+X programming models. Those extensions include supporting synchronous or asynchronous recovery modes, extending the MPI error-handler interface to select the recovery mode, adding a function to the API to explicitly trigger recovery in synchronous mode, and articulating assumptions in the fault and recovery model supported by Reinit. Further, we will present insights from applying Reinit recovery on MPI applications that use GPUs.

Georgis Georgakoudis, Ignacio Laguna

Lawrence Livermore National Laboratory  
georgakoudis1@llnl.gov, lagunaperalt1@llnl.gov

### MS33

#### Towards Access Pattern Aware Checkpointing For Kokkos Applications

Portability of resilience despite increasing heterogeneity is vital for the future of HPC systems. While there is an active push towards performance portability, relatively little attention has been dedicated to resilience portability. This talk focuses on the challenges and opportunities of adopting checkpoint-restart techniques (commonly used to provide resilience for HPC systems) in the context of performance portable abstractions and runtimes. Starting from the problem of sparse updates and implications it has on incremental checkpointing, it generalizes the discussion towards the broader question how application data access patterns can be identified and leveraged through memory views for the purpose of automating and accelerating checkpoint-restart at scale.

Nigel Tan

University of Tennessee Knoxville  
ntan1@vols.utk.edu

Bogdan Nicolae  
MCS Division  
Argonne National Laboratory  
bnicolae@anl.gov

Nicolas M. Morales  
Sandia National Labs  
nmmoral@sandia.gov

Keita Teranishi  
Sandia National Laboratories  
Livermore, California  
knieran@sandia.gov

Sanjukta Bhowmick  
University of North Texas  
sanjukta.bhowmick@unt.edu

Michela Taufer  
The University of Tennessee  
taufer@utk.edu

### MS34

#### **A Case Study on the Impact of Chapel within an Academic Computational Aerodynamic Laboratory**

Aerodynamic simulations involve solving Computational Fluid Dynamics (CFD) models on supercomputers. The resulting non-linear partial differential equations, comprising up to 10 billion unknowns, are solved using proprietary C/C++ or Fortran source codes. Scalability via domain decomposition is mandatory for industrial run times, typically using MPI libraries. When multi-physical phenomena are accounted, such as aero-icing, code development becomes more challenging since different packages constructed with different paradigms must interact. Orthogonal to these challenges are organizational ones. As R&D involves cascading university-based low Technology Readiness Level (TRL) advances into industrial high TRL workflows, the technical gap for graduate students with knowledge of precompiled-based tools handling few unknowns on single CPU machines is tremendous. At Polytechnique Montreal, we used the Chapel programming language so that students can spend time addressing scientific problems minimizing technical challenges. In this talk, a case study highlights how Chapel enabled a small group of graduate students write a world-class CFD software participating in international workshops alongside similar software, albeit with MPI libraries, developed by R&D professionals. The obtained results showcase the efficiency of the Chapel language not only in terms of HPC metrics, but also in terms of capability enhancement for university-based contributions to the aeronautical field.

Éric Laurendeau  
Polytechnique Montréal Department of Mechanical Engineering  
eric.laurendeau@polymtl.ca

### MS34

#### **Arkouda: Interactive Supercomputing for Data Analytics Made Possible by Chapel**

Exploratory data analysis (EDA) is a prerequisite for all data science, as illustrated by the ubiquity of Jupyter notebooks, the preferred interface for EDA among data scientists. The operations involved in exploring and trans-

forming the data are often at least as computationally intensive as downstream applications (e.g. machine learning algorithms), and as datasets grow, so does the need for HPC-enabled EDA. However, the inherently interactive and open-ended nature of EDA does not mesh well with current HPC usage models. Meanwhile, several existing projects from outside the traditional HPC space attempt to combine interactivity and distributed computation using programming paradigms and tools from cloud computing, but none of these projects have come close to meeting our needs for high-performance EDA. To fill this gap, we have developed a software package, called Arkouda, which allows a user to interactively issue massively parallel computations on distributed data using functions and syntax that mimic NumPy, the underlying computational library used in the vast majority of Python data science workflows. The computational heart of Arkouda is a Chapel interpreter that accepts a pre-defined set of commands from a client (currently implemented in Python) and uses Chapel's built-in machinery for multi-locale and multi-threaded execution. Arkouda has benefited greatly from Chapel's distinctive features and has also helped guide the development of the language.

Michael Merrill  
US Government  
mhmerrill@mac.com

### MS34

#### **Experiences with Chapel in Cosmology from Data Analysis to Simulations**

Addressing problems in cosmology require analyzing large amounts of data, either from simulations or observations. At the same time, extracting the maximal amount of information from the data often requires researchers to explore new algorithms, which are most easily developed and explored on small datasets on single machines. It is therefore very useful to have the ability to (relatively) seamlessly scale up algorithms from simple serial implementations to fully distributed parallel implementations. Chapel provides a promising answer to this problem. This talk will describe personal experiences using Chapel for research in cosmology, describing how Chapel has allowed us to easily scale up algorithms. While I will present performance numbers, I will also emphasize the productivity aspects of the language, an essential consideration for a researcher. For concreteness, I will focus on a single problem: the astrophysical simulation of ultralight dark matter, which requires solving coupled Schrodinger and Poisson equations. I will start by describing how Chapel's distributed-parallel features allowed us to simply extend local 1D FFTs (as implemented in FFTW) to a 3D distributed setting and describe its performance. I will then turn to how we implemented our simulation and associated analysis codes in Chapel, highlighting the relative closeness between the mathematical description of the algorithm and the fully distributed-parallel code. Based on our experiences, I will attempt to draw some general lessons on Chapel's suitability for scientific analyses.

Nikhil Padmanabhan  
Yale Department of Physics  
nikhil.padmanabhan@yale.edu

### MS34

#### **Chapel Overview and Future Opportunities and**

## Challenges for Chapel

A variety of applications have taken advantage of Chapel's ease of programming and high performance, some of which will be featured in subsequent talks in this session/minisymposium. This talk will provide an overview of some key concepts in Chapel that users have reported finding attractive such as its global view of computation and constructs that enable the specification of distributed- and shared-memory parallelism. These concepts will be illustrated with some example programs running on laptops to supercomputers. This talk will also feature future opportunities and challenges for Chapel such as a compiler rework to provide faster and interactive compilation, expanding Chapel's support for GPUs, and finalizing the language and standard library stabilization effort.

Michelle Strout  
Hewlett Packard Enterprise  
University of Arizona  
michelle.strout@hpe.com

Bradford Chamberlain  
Hewlett Packard Enterprise  
bradford.chamberlain@hpe.com

## MS35

### Optimizations of Lattice H-Matrix-Vector Multiplication for Modern Supercomputers

The Hierarchical matrices (H-matrices) are used to approximate the dense matrices that appear in the Boundary Element Method (BEM). Accelerating the Hierarchical Matrix Vector multiplication (HiMV) is important for fast solution of the simultaneous linear equations of BEM, but it has not been sufficiently studied compared to the dense matrix vector multiplication (GEMV) and the sparse matrix vector multiplication (SpMV). In this study, in order to speed up HiMV, we have optimized it for modern multi-core and many-core processors while considering the complexity of H-matrix. We also improved the scalability by developing a communication hiding method for the lattice H-matrix method, which is an extension of the H-matrix for distributed memory environments. Using H-matrices of various shapes appearing in the electrostatic field BEM analysis, HiMV was evaluated on a cluster with A64FX, Intel Xeon Phi Knights Landing and Intel Xeon Cascade Lake processors.

Tetsuya Hoshino  
The University of Tokyo  
hoshino@cc.u-tokyo.ac.jp

Akihiro Ida  
Japan Agency for Marine-Earth Science and Technology,  
Japan  
ida.a@jamstec.go.jp

Toshihiro Hanawa  
Information Technology Center  
The University of Tokyo  
hanawa@cc.u-tokyo.ac.jp

## MS35

### Adaptation of Explainable AI for Auto-Tuning on Accurate Matrix Multiplication Library

Explainable AI (XAI) is one of very important topics for not only AI field but also numerical computations. The

authors have been developing for Auto-tuning (AT) facility for numerical libraries in this two decades, especially to autotune parameters on numerical algorithms, selections of high-performance implementations, etc. In this presentation, the authors focus on XAI for AT to numerical libraries. Target numerical library is chosen for high accurate matrix multiplication library, which is developed by author's group. For more details, AI is adapted to select the best implementation between 11 implementations on the library, including CPU and GPU. LIME and SHAP are utilized for tools of XAI. To prepare training data, a GPU based-supercomputer, named the Supercomputer Flow Type II Subsystem in Information Technology Center, Nagoya University, is used. According to preliminary result, the tools explain well for behavior of learned models by random-forest regressor.

Takahiro Katagiri  
Information Technology Center  
Nagoya University  
katagiri@cc.nagoya-u.ac.jp

Shota Aoki  
Graduate School of Informatics  
Nagoya University  
aoki@hpc.itc.nagoya-u.ac.jp

Satoshi Ohshima, Toru Nagai  
Information Technology Center  
Nagoya University  
ohshima@cc.nagoya-u.ac.jp, nagai@cc.nagoya-u.ac.jp

## MS35

### Massively Parallel Eigensolvers based on Unconstrained Minimization

This presentation will show recent developments in unconstrained minimization schemes for the solution of eigenvalue problems in electronic structure calculations. These schemes employ a preconditioned conjugate gradient approach that avoids an explicit reorthogonalization of the trial eigenvectors, in contrast to typical iterative eigensolvers, therefore reducing communications and becoming an attractive approach for the solution of very large problems on massively parallel computers. We show results for a set of benchmark systems with an implementation of the unconstrained minimization in first-principles materials and chemistry codes that perform electronic structure calculations based on a density functional theory (DFT) approximation to the solution of the many-body Schrödinger equation. The results show that the unconstrained formulation, together with an appropriate preconditioner, offers good convergence properties and scales well on a large number of cores. We also show results for an implementation in the commonly used plane wave formulation of DFT, on mixed GPU/CPU computers such as the IBM Summit computer at the leadership computing facility at Oak Ridge National Laboratory.

Doru Thom A. Popovici  
Lawrence Berkeley National Laboratory  
dtpopovici@lbl.gov

Osni A. Marques  
Lawrence Berkeley National Laboratory  
Berkeley, CA  
oamarques@lbl.gov

Mauro Del ben, Andrew M. Canning

Lawrence Berkeley National Laboratory  
mdelben@lbl.gov, acanning@lbl.gov

### MS35

#### Performance Evaluation of Iterative Refinement for Symmetric Eigenvalue Problem on Supercomputer

This study considers a refinement algorithm for the eigenvalue decomposition of a real symmetric matrix. Various numerical algorithms have been developed for standard eigenvalue problems. The accuracy of approximate eigenvectors obtained using a numerical algorithm decreases with the gap between the associated eigenvalue and its closest one. Therefore, using iterative refinement methods is useful for obtaining the eigenvalue decomposition with high accuracy. Recently, Ogita and Aishima proposed an efficient refinement algorithm for improving approximate eigenpairs obtained using a numerical algorithm. Since that algorithm is based on Newton's method, that converges quadratically if an initial guess is moderately accurate. The refinement algorithm is constructed via four high-accuracy matrix multiplications. Thus, that may yield high-accuracy results faster than those obtained using multiple-precision arithmetic. In this study, we implemented that refinement algorithm on a supercomputer using high-accuracy matrix multiplications proposed by Ozaki et al. and Double-Double arithmetic. We improved the accuracy of approximate eigenpairs provided by PDSYEVD from ScaLAPACK. Numerical results are presented showing the efficiency of the refinement algorithm implemented on a supercomputer.

Yuki Uchino

Shibaura Institute of Technology  
mf20013@shibaura-it.ac.jp

Takeshi Terao  
R-CCS  
RIKEN  
takeshi.terao@riken.jp

Katsuhisa Ozaki  
Shibaura Institute of Technology / JST  
ozaki@sic.shibaura-it.ac.jp

Toshiyuki Imamura  
RIKEN Center for Computational Science  
imamura.toshiyuki@riken.jp

### MS36

#### Operator-Splitting Decompositions for Efficient Solution to the Emi Equations

Temporal operator-splitting is a general approach for decomposing time-dependent multi-physics problems into smaller constitutive chunks that can each be solved independently and then recombined to form an approximate solution. The general approach is typically most useful when the individual chunks function on different time-scales, each requiring specialized methods for solution. The extracellular-membrane-intracellular (EMI) model of cardiac electrophysiology—which aims to explicitly model electrical potentials of extracellular space, intracellular space, and across the cell membrane of cardiomyocytes—is a candidate multi-physics example where operator-splitting can show its benefits. Using a recently developed operator-splitting class within the deal.ii finite element library, we discuss the implementation of such multi-physics mod-

els in this framework and present a survey of computational results pertaining to different order operator-splitting schemes in electrophysiological modelling.

Kevin R. Green

Global Institute for Water Security  
University of Saskatchewan  
kevin.green@usask.ca

### MS36

#### An Operator-Splitting Approach to Solving Cell-Based Mathematical Models of Cardiac Tissue using Modern CPU Architectures

A number of pathologies related to the electrical activity in the heart can be studied using computer simulations of reaction-diffusion models. With the recent extracellular-membrane-intracellular (EMI) model, the geometry of each cell is resolved in the mesh, allowing for a more accurate representation of cardiac tissue on the cell scale. However, the EMI model requires a very fine mesh, and the linear systems arising from the diffusion process in the extracellular and the intracellular domains are ill-conditioned. In this talk, we present an improved operator-splitting method that decouples the intracellular and extracellular domains, such that each sub-problem becomes a classical elliptic partial differential equation. Using this operator-splitting method, the computing time scales linearly with the problem size. This operator-splitting method enables us to solve the linear systems efficiently on shared-memory parallel computers, and we demonstrate that we are able to solve a system with  $512 \times 256$  cardiac cells, solving linear systems with approximately  $2.5 \cdot 10^8$  degrees of freedom.

Kristian Gregorius Hustad

Simula research laboratory  
kghustad@simula.no

### MS36

#### Microscale Cardiac Electrophysiology on Exascale Supercomputers

The bidomain model of cardiac electrophysiology has made it possible to simulate entire human hearts, by homogenizing the heart's two billion myocytes and their interconnections and replacing them by only a few million model elements. This simplification has allowed enormous progress in the area of cardiac simulations with a large impact on cardiology itself. However, bidomain models cannot represent the electrical interactions between individual cells or small groups of cells. These interactions are thought to play a crucial role in arrhythmia in ageing hearts and in a swath of ischemic, infectious, hypertrophic, and inheritable cardiomyopathies. In order to understand these interactions and to learn to recognize their electric signatures, we must model the heart cell by cell. This problem is challenging due to its size and complexity, and will require exascale supercomputers to be solved on a whole heart. We are working on this problem with an international, interdisciplinary consortium involving computer scientists, applied mathematicians, biomedical engineers, and medical researchers. In this talk I will outline our motivation for this work and initial results of the project.

Mark Potse

Liryc institute of cardiac electrophysiology and modeling  
University of Bordeaux

mark@potse.nl

### MS36

#### Cardiac Electrophysiology Simulation: Past, Present, and Future

Cardiac electrophysiology simulation provides a non-invasive means to better diagnose and treat cardiac pathologies. At present, the most reliable model for cardiac electrophysiology simulation is arguably the bidomain model, which is based on a homogenization of myocardial tissue such that the intra- and extra-cellular domains are co-located. Historically, this homogenization was done mostly to make computations feasible despite some desire to model each of the heart's two billion muscle cells individually. The homogenization, however, obscures the cell membrane dynamics, thus removing the possibility of studying membrane channelopathies that contribute to cardiac pathologies such as long QT syndrome. As computing power has increased, so have the expectations from models and simulations. The extracellular-membrane-intracellular (EMI) model resolves the detailed characteristics of ion channels on individual cell membranes, thus making it highly suitable for studying channelopathies. The EMI model resolves the tissue-scale dynamics at about four orders of magnitude higher than the bidomain model. This increased resolution opens the door to studying a wider array of heart conditions but at a significant increase in computational cost. This presentation gives an overview of the current state of cardiac electrophysiology simulation, starting from its relatively humble beginnings and ending with a glimpse into what the future may hold for it.

Raymond J. Spiteri  
University of Saskatchewan, Canada  
Department of Computer Science  
spiteri@cs.usask.ca

### MS37

#### Accelerating Configuration Interaction Calculations for Atomic Nuclei

Many-Fermion Dynamics nuclear, or MFDn, is a Configuration Interaction (CI) code for ab initio calculations of the structure of atomic nuclei. It is a platform-independent Fortran 90 code using a hybrid MPI+X programming model, which performs efficiently on CPU-based HPC platforms using MPI between nodes and OpenMP for shared-memory parallelism within a node. A typical calculation involves constructing a large irregular sparse symmetric matrix, of which we obtain the lowest eigenvalues and corresponding eigenvectors using the Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) algorithm. Rapid convergence of this block iterative method is achieved by a suitable choice of starting guesses of the eigenvectors and an effective distributed preconditioner. We have recently updated MFDn in order to take advantage of accelerators in current and future HPC platforms using OpenACC, targeting Perlmutter at NERSC; in addition we are exploring OpenMP target offload as well. In this talk I describe some of the challenges of creating an efficient GPU implementation using directives, both for the matrix construction and for the iterative eigensolver. Furthermore, with increasing local compute-power, data movement and inter-node communication can easily become a bottleneck, in particular for distributed iterative solvers; we are therefore also implementing a matrix-free version of

the solver.

Pieter Maris  
Iowa State University  
pmaris@iastate.edu

### MS37

#### Parallel Simulation of Non-Equilibrium Dynamics of a Quantum Many-Body System

The dynamical properties of a non-equilibrium quantum many-body system can be obtained from the solution of the Kadanoff-Baym equations (KBEs), which are a set of nonlinear integral differential equations satisfied by a two-time Greens function. The numerical procedure for solving KBEs has a high computational complexity in both the floating point operations and memory footprint. We will describe a parallelization strategy for solving the KBEs efficiently within a small time window and using the solution snapshots to predict long time dynamics of the Greens function by performing a parallel dynamic mode decomposition.

Jia Yin  
Lawrence Berkeley National Laboratory  
jiayin@lbl.gov

### MS38

#### Designing Smart and Resilient Extreme-Scale Systems

Resilience is one of the critical challenges of extreme-scale high-performance computing (HPC) systems, as component counts increase, individual component reliability decreases, and software complexity increases. Building a reliable supercomputer that achieves the expected performance within a given cost budget and providing efficiency and correctness during operation in the presence of faults, errors, and failures requires a full understanding of the resilience problem. This talk provides an overview of recent achievements in developing a taxonomy, catalog and models that capture the observed and inferred fault, error, and failure conditions in current supercomputers and in extrapolating this knowledge to future-generation systems. It also describes the path forward in machine-in-the-loop operational intelligence for smart computing systems, leveraging operational data analytics in a loop control that maximizes productivity and minimizes costs through adaptive autonomous operation for resilience.

Christian Engelmann  
Oak Ridge National Laboratory  
engelmann@ornl.gov

### MS38

#### Characterizing Failures in HPC using Benfords Law

To better understand the potential impacts of failures on next-generation systems, significant effort has been devoted to collecting, characterizing and analyzing failures on current systems. These studies typically require large volumes of data and complex analysis in an attempt to identify statistical properties within the data to exploit. In this talk, we examine failures on a number of systems and demonstrate that for many systems, the time between failures obeys Benfords law. This law applies to a number of naturally occurring collections of numbers and states

that the leading digit is more likely to be small, for example a leading digit of 1 is more likely than 9. Additionally, we outline how we might utilize this property in the failure data to lower mitigation overheads.

Kurt Ferreira

Sandia National Laboratories  
kbferre@sandia.gov

### MS38

#### A Cooperative Compiler and Runtime Checkpoint/Restart Approach for Kokkos

Resilience is an imminent issue for next-generation platforms due to projected increases in node failures. Typically, Checkpoint/restart is used to save the temporary state of a program and to resume from where failures happened. However, it is not always trivial for users to add the Checkpoint/restart capability to their applications because there is no performance portable and out-of-box Checkpoint/restart framework for different HPC platforms. Kokkos is designed to provide a productive yet performance portable parallel programming model. In this talk, we explore the possibility of constructing a performance portable Checkpoint/restart framework built on top of the Kokkos programming system. Essentially, our key contributions are twofold: First, we propose a clang-based source-to-source translator that analyzes the use of Kokkos's View in Kokkos's parallel constructs and automatically generates a sequence of checkpoint API. Second, we propose a cooperative runtime and its API, which enables optimizing the frequency of checkpointing, thereby avoiding too-fine-grain checkpointing. Such a synergetic approach enables Kokkos programmers to focus on writing standard Kokkos programs. We will discuss the effectiveness and applicability of our approach using different Kokkos applications such as a 3D heat distribution applications and mini-applications from the Mantevo project.

Akihiro Hayashi

Georgia Tech  
ahayashi@gatech.edu

### MS38

#### Multi-Model Fault-Tolerant and Malleable MPI+x Parallel Programming with Transparent Checkpointing and Next-Generation RDMA Protocols

The authors detail multiple models of fault-tolerance with MPI+X (X=GPUs or threads) and discuss both explicit and transparent checkpointing as a collaborative aspect of supporting both malleability and fault-tolerance in such programs. Discussion of collaboration between the fault-tolerant MPI program and the fault-tolerant model, as well as between the fault-tolerant model(s) and the checkpoint system are covered. Aspects of malleability, under consideration in for MPI-5 are covered too, as complementary. Discussion of robust operation in coverage cloud-HPC environments will ephemeral/spot resources is included. Aspects of MPI semantics provide hints as to where transparent checkpointing and malleability can most appropriately cope with MPI message activity and handle both quiescent and non-quiescent state. Key contributions of this work include integration of CRAC+MANA from Northeastern with the UTC ExaMPI research implementation of MPI extended to support both the MPI Stages and FA-MPI models of fault-tolerant MPI. Implications for explicit checkpointing libraries like SCR are also discussed. The potential for integrating malleability with message re-routing

through advanced RDMA techniques is mentioned, particularly involve RVMA (aka, RDMA 2.0) as a further source of resilience and reconfigurability.

Anthony Skjellum

University of Tennessee at Chattanooga  
tony-skjellum@utc.edu

Gene Cooperman

Khoury College of Computer Sciences  
Northeastern University, Boston, MA  
gene@ccs.neu.edu

Ryan Grant

Queen's University  
ryan.grant@queensu.ca

Purushotham Bangalore

University of Alabama  
pvbangalore@ua.edu

### MS39

#### HP-Adaptive Finite Element Methods for Large-Scale Parallel Computations

hp-finite element methods – where both the mesh size and the polynomial degree used on every cell are chosen adaptively – have been known to have excellent convergence properties since the 1980s. Yet, they are not widely used, possibly because the large finite element libraries do not provide implementations. We have undertaken the task of providing comprehensive support for hp methods for parallel computations in the deal.II library where this support has already been available for sequential computations since around 2007. In this talk, we will discuss what algorithmic infrastructure is necessary for efficient, large-scale parallel support for hp-FEM. We will then illustrate our approach with performance and scalability data for all parts of realistic computations on up to many thousands of processors.

Wolfgang Bangerth, Marc Fehling

Colorado State University  
Department of Mathematics  
bangerth@colostate.edu, marc.fehling@colostate.edu

### MS39

#### AMR at Scale for Radiation and Combustion in Uintah

AMR has been used in the Uintah code in many different physically applications including phase field models and combustion. Its use in combustion has been in conjunction with Ray tracing radiation models so far. In this talk the scalability of the AMR radiation approach is shown on a variety of systems such as Fronterra Summit and Lassen. We hope to include full scale summit results and to compare GPU and CPU performance. At the same time the Uintah combustion users have been reluctant to use AMR for their standard combustion solver. We will explore the algorithmic issues they have raised and discussed how they may be resolved. Finally we will describe a Uintah combustion benchmark for exascale and discuss what may be the challenges an getting it to run on Aurora and Frontier.

Martin Berzins

Scientific Computing and Imaging Institute  
University of Utah

mb@sci.utah.edu

**MS39****Adaptive Mesh Refinement Simulations of Compressible Reacting Flows in Complex Geometry**

Compressible reacting flow systems, such as turbine combustors and internal combustion engine cylinders, exhibit complex multiscale behavior and their detailed simulation at the scale of realistic devices requires massive computing resources. Even with the availability of exascale computing hardware, such simulations at practically relevant time and length scales require advanced algorithms that can intelligently reduce the computing resources required while also retaining required accuracy in the solution. PeleC is a code for compressible reacting flows in complex geometries that features adaptive mesh refinement (AMR). AMR is used to capture and track critical time-dependent fine-scale flow features that arise from chemical reactions and their interaction with complex boundary structures. In this talk, we discuss the embedded boundary formulation used in conjunction with AMR in PeleC to represent practically relevant complex flow geometries, and the coupling strategy used to incorporate the combustion chemistry. We will demonstrate that PeleC, which leverages the AMReX library for block-structured AMR, is well-suited for modern, extreme-scale, heterogeneous compute platforms.

Marc T. Henry de Frahan, Jon Rood, Hariswaran Sitaraman, Bruce A. Perry, Shashank Yellapantula, Ray W. Grout, Marcus Day  
National Renewable Energy Laboratory  
marc.henrydefrahan@nrel.gov, jon.rood@nrel.gov,  
hariswaran.sitaraman@nrel.gov, bruce.perry@nrel.gov,  
shashank.yellapantula@nrel.gov, ray.grout@nrel.gov,  
marcus.day@nrel.gov

**MS39****A Family of MAC-Like Discretizations for Hierarchically Refined Meshes**

The Marker-and-Cell (MAC) discretization for Navier-Stokes equations is famously efficient and super-convergent for uniform grids. Practitioners of finite difference, finite volume, discontinuous Galerkin, and mixed-finite element methods all have ways of deriving the MAC scheme within their favored frameworks. But while these derivations coincide on a uniform grid, they diverge into different methods in the presence of the coarse-fine boundaries that occur in adaptively refined meshes, each with their own merits. Inspired by local codifferential discretizations in other work (which have appeared in the multipoint flux mixed finite element method of Wheeler & Yotov and the more recent augmented serendipity elements of Lee & Winther), we show how the MAC scheme can be framed as a mixed discretization of the vector Laplacian in which the codifferential is locally eliminated. From this starting point, we can generalize to AMR meshes: quadrees and octrees, but beyond that to any hierarchically refined mesh. The generalization retains the formal order of accuracy and sparsity of the original and respects the Helmholtz decomposition. We will also discuss the finite difference stencils generated by this approach, and the possibilities of combining this approach with higher order methods.

Tobin Isaac  
Georgia Tech

tisaac@cc.gatech.edu

**MS40****High Performance Bayesian Computing with Applications to Spatial/Spatiotemporal Data Analysis and Genome Wide Association Studies**

We present Bayesian algorithms for multi-core architectures for the analysis of spatial/spatiotemporal data and genome wide association studies (GWAS) data. To design Bayesian algorithms, we use the following guidelines: (1) Minimize communication amongst cores; (2) Break large datasets into many small pieces; (3) Transform the data to make computations faster. We employ these guidelines in three data analysis applications. The first application presents a multiscale decouple/recouple algorithm for the analysis of multivariate spatiotemporal data. This algorithm decouples the original dataset into many small datasets at multiple scales of resolution, analyzes these small datasets in parallel, and finally recouples these analyses into a coherent analysis at the original scale of resolution. The second application presents a Bayesian algorithm for the analysis of areal/regional spatial data where the selection of a best subset of predictors is of interest. This algorithm uses the spectral decomposition of the covariance matrix to transform the dependent variable and the predictors to the spectral domain, fits each possible model in parallel, and finally combines the results from the various models to choose the best subset of predictors. Related to the second application, the third application presents an algorithm for the identification of important predictors in GWAS where the number of possible predictors is on the order of hundreds of thousands to millions.

Marco Ferreira  
Associate Professor of Statistics, Virginia Tech  
marf@vt.edu

Jacob Williams, Erica Porter  
Virginia Tech  
jwilliams@vt.edu, ericamp@vt.edu

Mohamed Elkhoully  
Cairo University  
emoham7@vt.edu

Christopher Franck  
Virginia Tech  
chfranck@vt.edu

**MS40****Supercomputing-Driven Climatic Clustering and Longitudinal Analysis with Impacts on Food, Bioenergy, and Pandemics**

Predicted growth in world population will put unparalleled stress on the need for sustainable energy and global food production, as well as increase the likelihood of future pandemics. In this work, we identify high-resolution environmental zones in the context of a changing climate to predict longitudinal processes relevant to these challenges. We develop novel climatic similarity methods to find high-order relationships between climate zones at high geospatial resolution. We run these methods at massive scale using two of the world's top computing systems. We also perform unbiased, data-driven climatic analyses at unprecedented scale and accuracy compared to previous approaches. Our application runs for this work result in one of the largest

matrix product calculations ever performed ( $168.7 \times 10^{21}$  operations), as well as the highest operation rate within an application ever reported ( $9.37 \times 10^8$  operations/sec). The methods discussed herein are also applicable to correlation analyses in other diverse fields such as systems biology, ecology, materials science, carbon cycles, biogeochemistry, additive manufacturing, and zoonosis research.

Daniel A. Jacobson  
Oak Ridge National Laboratory  
jacobsonda@ornl.gov

#### MS40

##### **Fine Tuning the Grouping Approach to Parallelization of Statistics/Machine Learning Methods**

A common approach to parallelization in data science is to randomly group the rows of the data, apply a serial algorithm to each group, then apply averaging or voting to the groups. This approach can be highly effective, even producing superlinear speedup (Matloff, JSS, 2016). However, it raises an issue regarding tuning parameters. Since the latter depend on the number of data points, the best parameter values for the grouped approach may differ substantially from the best set for the ungrouped data. A case study will involve parallelizing the FOCI method of variable selection.

Norman Matloff  
University of California, Davis  
matloff@cs.ucdavis.edu

#### MS40

##### **Scaling Shrinkage Estimates of Large Covariance and Correlation Matrices**

Many real-world data applications require the estimation of correlation/covariance matrices, and when dealing with large datasets, empirical correlation/covariance matrices may be ill-conditioned. However, shrinkage-estimated covariance matrices are always positive definite, well conditioned (so that the inverse always exists), and show small mean squared error. Nevertheless, the latter requires access to the full data matrix and when both  $n$  and  $p$  are large this can be problematic. In particular, forming the covariance matrix may require enormous amounts of memory, far exceeding the capacity of a typical users workstation. Operating out of core is one possible solution, but because disk I/O speeds are slow, this is an unappealing solution for large datasets. Therefore, by exploiting in-house high-performance computing (HPC) capabilities we scaled our implementation of a shrinkage approach for the efficient estimation of correlation/covariance matrices. We apply our correlation/covariance matrix shrinkage approach to the uncertainty quantification for DL models by estimating an a priori abstaining classifier for multitask convolutional neural networks (MTCNN).

Alina Peluso  
ORNL  
peluso@ornl.gov

Drew Schmidt  
Oak Ridge National Laboratory  
schmidt@ornl.gov

Blair Christian  
ORNL  
christianjb@ornl.gov

George Ostrouchov  
Computer Science and Mathematics Division  
Oak Ridge National Laboratory  
ostrouchovg@ornl.gov

#### MS41

##### **Lossy Compression and Mixed Precision Strategies for Memory-Bound Linear Algebra**

The performance of sparse linear algebra is to a large extent constrained by the communication bandwidth, motivating the recent investigation of sophisticated techniques to avoid, reduce, and/or hide data transfers in-between processors and between processors and main memory. One promising strategy is to decouple the memory precision from the arithmetic precision, and compress the data before invoking communication operations. While this generally comes with a loss of information, the strategy can be reasonable when operating with approximate objects, e.g. preconditioners used in iterative methods or self-healing iterative solvers. We will present a memory accessor separating the arithmetic precision from the memory precision and mixed precision algorithms based on the strategy of employing lower precision formats for communication and memory access without impacting the final accuracy.

Hartwig Anzt  
University of Tennessee, U.S.  
hantz@icl.utk.edu

#### MS41

##### **Mixed-Precision Algorithm for Approximating Selected Eigenvalues and Eigenvectors of Symmetric and Hermitian Matrices**

As the new hardware is being equipped with powerful low-precision capabilities driven primarily by the needs of the burgeoning field of Artificial Intelligence (AI), mixed-precision algorithms are now showing far greater potential and renewed interest in scientific computing community. The multi-precision methods commonly follow approximate-iterate scheme by first obtaining the approximate solution from a low-precision factorization and solve. Then, they iteratively refine the solution to the desired accuracy that is often as high as what is possible with traditional approaches. While targeting symmetric and Hermitian eigenvalue problems of the form  $Ax = \lambda x$ , we revisit the SICE algorithm proposed by Dongarra et al. By applying the Sherman-Morrison formula on the diagonally-shifted tridiagonal systems, we propose an updated SICE-SM algorithm. By incorporating the latest two-stage algorithms from the PLASMA and MAGMA software libraries for numerical linear algebra, we achieved up to 3.6 times speedup using the mixed-precision eigensolver with the blocked SICE-SM algorithm for iterative refinement when compared with full double complex precision solvers for the cases with a portion of eigenvalues and eigenvectors requested.

Piotr Luszczek  
University of Tennessee  
luszczek@icl.utk.edu

Yaohung M. Tsai  
Innovative Computing Laboratory  
University of Tennessee  
ytsai2@icl.utk.edu

Jack J. Dongarra

University of Tennessee, Oak Ridge National Laboratory,  
USA  
dongarra@icl.utk.edu

#### MS41

##### **BLIS: From BLAS to Tensor Contraction, Mixed-Precision, and Beyond**

The BLIS project provides a portable framework for implementing high-performance BLAS operations such as matrix-matrix multiplication. Portability (in particular, performance portability) is supported by formulating level 3 BLAS operations as a set of 5 loops, written in C99, around an assembly-language microkernel, in addition to matrix packing operations to improve data locality. Because the only "real" work happens in these computational and packing microkernels, we can use this framework to go beyond the standard BLAS API. In particular, we discuss how BLIS enables low-complexity implementations of mixed-precision/mixed-domain linear algebra and extended operations such as tensor contraction and machine learning primitives.

Devin Matthews

Southern Methodist University  
damatthews@smu.edu

#### MS41

##### **HPAC: Evaluating Approximate Computing Techniques on HPC Openmp Applications**

As we approach the limits of Moores law, researchers are exploring new paradigms for future high-performance computing (HPC) systems. Approximate computing has gained traction by promising to deliver substantial computing power. However, due to the stringent accuracy requirements of HPC scientific applications, the broad adoption of approximate computing methods in HPC requires an in-depth understanding of the applications amenability to approximations. We develop HPAC, a framework with compiler and runtime support for code annotation and transformation, and accuracy vs. performance trade-off analysis of OpenMP HPC applications. We use HPAC to perform an in-depth analysis of the effectiveness of approximate computing techniques when applied to HPC applications. The results reveal possible performance gains of approximation and its interplay with parallel execution. For instance, in the LULESH proxy application approximation provides substantial performance gains due to the reduction of memory accesses. However, in the leukocyte benchmark approximation induces load imbalance in the parallel execution and thus limiting the performance gains.

Konstantinos Parasyris, Giorgis Georgakoudis, Harshitha Menon, James Diffenderfer, Ignacio Laguna, Daniel Osei-Kuffuor, Markus Schordan  
Lawrence Livermore National Laboratory  
parasyris1@llnl.gov, georgakoudis1@llnl.gov,  
harshitha@llnl.gov, diffenderfer2@llnl.gov, lagunaperalt1@llnl.gov, osei-kuffuor1@llnl.gov, schordan1@llnl.gov

#### MS42

##### **Extreme-Scale Grid Optimization with the Exascale Grid Optimization (ExaGO) Toolkit**

In this talk, we introduce Exascale Grid Optimization (ExaGO), an open-source package for large-scale power

grid optimization on parallel, distributed, and heterogeneous GPU architectures. We will discuss ExaGO's architecture, formulation, and features along with presenting its performance on large-scale grid optimization problems.

Shrirang Abhyankar

Pacific Northwest National Laboratory  
shri@pnnl.gov

Slaven Peles

PNNL

slaven.peles@pnnl.gov

Jingyi Wang

Lawrence Livermore National Laboratory  
7000 East Avenue, Livermore, CA 94550  
wang125@llnl.gov

#### MS42

##### **A Scalable Decomposition for Solving ACOPF with GPU Batch Solver**

We present the application of distributed control of alternating current optimal power flow, where a large problem is decomposed into many smaller nonlinear programs using a Lagrangian approach. The small nonlinear programs are solved by ExaTron, the implementation of a trust-region Newton algorithm for bound-constrained nonlinear programming problems, fully running on multiple GPUs. Without data transfers between CPU and GPU, our implementation achieves the elimination of a major performance bottleneck under a memory-bound situation, particularly when solving many small problems in batch. We discuss the design principles and implementation details for our kernel function and core operations. Different design choices are justified by numerical experiments. We present the numerical results that demonstrate computational performance of ExaTron on the Summit supercomputer at Oak Ridge National Laboratory, showing the linear scaling with respect to the batch size and the number of GPUs and more than 35 times speedup on 6 GPUs than on 40 CPUs available on a single node.

Kibaek Kim

Argonne National Laboratory  
kimk@anl.gov

Youngdae Kim

Computer Science Department  
University of Wisconsin Madison  
youngdae@cs.wisc.edu

#### MS42

##### **Towards Real-Time ACOPF on Emerging Architectures**

The ACOPF is an essential algorithm used to plan and optimize the dispatch of thousands of generators in the United States power grid. The ACOPF is a large, nonlinear optimization problem that needs to be re-computed several times as the energy demand and generation change over the day. As renewable generation becomes an increasingly important part of the energy mix, its rapid variability requires solving the ACOPF at a much higher frequency thus becoming a "real-time" optimization problem. In this "real-time" setting, there is a requirement for the algorithm to finalize in some pre-established amount of time and to produce a feasible solution. Our talk will review the ACOPF problem which consists of a set of equality

constraints representing the physics of the network, and a set of inequality constraints that represent the engineering limitations of the system. We will describe the roadblocks for real-time optimization and study how novel single-instruction multiple-dispatch (SIMD) architectures can be leveraged to overcome these hurdles.

Daniel A. Maldonado  
Argonne National Laboratory  
9700 S. Cass Avenue Lemont, IL 60439  
maldonadod@anl.gov

#### MS43

##### **Inverse-Deletion BFS - Revisiting Static Graph BFS Traversals with Dynamic Graph Operations**

Breadth-First Search (BFS) traversals appear in a wide range of applications and domains. BFS traversals determine the distance between key vertices and the remaining vertices in the network. The distance between the vertices often referred to as the number of hops, is the shortest path between a root and the remaining vertices in the graph. Given its applicability across multiple domains, BFS has received significant attention from both the theoretical and applied communities, including many algorithms for parallelizing BFS. BFS traversals are typically conducted on a snapshot of a graph (commonly referred to as a static graph). In this paper, we show a novel algorithm for executing a BFS traversal. While our algorithm also uses a static graph for its input, we show how to execute a BFS traversal by using dynamic graph operations. Specifically, in each BFS level, we remove the subset of the edges that are not necessary for the next phase. These edges do not impact finding the vertices in the next level and these reduce the number of random memory accesses. We show a top-down BFS variation of our new algorithm. While our implementation does not outperform state-of-the-art implementations its performance is competitive. Furthermore, it shows a novel way to implement BFS and opens up many research opportunities.

Oded Green  
NVIDIA  
Georgia Tech:School of Computational Science and Engineering  
ogreen@gatech.edu

#### MS43

##### **Graph Analytics in the Exascale Era**

Combinatorial algorithms in general and graph algorithms in particular play a critical enabling role in numerous scientific applications. The irregular memory access nature of these algorithms makes them one of the hardest algorithmic kernels to implement on parallel systems. To address the challenges, the co-design center on combinatorial algorithms, ExaGraph, was established to design and develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms chosen from a set of exascale applications, targeting the pre-exascale and exascale systems. Algebraic and combinatorial methods have a complementary role in the advancement of computational science and engineering, including playing an enabling role on each other. In this talk, we present a brief overview of the latest work on multi-graphics processing unit (GPU) systems for two prototypical graph problems: graph clustering and influence maximization, along with demonstrations on substantial gains in performance for both PNNL systems and the current No. 2 supercomputer, Summit.

We believe that several applications will benefit from the algorithmic and software tools developed by the ExaGraph team.

Mahantesh Halappanavar  
Pacific Northwest National Laboratory  
hala@pnnl.gov

Marco Minutoli  
Pacific Northwest National Lab  
marco.minutoli@pnnl.gov

Sayan Ghosh  
Pacific Northwest National Laboratory  
sayan.ghosh@pnnl.gov

Aydin Buluc  
Lawrence Berkeley National Laboratory  
abuluc@lbl.gov

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

Alex Pothen  
Purdue University  
Department of Computer Science  
apothen@purdue.edu

#### MS43

##### **Parallel Graph Algorithms by Blocks for Heterogeneous Systems**

Efficient graph analysis is crucial for many data analytics applications. However, developing high-performance parallel graph algorithms is a challenging task: Several well-known hardware and software-related challenges arise. Recent architectural advances and variety makes designing flexible graph kernels that can run well on various platforms more challenging. In this work, we propose a novel graph processing framework for modern shared-memory heterogeneous platforms. Our framework implements a block-based programming model. That allows a user to express a graph algorithm using host and device functors that operate on an ordered list of blocks (subgraphs). Our framework deploys these computations to all available resources in a heterogeneous architecture. One can implement a diverse set of graph algorithms in our framework, and we support graph computations that fit in host DRAM but not in GPU device memory. Our experimental results show that these achieve competitive performance compared to hand-optimized implementations. Based on our experiments in the median, our framework achieves 1.6 times to 5.7 times better performance than state-of-the-art graph processing systems, respectively.

Abdurrahman Yasar  
Georgia Institute of Technology  
School of Computational Science and Engineering  
ayasar@gatech.edu

Siva Rajamanickam, Jonathan Berry  
Sandia National Laboratories  
srajama@sandia.gov, jberry@sandia.gov

Umit V. Catalyurek  
Georgia Institute of Technology  
School of Computational Science and Engineering

umit@gatech.edu

#### MS44

##### **ComPort: Rigorous Testing Methods to Safeguard Numerics During Software Porting**

The ComPort project develops techniques to ensure that numerical results exhibited by codes continue to meet the intended science objectives even after porting to new environments that involve different compilers and/or the execution platforms. Our specific focus will be the increased use of GPUs and other accelerators in future platforms. After an initial study to obtain an in-depth understanding of how accelerators treat numerical expressions differently from traditional CPUs, the ComPort project will aim for techniques to root-cause code sections that introduce unacceptable degrees of numerical variability. Incisive test-generation methods driven by user-written higher-level specifications will help reveal such code regions with higher assurance. Automated expression repair through rewriting will help restore acceptable numerical behavior within observed operating ranges of values. ComPort also intends to incorporate these methods into performance portability layers such as Kokkos and Raja.

Ganesh Gopalakrishnan  
University of Utah  
ganesh@cs.utah.edu

Ignacio Laguna  
Lawrence Livermore National Laboratory  
ilaguna@llnl.gov

Ang Li  
Pacific Northwest National Laboratory  
ang.li@pnnl.gov

Pavel Panchekha  
University of Utah  
pavpan@cs.utah.edu

Cindy Rubio González  
University of California, Davis  
crubio@ucdavis.edu

Zachary Tatlock  
University of Washington, Seattle  
ztatlock@cs.washington.edu

#### MS44

##### **Testing Strategies in a Large Scale Scientific Code Experiences and Aspirations**

Maintaining correctness of Scientific Codes is important to ensure confidence in the results. In this talk I will discuss various methods of testing that I have encountered during my work with the Chroma Lattice QCD code, including algorithmic testing, regression testing, unit testing and discuss our current status and future aspirations.

Balint Joo  
Oak Ridge National Laboratory  
joob@ornl.gov

#### MS44

##### **Toward Automated Test Generation for Performance Portable Programs using Clang/llvm and**

#### Formal Methods

Automated test generation using formal methods has been widely applied to verifying the correctness of system software and hardware, but comprehensive test coverage for scientific applications relevant to domain scientists is difficult to achieve with current frameworks. Furthermore, increasing complexity and limited accessibility of leadership-class high performance computing (HPC) systems make the testing of emerging applications even more complicated. To address these challenges, we are going to develop and integrate high performance computing (HPC) runtime abstractions and emulation into automated test generation frameworks to make test generation faster, more comprehensive and more understandable to the domain scientists. In this talk, we will demonstrate our initial attempt to automate detection of error-prone segments of application program source written with Kokkos (performance portable parallel programming model) for successful automatic test case generation. We will discuss (1) formal specification of the Kokkos programming model and (2) adaptation of Clangs abstract syntactic tree (AST) to extract Kokkos abstractions and API calls in a program source and (3) the integration of these two techniques.

Shyamali Mukherjee, Noah Evans  
Sandia National Laboratories  
smukher@sandia.gov, nevans@sandia.gov

Alessandro Orso, Vivek Sarkar  
Georgia Institute of Technology  
orso@cc.gatech.edu, vsarkar@gatech.edu

Keita Teranishi  
Sandia National Laboratories  
Livermore, California  
knteran@sandia.gov

#### MS44

##### **Testing, Verifying, and Validating Openmp Implementations**

The validation and verification of OpenMP implementations is critical as it provides a mechanism that tests for an implementations conformance to the OpenMP standard, as well as highlights ambiguities in the OpenMP specification. Currently OpenMP is at Specification version 5.1 while the vendors are still at different stages of implementing 5.0 features. With difference in vendor support it is critical for applications to know 1) the status of the implementations, especially with respect to the offload features that allow computations to be performed on target devices such as GPUs; 2) the conformance of the implementations to the specification this is the correctness aspect –which becomes increasingly complex when multiple OpenMP features are exercised together. Through the SOLLVE V&V we have tried to provide the fundamental tests that check for conformance keeping in mind how applications would use these OpenMP features. This talk will focus on the how, what, and the interesting lessons learnt during the testing of different OpenMP implementations and the key takeaways from our experiences.

Swaroop Pophale  
Oak Ridge National Laboratory  
pophale@ornl.gov

Nolan Baker  
University of Delaware  
nolanb@udel.edu

David E. Bernholdt  
Oak Ridge National Laboratory  
Computer Science and Mathematics Division  
bernholdtde@ornl.gov

Sunita Chandrasekaran, Josh Davis, Thomas Huber,  
Sergio Pino Gallardo  
University of Delaware  
schandra@udel.edu, jhdavis@udel.edu, thuber@udel.edu,  
sergiop@udel.edu

#### MS45

##### Online and Offline Precision Tuning in Hardware Accelerators

In recent years, there has been considerable interest in application-specific or domain-specific hardware accelerators for computational kernels. The ability to define the hardware architecture to suit the numerical application at hand gives rise to considerable potential sources of energy and cost efficiency, including through the development of custom and mixed-precision arithmetics. In this talk, I will summarise two of our recent research contributions in this space to highlight the possibilities. In one approach, we have developed a hardware accelerator using redundant number representations that customises the precision of its computation on-the-fly, computing in an 'online' manner; I will present some analysis of the convergence of such computations. In another approach, we have utilised recent developments in satisfiability modulo theories to estimate the probabilistic impact of low precision computation on small computational kernels suitable for a hardware accelerator. I hope to end the talk with a summary of some of the key areas where I believe the NA community and the hardware accelerator community should prioritise working together over the next few years.

George A. Constantinides  
Imperial College London  
g.constantinides@imperial.ac.uk

#### MS45

##### Fluid Simulations Accelerated with 16 bit: Approaching 4x Speedup on A64FX by Squeezing ShallowWaters.jl into Float16

Most Earth-system simulations run on conventional CPUs in 64-bit double-precision floats, although the need for high-precision in the presence of large uncertainties has been questioned. Fugaku, currently the worlds fastest supercomputer, is based on A64FX microprocessors, which also support the 16-bit low-precision format Float16. We investigate the Float16 performance on A64FX with ShallowWaters.jl, the first fluid circulation model that runs entirely with 16-bit arithmetic. Techniques are implemented to address precision and dynamic range issues in 16 bit. The precision-critical time integration is augmented to include compensated summation to minimize rounding errors. Such a compensated time integration is as precise but faster than mixed-precision with 16 and 32-bit floats. As subnormals are inefficiently supported on A64FX the very limited range available in Float16 is  $6 \cdot 10^{-5}$  to 65504. We develop the analysis-number format Sherlogs.jl to log the arithmetic results during the simulation. The equations in ShallowWaters.jl are then systematically rescaled to fit into Float16, using 97% of the available representable numbers. Consequently, we benchmark speedups of 3.8x on A64FX with Float16 and 3.6x when adding a compensated time integration. ShallowWaters.jl shares essential algorithms

with large Earth-system models and therefore shows that 16-bit calculations are indeed a competitive way to accelerate Earth-system simulations on available hardware.

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

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

Matteo Croci  
Mathematical Institute  
University of Oxford  
matteo.croci@maths.ox.ac.uk

Peter D. Dueben  
European Weather Centre (ECMWF)  
peter.dueben@ecmwf.int

Tim Palmer  
University of Oxford, UK  
t.n.palmer@atm.ox.ac.uk

#### MS45

##### Reducing Data Movement in Mixed Precision LU Factorization with GPU Tensor Cores

Mixed precision accelerators such as GPU tensor cores can significantly improve both the performance and accuracy of the solution of linear systems, in particular in the context of iterative refinement based on mixed precision LU factorization. The simplest and safest implementation of LU factorization with tensor cores is to store the matrix in simple, rather than half, precision. However, this approach is expensive in terms of storage and data movement. We will present a new LU factorization algorithm that works on a matrix stored almost entirely in half precision, thereby halving storage and data movement, and achieving a speedup of nearly a factor two. A careful rounding error analysis allows us to do this without endangering the numerical stability of the factorization and thus preserving the fast convergence of iterative refinement.

Atef Dorai  
LIP6  
atef.dorai@lip6.fr

Roman Iakymchuk  
Sorbonne Université and Fraunhofer ITWM  
roman.iakymchuk@sorbonne-universite.fr

Florent Lopez  
Ansys  
florent.joan.lopez@gmail.com

Theo Mary  
Sorbonne Université, CNRS, LIP6  
theo.mary@lip6.fr

#### MS45

##### BLIS: Mixing, Matching, and Extending Precision

We report on efforts to support alternative floating point precisions, mixed precision, and/or mixed domain computation within the BLAS-like Library Instantiation Software (BLIS) framework. We describe a novel approach

to the high-performance implementation of double-double precision matrix-matrix multiplication (gemm) in terms of a given double-precision matrix-matrix multiplication (dgemm). This technique can be generalized to bootstrap a higher-precision gemm in terms of a given lower-precision gemms. The alternative object-based interface available for BLIS simplifies the calling sequences for accessing such BLAS-like functionality.

Greg Henry  
Intel Corporation  
greg.henry@intel.com

Devin Matthews  
Southern Methodist University  
damatthews@smu.edu

Maggie E. Myers  
The University of Texas at Austin  
myers@cs.utexas.edu

Devangi N. Parikh  
University of Texas at Austin  
devangi.parikh@utexas.edu

Robert A. van de Geijn  
The University of Texas at Austin  
Department of Computer Science  
rvdg@cs.utexas.edu

Field G. Van Zee  
Department of Computer Sciences  
The University of Texas at Austin  
field@cs.utexas.edu

## MS46

### An Overview of Particles in Amrex

Particle methods are used to model a variety of phenomena in scientific computing, including the dynamics of charged particle beams subject to electric and magnetic fields in accelerators, dark matter in an expanding universe, the solids phase in chemical looping reactors, and many more. The AMReX block-structured adaptive mesh refinement framework provides a set of tools for implementing such methods that is particularly tailored towards particles that interact with a nested hierarchy of mesh data. The addition of particles to structured mesh calculations introduces additional challenges stemming from the irregularity and dynamic nature of the data layout, memory access patterns, and connectivity information. In this talk, I will give an overview of the particle functionality provided in the AMReX framework and demonstrate how it can be used to build scalable, performance-portable adaptive mesh refinement applications that implement a variety of particle methods across a number of different scientific domains.

Andrew Myers  
Lawrence Berkeley National Lab  
atmyers@lbl.gov

## MS46

### Parallel Simulation of Cancer Cell Migration

In biological tissues, cells touch each other with minimal space in-between, limiting their relative motion. When cells migrate, they undergo a jamming transition from a fluid-like to a solid-like state, so called glassy dynamics.

Previous studies have revealed that jamming transition is governed by active motility, interaction forces, cell shape and applied stress. Complementing experimental techniques, numerical modeling and simulation offer significant benefits towards gaining insights into biological processes and predicting outcomes. In this talk I describe a numerical model for studying jamming and glass transitions that account for both cell shapes and motility. A hybrid vertex and deformable polygon method are introduced that accounts for the viscoelasticity of the cell membrane. The model is implemented using the particle data structure in the AMReX software framework, where each particle represents one cell with multiple vertices defining its boundary; the intra- and inter-cellular forces are represented as interactions between data members of the same particle vs of different particles, respectively, and the AMReX neighbor list construct is used to minimize the number of interactions that must be evaluated. Exploiting the AMReX framework to achieve performance portability, this model can run efficiently on high-performance computer systems with heterogeneous architectures.

Dung Nguyen  
Seattle Pacific University  
nguyend16@spu.edu

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

Emily Bogle, Amanda R. Harkleroad, Wai Lau  
Seattle Pacific University  
boglee@spu.edu, harkleroad@spu.edu, lau@spu.edu

Victor Zendejas Lopez  
UC Berkeley  
vzendejasl@berkeley.edu

## MS46

### CFD-Dem Simulations of Reacting Fluidized Beds with MFIx-Exa

MFIx-Exa is a CFD-DEM HPC code under development for simulating chemically reacting multiphase flows and optimizing modern complex geometries for chemical looping reactors (CLRs). Specifically, we model the fluid using a low Mach number formation with a multicomponent ideal gas state equation, and we discretize the particle-unresolved fluid equations using an embedded boundary (EB) aware Godunov scheme with an approximate projection. We represent particles by a soft-sphere spring-dashpot model and evolve them using a forward Euler method with subcycling. The fluid and particle models are coupled through a volume fraction field and interphase mass, momentum, and energy transfer. These coupling terms are computed at the p-th particle position, mapping continuous fluid properties to the particle location through an interpolation operator. The resulting particle properties are transformed into continuous fields through a unit normal transfer kernel. We present numerical results simulating a CLR prototype and benchmark it against experimental data to validate the approach. Besides the prototype's simplifications and reduced scale, these results represent a milestone towards realistic CLR simulations.

Roberto Porcu  
National Energy Technology Laboratory  
roberto.porcu@netl.doe.gov

Jordan Musser  
National Energy Technology Laboratory (NETL)  
jordan.musser@netl.doe.gov

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

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

William Fullmer  
National Energy Technology Laboratory  
Leidos Research Support Team  
william.fullmer@netl.doe.gov

Deepak Rangarajan  
National Energy Technology Laboratory  
deepak.rangarajan@netl.doe.gov

## MS46

### A Scalable Particle-In-Cell Algorithm For Neutrino Quantum Kinetics

Neutron star mergers, or kilonovae, are one of the most intense subjects of astrophysics research in recent years after the LIGO collaboration observed them to emit strong gravitational wave radiation. We have yet to understand many of the details of the merger dynamics, and one of the remaining open questions is to what extent neutrino flavor transformations determine nucleosynthesis for the heavy elements these events produce. Answering this question requires sophisticated computing resources to solve the underlying nonlinear equations for neutrinos in six-dimensional phase space. To achieve this, we developed Emu, a particle-in-cell simulation code that solves the neutrino quantum kinetics equations with arbitrary angular resolution. Emu represents the neutrino distribution using a set of computational particles, each having a unique position, momentum, and quantum state. For flexibility, Emu's C++ kernels for evolving each particle's quantum state are all symbolically generated using Sympy. Emu is built on top of the AMReX computational framework to take advantage of AMReX's scalable, domain-distributed particle-mesh routines, and Emu is performance portable on CPUs and GPUs. Emu thus enables detailed multidimensional studies of the neutrino fast flavor instability, resolved in space, angle, and time. We will discuss Emu's design and scaling on modern supercomputing platforms along with our ongoing studies of the fast flavor instability in neutron star mergers.

Don E. Willcox  
Lawrence Berkeley National Laboratory  
dewillcox@lbl.gov

Sherwood Richers  
University of California, Berkeley  
USA  
srichers@berkeley.edu

Andrew Myers  
Lawrence Berkeley National Lab

atmyers@lbl.gov

## MS47

### Solving Mixed-Integer Nonlinear Programs in the ARPA-e Grid Optimization Competition

The talk will cover recent work on solving large-scale Mixed-Integer NonLinear Programs (MINLPs). We will discuss ideas such as automatic projection of auxiliary variables, building convex outer relaxations as well as convex inner restrictions, and showcasing how they both play an important role in providing feasible solutions and optimality guarantees for challenging MINLPs. We will also discuss open-source software used to implement these ideas while exploiting parallelization in computing clusters. Numerical results on the ARPA-e Grid Optimization Competition Challenge 2 (<https://gocompetition.energy.gov>) will be presented.

Hassan L. Hijazi  
Los Alamos National Laboratory  
PO Box 1663, MS B284, Los Alamos, NM 87545  
hlh@lanl.gov

## MS47

### A Hierarchical Interior-Point Method for Doubly-Bordered Block-Diagonal LPs from Energy System Modeling

The current shift of energy systems in several European countries towards sustainable, renewable technologies comes with a massive decentralization and a concomitant increase in the scale of realistic energy models. The resulting large-scale linear programs (LPs) often contain several hundred million constraints and variables. These models can take prohibitively long to solve for state-of-the-art commercial solvers and even exceed the capacities of a single desktop machine. Against this backdrop, the research project UNSEEN has been launched to develop methods for solving currently intractable energy optimization problems that exhibit a block-diagonal structure with linking constraints and linking variables. Our solver PIPS-IPM++ implements a parallel interior-point method (IPM) and exploits the problem structure to parallelize the linear algebra kernel. However, PIPS-IPM++'s performance crucially correlates with the number of linking constraints and variables present in the model; linking parts form a sequential bottleneck in our algorithm. We present the hierarchical approach, a nested Schur complement strategy that exploits the constraints between two consecutive blocks. We demonstrate how PIPS-IPM++ scales with the dimension on several real-world energy system problems, solved with the supercomputer at JSC (Juelich) and compared against commercial solvers.

Nils-Christian Kempke, Charlie Vanaret  
Technische Universität Berlin  
kempke@zib.de, vanaret@math.tu-berlin.de

## MS47

### A Distributed-Memory Optimization Solver for Two-Stage Stochastic Programming Problems

We present a scalable optimization algorithm and its parallel implementation for large-scale two-stage stochastic programming problems. The application includes the security constrained optimal power flow (SCOPF) models routinely used in electrical power grid operations. Such

problems are often nonsmooth, nonconvex in two-stage formulation and can be prohibitively expensive to solve on an industrial scale with the traditional methods. To overcome these challenges, the distributed algorithm proposed in this work solves the first-stage master problem with local quadratic approximation of the second-stage optimization solution. It is shown that under reasonable assumptions and smoothing, global convergence can be achieved. The second-stage subproblems, which require highly varying execution times, are scheduled asynchronously and solved in parallel. This is crucial in achieving good load balancing and parallel efficiency. The algorithm draws similarity to bundle methods, in avoiding line search to reduce the computation time. The step size is instead controlled by the quadratic approximation coefficient. The asynchronous scheduling is implemented in C++ via MPI non-blocking calls to overlap computations with communication. The convergence and scaling results are provided through numerical experiments, in synthetic and SCACOPF problems, conducted on Summit and Lassen supercomputers. We observe quick converging behavior and high parallel efficiency.

Jingyi Wang  
Lawrence Livermore National Laboratory  
7000 East Avenue, Livermore, CA 94550  
wang125@llnl.gov

#### MS48

##### **GPU Support in SUNDIALS Time Integrators and Nonlinear Solvers**

SUNDIALS is a suite of robust and scalable time integrators and nonlinear solvers utilized in a wide variety of applications on computing systems ranging from desktop machines to super computers. With the shift to heterogeneous architectures in exascale systems where most of the computational power comes from GPUs there has been a significant development effort in SUNDIALS to support hardware from various vendors and a variety of programming models as well as provide interfaces to GPU-enabled algebraic solvers. In this talk we will overview our GPU porting strategy leveraging SUNDIALS object-oriented design to offloading computations while keeping algorithmic control on the CPU, summarize the current GPU capabilities of SUNDIALS, and highlight applications uses and recent results on GPU systems. 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-826702.

David J. Gardner, Cody J. Balos  
Lawrence Livermore National Laboratory  
gardner48@llnl.gov, balos1@llnl.gov

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

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

#### MS48

##### **Advances in GPU Methodologies in AMReX**

The AMReX software framework for massively parallel block-structured AMR applications has undergone exten-

sive improvements to run efficiently on GPU supercomputers, especially the DOE exascale systems Perlmutter, Frontier and Aurora. The latest generation of computing technologies has led to additional studies in performance and algorithmic design with a focus on usability and scientific achievement. These advancements have demonstrated substantial gains across the AMReX suite of applications, including WarpX, Nyx, Castro, Pele, MFix-Exa and AMR-Wind. This talk will give an overview of recent AMReX advancements in GPU design and implementation. Topics include advancements in porting to AMD and Intel software frameworks, advances and remaining deficiencies in GPU performance and new technologies explored to enhance AMReXs capabilities and prepare for the next-generation of scientific research.

Kevin N. Gott  
LBNL  
kngott@lbl.gov

Weiqun Zhang  
Lawrence Berkeley National Laboratory  
Center for Computational Sciences and Engineering  
weiqunzhang@lbl.gov

Andrew Myers  
Lawrence Berkeley National Lab  
atmyers@lbl.gov

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

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

#### MS48

##### **Development of Unstructured Mesh Pic Applications for Performant Execution on Gpu Accelerated Supercomputers**

Multiscale simulations using the Particle-in-Cell (PIC) method couple particle motion to fields described by PDEs discretized over meshes of the domain. Such methods are commonly used for plasma physics simulations in fusion devices (ITER, DIII-D, etc.). Given the complexity of the geometric models of these devices, unstructured mesh discretizations are key to efficiently capturing the model feature geometry, often using anisotropic mesh gradations, while controlling mesh quality and overall size. Efficiently developing and running these simulations on leadership class GPU accelerated systems is a challenge that requires a flexible infrastructure that can be tuned to different system architectures by computational scientists and easy to use, stable, interfaces that application developers use to implement physics methods. PUMIPic, the parallel unstructured mesh infrastructure for PIC, is a C++, Kokkos based, library that meets these needs. An overview of PUMIPic and its methods will be presented along with progress on developing and verifying the edge plasma physics application, XGCM, and the impurity transport application, GITRM.

Mark S. Shephard  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

Cameron W. Smith  
 Scientific Computation Research Center  
 Rensselaer Polytechnic Institute  
 smithc11@rpi.edu

#### MS49

##### **Sphynx: The First Graph Partitioner on AMD GPUs**

Sphynx: Graph partitioning is an essential step of large-scale scientific computations to distribute the work among processing elements efficiently. New accelerators used in leadership-class supercomputers have led scientific application developers to port their libraries to those architectures. Sphynx was proposed last year to address the need of a graph partitioner which is portable across the different architectures, and it is the first multi-GPU graph partitioner useable in an MPI+CUDA setting. Sphynx relies on the Trilinos scientific software library and thanks to the portability provided by the Kokkos package, Sphynx is now the first graph partitioning tool to run on (multiple or single) AMD GPUs. In this talk, we present the performance results of Sphynx obtained on an AMD system, which is a small replica of the first exascale system, Frontier.

Seher Acer

Oak Ridge National Laboratory  
 acers@ornl.gov

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

#### MS49

##### **Sparse Tensor Partitioning for Scalable Distributed CPD-ALS**

Sparse tensor decomposition is a fundamental tool widely used in the analysis of multi-dimensional data. The canonical polyadic decomposition (CPD) is one of the most popular decomposition methods and commonly found by the CPD-ALS algorithm. High computational and memory costs of CPD-ALS necessitate the use of a distributed-memory-parallel algorithm for efficiency. MTTKRP operation, which is the bottleneck of the CPD-ALS, requires sparse expand and sparse reduce communications as well as two dense collective communications. MTTKRP becomes bandwidth bound for large decomposition ranks, whereas it becomes latency bound for small decomposition ranks and large number of processors. Several successful hypergraph-partitioning-based tensor partitioning methods have been proposed for reducing the bandwidth requirement. However, the literature is lacking methods and models, which directly target at optimizing the latency requirement of CPD-ALS, although this is a very important aspect for scaling on large number of processors. In this talk, we will present our recently proposed framework and partitioning for reducing the latency overhead of the sparse expand and reduce communications, thus increasing the scalability of CPD-ALS significantly.

Cevdet Aykanat, Ozan Karsavuran, Nabil Abubaker  
 Bilkent University  
 aykanat@cs.bilkent.edu.tr,  
 ozan.karsavuran@cs.bilkent.edu.tr,

nabil.abubaker@bilkent.edu.tr

#### MS49

##### **Hierarchical Partitioning for GPU-Based High Performance Computational Platforms**

This talk will discuss the extension of our existing PuLP/XtraPuLP graph partitioners onto GPU-based distributed computational systems. Two primary challenges and our solutions will be discussed. First, we will discuss our implementation of an efficient graph data structure for coarsening, a key processing step of multilevel partitioning. We use this data structure as part of a novel, fast, and distributed coarsening phase. This efficient coarsening phase allows us to better exploit the massive concurrency available on our target systems, without sacrificing partitioning edge cut quality or balance. Second, we extend the sub-routines within PuLP/XtraPuLP to allow for solutions to hierarchical formulations of the partitioning problem. We demonstrate considerable improvements in terms of speed and efficiency relative to our preliminary efforts. The end goal of this work is to accelerate large-scale climate simulations on leadership-class supercomputers.

Christopher Brissette, George M. Slota  
 Rensselaer Polytechnic Institute  
 brissc@rpi.edu, slotag@rpi.edu

#### MS49

##### **Recent Advances in Shared-Memory (Hyper)Graph Partitioning**

(Hyper)graph partitioning is an important preprocessing step for optimizing data placement and minimizing communication volume in many high-performance computing applications. It asks for a partition of the vertex set of a (hyper)graph into  $k$  blocks of bounded size while minimizing an objective function defined on the (hyper)edges that connects two or more blocks. During the last decades, sequential high-quality partitioners have evolved into sophisticated frameworks, while existing parallel approaches use comparatively weak components that are easier to parallelize. However, to cope with ever growing problem sizes and modern architectures, it has become increasingly important to develop efficient parallel partitioning algorithms whose solution quality is competitive with existing sequential systems. In recent years, research has begun to address this issue by parallelizing the core techniques of their sequential high-quality counterparts without compromising solution quality. In this talk, we will cover recent developments in shared-memory (hyper)graph partitioning and show that this class of new systems achieves good speedups on (hyper)graphs with more than 100 million nodes and edges and significantly outperform previously developed parallel partitioners in terms of solution quality and running time, while producing solutions with comparable quality than their sequential counterparts.

Tobias Heuer

Karlsruhe Institute of Technology  
 tobias.heuer@kit.edu

#### MS50

##### **Composability, Flexibility, and Ease of use for CliMAs Next Generation Earth System Model**

The Climate Modelling Alliance (CliMA) project is developing a full Earth System Model (ESM), entirely in the

high-level, dynamic Julia language. The CliMA model has been developed to target both CPU and GPU architectures, using a common codebase. The high-level application programming interface (API) facilitates modularity and composition of differential operators and the definition of flexible discretizations. This, in turn, is coupled with the low-level API that supports different data layouts, specialized implementations and flexible models for threading, to better face high-performance optimization, data storage, and scalability challenges on modern HPC ecosystems. This talk will describe the design of the software and numerical aspects comprising the ClimaCore.jl low-level interface and its high-level counterparts.

Valeria Barra  
University of Colorado Boulder  
valeria@caltech.edu

Simon Byrne  
California Institute of Technology, U.S.  
simonbyrne@caltech.edu

Jake Bolewski, Sriharsha Kandala  
California Institute of Technology  
jbolewsk@caltech.edu, skandala@caltech.edu

Daniel Z. Huang  
Stanford  
zhengyuh@stanford

Kiran Pamnany  
California Institute of Technology, U.S.  
kpamnany@caltech.edu

Toby Bischoff  
California Institute of Technology  
tobias@caltech.edu

Jeremy E. Kozdon  
Department of Applied Mathematics  
Naval Postgraduate School  
jekozdon@nps.edu

## MS50

### FEniCSx - a Modular PDE Framework: Developing for Heterogeneous Computing

In the latest development version of the popular FEniCS Finite Element framework, we have increased the flexibility of problem composition. In general, meshing, FEM kernel generation, assembly, and solve phases can now be considered as pluggable units. In this context, we have investigated performance portability of the central kernel generation and assembly phase with SYCL, a C++ library that allows us to easily target both CPU and GPU. With assembly "on device", this gives us the flexibility to call device-specific solvers, for example, amgX, and experiment with matrix-free methods. Obtaining the best performance requires attention to detail at the lowest level, as varying memory access and loop transformations can make large timing differences, especially for higher-order kernels. Despite recent improvements in memory capacity, GPU implementations are still limited by the amount of memory available and data transfer rates between devices. We will give an overview of the approaches we have taken to obtain the best performance we can from the available hardware. We will also discuss how carefully arranging memory transfer and allocations can reduce latency and increase

throughput in different accelerators.

Jørgen Dokken, Igor Baratta  
Department of Engineering  
University of Cambridge  
jsd55@cam.ac.uk, ia397@cam.ac.uk

Chris N. Richardson  
University of Cambridge  
chris@bpi.cam.ac.uk

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

## MS50

### Modularity and Composability for High Performance Computing via PETSc

The Portable, Extensible Toolkit for Scientific Computation (PETSc) emerged as a linear algebra for high-performance computing (HPC) library. The composability concept in PETSc relies on the idea that different solvers can play different roles in an algorithmic sequence. For example, the conjugate gradient can be either a direct solver or a preconditioner. Based on this assertion, the user interface design allows different compositions to ensure the solvability of complex engineering problems or advance linear algebra research. To facilitate the interface to different discretizations yielding a range of matrix structures, PETSc allows for modular constructions of discrete operators. Implicit time-steppers which require linear algebraic treatments interface easily with linear solver objects (KSP) or non-linear ones (SNES). Across HPC environments, and more so on nascent heterogeneous architectures, scalability and portability are significant concerns. For example, the convergence of iterative solvers is subject to variations in the number of iterations or computing precision. PETSc's advanced logging and testing infrastructure tracks numerical behavioral changes and assures portability. To enforce scalability, PETSc provides data management (DM) interfaces that the user can control within the solver, at runtime or delegate to PETSc to decide optimal splits for performance.

Oana Marin  
Argonne National Laboratory  
oanam@anl.gov

## MS50

### Composition and DSLs for Finite Element Computation: Pros and Cons

The Firedrake system uses a multi-layered set of DSLs and domain-specific optimising compilers to generate (hopefully) high performance code for finite element assembly. The goal is to deliver expert-level performance to all users, who may not be experts in the necessary low level optimisations, and data structure design. In this talk I will discuss some of the design choices, optimisations and algorithmic transformations that they enable, and pros and cons of this generative approach.

Lawrence Mitchell  
Durham University

lawrence.mitchell@durham.ac.uk

georg.hager@fau.de

MS51

### perf-taint: Extracting Clean Performance Models from Tainted Programs

Performance models are well-known instruments to understand the scaling behavior of parallel applications. They express how performance changes as key execution parameters, such as the number of processes or the size of the input problem, vary. Besides reasoning about program behavior, such models can also be automatically derived from performance data. This is called empirical performance modeling. While this sounds simple at the first glance, this approach faces several serious interrelated challenges, including expensive performance measurements, inaccuracies inflicted by noisy benchmark data, and overall complex experiment design, starting with the selection of the right parameters. The more parameters one considers, the more experiments are needed and the stronger the impact of noise. In this paper, we show how taint analysis, a technique borrowed from the domain of computer security, can substantially improve the modeling process, lowering its cost, improving model quality, and help validate performance models and experimental setups.

Marcin Copik  
ETH Zurich  
marcin.copik@inf.ethz.ch

MS51

### Computational Waves in Parallel Programs and Their Impact on Performance Modeling

Analytic, first-principles performance modeling of nontrivial parallel programs is hampered by the observation that simply adding runtime predictions for different program phases is often inaccurate. One of the reasons for this is that even perfectly load-balanced parallel code can get out of lockstep unless frequent synchronization is enforced. In this talk we show under which circumstances this "desynchronization" can occur and how it can lead to surprising effects, first and foremost an automatic overlap of communication overhead with useful work. Execution bottlenecks such as memory bandwidth and the behavior of a parallel code under external or internal noise are both crucial for the understanding of the dynamics of desynchronization: While bottleneck-free programs tend to synchronize automatically, bandwidth-bound programs appear to be unstable against noise and move into a state that can be described as a "computational wave," where different processes arrive at the same time step at different wall times. The amplitude of such waves is influenced by the communication topology of the code and connected to the velocity of "idle waves," propagating phases of idleness caused by one-off noise injections.

Ayesha Afzal  
Friedrich-Alexander University Erlangen-Nürnberg  
Erlangen National High Performance Computing Center  
ayasha.afzal@fau.de

Georg Hager  
Erlangen National High Performance Computing Center

MS51

### The Price Performance of Performance Models

To understand the scaling behavior of HPC applications, developers often use performance models. A performance model is a formula that expresses a key performance metric, such as runtime, as a function of one or more execution parameters, such as core count and input size. Performance models offer quick insights on a very high level of abstraction, including predictions of future behavior. In view of the complexity of today's applications, which often combine several sophisticated algorithms, creating performance models manually is extremely laborious. Empirical performance modeling, the process of learning such models from performance data, offers a convenient alternative, but comes with its own set of challenges. The two most prominent ones are noise and the cost of the experiments needed to generate the underlying data. In this talk, we will review the state of the art in empirical performance modeling and investigate how we can employ machine learning and other strategies to improve the quality and lower the cost of the resulting models.

Alexandru Calotoiu  
ETH Zurich  
alexandru.calotoiu@inf.ethz.ch

Alexander Geiss, Benedikt Naumann, Marcus Ritter,  
Felix Wolf  
Technical University of Darmstadt  
alexander.geiss1@tu-darmstadt.de,  
benedikt.naumann@tu-darmstadt.de, marcus.ritter@tu-darmstadt.de, felix.wolf@tu-darmstadt.de

MS51

### Extra-P Meets Hatchet: Towards Modeling in Performance Analytics

As the scale of HPC systems increases, so does the volume of performance data generated by various profiling tools. These tools capture multiple performance metrics at various levels of parallelism (e.g., threads and processes) and from different sources (e.g., CPU and GPU) resulting in performance profiles that cannot be analyzed by a simple visual inspection. Hatchet is a Python-based tool that was developed to address this challenge. It brings the power of modern data science into performance analysis by allowing filtering, pruning, and aggregating performance data programmatically. In our ongoing work, we aim to supplement the performance analytics functionality enabled by Hatchet with automated modeling capabilities. Extra-P is a unique and powerful tool for automated performance modeling. It generates human-readable models that help uncover scalability bugs and provide insight into code performance at scale. In this talk, we present our recent effort to integrate Extra-P into Hatchet. With only a few lines of code it is now possible to create insightful models for complete sequences of GraphFrames in Hatchet. We aim to make this functionality accessible to users and further explore various workflows on top of it.

Sergei Shudler  
Lawrence Livermore National Laboratory

shudler1@llnl.gov

## MS52

### Overview of Excalibur Project Neptune

The talk will describe ExCALIBUR project NEPTUNE, list contributors, and describe the aims of the project as set in the abstract for the minisymposium. Background to the other presentations in the session(s) will be given.

Wayne Arter

United Kingdom Atomic Energy Authority, UK  
wayne.arter@ukaea.uk

## MS52

### Numerical and Computational Performance of Spectral Element Methods for Prototype Fusion Problems

High-order methods are increasingly being viewed as an enabling technology for bespoke high-fidelity modelling of challenging physical systems. The use of high-order polynomials within elements, when compared to classical linear and second-order methods, offers advantages from both a numerical analysis and implementation perspective. The resolution power of high-order discretisations is far greater than at lower orders, with favourable dispersion and diffusion characteristics meaning that complex spatial structures and phenomena spanning multiple length scales can be accurately resolved and tracked across long timescales. This accuracy is critical in highly sensitive and complex systems, such as fluid dynamics, where high-order methods are amongst the most promising methods for next generation numerical simulations of high-fidelity problems at the exascale. This talk will outline the developments currently being made in the algorithmic and software improvements required to enable high-order simulations to be performed for the problems of interest to the NEPTUNE community. In particular, we will present current developments in curvilinear mesh generation for fusion-related geometries, and on the computationally-efficient high-order implementation of prototype systems related to heat transport in highly anisotropic systems.

David Moxey

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

Bin Liu

Department of Aeronautics  
Imperial College London  
bin.liu@imperial.ac.uk

Chris Cantwell, Spencer Sherwin

Imperial College London  
c.cantwell@imperial.ac.uk, s.sherwin@imperial.ac.uk

## MS52

### Exploring A Domain Specific Language For Performance Portable Particle Based Algorithms in Julia

Here we present an exploratory Domain Specific language (DSL) for describing particle based algorithms. Particle based methods are widely used in many application areas and, in particular, plasma modelling. As part of the ExCALIBUR NEPTUNE project we are interested in the application of Particle In Cell (PIC) methods as a simu-

lation technique for modelling plasma in the edge regions of fusion reactors. Implementing a PIC algorithm for new a hardware architecture requires substantial knowledge in both the physical domain and the computer science domain. However, the Julia programming language has recently emerged as a single source language that offers a solution to the two-language problem in numerical computing. In this work we investigate the implementation and extension of an existing abstraction for particle based operations within the Julia language. This abstraction provides a separation of concerns that allows efficient implementation of particle based operations within PIC algorithms without requiring detailed knowledge of the target hardware. We provide an overview of the particular algorithms this DSL is designed to facilitate alongside performance results. In particular we discuss applicability and efficiency of this Julia based approach on the hardware architectures of interest in exascale computing.

Will R. Saunders, James Cook

UKAEA

will.saunders@ukaea.uk, james.cook@ukaea.uk

## MS52

### Speeding-Up and Improving Neptune Simulations with Novel Preconditioners

At the heart of the NEPTUNE Programme is the solution of coupled partial differential equations (PDEs) that exhibit challenging properties when trying to solve them. Following discretisation of the PDE, a large linear system  $Ax = b$  must be solved or a series of such large linear systems: the size of the systems to be solved mean that iterative methods are preferable over direct factorisation methods but, without preconditioning to improve rate of convergence, the number of iterations needed to solve the system is too large and it is necessary to stop the iterative process before the required levels of accuracy are met. This loss in accuracy can have a detrimental effect on the convergence properties of any overarching iterative methods that are, for example, dealing with the non-linearity of the PDEs being solved. By using a preconditioner that is suitable for exascale computations, takes into account that the matrix  $A$  is not normally available and, instead, the matrix-vector multiplication is provided via a function, is relatively cheap to initialise and apply, and significantly speeds up the convergence towards the solution of  $Ax = b$ , simulations should become faster to run and have improved accuracy. In this talk, we will summarise the approaches that we have trialled, for example, the novel Markov Chain Monte Carlo Matrix Inverse and Implicit-Factorisation preconditioners, and show the successes and the pitfalls of these approaches.

Sue Thorne

Rutherford Appleton Laboratory  
sue.thorne@stfc.ac.uk

Maksims Abalenkovs

Hartree Centre

UKRI Science and Technology Facilities Council  
maksims.abalenkovs@stfc.ac.uk

Vassil Alexandrov

Hartree Centre, STFC

UK

vassil.alexandrov@stfc.ac.uk

Anton Lebedev

Hartree Centre STFC, UK  
anton.lebedev@stfc.ac.uk

Emre Sahin  
Hartree Centre, STFC UK  
emre.sahin@stfc.ac.uk

### MS53

#### Improving the Speed and Quality of Parallel Graph Coloring

Graph coloring assigns a color to each vertex of a graph such that no two adjacent vertices get the same color. It is a key building block in many applications. In practice, solutions that require fewer distinct colors and that can be computed faster are often preferable. Various coloring heuristics exist that provide different quality versus speed tradeoffs. The highest-quality heuristics tend to be quite slow. To improve performance, several parallel implementations have been proposed. In this talk, I will describe two improvements to the widely used LDF heuristic. First, I will present shortcutting approaches to increase the parallelism by non-speculatively breaking data dependencies. Second, I will present color reduction techniques to boost the solution quality. On 18 graphs from diverse domains, the shortcutting techniques yield 2.5 times more parallelism on average, and the color-reduction techniques result in up to 20% fewer colors needed. My deterministic CUDA implementation running on a Titan V is 2.9 times faster and uses as few or fewer colors as the best GPU codes from the literature.

Ghadeer Alabandi, Martin Burtscher  
Texas State University  
gaa54@txstate.edu, burtscher@txstate.edu

### MS53

#### GPU Accelerated Online Search for Gravitational Waves

Gravitational waves (GWs) open a completely new window for us to learn about the universe. In this talk, a Graphics Processing Unit (GPU) accelerated time-domain low-latency method to search for GWs from coalescing binaries of compact objects will be introduced. The objective of this research is to facilitate fast detection of GWs with a very low delay to allow prompt electromagnetic follow-up observations. Our optimizing methods on GPU can achieve significant performance improvement by matching the application, parallel program and manycore processors well together. The proposed method and the corresponding SPIIR (Summed Parallel Infinite Impulse Response) data processing pipeline have been successfully used in LIGO's GW events detection.

Zhihui Du  
New Jersey Institute of Technology  
zd4@njit.edu

### MS53

#### Kernel Specialization for Faster Segmented Sort on GPUs

Segmented sort orders many variable-length arrays independently. Variability of array sizes obstructs efficient use of GPU resources. This talk introduces a new segmented sort algorithm that addresses this issue. The algorithm includes two stages. The first stage consists of efficient

three-way segments partitioning on GPU. The second stage relies on this partitioning to apply specialized kernels for different classes of segments. This talk highlights implementation details and performance characteristics of the new segmented sort that is implemented as part of CUB.

Georgii Evtushenko  
NVIDIA, US  
evtushenko.georgy@gmail.com

### MS53

#### Delegate Centric Top-K Computation on GPUs

Recent top-k computation projects explore the possibility of revising various sorting algorithms to answer top-k queries on GPUs. These endeavors, unfortunately, perform significantly more work than needed. This talk will present a Delegate-centric top-k system on GPUs that can reduce the top-k workloads significantly. Particularly, we will introduce a comprehensive design of the delegate-centric concept, including maximum delegate, top-k delegate-based filtering, and  $\beta$  delegate mechanisms to help reduce the workload for top-k by more than 99%. We will also introduce four key designs to enable fast multi-GPU.

Hang Liu  
Stevens Institute of Technology  
hliu77@stevens.edu

### MS54

#### Simulating Next-Gen Dataflow Architectures for HPC

The slowing of Moore's Law has led to an explosion of accelerator designs from both industry and academia. Many of these emerging architectures promise to provide savings in energy efficiency, area, and latency but are domain specific. Moreover, developing a simulation framework capable of exploring a design space that includes a model for each of these accelerators is infeasible. One promising family of generalizable technologies are dataflow architectures, which diverge from the traditional sequentially executed instruction model into one that reflects the inherent instruction-level parallelism in a program. These dataflow accelerators can offer better programmability and more flexibility than domain specific hardware and, in a simulation framework, be used to emulate the behavior of a range of accelerators. This talk will cover some of the research and engineering that has gone into integrating a dataflow model inside of the Structural Simulation Toolkit with a focus on AI/ML and linear algebra applications.

Clayton Hughes, Gwendolyn Voskuilen, Arun Rodrigues  
Sandia National Laboratories  
chughes@sandia.gov, grvosku@sandia.gov,  
afrodri@sandia.gov

Simon D. Hammond  
Scalable Computer Architectures  
Sandia National Laboratories  
sdhammo@sandia.gov

### MS54

#### COMET: A Compiler Framework for Next-Generation Heterogeneous Systems

As Moore's Law is coming to an end, simple technology scaling cannot be relied on for performance gain, and new

technologies and computing paradigms must be developed to achieve application performance, which is leading to an explosion of accelerator designs from both industry and academia. Leveraging the new types of compute resources will require new methods and tools to not only achieve high-performance on current heterogeneous systems but also interface with hardware performance modeling and simulation tools to estimate application performance on future specialized accelerators. Compiler technologies have considerably evolved during the last decades and now support sophisticated code analysis and translation methodologies as well as efficient code generation for target heterogeneous systems. In the context of co-design, a compiler-based tool is desirable because it offers the possibility of automatically generate code that leverages new hardware concepts without or with little code modifications in the applications. In this talk, I present COMET tensor algebra compiler targeting quantum chemistry and graph analytics applications. COMET provides an opportunity to perform hardware/software co-design and design space exploration efficiently and to assess the performance of the entire application, instead of only the innermost kernel.

Gokcen Kestor, Ruiqin Tian, Luanzheng Guo  
Pacific Northwest National Laboratory  
gokcen.kestor@pnnl.gov, ruiqin.tian@pnnl.gov,  
lenny.guo@pnnl.gov

Geonhwa Jeong, Tushar Krishna  
Georgia Institute of Technology  
geonhwa.jeong@gatech.edu, tushar@ece.gatech.edu

Roberto Gioiosa  
Pacific Northwest National Laboratory  
roberto.gioiosa@pnnl.gov

#### MS54

##### Streaming Benchmarks on the Graphcore Intelligence Processing Unit

The Graphcore IPU is a new massively parallel processor with architecture designed to optimize AI computations. In this talk we will discuss implementing benchmarks originally developed for traditional compute methods on IPUs. Code for IPUs is developed using the Poplar SDK, which was co-developed by Graphcore with the IPU and supports both direct programming with C++ and integration with standard machine learning frameworks. To take full advantage of the architecture of the IPU, we implement the benchmarks with maximal parallelization and seek to limit communication between pairs of processes in locations with relatively higher communication costs. Finally, we show performance results and comparisons to other implementations of the benchmarks.

Alexandra Porter  
Stanford University  
amporter@cs.stanford.edu

Cannada A. Lewis  
Sandia National Laboratories  
canlewi@sandia.gov

#### MS54

##### Co-Designing Data Flow Accelerators

Several emerging accelerators have characteristics that are attractive for use cases in ML for science domain, traditional machine learning use cases, and in scientific simula-

tions. As these accelerators are being developed, co-design of algorithms, architectures, and applications could lead to better hardware and software. This talk will highlight several recent work that is focused on this co-design effort. We will introduce a suite of mini-applications that will help in the co-design effort called Mantevo-DF focused on data flow focused use cases for scientific simulations. We will also introduce machine learning use cases that are important for ML for science use cases and their uses on modern accelerator hardware. Finally, we will set the stage for the rest of this three part mini-symposium by describing the open challenges in co-design efforts and how the rest of the talks connect to these co-design efforts.

Sivasankaran Rajamanickam  
Sandia National Laboratories  
srajama@sandia.gov

#### MS55

##### Multi-Scale Modeling in Storm Surge Applications with an Eye Towards HPC Environments

Coastal hazards related to strong storms are one of the most frequently recurring and widespread hazards to coastal communities today. In particular, storm surge, the rise of the sea surface responding 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. This talk will focus on the description and approach to the optimization of coastal protection in the presence of climate change impacts. This includes a multimodel approaches, the use of multiscale techniques to represent difficult to resolve features, and reduce order models for uncertainty quantification.

Kyle T. Mandli  
Columbia University  
Applied Physics and Applied Mathematics  
kyle.mandli@columbia.edu

#### MS55

##### High-Order Matrix-Free GPU-Accelerated Linear Solvers with Applications to Fluid Flow

In this talk, we will present a GPU-accelerated incompressible flow solver using a high-order finite element discretization. As in most projection-type methods, solving a Poisson problem for the pressure at every time step represents the dominant computational cost. We will present matrix-free linear solvers for the ill-conditioned linear systems that result from high-order finite element discretizations. At high orders, the cost of assembling the system matrix becomes infeasible, necessitating the use of matrix-free preconditioners that can be built without access to the matrix entries. In this talk, we will consider several matrix-free approaches, including *hp*-multigrid and low-order refined preconditioning. Both the application of the preconditioners and the matrix-free operator evaluation are performed on the GPU.

Will Pazner  
Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory

pazner1@llnl.gov

### MS55

#### Extensible Computational Fluid Dynamics in Julia with Trixi.jl

In recent years, the Julia programming language has made considerable advances in the scientific computing community and positions itself as a viable alternative to established high-performance computing languages such as Fortran and C/C++. Its paradigm of multiple dispatch with just-in-time compilation allows one to easily reuse existing code and surgically extend its functionality without compromising the computational performance. In this talk, we will present some recent results on efficient implementations of entropy stable and kinetic energy preserving discontinuous Galerkin schemes used for the discretization of the compressible Euler equations. These implementations are evaluated in two open source codes: Trixi.jl, an adaptive high-order simulation framework written in Julia, and FLUXO, a Fortran-based simulation tool for computational fluid dynamics. We will discuss several implementation techniques for such summation-by-parts flux differencing operators and evaluate their impact on the execution speed. A particular focus will be on demonstrating how Julia's extensibility features can be used to achieve these efficiency improvements without having to change the existing implementation, and how the resulting performance is on par with a mature Fortran code.

Michael Schlottke-Lakemper  
University of Stuttgart  
m.schlottke-lakemper@hirs.de

Hendrik Ranocha  
King Abdullah University of Science and Technology (KAUST)  
mail@ranocha.de

Gregor Gassner  
University of Cologne  
ggassner@uni-koeln.de

### MS55

#### Exahype 2 Finally Solve Some Astrophysics Problems

ExaHyPE is an engine to simulate hyperbolic equation systems given in first order formulation. A fundamental idea behind ExaHyPE - currently written in its second generation - is the strict separation of what, how, when, where: Users specify their physics via a few PDE terms, while the engine decides how to translate these terms into a numerical scheme, when to launch the individual steps of this scheme, and where to execute the computation's steps. The focus on particular numerical schemes, i.e. this particular software design, distinguishes ExaHyPE from other solutions in the field and facilitates a tight co-orchestration and tuning of all engine ingredients - the AMR, the timestepping, the I/O, and so forth. We give a tour de force through ExaHyPE, which can be controlled exclusively through a Python (code generation) interface, and discuss some particular HPC techniques under the hood: The underlying domain decomposition, ExaHyPE's task formalism, and its ability to deploy tasks to accelerators. Eventually, we will demonstrate how to use the code to simulate a binary black hole merger.

Tobias Weinzierl

Durham University  
tobias.weinzierl@durham.ac.uk

### MS56

#### Asynchronous Truncated Multigrid-Reduction-in-Time (AT-MGRIT)

In this talk, we present the new "asynchronous truncated multigrid-reduction-in-time" (AT-MGRIT) algorithm for introducing time parallelism to the solution of discretized time-dependent problems. The new algorithm is based on the multigrid-reduction-in-time (MGRIT) approach, which, in certain settings, is equivalent to another common multilevel parallel-in-time method, Parareal. In contrast to Parareal and MGRIT that both consider a global temporal grid over the entire time interval on the coarsest level, the AT-MGRIT algorithm uses truncated local time grids on the coarsest level, each grid covering certain temporal subintervals. These local grids can be solved completely independent from each other, which reduces the sequential part of the algorithm and, thus, increases parallelism in the method. Here, we study the effect of using truncated local coarse grids on the convergence of the algorithm, and present parallel results on challenging nonlinear problems.

Jens Hahne  
Bergische Universität Wuppertal  
jens.hahne@math.uni-wuppertal.de

Stephanie Friedhoff  
University of Wuppertal  
friedhoff@uni-wuppertal.de

Ben Southworth  
Los Alamos National Laboratory  
southworth@lanl.gov

### MS56

#### Analysis of a Three-Level Variant of Parareal

In this talk, we present a three-level variant of the parareal algorithm that uses three propagators at the fine, intermediate and coarsest levels. The fine and intermediate levels can both be run in parallel; only the coarsest level propagation is completely sequential. We compare our algorithm with three-level MGRIT, and we present a convergence analysis that uses parareal-type assumptions, i.e., those that involve Lipschitz constants on the propagators. We present numerical experiments to illustrate how sharp the estimates are for various time dependent problems.

Felix Kwok  
Université Laval  
felix.kwok@mat.ulaval.ca or felixk80@gmail.com

### MS56

#### A Unified Analysis Framework for Iterative Parallel-in-Time Algorithms

Parallel-in-time (PinT) methods have received renewed attention in the last two decades, mainly motivated by the advent of new supercomputing resources that provide more and more concurrency. Four of the most widely studied methods are Parareal (Lions, Maday, Turinici), PFASST (Minion and Emmett), MGRIT (Falgout, Friedhoff, Kolev, MacLachlan, Schroder) and a specific form of Space-Time Multigrid (Gander and Neumueller). All those algorithms

have a common aspect: they are iterative PinT "across-the-steps" methods, i.e they represent the time sequential solution process as a large linear or nonlinear system and solve it by iterating on all time steps simultaneously. While various convergence analyses exist for each algorithm separately, it is difficult to relate them and compare convergence of these iterative PinT methods when applied to various model problems, and in applications. In this talk, we present a new approach that lets us analyze the convergence of these four iterative algorithms in a common framework. Following an idea of Hairer and Gander already used to analyze Parareal convergence, this framework is based on an abstract view of each iterative PinT algorithm and provides error bounds for the Dahlquist equation, the fundamental time-dependent test problem. Furthermore, it provides understanding of the different convergence mechanisms in each algorithm and its level of abstraction could eventually lead to novel PinT algorithms.

Thibaut Lunet

Hamburg University of Technology  
thibaut.lunet@tuhh.de

Martin J. Gander  
University of Geneva  
martin.gander@unige.ch

Daniel Ruprecht  
Institute of Mathematics  
Hamburg University of Technology  
daniel.ruprecht@fu-berlin.de

Robert Speck  
Juelich Supercomputing Centre  
Forschungszentrum Juelich GmbH  
r.speck@fz-juelich.de

#### MS56

##### **Assignment of Idle Processors to Spatial Redistributed Domains for Multigrid Reduction in Time**

Parallel-in-time approaches for time-dependent PDEs attract much attention in the context of massively parallel environments. One such approach, multigrid reduction in time (MGRIT), extracts temporal parallelism by applying the reduction-based multigrid method for the temporal domain. To make efficient use of this parallelism, some processors need to be idle on a temporal coarse level of MGRIT. This talk considers an optimization of the implementation by assigning these idle processors to the redistributed spatial domain on coarse levels. It accelerates coarse-level spatial solvers and promotes early switching to parallelization-in-time by reducing its overhead. Although the spatial redistribution is contrary to the motivation of parallel-in-time approaches, it is expected to be effective when we need to balance spatial and temporal parallelism. This is because parallelization-in-time is already usually switched on when parallelization-in-space starts to saturate, so one can still benefit from additional parallelism. Numerical experiments show the effectiveness of the spatial redistribution compared with pure MGRIT, which assigns the best spatial and temporal parallelism at specific parallelism.

Ryo Yoda

University of Tokyo  
ryo.yoda@mist.i.u-tokyo.ac.jp

Matthias Bolten

University of Wuppertal  
bolten@math.uni-wuppertal.de

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

Akihiro Fujii  
Kogakuin University  
fujii@cc.kogakuin.ac.jp

#### MS57

##### **Practical, Out-of-Memory Sparse MTTKRP on Massively Parallel Architectures**

We propose a novel frame-work for massively parallel execution of fundamental tensor decomposition operations. Our new Blocked Linearized CoOrdinate (BLCO) format enables out-of-memory computation of the matricized tensor times Khatri-Rao product (MTTKRP) for every tensor dimension, using a unified implementation that works on a single tensor copy. To address the substantial synchronization cost on massively parallel architectures, we introduce a hierarchical on-the-fly algorithm to efficiently discover and resolve update conflicts across threads without keeping any auxiliary information or storing non-zero elements according to specific mode orientations. As a result, our BLCO-based MTTKRP not only handles out-of-memory tensors, but also delivers superior in-memory performance compared to mode-specific approaches. On the latest NVIDIA A100 GPU, BLCO achieves 2.9 geometric-mean speedup over the state-of-the-art mixed-mode compressed sparse fiber (MM-CSF) on a range of real-world sparse tensors.

Andy Nguyen, Jee W. Choi  
University of Oregon  
andyn@cs.uoregon.edu, jeecc@uoregon.edu

#### MS57

##### **A Distributed Memory Implementation of CP-POPT-GDGN: An All-at-Once Decomposition Algorithm for Count Tensors**

Data sets from many applications can be represented as count tensors. As such tensors grow larger and increasingly elaborate, more sophisticated algorithms are required to enable accurate CP decompositions of these tensors. Recent work has demonstrated that CP-POPT-GDGN, an allatonce algorithm for count tensor decomposition that utilizes a generalized damped Gauss-Newton method, typically achieves more accurate decompositions than earlier count tensor decomposition algorithms, including CP-APR. In this work, we consider a distributed memory implementation of CP-POPT-GDGN, and discuss its scalability. This implementation allows us to compute CP decompositions in an amount of time comparable to that required by a distributed memory implementation of CP-APR, when using the same computing resources.

Teresa Ranadive  
Laboratory for Physical Sciences

tranadive@lps.umd.edu

### MS57

#### Fully Constrained PARAFAC2 with AO-ADMM

The PARAFAC2 model is a powerful method to analyse multiway measurements with variations across one mode. The model captures latent factors and allows the factors of one mode to evolve across another, which has proven beneficial for applications in several domains, such as chemometrics, neuroscience and data mining. For example, the PARAFAC2 model can accurately describe chromatographic data despite retention shifts and shape changes of the temporal profiles across samples. A challenge with PARAFAC2, however, is the constant cross-product constraint, which is necessary to ensure a unique decomposition. To impose this constraint, the evolving factor matrices are traditionally modelled implicitly as the product of a shared coordinate matrix and a set of orthogonal basis matrices. However, this implicit handling makes it difficult to impose constraints on the evolving factor matrices. In this work, we instead reformulate the regularised PARAFAC2 problem in a form suited for alternating optimisation (AO) with the alternating direction method of multipliers (ADMM), thus making it possible to fit PARAFAC2 models with any proximable constraint on all modes. Using simulated datasets, we show that the proposed scheme is accurate, flexible in terms of constraints and computationally efficient. Finally, we use our method in a real-world chromatography application and show that constraining all modes to be non-negative improves the interpretability of the extracted components.

Marie Roald

Simula Metropolitan Center for Digital Engineering  
Oslo Metropolitan University  
mariero@simula.no

Carla Schenker

Simula Metropolitan Center for Digital Engineering  
carla@simula.no

Rasmus Bro

Copenhagen University  
rb@life.ku.dk

Jeremy Cohen

Universite de Rennes, INRIA, CNRS, IRISA  
jeremy.cohen@irisa.fr

Evrin Acar

Simula Metropolitan Center for Digital Engineering  
evrim@simula.no

### MS57

#### A New Alternating Optimization Algorithm for CP Decomposition

CP decomposition (CPD) is prevalent in chemometrics, signal processing, data mining and many more fields. While a plethora of algorithms have been proposed to compute the CPD, alternating least squares (ALS) still remains one of the most widely used algorithm for computing the decomposition. Recent works have shown that ALS suffers from slow convergence when the CP factors become collinear (also known as the swamp phenomenon). Several attempts have been made to overcome this issue through incorporating techniques like line search, pivoting in the ALS algo-

rithm. We propose a new algorithm which minimizes a different objective function via alternating optimization leading to a sequence of better conditioned iterates of the CP factors, alleviating the swamp phenomenon. Alternating optimization of this new objective leads to simple updates to the factor matrices with the same asymptotic computational cost as ALS. We show that a subsweep of this algorithm can achieve a superlinear convergence rate for CPD with known rank and verify it experimentally. We further introduce a formulation which generalizes our approach to interpolate between updates corresponding to the ALS and the new algorithm to manage the tradeoff between stability and fitness of the decomposition for approximation problems. Our experimental results show that for approximate CPD, this algorithm and its variants converge to a better conditioned decomposition with a small change in fitness as compared to the ALS.

Navjot Singh

University of Illinois at Urbana Champaign  
navjot2@illinois.edu

Edgar Solomonik

University of Illinois at Urbana-Champaign  
solomon2@illinois.edu

### MS58

#### Automatic Application Performance Data Collection with Caliper and Adiak

Most performance analysis, optimization, and modeling tasks rely on accurate performance measurements, but collecting this data is often a time-consuming process that requires specialized tools and expert users. This talk presents two software libraries, Caliper and Adiak, that simplify and automate performance data collection. Caliper embeds performance profiling capabilities directly into HPC codes, so that performance measurements can be enabled at runtime for any program run. Equally important, Adiak collects program metadata - what was the machine, software environment, and program configuration for a given run? Together, these tools enable new, automated workflows for large performance comparison studies or long-term performance regression testing. We demonstrate use cases and analysis techniques for automated performance data recording involving large collections of runs.

David Boehme

Lawrence Livermore National Laboratory  
boehme3@llnl.gov

### MS58

#### Machine Learning Enabled Scalable Performance Prediction of Scientific Codes

We present the Analytical Memory Model with Pipelines (AMMP) of the Performance Prediction Toolkit (PPT). PPT-AMMP takes high-level source code and hardware architecture parameters as input and predicts the runtime of that code on the target hardware platform. PPT-AMMP transforms the code to an (architecture-independent) intermediate representation, then (i) analyzes the basic block structure of the code, (ii) processes architecture-independent virtual memory access patterns that it uses to build memory reuse distance distribution models for each basic block, and (iii) runs detailed basic-block level simulations to determine hardware pipeline usage. PPT-AMMP uses machine learning and regression techniques to build the prediction models based on small instances of the in-

put code, then integrates into a higher-order discrete-event simulation model of PPT running on Simian PDES engine. We validate PPT-AMMP on four standard computational physics benchmarks and present a use case of hardware parameter sensitivity analysis to identify bottleneck hardware resources on different code inputs. We illustrate a further use case of PPT-AMMP by predicting the performance of the scientific application code SNAP. To this end, we analyze multi-variate regression models that accurately predict the reuse profiles and the basic block counts. We validate predicted SNAP runtimes against actual measured times.

Stephan Eidenbenz, Nandakishore Santhi  
Los Alamos National Laboratory  
eidenben@lanl.gov, nsanthi@lanl.gov

### MS58

#### Performance Modeling of Graph Processing Workloads

Graph processing workloads are increasing in popularity and complexity: many problems ranging from bioinformatics to text analysis, and from social network analysis to logistics and optimizations use graphs as primary data models. With the increase in analysis complexity and data size, graph processing performance becomes problematic: execution time increases significantly, yet unpredictably. For the past 15+ years, using massive parallelism - often in the form of GPU-like accelerators - has been a successful approach for many applications. However, parallel graph processing workloads do not always provide the expected performance benefits; moreover, the performance is heavily dependent on the graph structure, system architecture, and the algorithm itself. In this talk, we will discuss several methods for performance modeling for parallel graph processing workloads on GPUs. We will contrast and compare analytical and machine-learning models, and discuss the pro's and con's of both approaches. We will show results on multiple GPU systems, different applications, and a large number of datasets.

Ana Lucia Varbanescu  
University of Amsterdam  
a.l.varbanescu@uva.nl

Merijn Verstraaten  
University of Amsterdam  
Netherlands e-Science Center  
m.verstraaten@esciencecenter.nl

### MS59

#### Code Generation for Matrix-Free Iterative Linear Solvers

Fast matrix-free iterative linear solvers play an important role in high-performance simulations on modern computer architectures because they allow to overcome the bandwidth limitation of iterative solvers based on assembled matrices. High compute intensity at reduced computational complexity can be achieved in the context of high-order finite element discretizations by exploiting tensor product structure of the finite element basis. Robustness of iterative solvers with respect to the mesh size as well as model coefficients is achieved by a two-level preconditioner combining a local matrix-free preconditioner on the high-order part with global matrix-based preconditioners in a low-order subspace. In this talk we focus on code-generating all required components of such precon-

ditioners for high-order discontinuous Galerkin discretizations. These components are matrix-free operator application, matrix-free block SSOR smoothers as well matrix generation for the low-order subspace. For all components it is possible to generate high-performance code specifically tailored to the current kernel and hardware through SIMD vectorization and code transformations based on auto-tuning. As an example application we consider viscous fingering in miscible displacement with operator splitting techniques.

Peter Bastian  
Interdisciplinary Center for Scientific Computing  
University of Heidelberg  
peter.bastian@iwr.uni-heidelberg.de

René Heß  
IWR, Heidelberg University  
rene.hess@iwr.uni-heidelberg.de

### MS59

#### Transforming (not just) DG-FEM Array Expressions (not just) on GPUs

While domain-specific languages for finite-element-based PDE solvers have enjoyed considerable success, composing them with supporting computations (e.g. chemistry, neural networks, non-variational discretization schemes) has proved challenging. We demonstrate a fully generic array package with numpy-like semantics that is able to represent and transform the action of complex discontinuous Galerkin (DG) FEM operators for near-roofline performance while straightforwardly incorporating non-FEM computations on the degrees of freedom. Our array package (pytato) can straightforwardly be extended to transform other workloads and allows code generation into an imperative, scalar, polyhedrally-based Intermediate Representation (IR), based on our Loopy tool for loop transformations. With our implementation, we also provide a transformation pipeline tuned for a general DG-FEM operator evaluation. Based on the metadata attached to various IRs in the lowering pipeline, optimization passes such as kernel fusion, loop fusion/distribution, common sub-expression elimination and loop tiling get triggered.

Kaushik Kulkarni  
University of Illinois at Urbana-Champaign  
kgk2@illinois.edu

Andreas Kloeckner  
University of Illinois  
andreask@illinois.edu

### MS59

#### Robust Multigrid for Stokes with Extreme Viscosity Variation using Firedrake

In this talk, I will present augmented Lagrangian Schur complement preconditioners and robust multigrid methods for incompressible Stokes problems with extreme viscosity variations, with a focus on its implementation using the open source library Firedrake. Such Stokes systems arise, for instance, upon linearization of nonlinear viscous flow problems, and they can have severely inhomogeneous and anisotropic coefficients. The augmented Lagrangian formulation for the incompressibility constraint makes the Schur complement easier to approximate, but results in a nearly singular (1,1)-block in the Stokes system. To cope with the near-singularity of the (1,1)-block, we use a multigrid

scheme with discretization-dependent smoother and transfer operators. We confirm the robustness of the multigrid scheme and the overall efficiency of the solver using numerical examples with scalar and with anisotropic fourth-order tensor viscosity arising from linearization of a viscoplastic constitutive relation. With Firedrake's parallelism, we present scalability results using up to 28,672 parallel tasks for problems with up to 1.6 billion unknowns and a viscosity contrast up to ten orders of magnitude.

Yu-hsuan Shih

Courant Institute of Mathematical Science  
New York University  
shihyh@cims.nyu.edu

Georg Stadler

Courant Institute for Mathematical Sciences  
New York University  
stadler@cims.nyu.edu

Florian Wechsung

Courant Institute of Mathematical Science  
wechsung@nyu.edu

## MS59

### **An Automated Framework for Fully Matrix-Free, Hybridised, Compatible, High Order Finite Element Methods**

One way to achieve high accuracy for simulations of physically complex problems is to discretise the partial differential equations with higher order, compatible finite element methods (FEM). Since higher order, compatible approximations result in big equation systems, performance optimisations become crucial. A speedup of the solve of the corresponding equation system can be achieved by loosening the global coupling of the FEM with a hybridization preconditioner, such that it is sufficient to execute expensive operations on smaller matrices. The preconditioner is defined through local linear algebra operations on FEM tensors and is represented in the domain specific language Slate in Firedrake. The local linear algebra operations need further optimisations in order to achieve high performance for high order, compatible FEM, because not only the size of the global tensors, but also of the local tensors is considerably high. By employing locally matrix-free methods high storage requirements/data movement can be avoided in favor of more FLOPS. A high FLOPS to data ratio is advantageous for high performance on recent computer architectures. I will talk about performance optimisations in form of fully matrix-free, hybridised, compatible, high order FEM in Firedrake and their portability between different problems and across architectures through automatic code generation.

Sophia Vorderwuelbecke

Imperial College London  
s.vorderwuelbecke18@imperial.ac.uk

David Ham

Department of Mathematics  
Imperial College London  
david.ham@imperial.ac.uk

Colin J. Cotter

Imperial College London  
Department of Mathematics

colin.cotter@imperial.ac.uk

## MS60

### **Partition Improvement for High-Order Unstructured Mesh and Particle-In-Cell Simulations**

One of the most common techniques for partitioning unstructured meshes is to construct a graph abstraction of the mesh and perform a graph partitioning method. When these partitioning methods do not create a sufficiently good partition of the mesh, multicriteria diffusive partition improvement methods can be employed to further improve performance in the target application. In this presentation, we explore alternative graph constructions of the unstructured mesh to specifically target the application requirements allowing better diffusive partition improvements. Diffusive load balancing is provided by the EnGPar library that employs a multihypergraph to represent the unstructured mesh. We present results of EnGPar partitioning the unstructured mesh for high-order finite element simulations and dynamic load balancing of particles in distributed unstructured mesh particle-in-cell simulations.

Gerrett Diamond

Rensselaer Polytechnic Institute  
diamog2@rpi.edu

Cameron W. Smith

Scientific Computation Research Center  
Rensselaer Polytechnic Institute  
smithc11@rpi.edu

Mark S. Shephard

Rensselaer Polytechnic Institute  
Scientific Computation Research Center  
shephard@rpi.edu

## MS60

### **An MPI-Based Algorithm for Mapping Complex Networks onto Hierarchical Architectures**

Processing massive application graphs on distributed memory systems requires to map the graphs onto the system's processing elements (PEs). This task becomes all the more important when PEs have non-uniform communication costs or the input is highly irregular. Typically, mapping is addressed using partitioning, in a two-step approach or an integrated one. Parallel partitioning tools do exist; yet, corresponding mapping algorithms or their public implementations all have major sequential parts or other severe scaling limitations. We propose a parallel algorithm that maps graphs onto the PEs of a hierarchical system. Our solution integrates partitioning and mapping; it models the system hierarchy as an implicit labeled tree. The vertices of the application graph are labeled as well, and these vertex labels induce the mapping. The mapping optimization follows the basic idea of parallel label propagation, but we tailor the gain computations of label changes to quickly account for the induced communication costs. Our MPI-based code extends the partitioning library ParHIP. To evaluate our algorithm's implementation, we perform comparative experiments with complex networks in the million- and billion-scale range. In general our mapping tool shows good scalability on up to a few thousand PEs. Compared to other MPI-based competitors, our algorithm achieves the best speed to quality trade-off and our quality results are even better than non-parallel mapping

tools.

Maria Predari, Charilaos Tzova  
Humboldt-Universität zu Berlin  
maria.predari@informatik.hu-berlin.de,  
charilaos.tzovas@informatik.hu-berlin.de

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

Henning Meyerhenke  
Humboldt-Universität zu Berlin  
meyerhenke@hu-berlin.de

## MS60

### Acyclic Partitioning for Mapping and Scheduling

With the disproportional improvements of computational and communication efficiency in computing architectures, data movement remains as the most important challenge in efficient execution. A good trade-off between load balance and data locality is necessary when scheduling computations. Directed Acyclic Graphs (DAGs) have been the de-facto abstraction to model computational tasks and their dependencies. Acyclicity provides many benefits on the ordering and mapping of vertices of the DAG (i.e., tasks) to computational resources. Thus, it is desirable to preserve acyclicity at different granularities of computation. In this talk, we show how acyclic partitioning of DAGs, i.e. partitioning in which the inter-part edges between the vertices from different parts should preserve an acyclic dependency structure among the parts, can be explored to minimize redundant data movement in two different computational models: on a single processor with two-level memory setting and on distributed memory setting. We demonstrate the usefulness of our acyclic partitioning using a high-performance distributed quantum circuit simulation application. To address the limitations of the graph model, we also develop and analyze the performance of acyclic partitioning approaches for directed hypergraphs and discuss and evaluate the scenarios where each can be preferable.

Yusuf Ozkaya  
mit Catalyurek  
my.ozkaya@gmail.com

mit Catalyurek  
Georgia Institute of Technology  
umit@gatech.edu

## MS60

### Lightning Fast Streaming Graph Partitioning for Modern Architectures

Partitioning a graph into balanced blocks such that few edges run between blocks is a key problem for large-scale distributed processing. If the topology of a distributed system is known, it is possible to further optimize the communication costs by mapping blocks onto processing elements. A current trend for partitioning huge graphs are streaming algorithms, which are very fast and have a small memory footprint. These algorithms receive nodes, one at a time, and their neighborhood and directly make a block assignment decision which can not be undone. Up to now, no streaming algorithm maps nodes to processing elements of distributed systems taking into account the different communication speeds of distinct communication links. In this talk, we present a number of improvements to streaming

graph partitioning algorithms. We present a streaming multi-recursive partitioning scheme that performs recursive multi-sections on the fly without knowing the overall input graph. This enables us to adopt the scheme to be able to take different communication speeds of a hierarchical compute system into account. Our approach has a considerably lower running time complexity in comparison with state-of-the-art one-pass partitioning algorithms designed for graph partitioning. Our algorithms are up to two orders of magnitude faster than competing algorithms when partitioning inputs into large amounts of blocks while also computing better partitions w.r.t. a process mapping objective.

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

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

## MS61

### Fast High Order GPU Simulation Support for Exascale Applications

MFEM is an open-source, lightweight, flexible and scalable C++ high-order finite element methods library targeting portability, scalability and high-performance computing. This talk will present the GPU research and development activities of the MFEM library within the U.S. Department of Energy Exascale Computing Project co-design Center for Efficient Exascale Discretization. We will discuss the GPU developments strategies in several components and our work on performance optimizations for large-scale architectures, targeting high-order finite-element algorithms for high-order applications on both NVIDIA and AMD GPU-accelerated platforms. We will report performance and capability improvements in several benchmarked proxies, miniapps and for CEED's applications on both systems.

Jean-Sylvain Camier  
Lawrence Livermore National Laboratory  
camier1@llnl.gov

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

Veselin Dobrev  
Lawrence Livermore National Laboratory  
dobrev1@llnl.gov

Yohann Dudouit  
LLNL  
dudouit1@llnl.gov

## MS61

### Toward Performance-Portable PETSc for GPU-Based Exascale Systems

The Portable Extensible Toolkit for Scientific computation (PETSc) library provides scalable solvers for nonlinear time-dependent differential and algebraic equations and for numerical optimization. The PETSc design for performance portability addresses fundamental GPU accelerator

challenges and stresses flexibility and extensibility by separating the programming model used by the application from that used by the library, enabling application developers to use their preferred programming model, such as Kokkos, RAJA, SYCL, HIP, CUDA, or OpenCL, on upcoming exascale systems. We will describe some of the major challenges GPUs pose for scientific software libraries, discuss design decisions we have made to address these, provide a blueprint for using GPUs from PETSc-based codes, and present case studies emphasizing the flexibility and high performance achieved on current GPU-based systems. Additionally, we will discuss recent developments in PETSc's communication module, PetscSF, that enable flexibility and scalable performance on GPU-based systems.

Richard T. Mills

Argonne National Laboratory  
rtmills@anl.gov

Mark Adams

Lawrence Berkeley National Laboratory  
mfadams@lbl.gov

Satish Balay

Argonne National Laboratory  
balay@anl.gov

Jed Brown

University of Colorado Boulder  
jed@jedbrown.org

Alp Dener

Argonne National Laboratory  
Mathematics and Computer Science Division  
adener@anl.gov

Matthew G. Knepley

University at Buffalo  
Department of Computer Science and Engineering  
knepley@gmail.com

Scott Kruger

Tech-X Corporation  
kruger@txcorp.com

Hannah Morgan

University of Chicago  
hmmorgan@uchicago.edu

Todd Munson

Argonne National Laboratory  
Mathematics and Computer Science Division  
tmunson@mcs.anl.gov

Karl Rupp

TU Wien  
me@karlrupp.net

Barry Smith

Center for Computational Mathematics  
Flatiron Institute  
bsmith@petsc.dev

Stefano Zampini

King Abdullah University of Science and Technology  
stefano.zampini@kaust.edu.sa

Hong Zhang, Junchao Zhang

Argonne National Laboratory  
hongzhang@anl.gov, jc Zhang@anl.gov

**MS61**

**Automatic Discovery of Implementation Rules for Fast GPU + MPI Operations**

Developing a high-performance implementation of a distributed computational kernel for high-performing computing is increasingly challenging. Systems are composed of heterogeneous computational resources, and limited communication performance demands an asynchronous application design. Even if high-performance computation and communication libraries are available, the challenge becomes the best coordination of the provided operations to create an optimal result. This work presents a system that automatically generates design rules for a high-performance implementation of a compound operation provided as a dependence graph. The system searches among valid schedules to determine the fastest arrangement of operations. A post-processing step on the results of the search yields interpretable design rules. The fast implementation can be used directly, or experts can use the design rules to create a high-performance implementation.

Carl Pearson

Sandia National Laboratory  
cwpears@sandia.gov

**MS61**

**Accelerating a Sparse Eigensolver using OpenACC**

We describe the effort to accelerate a sparse eigensolver for nuclear many-body calculations on multiple GPUs. Our work involves modifying a previously developed hybrid MPI/OpenMP implementation of an eigensolver written in FORTRAN90 with OpenACC directives. The architectural differences between a CPU and a GPU device results in some differences between the way OpenACC directives are inserted and the way OpenMP directives were used. We point out these differences and demonstrate significant speedup achieved on GPUs compared to the on-node performance of a many-core CPU. We also show that the overall performance improvement of the eigensolver on multiple GPUs is more modest due to the communication overhead among different MPI ranks.

Chao Yang

Lawrence Berkeley National Lab  
cyang@lbl.gov

Pieter Maris

Iowa State University  
pmaris@iastate.edu

Brandon Cook

Lawrence Berkeley National Laboratory  
bgcook@lbl.gov

Dossay Oryspayev

Brookhaven National Laboratory  
doryspaye@bnl.gov

**MS62**

**An Inertia-Free Approach for Nonlinear Optimiza-**

**tion on GPU**

We present a filter line-search interior-point solver that does not require inertia information of the KKT system. The proposed approach performs curvature tests along the search step to detect negative curvature and correct it to guarantee global convergence. This feature enables the use of a wide range of linear algebra strategies and libraries, which is essential to tackle large-scale problems on modern computing architectures. Our numerical tests demonstrate the inertia-free approach on multiple CPUs (MPI) and on GPU.

Nai-Yuan Chiang  
Lawrence Livermore National Laboratory  
7000 East Avenue, Livermore, CA 94550  
chiang7@llnl.gov

Cosmin G. Petra  
Lawrence Livermore National  
Center for Advanced Scientific Computing  
petra1@llnl.gov

**MS62****A Hybrid Direct-Iterative Method for Solving Kkt Linear Systems**

We propose a method to solve sparse indefinite linear systems that arise at each iteration of an interior method for solving constrained optimization problems. The current gold standard for solving these systems is  $LDL^T$  factorization. However,  $LDL^T$  requires pivoting during factorization, which substantially increases communication cost. Our novel method solves a large indefinite system by solving multiple smaller positive definite systems, using an iterative solve for the Schur complement and an inner direct solve (via Cholesky factorization) within each iteration. Cholesky is stable without pivoting, thereby reducing communication and allowing reuse of the symbolic factorization. We show that on large systems, our method can efficiently utilize GPUs and outperform  $LDL^T$  factorization of the full system.

Shaked Regev  
Stanford University  
781 Escondido Road, Apt 416, Stanford, CA (USA) 94305  
sregev@stanford.edu

**MS63****Computational Optimization of Industrial High-Performance Data Analytics Applications**

High-performance data analytics applications are becoming increasingly common in industrial applications, but their deployment in a production setting on cluster-scale computers is often bottlenecked by shared memory architectures which are not optimized for the quantity of data throughput that such applications require. By analyzing an application of high-throughput electricity market simulation used by an international commodities trading company, we will show how memory bottlenecks can adversely affect the performance of high-throughput applications. To mitigate these costly bottlenecks, we will present a framework using LIKWID Performance Tools to analyze the computational nature of such applications, and strategies for achieving optimal throughput using the meta-scheduler GREASY to adjust the level of parallelism and introduce CPU affinity to ease memory congestion and improve cache

locality.

Timothy Holt  
Università della Svizzera italiana (USI)  
timothy.holt@usi.ch

Juraj Kardoš  
Università della Svizzera italiana  
Institute of Computational Science  
juraj.kardos@usi.ch

Olaf Schenk  
Università della Svizzera italiana  
olaf.schenk@usi.ch

**MS63****Yin and Yang: Balancing Cloud and HTC Workloads**

The NSF-funded Chameleon project ([www.chameleoncloud.org](http://www.chameleoncloud.org)) provides an OpenStack-based bare metal reconfigurable cloud that supports systems research, education, and emergent application projects. The facility has been in existence for 6+ years, and has served 6,000 users working primarily on systems experimentation. These types of experiments require interactive access; demand tends align with the academic calendar and conference deadlines which means that during certain periods the infrastructure is underutilized. High Throughput Computing (HTC) jobs consist of independent tasks that do not require interactivity and are resilient to resource loss. Cloud and HTC workloads thus have complementary characteristics that could be leveraged and combined to ensure the best possible utilization of resources and thus the lowest possible price for interactive usage. This talk will describe the mechanisms developed by the Chameleon project to enable Open Science Grid (OSG) workloads to run on the system; these workloads are invisible to Chameleon users, do not impact their usage of the system, and utilize only the resources that are not used for experimentation. I will present the architecture and implementation of the system, the methods and policies we use to ensure efficient usage, and share preliminary results from an actual deployment.

Kate Keahey  
Argonne National Laboratory  
keahey@anl.gov

Zhuo Zhen  
University of Chicago  
zhenz@uchicago.edu

**MS63****The Open Science Pool An OSG DHTC Service for the S&E Community**

Researchers who harness the capacity of the OSG operated Open Science Pool (OSPool) are likely to have access to accounts on campus clusters, allocations at national facilities, or CloudBank accounts. When they place their HTC workload on an Access Point provided by the Partnership to Advance Throughput Computing (PATH) project, they would like to steer the execution of some or all the jobs in the workload to these resources. Doing so improves the throughput of the workload and can address the needs of jobs with special resource requirements like large memory or long execution times. PATH has been developing tech-

nologies and services that allow a researcher to Bring Your Own Resource (BYOR) to the Access Point and to direct the matching of jobs to resources. We will present the capabilities available today and outline our future plans to support BYOR for HTC workloads. Provisioning and scheduling of BYORs is a major challenge as mistakes can be extremely costly to the researcher.

Miron Livny  
UW-Madison/CHTC  
miron@cs.wisc.edu

### MS63

#### Managing High-Throughput Application Workloads: Findings and Recommendations

High-throughput applications, in which application workload consists of a large ensemble of self-contained tasks and application performance is measured by the number of tasks completed per unit of time, are vital for scientific discovery. Managing high-throughput application workloads is not an easy task: it requires automation services that include workflow capabilities for managing all the tasks and all the input/output files across different runs of a high-throughput applications ensemble. These increasing complexities hinder the ability of scientists to generate robust science, which we define as the capacity of high-throughput applications to scale in performance, to exhibit trustworthiness, and to assure reproducibility. This talk presents recommendations for designing and implementing software systems for robust science across critical high-throughput applications. The findings were collected through a virtual mini-workshop in May 2021 called a Virtual World Cafe (VWC). In the VWC, we engaged communities of software developers to share state of the art recommendations through structured conversational processes: participants were distributed across several breakout sessions in an online meeting, with participants switching sessions periodically and getting introduced to the previous discussion at their new session by a session lead.

Michela Tauffer  
The University of Tennessee  
tauffer@utk.edu

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

### MS64

#### Multi-Scale Methods for Machine Learning

Many modern applications require us to very quickly find relevant results from an enormous output space of potential candidates, for example, finding the best matching product from a large catalog or suggesting related search phrases on a search engine. The size of the output space for these problems can be in the millions to billions. Moreover, observational or training data is often limited for many of the so-called long-tail of items in the output space. Given the inherent paucity of training data for most of the items in the output space, developing machine learning models that perform well for spaces of this size is a contemporary challenge. In this talk, I will present a multi-scale machine learning framework called Prediction for Enormous and Correlated Output Spaces (PECOS). PECOS proceeds by first building a hierarchy over the output space using unsupervised learning, and then learning a machine learning

model that makes predictions at each level of the hierarchy. Finally, the multi-scale predictions are combined to obtain an overall predictive model. This leads to an inference method that scales logarithmically with the size of the output space. A key to obtaining high performance is parallel computation, and leveraging sparsity to develop highly efficient sparse matrix routines.

Inderjit S. Dhillon  
University of Texas at Austin  
inderjit@cs.utexas.edu

### MS64

#### Sparse Precision Matrix Estimation for Large-Scale Datasets

The L1-regularized Gaussian maximum likelihood method is a common approach for sparse precision matrix estimation but one that poses a computational challenge for high-dimensional datasets. We present a novel method for performing large-scale sparse precision matrix estimation utilizing the block structures and parallelism in the underlying computations for added performance and scalability. Our numerical examples and comparative results with various modern open-source packages reveal that the proposed algorithm provides orders of magnitude faster runtimes and scales to significantly higher dimensional datasets.

Aryan Eftekhari  
Università della Svizzera italiana  
aryan.eftekhari@gmail.com

Matthias Bollhöfer  
TU Braunschweig  
m.bollhoefer@tu-bs.de

Olaf Schenk  
Università della Svizzera italiana  
olaf.schenk@usi.ch

### MS64

#### Handling Runtime Variability at Scale: Elasticity for Continuously Changing Graphs

Today's massive data are increasingly represented as graphs. Trillions of edges are common resulting in a growing importance for distributed graph systems. Most large graphs do not appear overnight and remain static. Instead, they are built up over time from continuous processes, e.g., users visiting web pages, product purchases, vehicles traveling on roads, and more. Such changing graphs are called dynamic and their algorithms operate on batches of changes instead of re-computing from scratch. An understudied but crucial property of dynamic graph systems is their high runtime variability, arising from three main areas: the graph change rate, e.g., web traffic increases during shopping holidays; the per-batch processing time, which can vary from constant to the full cost of re-computation; and the user query rate, e.g., hotspots on popular queries. High variability limits dynamic graph use for interactive applications. ELGA is a recent high-performance elastic, dynamic graph system. In this talk, we extend ELGA to address per-batch variability. We develop heuristics to predict batch runtime and then preemptively, automatically, scale out. We reactively adjust after observing iteration runtimes. This reduces runtime variability, enabling previously out-of-reach interactive applications on massive, dynamic graphs. SNL is managed and operated by NTESS under DOE NNSA contract DE-

NA0003525.

Kasimir Gabert

Georgia Institute of Technology  
kasimir@gatech.edu

Ali Pinar

Sandia National Labs  
apinar@sandia.gov

Umit V. Catalyurek

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

## MS64

### Should You Go with the Flow ? A New Tensor Algebra for Neural Networks

Big, high-dimensional information often involves multi-dimensional correlations that may remain latent by virtue of traditional matrix-based learning algorithms. In this study, we propose a tensor neural network framework that offers an exciting new paradigm for supervised machine learning. The tensor neural network framework is based upon the notion of tensor-tensor product, an algebraic formalism to multiply tensors, which inherits mimetic matrix properties. We further extend the framework to accommodate dynamical graphs in the form of tensor Graph Convolutional Neural Networks. We show that the proposed tensor neural network formalism is a natural high-dimensional extension of conventional neural networks. Lastly, we close with numerical experiments demonstrating the intrinsic advantages of the proposed tensorial architectures.

Lior Horesh

IBM Research  
lhoresh@us.ibm.com

Misha E. Kilmer

Mathematics Department  
Tufts University  
misha.kilmer@tufts.edu

Elizabeth Newman

Emory University  
elizabeth.newman@emory.edu

Shashanka Ubaru

IBM T. J. Watson Research Center  
shashanka.ubaru@ibm.com

Osman Asif Malik

University of Colorado, Boulder  
osman.malik@colorado.edu

Haim Avron

Tel Aviv University  
haimav@tauex.tau.ac.il

## MS65

### Dynamic and Intelligent Workflows with Eflows4hpc

Distributed computing infrastructures are evolving from traditional models to environments that involve sensors, edge devices, instruments, etc, and, as well, high-end computing systems such as clouds and HPC clusters. A key

aspect is how to describe and develop the applications to be executed in such platforms. What is more, the use of data analytics and artificial intelligence in general is on high demand in current HPC applications. However, the methodologies to develop such workflows that integrate HPC simulations and data analytics is not well integrated. eFlows4HPC project has recently started with the goal of providing workflow software stack and an additional set of services to enable the integration of HPC simulations and modelling with big data analytics and machine learning in scientific and industrial applications. The project will demonstrate its advances through three application Pillars with high industrial and social relevance: manufacturing, climate and urgent computing for natural hazards; these applications will help to prove how the realization of forthcoming efficient HPC and data-centric applications can be developed with new workflow technologies. The talk will present the motivation, challenges and project workplan.

Rosa M. Badia

Barcelona Supercomputing Center / CSIC  
rosa.m.badia@bsc.es

## MS65

### Investigations on the use of Hashing for Parallel Graph and Hypergraph Processing

Graphs and hypergraphs are widely used as models in many data science and scientific computing applications. This is so as these two discrete mathematical objects correspond naturally to sparse matrices and tensors. We investigate the use of hashing for implementing operations on graphs and hypergraphs in shared memory systems, with an eye towards linear and multi-linear algebraic operations on matrices and tensors.

Somesh Singh

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

Bora Ucar

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

## MS65

### Sparse Matrix Factorization by the Accelerator-First Principle

Within the last decade, heterogeneous computing has become mainstream. These days, fueled by the ML/AI revolution, its most common embodiment is the accelerated computing system, i.e. a system combining (a) CPU(s) with one or more compute accelerators. Most commonly, that is a GPU due to their massive FLOP count and extensive software stack. However, the number of alternative accelerators is still growing - leading to a classical chicken-egg problem: why would I buy a new compute accelerator if the software support is still scarce? In this talk, we consider sparse matrix factorization, a crucial building block of scientific applications. We analyze state-of-the-art implementations using accelerators and point out their issues. By re-thinking the organization of modern multifrontal codes with n 'accelerator-first' mindset, we design METAPACK, a system that automatically synthesizes sparse matrix factorization packages for various compute accelerator architectures. We discuss the components at play, from frontend to a novel compiler pass, and how systems research have affected the design. Lastly, we use first computational results to discuss the feasibility of the approach and future

of software synthesis in the HPC landscape.

Daniel Thürck  
NEC Laboratories Europe GmbH  
daniel.thuerck@neclab.eu

### MS65

#### Beyond a Hybrid Parallel GraphBLAS

GraphBLAS is a recent standard for writing graph algorithms in the language of linear algebra. While it has an accepted C API, we previously presented an alternative specification in line with standard C++ making use of its more recent features. Additionally, we added the notion of performance semantics and backends. An auto-vectorising and auto-parallelising implementation runs on both shared- and distributed-memory architectures, on Intel and ARM. In this talk we refine the notion of performance semantics, present a mechanism that ensures interoperability with existing parallel frameworks such as MPI or Spark, and explore how GraphBLAS relates to existing programming models that have been highly successful within data science, such as MapReduce or vertex-centric programming. In particular, we shall demonstrate how with minor extensions that remain truthful to the idea of programming using algebraic concepts, a thus-extended GraphBLAS in fact generalizes some of these earlier well-known programming models. Whether such an ‘algebraic programming model’ could form a basis of workloads beyond just graph computing depends also on its performance, which we shall present last.

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

### MS66

#### Computing with Biologically Inspired Dynamics

Some neuromorphic devices are most aptly described as neuronal emulators rather than simulators, as they replicate the electrodynamics of cortical cells rather than numerically calculating solutions to the associated differential equations. Coupled with the discrete, sparse nature of spike-based communication, the advantages of such systems become apparent in their emulation speed and power consumption. Being designed with a high degree of configurability, many of these devices support a wide range of applications. Here, we focus on two of these. First, we discuss the question of credit assignment for precise spike timing and show how synaptic plasticity can implement exact gradient descent in spiking networks. When applied to, for example, hierarchical neuronal networks, this enables exceptionally fast and efficient neuromorphic solutions to feedforward inference problems such as pattern recognition. Second, we consider the interpretation of neuronal spikes as samples from some underlying distribution. With appropriate neuronal parametrization, Hebbian synaptic plasticity can be used to effectively learn arbitrary shapes of this distribution. We show how a combination of cortical dynamics including spikes, short-term synaptic plasticity and background oscillations enable fast sampling from deep, multimodal distributions, with applications ranging from generative modeling to the representation of entangled quantum systems.

Mihai A. Petrovici  
Universitat Heidelberg  
mihai.petrovici@unibe.ch

Andreas Baumbach  
Heidelberg University  
andreas.baumbach@kip.uni-heidelberg.de

### MS66

#### Developing Spike-Based Graph Algorithms for Neuromorphic Processors

Applications for neuromorphic hardware recast computational tasks in terms of discrete, temporal pulses (or ‘spike’), transmitted between multi-level dynamical systems (neurons) via weighted connections (synapses). The interplay of neuronal and synaptic parameters, combined with external stimuli allow for complex spiking dynamics. We leverage the structural similarities between a weighted synaptic network, and general network models to design and implement graph-based applications. In this talk we will present a survey of several applications that have been developed to run on near-term neuromorphic processors. We focus on analysis of undirected graphs (community detection, centrality measure evaluation) and dynamical processes on undirected graphs (epidemic models and percolation modeling).

Kathleen Hamilton, Bill Kay  
Oak Ridge National Laboratory  
hamiltonke@ornl.gov, kaybw@ornl.gov

Tiffany Mintz  
AMD  
drtmmintz@gmail.com

Prasanna Date, Catherine Schuman, Travis Humble  
Oak Ridge National Laboratory  
datepa@ornl.gov, schumancd@ornl.gov, hum-  
blets@ornl.gov

### MS66

#### Hand Tuned Spiking Networks for Composition

This work focuses on composing small spiking neural network to create larger neural networks. The motivation is to ease the burden on training, by either hand-tuning networks or training small networks with machine learning and then composing them for larger networks. We present results in both categories.

James S. Plank, ChaoHui Zheng, Charles Rizzo, John Carmack  
University of Tennessee  
jplank@utk.edu, czheng4@utk.edu, crizzo@utk.edu, jcar-  
mac3@utk.edu

### MS66

#### Assessing a Neuromorphic Platform for Scientific Random Sampling

Recent advances in neuromorphic algorithm development have shown that neural inspired architectures can efficiently solve scientific computing problems including graph decision problems and partial-integro differential equations (PIDEs). The latter requires the generation of a large number of samples from a stochastic process. While the Monte Carlo approximation of the solution of the PIDEs converges with an increasing number of sampled neuromorphic trajectories, the fidelity of samples from a given stochastic process using neuromorphic hardware requires verification. Such an exercise increases our trust in this

emerging hardware and works toward unlocking its energy and scaling efficiency for scientific purposes such as synthetic data generation and stochastic simulation. In this talk, we focus our verification efforts on a one-dimensional Ornstein-Uhlenbeck stochastic differential equation. We sample trajectories of the stochastic process across a variety of parameters on an Intel 8-Loihi chip Nahuku neuro-morphic platform. Using relative entropy as a verification measure, we demonstrate that the random samples generated on Loihi are, in an average sense, acceptable.

Darby Smith, William Severa  
Sandia National Laboratories  
jsmit16@sandia.gov, wmsever@sandia.gov

## MS67

### Designing Task-Based Libraries in Legion

Scalable software development depends on having good facilities for developing reusable libraries. In the case of parallel, and particularly distributed, programming there is a question of how to write libraries that can deal with unknown distributions of application data, both on input and output. This talk will discuss some motivating examples of the problem and solutions developed in the context of the Legion distributed programming system.

Alex Aiken  
Stanford University  
aaiken@stanford.edu

## MS67

### The Chunks and Tasks Programming Model

We present Chunks and Tasks, a programming model that makes it easier to parallelize dynamic hierarchical algorithms on computer clusters with distributed memory. Contrary to conventional programming models there are no explicit communication calls in user code and the library takes care of the distribution of both work and data. We showcase what can be achieved with the model by looking at the example of parallelization of sparse matrix operations such as matrix-matrix multiplication and inverse factorization. We present recent changes to the Chunks and Tasks C++ interface that simplifies the handling of hierarchical data for the user, by moving responsibilities to the runtime library.

Emanuel H. Rubensson  
Uppsala University  
emanuel.rubensson@it.uu.se

Elias Rudberg  
Division of Scientific Computing, Department of I.T.  
Uppsala University  
elias.rudberg@it.uu.se

Anastasia Kruchinina  
Uppsala universitet  
anastasia.kruchinina@it.uu.se

Anton G. Artemov  
Uppsala University  
anton.artemov@it.uu.se

## MS67

### Interactions Between Task-Based Runtime Sys-

### tems and the Communication Layer

With the increasing complexity of heterogeneous supercomputers, task-based runtime systems have emerged to more easily take benefit of the full computation power of these machines. For distributed applications, these runtime systems can handle communications between nodes in an implicit way by automatically detecting them. This presentation will focus on the possible interactions between the task-based runtime system and the communication layer (usually a third-party library dedicated to high performance communications). Since the runtime system and the communication library have both informations about their internal states, there are several opportunities to make them more collaborate. I will present our last research results, covering both positive and negative interactions: for instance how broadcasts can be done efficiently or the possible interferences when executing in parallel computations and communications.

Philippe Swartvagher  
Inria Bordeaux Sud-Ouest  
philippe.swartvagher@inria.fr

## MS68

### Multi-GPU Implementation of Sparse Triangular Solve and Its Performance Analysis

Solving large-scale sparse triangular linear systems is a critical computation component for many ECP and SciDAC application codes. However due to low-arithmetic intensity and sequential nature of sparse triangular solve (SpTRSV), optimizing its performance on GPU-based HPC machines is a tremendously difficult problem. This talk presents recent efforts for improving multi-GPU sparse solves in the SuperLU\_DIST software. First, we improve the data structure, parallelization strategy and memory transfer for single-GPU SpTRSV on both NVIDIA and AMD GPUs. Second, we developed a NVSHMEM/ROCSHMEM-based multi-GPU SpTRSV code to minimize communication costs. Third, we developed a critical-path-based performance model to predict and analyze the multi-GPU SpTRSV performance.

Yang Liu  
Lawrence Berkeley National Laboratory  
liuyangzhuan@lbl.gov

Nan Ding  
lawrence berkeley national laboratory  
nanding@lbl.gov

Samuel Williams  
Lawrence Berkeley National Laboratory  
swwilliams@lbl.gov

Xiaoye S. Li  
Computational Research Division  
Lawrence Berkeley National Laboratory  
xsli@lbl.gov

## MS68

### Experiences with Quasi-Newton Methods for Numerical Optimization on GPUs

In this talk, I will discuss our experiences with using the Toolkit for Advanced Optimization to solve unconstrained and bound constrained numerical optimizations problems using GPUs. We focus on first order methods that only

require objective and gradient evaluations and in particular the conjugate gradient and quasi-Newton methods.

Alp Dener, Todd Munson  
Argonne National Laboratory  
Mathematics and Computer Science Division  
adener@anl.gov, tmunson@mcs.anl.gov

#### MS68

##### Scientific Data Refactoring and Compression on the GPU with MGARD

Rapid growth in scientific data and a widening gap between computational speed and I/O bandwidth make it increasingly infeasible to store and share all data produced by scientific simulations. Hierarchical data representations hold promise as a solution to this problem, allowing for flexible conversion between different fidelities so that, for example, data can be created at high fidelity and then transferred or stored at lower fidelity via logically simple and mathematically sound operations. The effective use of these representations has been hindered by the relatively high costs of creating, accessing, reducing, and otherwise operating on them. In this talk, we describe MGARD, a multi-level data reduction framework, and present a collection of highly optimized data refactoring GPU kernels that enable efficient creation and manipulation of data in MGARD's hierarchical representation. We demonstrate that our optimized design can achieve up to 250 TB/s aggregated data refactoring throughput—83% of theoretical peak—on 1024 nodes of the Summit supercomputer.

Jieyang Chen, Ben Whitney  
Oak Ridge National Laboratory  
chenj3@ornl.gov, whitneybe@ornl.gov

#### MS68

##### Accelerating Hypre with GPUs: Strategies and Performance

With the increasing inclusion of accelerators into current and future high-performance computers, it has become important to enable parallel mathematical libraries to take advantage of the increased performance potential of GPUs. The hypre software library provides high performance preconditioners and solvers for the solution of large sparse linear systems on massively parallel computers with focus on algebraic multigrid (AMG) methods. However, their implementation can be complex. This talk will discuss strategies in achieving portability, new developments in hypre and performance results on GPUs.

Ulrike Meier Yang  
Lawrence Livermore National Laboratory  
yang11@llnl.gov

#### MS69

##### Hippylib-Muq: Scalable Markov Chain Monte Carlo Sampling Methods for Large-Scale Bayesian Inverse Problems Governed by PDEs

Bayesian inference provides a systematic framework to quantify and reduce uncertainties in inverse problems governed by partial differential equations (PDEs), but its solution for large-scale inverse problems is computationally intractable. We present in this paper an extensible and scalable software framework to overcome the prohibitive nature of Bayesian inverse problems, providing unprecedented access to state-of-the-art algorithms for large-scale Bayesian

inverse problems governed by PDEs forward models. The software integrates two complementary open-source software packages, hIPPYlib and MUQ. hIPPYlib implements powerful large-scale gradient/Hessian-based inverse solvers in an environment that can automatically generate needed derivatives, but it lacks full Bayesian capabilities. MUQ provides a spectrum of powerful Bayesian inversion models and algorithms, but expects forward models to come equipped with gradients/Hessians to permit large-scale solution. By combining these two complementary libraries, we create a robust, scalable, and efficient software framework that realizes the benefits of each to tackle complex large-scale Bayesian inverse problems across a broad spectrum of scientific and engineering disciplines.

Ki-Tae Kim  
University of California Merced  
kkim107@ucmerced.edu

Noemi Petra  
University of California, Merced  
npetra@ucmerced.edu

Umberto Villa  
Electrical and Systems Engineering  
Washington University in St Louis  
uvilla@wustl.edu

Matthew Parno  
Dartmouth College  
matthew.d.parno@dartmouth.edu

Youssef M. Marzouk  
Massachusetts Institute of Technology  
ymarz@mit.edu

Omar Ghattas  
The University of Texas at Austin  
omar@oden.utexas.edu

#### MS69

##### Bridging the Gap: Complex Models and Advanced Uncertainty Quantification

Mathematical models of complex real-world phenomena result in computational challenges, necessitating advanced software and High Performance Computing systems. Uncertainty quantification (UQ) on such models is even more challenging, since uncertainties essentially increase dimensionality of the problem at hand. We propose a new approach to coupling arbitrary model codes with advanced uncertainty quantification packages, making it easy to apply advanced UQ methods on challenging models. Further gains are portability of models, separation of concerns between model experts and UQ developers, and the availability of standardized UQ benchmarks. Finally, we show a demonstration of the approach by applying a massively parallelized Multilevel Markov Chain Monte Carlo method on a shallow water PDE modelling the 2011 Tohoku tsunami, determining its source from buoy data.

Linus Seelinger  
Heidelberg University  
Interdisciplinary Center for Scientific Computing (IWR)  
mail@linusseelinger.de

Anne Reinartz  
Durham University  
anne.k.reinartz@durham.ac.uk

Tim J. Dodwell  
University of Exeter  
College of Engineering, Mathematics & Physical Sciences  
t.dodwell@exeter.ac.uk

### MS69

#### Customizable Physics Kernels in the Glacier Flow Modeling Package Icepack

In this talk, I'll describe some of the features of the glacier flow modeling package *icepack*. Icepack is built on the finite element modeling library Firedrake, which uses the Unified Form Language to specify variational forms of partial differential equations and implements a compiler from this specification to efficient low-level C code. Firedrake assumes that end users have a concrete idea in mind of what problems they wish to solve, but for domain scientists, the real service that a software package like icepack can provide is the flexibility to try many different problem formulations. For example, two key problems at the forefront of glaciology are (1) how fracture and damage to the ice affects large-scale flow and (2) how the state of the hydrological system beneath a glacier affects sliding, both of which require adding new physics to the system and including more physical fields in the problem formulation. In this talk I'll describe how physics models are specified in icepack and how end users can alter components of these physics models, even to the point of adding new physical fields.

Daniel Shapero  
Polar Science Center  
Applied Physics Lab  
shapero@uw.edu

### MS69

#### Prototyping Parallel-in-Time Integration Methods with pySDC

The efficient use of modern supercomputers has become one of the key challenges in computational science. For the numerical solution of time-dependent processes, time-parallel methods have opened new ways to overcome both strong and weak scaling limits. With the "parallel full approximation scheme in space and time" (PFASST), multiple time-steps can be integrated simultaneously. Based on spectral deferred corrections (SDC) methods, PFASST uses a space-time hierarchy with various coarsening strategies to maximize parallel efficiency. However, since the integration of SDC or PFASST into an existing application code is by far not straightforward and the potential gain is typically uncertain, the Python prototyping framework pySDC enables users to rapidly test new ideas in this field and to implement first toy problems more easily and more quickly. Yet, the main focus of this framework is time-integration of PDEs and in order to keep this focus, spatial discretization, solvers and (ideally) parallelization are left to established software frameworks like PETSc or FEniCS. In this talk we report on the pros and cons of this separation, the results achieved so far, current roadblocks and our ever-growing todo list. We will also shed some light on potential collaboration opportunities with other research software developers or potential users.

Robert Speck  
Juelich Supercomputing Centre  
Forschungszentrum Juelich GmbH

r.speck@fz-juelich.de

### MS70

#### Programming Heterogeneous Hpc Platforms with the Starpu Task-Based Runtime System

The StarPU task-based runtime system for heterogeneous platforms is developed by INRIA Team STORM in Bordeaux, France. It offers heterogeneous task scheduling capabilities on computing nodes equipped with accelerators, while presenting application developers with its programmer-friendly sequential task flow programming model. This talk will give an overview of StarPU's model and specific characteristics, together with an update on latest developments.

Oliver Aumage  
INRIA  
olivier.aumage@inria.fr

### MS70

#### Multi-Hetero Accelerated Computing: System and Application

GPU is the most promised accelerating device for HPC today and there are a number of GPU-clusters, especially for top-level systems in the world, for high absolute performance as well as its high performance per power consumption. However, the GPU solution is not perfect for some applications, especially for dynamically changing the behavior and degree of parallelism during the execution. There are several more features which cannot be covered perfectly by GPU. On the other hand, FPGA (Field Programmable Gate Array) is rapidly featured not just for embedded system but also for HPC use. FPGA has almost opposite characteristics from GPU (non-SIMD, hardware programmability, autonomous, equipped with high speed interconnection, etc.), so we are challenging to couple FPGA with GPU for compensation. We call this paradigm as CHARM (Cooperative Heterogeneous Acceleration with Reconfigurable Multi-devices), and implemented the world first GPU-FPGA combined multi-hetero cluster, Cygnus at Center for Computational Sciences in University of Tsukuba. In this talk, I focus on the basic concept of CHARM and supercomputer Cygnus as well as several supporting system environment for FPGA-FPGA direct interconnection network, intra-node GPU-FPGA DMA, a couple of programming method for CHARM, and the actual application which achieves a great performance improvement by this concept.

Taisuke Boku  
Center for Computational Sciences, University of Tsukuba  
taisuke@cs.tsukuba.ac.jp

### MS70

#### Designing an Intelligent Runtime System for Extremely Heterogeneous Computing

Across embedded, mobile, enterprise, and HPC systems, computer architectures are becoming more heterogeneous and complex. This complexity is causing a crisis in programming systems and performance portability. Several programming systems are working to address these challenges, but the increasing architectural diversity is forcing software stacks and applications to be specialized for each architecture. This talk argues that a more agile, proactive, and intelligent runtime system is essential to increase

performance portability and improve user productivity. To this end, this talk introduces a new runtime system called IRIS. IRIS enables programmers to write portable and flexible programs across diverse heterogeneous architectures from edge to exascale by orchestrating multiple programming platforms in a single execution and programming environment. This talk will present the challenges we have faced while designing IRIS and how we addressed the challenges.

Jungwon Kim  
Oak Ridge National Laboratory  
kimj@ieee.org

### MS70

#### Challenges on Global Task-Based Programming Model for Fugaku and Beyond

Recently, task-based parallel programming model is attracting attention from many researchers working on exascale computing. To realize exascale systems, there are many challenges and issues including architectures and programming models to exploit billions of parallelism. There are two major different approaches for exascale, manycore-based and accelerator-based. OpenMP task is a key to make use of manycore efficiently, and is to be evolved for exploiting global and large parallelisms and integration with communication layers for large scale system such as Fugaku. On the other hand, task-based programming model is also useful to combine accelerator devices by offloading functions to the devices as separate tasks. Future systems beyond exascale are expected to be large-scale, heterogeneous and complex integrating many accelerators devices so that these should be programmed by global task-based programming model. We are working on global task-based parallel programming by extending the PGAS programming language XcalableMP to tackle these challenges. In this talk, challenges for global task-based parallel programming models and languages for exascale system and beyond will be addressed with our project.

Mitsuhsa Sato  
RIKEN Advanced Institute for Computational Science  
msato@riken.jp

### MS71

#### Parallel Tensor Train Rounding using Gram SVD

Tensor Train (TT) is a low-rank tensor representation consisting of a series of three-way cores whose dimensions specify the TT ranks. Formal tensor train arithmetic often causes an artificial increase in the TT ranks. Thus, a key operation for applications that use the TT format is rounding, which truncates the TT ranks subject to an approximation error guarantee. Truncation is performed via SVD of a highly structured matrix, and current rounding methods require careful orthogonalization to compute an accurate SVD. We propose a new algorithm for TT rounding based on the Gram SVD algorithm that avoids the expensive orthogonalization phase. Our algorithm performs less computation and can be parallelized more easily than existing approaches, at the expense of a slight loss of accuracy. We demonstrate that our implementation of the rounding algorithm is efficient, scales well, and consistently outperforms the existing state-of-the-art parallel implementation in our experiments.

Hussam Al Daas  
INRIA

hussam.al-daas@stfc.ac.uk

Grey Ballard  
Wake Forest University  
ballard@wfu.edu

Lawton Manning  
Wake Forest University  
Winston-Salem, North Carolina  
mannlg15@wfu.edu

### MS71

#### FIST-HOSVD: Fused Inplace Sequentially Truncated Higher Order Singular Value Decomposition

Several novel methods of improving the memory locality of the Sequentially Truncated Higher Order Singular Value Decomposition (ST-HOSVD) algorithm for computing the Tucker decomposition are presented. We show how the two primary computational bottlenecks of the ST-HOSVD can be fused together into a single kernel to improve memory locality. We then extend matrix tiling techniques to tensors to further improve cache utilization. This blocked based approach is then leveraged to drastically reduce the auxiliary memory requirements of the algorithm. Our approach's effectiveness is demonstrated by comparing benchmark results between the traditional ST-HOSVD kernels and our single fused kernel. We then compare single-node CPU runtime results of a ST-HOSVD implementation utilizing our optimizations to TuckerMPI. We demonstrate  $\sim 2\times$  speedup over the existing state-of-the-art for dense high rank tensors, whilst increasing the problem size that can be computed for a given memory allocation by  $\sim 3\times$ . Finally, we demonstrate our approach's effectiveness with a motivating example of compressing Homogeneous Charge Compression Ignition (HCCI) and Stat Planar (SP) simulation datasets.

Benjamin Cobb  
Georgia Tech  
bcobb33@gatech.edu

Eric Phipps  
Sandia National Laboratories  
Center for Computing Research  
etphipp@sandia.gov

Hemanth Kolla  
Sandia National Laboratories  
hnkolla@sandia.gov

mit Catalyurek  
Georgia Institute of Technology  
umit@gatech.edu

### MS71

#### Parallel Randomized Algorithms for Tucker Decompositions

The Tucker decomposition is a low-rank tensor decomposition generalizing the matrix SVD to higher dimensions. Traditional algorithms used to compute Tucker decompositions can be computationally expensive as they involve computing the SVD of mode unfolding, which can be quite large. Existing parallel and randomized algorithms have reduced this cost, but computational challenges remain. We propose new parallel randomized algorithms to address these challenges. Using randomized matrix tech-

niques, we accelerate a distributed-memory implementation of the Tucker decomposition to obtain an efficient algorithm that avoids communication. Specifically, we employ a new sketching method that exploits Kronecker structure to accelerate a key computation. We also present probabilistic analysis of the error resulting from the algorithm, as well as numerical results demonstrating the computational benefits of this approach.

Rachel Minster

North Carolina State University  
minsterr@wfu.edu

Zitong Li

Wake Forest University  
Computer Science Department  
liz20@wfu.edu

Grey Ballard

Wake Forest University  
ballard@wfu.edu

### MS71

#### Multi-Start Poisson Tensor Factorization

The Poisson Tensor Factorization (PTF) is a special case of the CP decomposition applied to low-rank, multilinear structure via Poisson maximum likelihood estimation. Owing to the nonlinearity and non-convexity of the underlying problem, it is important to use a multi-start strategy to improve the chance of recovering the maximum likelihood estimator rather than suboptimal local solutions. We compare scalability-accuracy trade-offs among state-of-the-art methods for solving the PTF problem on real problems in the multi-start context.

Jeremy Myers

The College of William and Mary  
jmmyers01@email.wm.edu

Daniel M. Dunlavy

Sandia National Laboratories  
Computer Science and Informatics  
dmdunla@sandia.gov

### MS72

#### Update on the Intel Onemkl Library: Performance, Interfaces, and Portability

Today's applications are more diverse than ever, as is the hardware used to run and accelerate them. This causes a challenge for developers, as they must juggle multiple code bases, tools, programming languages, and workflows. The oneAPI initiative is designed to help overcome these challenges, by offering a unified, standards-based programming model that allows acceleration of the code on various processing architectures. The Intel oneAPI Math Kernel Library (oneMKL) product is the Intel product implementation of the oneMKL Data Parallel C++ (DPC++) specification (with DPC++ interfaces) as well as similar functionality with C and Fortran interfaces, and is highly optimized for Intel CPU and GPU hardware. In this talk, we will look at recent developments in the library and also discuss the oneMKL interfaces project. The oneMKL interfaces are an open-source implementation of the oneMKL specification, which works with multiple hardware backends using device-specific libraries underneath.

Sarah Knepper

Intel Corporation

sarah.knepper@intel.com

### MS72

#### Modern Mixed-Precision Methods in Portable C++ for Accelerated Hardware Platforms

Mixed precision methods are growing increasingly popular for numerical linear algebra. However, developing and testing these methods often requires repeatedly implementing the same routine with different combinations of datatypes. Furthermore, adding support for hardware accelerators often requires reimplementing codes for the differences in API for vendor libraries. Fortunately, C++ features, such as templating, allow for implementing a method once in a generic manner without significant loss of performance or readability.

Neil Lindquist

University of Tennessee  
nlindqu1@icl.utk.edu

### MS72

#### Design, Usage and Features of the Tlapack Library

The sustainability of legacy software libraries, in particular legacy numerical libraries, has been seriously deteriorated by several recent radical changes in microprocessor and system designs. These have rendered most legacy software infrastructure obsolete. The Linear Algebra PACKage (LAPACK) is a community standard for dense linear algebra and has been adopted and supported by a large community of users, computing centers, and high-performance computing (HPC) vendors. The  $\text{iT}_\ell$ LAPACK library is one proposal to close the gap between LAPACK and the new and emerging computing platforms. It uses C++ templates to provide precision-neutral algorithms, i.e., which work in single, double, half, and multiprecision types, and, in some cases, allow for mixed-precision. In the long term, we hope that some of the framework of  $\text{iT}_\ell$ LAPACK could be used to develop new implementations of state-of-the-art numerical linear algebra libraries. In this talk, I will present  $\text{iT}_\ell$ LAPACK, its design, features, and some examples of usage. I will talk about the test suite that has been developed jointly.

Wesley S. Pereira

Department of Mathematical and Statistical Sciences  
University of Colorado Denver  
wesley.pereira@ucdenver.edu

### MS72

#### P1673:A Proposal for a C++ Standard Linear Algebra Library

Linear algebra is a fundamental building block for a large range of applications, which had been standardized as far back as the early 90s through the Fortran BLAS interface. However, despite its importance linear algebra functionality never made it into the C++ standard, even though functionality of similar complexity - such as sorting - has been part of the standard library for a long time. Partly that is because BLAS was primarily used in high performance computing, which only represents a small fraction of the overall C++ user base. Consequently, linear algebra was considered a niche feature by the committee. The increasing use of high performance computing in the commercial sector, together with the advent of AI and Machine

Learning changed that calculus, putting linear algebra into the list of strategic exploration topics for the C++ standard. P1673 is a Linear Algebra proposal based on the existing BLAS standard. It builds on top of mdspan - multi-dimensional arrays for C++ - and extends the BLAS capabilities by modern capabilities such as mixed precision calculations, flexibility in data layout, and a design which is easily extended to support batched BLAS. This talk will present the design, demonstrate how it improves upon the existing BLAS standard, and discuss some options for future extension.

Christian Trott  
Sandia National Laboratories  
crtrott@sandia.gov

### MS73

#### Automatic Code Generation from High Level Mathematical Abstractions

The hardware for scientific computing evolved significantly in the past decade, especially with many core architectures like GPGPU's. Programming the numerical methods by the modellers on these complex machines need thorough understanding of its architecture, programming language and maintaining different code bases for each architecture. This is not sustainable, as huge time needs to be invested and capability enhancement becomes nearly impossible. Separation of concerns approach using Domain Specific Languages (DSL's) is a sensible option to tackle these issues. In this the numerical modeller writes a single source code using the DSL's Application Programming Interface (API) and the computer scientist's work on optimization for different architectures using source-to-source translation. However, the DSL's API's tend to be very specific and programming using DSL's is error prone. Code generation approach is a solution to overcome some of the complexities caused by ever changing architectures and programming paradigms. In this talk, experiences in developing automatic code generation frameworks for structured finite difference methods (OpenSBLI, <https://opensbli.github.io/>) and finite volume methods from high level mathematical abstractions and symbolic manipulations will be presented. Algorithmic changes that are impossible to implement by hand can be easily generated using these techniques and the speedup one can achieve will also be presented.

Satya Pramod Jammy  
SRM University AP, Andhra Pradesh  
satyapramod.j@srmmap.edu.in

### MS73

#### Evolutionary Re-Engineering of Multi-Physics Industrial HPC Applications with OP-DSLs

The rapid evolution of High Performance Computing architectures has resulted in highly complex systems with massively parallel heterogeneous processors, deep and multiple memory hierarchies and interconnects. As a result, maintaining performance as platforms change has become increasingly difficult. On the one hand, open standards have been slow to catch up with supporting new hardware, and for many real applications have not provided the best performance achievable from these systems. On the other hand, proprietary solutions have only targeted narrow vendor-specific devices resulting in a proliferation of parallel programming models and technologies. The only practically viable approach to addressing the above issue, is through the development of appropriate application-

oriented, high-level programming abstractions such as Domain Specific Languages (DSLs). In this talk I will present how two of the earliest DSLs developed in the UK, OP2 and OPS, have been able to build on multi-layered abstractions techniques to evolve production-grade HPC applications. I will detail how simple source-to-source and automatic code-generation techniques using maintainable software technologies have enabled us to develop a framework that has remained agile in delivering performance portability in the face of nearly a decade of hardware and parallel programming innovations. I will also discuss lessons learnt and my outlook for our DSLs in the run-up to deploying exa-scale systems.

Gihan Mudalige  
University of Warwick  
g.mudalige@warwick.ac.uk

### MS73

#### Teaching Weather(ed) Code New Languages: DSLs and Source-to-Source Translation for ECMWF's Integrated Forecasting System

The need to fully exploit the advantages of modern accelerators in HPC architectures plays a vital role in achieving the next step change in predictive skills for weather and climate models. However, the diversification of HPC hardware continues and incurs a recurring need to adapt and transform existing model software to new programming models, which is becoming increasingly impractical using a single code base. All current strategies tackling this issue involve some form of abstraction that separates numerical algorithms from hardware-specific implementation. ECMWF is aiming to achieve both compatibility and performance portability for accelerators and future architectures through a combination of DSLs, domain-specific libraries and bespoke source-to-source translation tools that allow an incremental adaptation of the code base alongside scientific development. In this talk we will describe algorithmic components of the Integrated Forecasting System (IFS) and highlight key elements of our roadmap for adapting them to GPU architectures. We discuss how a combination of flexible data structures, an extended use of library interfaces and the use of source-to-source translation tools is envisaged to allow the incremental adaptation of the code base. The second part will focus on Loki, an in-house developed tool that allows bespoke source-to-source transformations to be devised to adapt individual model components to novel programming paradigms or community-driven DSLs.

Balthasar Reuter, Michael Lange, Olivier Marsden  
ECMWF  
balthasar.reuter@ecmwf.int, michael.lange@ecmwf.int,  
olivier.marsden@ecmwf.int

### MS73

#### Boast: A Computing Kernel Metaprogramming and Autotuning Framework for Heterogeneous Architectures

In the ever-changing High-Performance Computing environment, porting and tuning applications to new platforms is of paramount importance but tedious and costly in terms of human resources. While solutions start to emerge to abstract most of the underlying programming models, the tuning of computing kernels is the work of highly trained specialists. An efficient computing kernel is often the combination of insights and techniques specific

to the targeted architecture. The performance portability, and sometimes compatibility, of these specialized versions are often poor. One way to tackle this problem is by using meta-programming approaches allowing the specialist to express optimization strategies and techniques in an orthogonal manner. Coupled with a runtime supporting many programming models found in HPC, different versions can be generated and benchmarked. The design space can then be explored to find the most suitable version for the target architecture. I will discuss how I developed and used BOAST, an autotuning framework, to address these challenges in several Heterogeneous HPC applications, using accelerators (OpenCL, CUDA, HIP) or vector architectures (AVX, NEON).

Brice Videau

Argonne National Laboratory  
bvideau@anl.gov

#### MS74

##### A Numerical Study of a Micro-Macro MGRiT Method for Scale-Separated SDEs

Time-parallel methods can significantly reduce wall clock time for accurate numerical solutions by parallelising across the time dimension. We present and test a micro-macro version of a two-level MGRiT method, in which the fine propagator is based on a (high-dimensional, slow-fast) stochastic microscopic model, and the coarse propagator is based on a low-dimensional approximate effective dynamics at slow time scales. At the microscopic level, we use an ensemble of Monte Carlo particles, whereas the approximate coarse propagator uses the (deterministic) FokkerPlanck equation for the slow degrees of freedom. The required coupling between microscopic and macroscopic representations of the system introduces several design options, specifically on how to generate a microscopic probability distribution consistent with a required macroscopic probability distribution and how to perform the coarse-level updating of the macroscopic probability distribution in a meaningful manner. We numerically study how these design options affect the efficiency of the algorithm in a number of situations.

Ignace Bossuyt, Stefan Vandewalle, Giovanni Samaey  
KU Leuven

ignace.bossuyt1@kuleuven.be, stefan.vandewalle@cs.kuleuven.be, giovanni.samaey@kuleuven.be

#### MS74

##### Time-Parallel Flow Estimation

Deformable image registration is a key technology in medical imaging; there the goal is to compute a meaningful spatial correspondence between two or more images of the same scene. One approach is to use an optimal control formulation to compute a stationary velocity field that parameterize the deformation map. The same methods can be used to estimate the motion of contrast agents from 3d ultrasound images. In the talk I'll introduce the application problem and discuss computational techniques for its solution, with a focus on using parallelization in time to reduce the time-to-solution.

Sebastian Götschel

TU Hamburg  
sebastian.goetschel@tuhh.de

Michael Minion

Lawrence Berkeley National Lab  
mlminion@lbl.gov

#### MS74

##### Fast Multigrid Reduction-in-Time for Advection Problems

Parallel-in-time integration of partial differential equations (PDEs) is an increasingly important topic in the context of modern, massively parallel computing architectures. Many parallel-in-time algorithms have been shown highly efficient for diffusion-dominated PDEs, but often are inefficient or even divergent for advection-dominated PDEs. We consider the application of the two-level Parareal algorithm, and its multilevel generalization of multigrid reduction-in-time (MGRIT) to linear advection problems. We focus on odd-order, finite-difference semi-Lagrangian discretizations of such problems, for which we present novel semi-Lagrangian-like coarse-grid operators that yield fast multilevel time integration. The coarse-grid operators occur in split form, consisting of a standard semi-Lagrangian discretization combined with a dissipative correction.

Hans De Sterck

University of Waterloo, Canada  
hdesterck@uwaterloo.ca

Robert D. Falgout

Lawrence Livermore National Laboratory  
falgout2@llnl.gov

Oliver A. Krzysik

School of Mathematics  
Monash University  
oliver.krzysik@monash.edu

#### MS74

##### Diagonal SDC Preconditioning

The spectral deferred correction method (SDC) is an iterative solver for time-dependent differential equations and can be interpreted as a preconditioned Picard iteration for the collocation problem. The key to an efficient iteration process is the choice of the preconditioner that defines the speed of convergence. While the de-facto standard choice is a fast-converging, yet serial preconditioner, our goal is to find a fast-converging, diagonal and therefore parallel one. Depending on the parameters of the equation at hand, we first try to find optimal diagonal preconditioners for Dahlquist's equation. Using spectral methods, we then apply these results to partial differential equations with a linear stiff part that we treat in a decoupled and implicit way with the previously optimized parallel preconditioners, whereas nonlinear parts are treated explicitly in real space. In this talk we present different approaches to find such optimal diagonal preconditioners, including mathematical optimization and reinforcement learning. We shed light on the pros and cons of these different strategies and show results for linear and nonlinear PDEs.

Ruth Schöbel

Juelich Supercomputing Centre  
Forschungszentrum Juelich GmbH  
r.schoebel@fz-juelich.de

Jan Ebert, Jenia Jitsev, Stefan Kesselheim  
Forschungszentrum Jülich GmbH  
ja.ebert@fz-juelich.de, j.jitsev@fz-juelich.de, s.kesselheim@fz-juelich.de

Robert Speck  
 Juelich Supercomputing Centre  
 Forschungszentrum Juelich GmbH  
 r.speck@fz-juelich.de

Michael Minion  
 Lawrence Berkeley National Lab  
 mlminion@lbl.gov

## MS75

### Vectorized Implicit Time Discretion

Numerical software to approximately solve partial differential equations (PDE) is of central importance in many fluid dynamics application. Constantly increasing model details require almost unlimited computer power, but the way increased hardware performance is achieved has changed fundamentally in recent years. To use modern hardware efficiently new dedicated numerical methods are necessary. Particular challenges are posed by the increasing parallelism, especially the instruction level parallelism, as well as the ever increasing gap between memory and computing speed. We will discuss a novel approach to increase the arithmetic intensity and to make full use of instruction level parallelism, like AVX512. The idea is to reformulate the space-time problem such yields a matrix equation, similar to the approaches taken in parallel in time (PiT) methods. The fundamental difference to PiT methods is that we do not aim at using more nodes, but to increase the local efficiency on the individual nodes. The arising matrix equation is then solved efficiently using block-Krylov methods. We will discuss how to efficiently implement and vectorize block-Krylov methods and investigate how the increase performance of a single Krylov iteration and the decreased convergence rate (due to our reformulation) can be balanced to achieve the best over-all performance.

Christian Engwer, Nils-Arne Dreier  
 INAM, University of Münster  
 christian.engwer@uni-muenster.de, n.dreier@uni-muenster.de

## MS75

### Flexible and Generic Software Design for Lattice Boltzmann Methods on CPU and GPU with Walberla and Lbmpy

As more and more heterogeneous hardware is used in high-performance computing like GPUs as accelerators, it becomes a demanding challenge to provide compute kernels optimised explicitly for each architecture. Furthermore, in the context of lattice Boltzmann methods, a wide variety of collision models, stencils, boundary conditions, and equilibria formulations exist, making the situation even more demanding. In this talk, the code generator lbmpy is introduced to tackle the problem. It allows defining lattice Boltzmann methods on a symbolic mathematical description. From this description, highly-optimised compute kernels, communication schemes, and boundary conditions automatically adapted to the model's description are generated. The generated compute kernels can run independently on a single machine, or they can be combined with the massively parallel multiphysics framework waLberla to use domain decomposition and allow for extensive MPI-parallel simulations.

Markus Holzer  
 FAU

markus.holzer@fau.de

## MS75

### A Moment Limiter for the DG Method on Adaptively Refined Triangular Meshes

We will present a second-order limiter for the discontinuous Galerkin method on nonconforming triangular meshes that arise in adaptive computations. It limits the linear solution coefficients (or moments) by reconstructing the slopes along a set of directions in which the moments decouple. We perform the reconstruction of the slopes on a compact stencil consisting of eight elements. We propose a simple algorithm to update the reconstruction stencil of elements in an adaptively refined triangular mesh. Our algorithm is implemented entirely on the graphics processing unit (GPU) and avoids race conditions. We provide numerical experiments to validate the robustness of the limiter in the presence of discontinuities. Finally, we perform wall clock studies to analyze the performance of the proposed limiter for the computational cost involved in setting up the reconstruction neighborhood and executing the limiter.

Lilia Krivodonova, Krishna Dutt  
 University of Waterloo  
 lgk@uwaterloo.ca, kdkrishnadutt@uwaterloo.ca

## MS76

### Automatic High-Performance Kernel Generation for Accelerators : Specialization to Generalization

The goal of achieving lower latency, higher throughput, and higher energy efficiency for compute-intensive domains (e.g., Deep Learning) has led to a meteoric increase in domain-specific hardware accelerators. However, it requires significant efforts to develop both hardware and software stacks for each domain. An even bigger problem is the lack of flexibility and scalability of the hardware and software stack across multiple domains. As a result, there is a great interest in building programmable and flexible accelerators. In this talk, I present two works starting with closely looking at the "Vyasa" compiler generating high-performance code for the specific accelerator, i.e., for Xilinx AI Engine, and the "PolyEDDO" framework targetting code generation for various accelerators using hardware abstractions.

Prasanth Chatarasi  
 IBM  
 prasanth@ibm.com

## MS76

### Understanding the Design Space of Sparse/Dense Multiphase Dataflows for Mapping Graph Neural Networks on Spatial Accelerators

Graph Neural Networks (GNNs) have garnered a lot of recent interest because of their success in learning representations from graph-structured data across several critical applications in cloud and HPC. Owing to their unique compute and memory characteristics that come from an interplay between dense and sparse phases of computations, the emergence of reconfigurable dataflow (aka spatial) accelerators offers promise for acceleration by mapping optimized dataflows (i.e., computation order and parallelism) for both phases. The goal of this work is to characterize and understand the design-space of dataflow choices for running GNNs on spatial accelerators in order for the

compilers to optimize the dataflow based on the workload. Specifically, we propose a taxonomy to describe all possible choices for mapping the dense and sparse phases of GNNs spatially and temporally over a spatial accelerator, capturing both the intra-phase dataflow and the inter-phase (pipelined) dataflow. Using this taxonomy, we do deep-dives into the cost and benefits of several dataflows and perform case studies on implications of hardware parameters for dataflows and value of flexibility to support pipelined execution.

Raveesh Garg, Eric Qin  
Georgia Institute of Technology  
raveesh.g@gatech.edu, ecqin@gatech.edu

Francisco Muñoz-Martínez  
Universidad de Murcia  
francisco.munoz2@um.es

Robert Guirado  
Universitat Politècnica de Catalunya  
rguirado@ac.upc.edu

Akshay Jain  
Neutron  
akshay.jain@neutron.com

Sergi Abadal  
Universitat Politècnica de Catalunya  
abadal@ac.upc.edu

José Abellán  
Universidad Católica de Murcia  
jlabellan@ucam.edu

Manuel Acacio  
Universidad de Murcia  
meacacio@um.es

Eduard Alarcon  
Universitat Politècnica de Catalunya  
eduard.alarcon@upc.edu

Sivasankaran Rajamanickam  
Sandia National Laboratories  
srajama@sandia.gov

Tushar Krishna  
Georgia Institute of Technology  
tushar@ece.gatech.edu

## MS76

### FPGA/DNN Codesign and Acceleration

In this talk, I will introduce a neural network and hardware implementation co-search methodology, named NAIS, to pursue aggregated solutions of high accuracy DNN designs and efficient hardware deployments simultaneously. To enable a comprehensive co-search framework, there are three indispensable components: 1) efficient hardware accelerator design (e.g. FPGA); 2) hardware-aware neural architecture search (NAS); 3) automatic design tools to quickly deploy DNNs to hardware platforms. I will discuss each component and the integrations to support an efficient and optimal NAIS implementation.

Callie Hao  
Georgia Institute of Technology

callie.hao@gatech.edu

## MS76

### Sparseloop: An Analytical Design Space Exploration Tool for Sparse Tensor Accelerators

Sparse tensor algorithms are critical to many emerging workloads (DNNs, data analytics, recommender systems, graph algorithms, etc.). As a result, recently, many sparse tensor accelerators and systems have been proposed to improve efficiency and performance for sparse tensor algorithms. These designs often involve complex tradeoffs between compression formats, hardware support that exploits sparsity, and on-chip dataflows. Each design therefore represents a specific design point in the large space. As more and more distinct solutions are proposed, we as a community need a better tool to model and search the large design space for various sparse tensor algorithms. In this talk, I will introduce Sparseloop, a tool that analytically models the impact of different dataflows, data representations, and associated hardware support for sparsity in sparse tensor accelerators. Sparseloop builds on prior dense tensor accelerator modeling tools and serves as an efficient tool for architects to explore the large design space for sparse tensor accelerators. Sparseloop is fast, accurate, and flexible for design space exploration. Each modeling takes less than 0.1 seconds to generate results, which is orders of magnitude faster than cycle-accurate simulators, and Sparseloop produces less than 1% error compared to custom hardware analytical models. I will also present case studies of using Sparseloop to study various sparsity optimizations, such as hardware support for structured sparsity.

Po-An Tsai  
NVIDIA  
poant@nvidia.com

## MS77

### Formulation and Solvers for Kinetic-Fluid Models for Plasma Simulation

We review recent progress on the formulation and numerical solution of kinetic-fluid models and algorithms for kinetic plasma simulation. The models we are interested in combine fluid and kinetic descriptions, either from a physics perspective (e.g., hybrid methods) or from an algorithmic one (e.g., hierarchical multi-model solvers), or both. Kinetic representations are characterized by an augmented (in principle, up to 6D) phase-space representation, which includes a configuration space and a velocity space, the latter discretized either with particles or on a mesh. We will touch upon different strategies for the efficient nonlinear solution of these systems such as linear moment-based preconditioning, and nonlinear Picard-based iterative solvers exploiting high-order/low-order (HOLO) hierarchical decompositions. We will demonstrate the effectiveness of our formulations with various applications of interest.

Luis Chacon, Guangye Chen, Adam Stanier, William T. Taitano  
Los Alamos National Laboratory  
chacon@lanl.gov, gchen@lanl.gov, stanier@lanl.gov, taitano@lanl.gov

## MS77

### Coupling Kinetic PIC to Multi-Fluid Plasma Mod-

### els using Compatible Discretizations

This talk investigates coupling between a particle-in-cell (PIC) plasma model and a multi-fluid continuum treatment. A central obstacle to this is evolving the stiff time scales that arise, while maintaining critical properties, like no magnetic monopoles and a discrete variant of Gauss' law. Our approach is to develop a DG discretization of the fluid equations coupled to the particles through collisional interactions, and indirectly through current coupling into Maxwells equations. For the later coupling, a compatible discretization is used to evolve Maxwells equations. This talk will demonstrate how these discretizations can be coupled while preserving the Gauss law involutions, and implicitly integrating the implicit modes, including the plasma and cyclotron frequencies. A discussion of the parallel solution of the systems will be included to demonstrate scalability.

Eric C. Cyr  
Computational Mathematics Department  
Sandia National Laboratories  
eccyr@sandia.gov

### MS77

#### Coupling Particle-in-Cell and Continuum Methods in the Global Total-f Gyrokinetic Magnetic Fusion Code XGC

The X-point gyrokinetic code (XGC) is a 5D total-f gyrokinetic particle-in-cell (PIC) code used to study kinetic transport phenomena in magnetic fusion plasmas. In addition to PIC methods, the code evaluates dissipative operations such as Coulomb collisions and heat sources/sinks on a 5D grid at each fixed time step. This requires a PIC-continuum coupling in the form of a mapping between the marker particles and the 5D phase-space grid. If the conservation errors in the mapping are undesirably large they can cause a significant error accumulation in long-time simulations of plasmas with steep density and temperature gradients. Here we discuss a novel mapping technique in 2D velocity-space, based on the calculation of a pseudo-inverse, to exactly preserve moments up to the order of the chosen discretization space and to minimize the errors. This new particle to/from velocity grid interpolation method has been implemented in XGC, and we demonstrate its effectiveness. The mapping relies on a particle resampling technique which is used to create, annihilate or redistribute particles in configuration space while preserving a desired number of moments. We will also discuss details of the resampling and how the new mapping could be used to reduce the error when coupling XGC to other codes.

Albert Mollén  
Princeton Plasma Physics Laboratory  
amollen@pppl.gov

Varis Carey  
Theodon Consulting  
varis.carey@theodonconsulting.com

Mark Adams  
Lawrence Berkeley National Laboratory  
mfadams@lbl.gov

Matthew G. Knepley  
University at Buffalo  
Department of Computer Science and Engineering  
knepley@gmail.com

Robert Hager  
Princeton Plasma Physics Laboratory  
rhager@pppl.gov

Julien Dominski  
PPPL  
jdominsk@pppl.gov

Choong-Seock Chang  
Princeton Plasma Physics Laboratory  
cschang@pppl.gov

### MS78

#### Fast Parallel Solver for the Space-Time IgA-Dg Discretization of the Anisotropic Diffusion Equation

We consider the space-time discretization of the anisotropic diffusion equation, using an isogeometric analysis (IgA) approximation in space and a discontinuous Galerkin (DG) approximation in time. Drawing inspiration from a former spectral analysis of space-time operators in [Benedusi, 2018], we propose for the resulting space-time linear system a parallel multigrid preconditioned GMRES method, which combines a preconditioned GMRES with a standard multigrid acting only in space. The performance of the proposed parallel solver is illustrated through numerical experiments, showing its competitiveness in terms of iteration count, run-time and parallel scaling. More general parallel space-time multigrid techniques applied also to non-linear problems, will be also discussed.

Pietro Benedusi  
Institute of Computational Science  
USI Lugano  
pietro.benedusi@usi.ch

Paola Ferrari  
University of Insubria  
pferrari@uninsubria.it

carlo garoni  
Tor Vergata  
carlogaroni@hotmail.it

Rolf Krause  
Università della Svizzera italiana, Switzerland  
rolf.krause@usi.ch

Stefano Serra-Capizzano  
University of Insubria  
s.serracapizzano@uninsubria.it

### MS78

#### Space-Time Block Preconditioning

Parallel-in-time (PinT) methods have become increasingly popular in the simulation of time-dependent numerical PDEs, owing to their potential for providing additional speedup when more classical spatial parallelism saturates. In particular, all-at-once schemes attempt to achieve parallelisation in time by simultaneously solving the whole system which arises from the monolithic space-time discretisation of the target PDEs. In this talk, we consider a novel take on the all-at-once solution of systems of PDEs, and introduce the concept of space-time block preconditioning. The proposed method takes motivation from well-developed spatial block preconditioning techniques used

in sequential time-stepping or steady-state problems, and extends such principles to the whole space-time setting. The resulting preconditioning procedure has the advantage of making the system more amenable to be solved via PinT methods, by reducing the target system to (parallel-)integrate in time to single-variable PDEs. To showcase its efficiency and flexibility, the method is tested on problems in incompressible fluid dynamics and magnetohydrodynamics. Results indicate good scalability under temporal and spatial mesh refinement in terms of number of iterations to convergence for both the linear and nonlinear solvers. Moreover, the total number of preconditioner applications compares well with sequential time-stepping, showing only minimal overhead in this sense.

Federico Danieli

University of Oxford  
federico.danieli@trinity.ox.ac.uk

## MS78

### ParaRKC: A Parallel Runge-Kutta-Chebyshev Method

The Parareal parallel-in-time algorithm [Lions et al., "A parareal in time discretization of PDEs, 2001] has been successfully applied to numerous engineering problems over the last two decades. It allows for the concurrent solution of initial value problems, independent of the choice of time integrator. Often, it is employed in combination with implicit time integration schemes to solve semi-discretized parabolic partial differential equations. An alternative to implicit integrators is given by Runge-Kutta-Chebyshev (RKC) methods [Sommeijer et al., "RKC: An explicit solver for parabolic PDEs", 1998]. Due to their explicit structure, they avoid matrix inversions and may therefore save computational effort compared to implicit schemes. Recently it was noted that a Parareal-like update may be used to parallelize any recursion relation admitting solutions based on different accuracies [Gander et al., "ParaStieltjes: Parallel computation of Gauss quadrature rules using a Parareal-like approach for the Stieltjes procedure", 2021]. Using this idea we propose a new approach (ParaRKC) combining Parareal and RKC type methods. The new method parallelizes over the internal stages of the RKC method, making use of a Parareal-like update and a coarse spatial discretization for the coarse propagator. We present results on the performance of ParaRKC and a comparison with Parareal using only a coarsening in space.

Peter Förster

TU Darmstadt  
peter.foerster@tu-darmstadt.de

Sebastian Schöps

Theorie Elektromagnetischer Felder (TEMF) and GSCE  
Technische Universität Darmstadt  
schoeps@gsc.tu-darmstadt.de

Iryna Kulchytska-Ruchka

Graduate School of Computational Engineering  
TU Darmstadt  
kulchytska@gsc.tu-darmstadt.de

Martin J. Gander

University of Geneva

martin.gander@unige.ch

## MS78

### Comparing the PFASST, MGRIT and Parareal Methods

PFASST, MGRIT and Parareal are three of the most popular parallel-in-time methods. However, these three methods have not been carefully compared in the literature. One reason for this is that the performance of each method depends on many factors, including the PDE being solved, the implementation specifications, and many method-specific parameters. In this work, we compare these three methods with the goal of investigating the performance of the algorithms rather than implementations. We have implemented all three methods in the LibPFASST library to reduce the effect on performance of using different software libraries. We distill the parameter space down to those that have the greatest impact on the reduction of the error: for MGRIT and Parareal, we vary the time discretization scheme and for PFASST, we vary the number of SDC sweeps. We conduct experiments on the Cori supercomputer at NERSC using two test problems. We look at three important results for each method: error versus wall-clock time, error versus number of time steps, and strong scaling.

Jordi Wolfson-Pou

Lawrence Berkeley National Laboratory  
jwolfsonp@lbl.gov

## MS79

### Moccuda: Running Cuda Codes on Fugaku

In recent years, the fastest supercomputers have been dominated by GPU-accelerated node architectures, and hence a majority of scientific applications have been ported to the CUDA programming paradigm and Nvidia's performance libraries (e.g. cuBLAS), while improvements to the original CPU-version received less attention. Furthermore, the proliferation of Deep Learning (DL) resulted in frameworks which are implemented with a GPU-first or GPU-only design philosophy and the CPU version only serves for unoptimized debugging purposes. These trends cause a non-trivial challenge to novel architectures, such as RIKEN's Arm/CPU-based Supercomputer Fugaku. To assist the early porting efforts, we designed, and present in this talk, a framework to natively execute CUDA code on CPUs. However, manually porting existing GPU-first applications and frameworks to Fugaku is a tedious (at best) or impossible (at worst) task, especially when it comes to highly-complex DL frameworks like Tensorflow or Pytorch with specific performance goals in mind. Hence, we explore an alternative approach and present the application of a virtual GPU plus re-implementations of CUDA libraries for CPU, to execute those CUDA applications natively. We demonstrate a proof-of-concept of our idea and show results of running Resnet50 with Pytorch, among other examples, on Supercomputer Fugaku.

Jens Domke

RIKEN-CCS  
jens.domke@riken.jp

## MS79

### Generating GPU Kernels from Chapel's Features

## for Parallelism And Locality

Chapel is a portable parallel programming language. Since the advent of using GPUs in HPC applications, targeting GPUs natively has been one of the most sought for features by Chapel users. Chapel's high-level language constructs, such as data-parallel *forall* loops that specify parallelism and *on* statements that specify locality, offer a great potential in writing portable applications that can execute on CPU and/or GPU efficiently. With this vision, the Chapel team at HPE has been making significant progress in designing and implementing support for GPUs, particularly in the September 2021 release of Chapel 1.25. In this talk, we will cover basic parallel concepts of the Chapel programming language focusing on those that are used for GPU programming. We will outline Chapel's vision of how GPUs can be targeted using those concepts in a portable manner. Then, we will dive deeper into how the Chapel compiler transforms data-parallel *forall* loops into GPU kernels and launches them. Finally, we will discuss near- and long-term plans for Chapel's native GPU support.

Engin Kayraklioglu, Andy Stone, David Iten, Sarah Nguyen, Bradford L. Chamberlain, Michael Ferguson  
Hewlett Packard Enterprise  
engin@hpe.com, andy.stone@hpe.com,  
david.iten@hpe.com, sarahelizabeth.nguyen@hpe.com,  
bradford.chamberlain@hpe.com,  
michael.ferguson@hpe.com

Michelle Strout  
Hewlett Packard Enterprise  
University of Arizona  
michelle.strout@hpe.com

### MS79

#### Using Macros to Assemble and Generate Specialized Code for Multiple Devices from a Single Source

Scientific software used on high performance computing (HPC) platforms is in a phase of transformation because of a combined increase in the heterogeneity and complexity of models and hardware platforms. Because separate implementations for different platforms is not an option, the computational science community has been looking for mechanisms to express code through abstractions that can be specialized for different platforms. Some approaches have met success through the use of template meta-programming in C++. However, these approaches are inherently inaccessible to non C++ codes. In this presentation, we describe a language agnostic methodology that mimics the behavior of templates as applied in the abstractions described above. We also include an example of code refactoring utilizing this approach.

Thomas Klosterman, Anshu Dubey  
Argonne National Laboratory  
tklosterman@anl.gov, adubey@anl.gov

### MS79

#### Transformation from DSL to Source Code for the Multi-Physics Application Flash-X

We propose a code transformation toolchain for Flash-X that realizes the "recipe-to-source" code transformation for Flash-X simulations and that is one crucial component in order to achieve performance portability. We de-

sign a high-level language to express operations of simulations in so-called "recipes," which are given as input to the toolchain. The tools of the code transformation pipeline include tree-based source code representation techniques and code orchestration and generation based on control flow graphs. The generated source code utilizes a new runtime, developed for Flash-X, that orchestrates dynamic and asynchronous data movement and task execution.

#### Johann Rudi

Argonne National Laboratory  
Mathematics and Computer Science  
jrudi@anl.gov

#### Jared O'Neil

Argonne National Laboratory  
joneal@anl.gov

#### Mohamed Wahib

AIST-TokyoTech Open Innovation Lab, Japan  
mohamed.attia@riken.jp

#### Anshu Dubey

Argonne National Laboratory  
adubey@anl.gov

### MS80

#### Overview of Some Popular Machine Learning Frameworks for Data Parallelism

Data parallelism is used in machine learning to improve the performance of model optimization and evaluation. In this scheme, the data is split into multiple workers/nodes, and every worker optimizes a replica of a model with its piece of data. In the process of optimization, every worker computes the loss function and gradients, but before updating the model parameters, gradients are averaged over all workers through the collective operation. Although the idea of data parallelism is quite simple, the practical implementation of it is complicated and, without proper handling of communications, may result in poor performance. We do an overview of several popular data parallelism frameworks and study their scaling with a number of nodes.

#### Denis Boyda

Argonne National Laboratory  
dboyda@anl.gov

### MS80

#### Scaling Data Parallel Training with PyTorch

Data parallelism has emerged as a popular solution for distributed training due to its broad applicability. As of v1.11, PyTorch offers three main data-parallel training paradigms, namely, DistributedDataParallel (DDP), Pipelined Data Parallelism (PDP), and FullyShardedDataParallel (FSDP). DDP replicates the model on every device to independently generate gradients and then communicates those gradients at each iteration to keep model replicas consistent. To accelerate training, DDP organizes gradients into buckets and overlaps backward computation with gradient communication. With the increase of model size, a single device might no longer host the entire states of the model and the optimizer replicas due to device memory capacity. To reduce per-device memory consumption, PDP partitions the model across multiple devices on the same machine and applies data parallelism across pipelines; and FSDP shards the states across all devices and only fetch states on-demand before usage. This presentation

covers technical details and comprehensive comparisons of the three data-parallel training paradigms.

Shen Li

Facebook

shenli@fb.com

**MS80**

### Data Parallel Training in DeepSpeed with ZeRO (Zero Redundancy Optimizer)

Large deep learning models offer significant accuracy gains, but training billions to trillions of parameters is challenging. Existing solutions such as data and model parallelisms exhibit fundamental limitations to fit these models into limited device memory, while obtaining computation, communication and development efficiency. In this talk we introduce a novel solution, Zero Redundancy Optimizer (ZeRO), to optimize memory, vastly improving training speed while increasing the model size that can be efficiently trained. ZeRO eliminates memory redundancies in data- and model-parallel training while retaining low communication volume and high computational granularity, allowing us to scale the model size proportional to the number of devices with sustained high efficiency. In fact, ZeRO can scale to trillions of parameters on modern GPU clusters and can be combined with heterogeneous memory technology to scale to tens of trillions of parameters without compromising efficiency or throughput scalability. Unlike other forms of parallelism, ZeRO enables data scientists to scale their model without requiring model code refactoring. A robust implementation of ZeRO and its extended family of technologies have been implemented in the DeepSpeed Library, and is used extensively by the DL community for large model training.

Samyam Rajbhandari

Microsoft

samyamr@microsoft.com

**MS80**

### Improving Large Scale Data Parallel Training Performance of Horovod

Horovod is a popular software library for communication in data parallel deep learning applications. While originally written by Uber to enable straightforward data parallel conversion of TensorFlow models to run across a handful of server nodes, the library has grown to support all major frameworks and has been used to run deep learning workloads on some of the worlds largest supercomputing clusters. Additionally, many new features and enhancements have been incorporated into the library to continually improve performance and usability. This talk will provide an overview of Horovods framework agnostic design and the implementation of several performance improvements to the original library to improve scalability, illustrated with real-world use cases.

Joshua Romero

NVIDIA

josh@nvidia.com

**MS81**

### Parallel Memory-Efficient Computation of Symmetric Higher-Order Joint Moment Tensors

The decomposition of higher-order joint cumulant tensor of spatial-temporal data sets is useful in analyzing multi-

variate non-Gaussian statistics with a wide variety of applications (e.g. anomaly detection, independent component analysis, dimensionality reduction). Computing the cumulant tensor often requires computing the joint-moment tensor of the input data first, which is very expensive using a naive algorithm. The current state-of-the-art algorithm for computing joint moment tensors takes advantage of the symmetric nature of a moment tensor by dividing it into smaller cubic tensor blocks and only compute the blocks with unique values and thus reducing computation. We propose an improvement over this algorithm by posing its computation as matrix operations, specifically Khatri-Rao products and standard matrix multiplications. Because this approach is much more cache efficient, we expect considerable speedup in a single processor. We implemented our algorithm in Julia and in Matlab and compared it against the state-of-the-art approach. The results show a speedup of up to 10x.

Zitong Li

Wake Forest University

Computer Science Department

liz20@wfu.edu

Hemanth Kolla

Sandia National Laboratories

hnkolla@sandia.gov

Eric Phipps

Sandia National Laboratories

Center for Computing Research

etphipp@sandia.gov

**MS81**

### Structured Matrix Approximations via Tensor Decompositions

We provide a computational framework for approximating a class of structured matrices (e.g., block Toeplitz, block banded). Our approach has three steps: map the structured matrix to tensors, use tensor compression algorithms, and map the compressed tensors back to obtain two different matrix representations — sum of Kronecker products and block low-rank format. The use of tensor decompositions enable us to uncover latent structure in the matrices and lead to computationally efficient algorithms. The resulting matrix approximations are memory efficient, easy to compute with, and preserve the error due to the tensor compression in the Frobenius norm. While our framework is quite general, we illustrate the potential of our method on structured matrices from three applications: system identification, space-time covariance matrices, and test matrices from SuiteSparse collection. See preprint: <https://arxiv.org/abs/2105.01170>.

Arvind Saibaba

North Carolina State University

saibab@ncsu.edu

Misha E. Kilmer

Mathematics Department

Tufts University

misha.kilmer@tufts.edu

**MS81**

### Robust Approximation of Tensor Networks, and Its Applications in Quantum Chemistry

Approximation of a tensor network by approximating (e.g.,

factorizing) one or more of its constituent tensors can be improved by canceling the leading-order error due to the constituents' approximation. The utility of such robust approximation is demonstrated for robust canonical polyadic (CP) approximation of a (density-fitting) factorized 2-particle Coulomb interaction tensor. The resulting algebraic (grid-free) approximation for the Coulomb tensor, closely related to the factorization appearing in pseudospectral and tensor hypercontraction approaches, is efficient and accurate, with significantly reduced rank compared to the naive (non-robust) approximation. Application of the robust approximation to the particle-particle ladder term in the coupled-cluster singles and doubles reduces the size complexity from  $\mathcal{O}(N^6)$  to  $\mathcal{O}(N^5)$  with robustness ensuring negligible errors in chemically-relevant energy differences using CP ranks approximately equal to the size of the density-fitting basis. Preliminary applications to other electronic structure methods will also be discussed.

Edward F. Valeev, Karl Pierce  
Virginia Tech  
efv@vt.edu, kmp5@vt.edu

### MS81

#### Memory-Efficient Tensorized Embedding Layers for Neural Networks

Tensor-trains have been proposed to reduce the memory requirements of embedding tables in industrial-scale deep neural networks, including recommendation models and graph neural networks. This talk summarizes the tensor-trains use cases and implementation challenges when training large embeddings for such models. In this context, we are developing a new prototype infrastructure of performance-optimized kernels, TT-EmbeddingBag, that aims to reduce model sizes while preserving model accuracy and incurring only modest increases in training-time over uncompressed baselines.

Chunxing Yin, Richard Vuduc  
Georgia Institute of Technology  
cyin9@gatech.edu, richie@cc.gatech.edu

### MS82

#### Robust Identification of Differential Equations from Noisy Data

Identifying unknown differential equations from given discrete time dependent data is a challenging problem. Noisy data make such identification particularly challenging. In this talk, we present robust methods against a high level of noise which approximate the underlying noise-free dynamics well. This approach is fundamentally based on numerical PDE techniques, and we introduce successively denoised differentiation and utilize subspace pursuit time evolution error for PDE identification.

Yuchen He  
Georgia Tech  
royarthur@gatech.edu

Sung Ha Kang  
Georgia Inst. of Technology  
Mathematics  
kang@math.gatech.edu

Wenjing Liao  
Georgia Institute of Technology

wlia060@gatech.edu

Hao Liu  
Hong Kong Baptist University  
haoliu@hkbu.edu.hk

Yingjie Liu  
School of Mathematics, Georgia Inst of Tech  
yingjie@math.gatech.edu

### MS82

#### Massively Parallel Solvers for Buoyancy-Driven Convection in Stokes Flow

Challenges in Earth mantle convection motivate the development of massively parallel solvers for advection dominated transport in viscous fluids. We present our hierarchical hybrid grids (HHG) based approach that couples a geometric multigrid solver for the Stokes equations to an Eulerian-Lagrangian method for the advection-diffusion equation on block-structured tetrahedral meshes. Here, we use a matrix-free approach, where the stencil-based kernels are partly generated from an abstract description in Python. The parallel performance of the implementation is demonstrated for Stokes systems with trillions ( $> 10^{12}$ ) of unknowns via the notion of textbook multigrid efficiency, and for advection dominated transport problems that are simulated with more than ten billion particles ( $> 10^{10}$ ). Both approaches are run on up to 147,456 parallel processes, and coupled via a predictor-corrector scheme to simulate buoyancy-driven convection in Stokes flow.

Nils Kohl  
University of Erlangen-Nuremberg  
nils.kohl@fau.de

### MS82

#### Towards Data-Driven Inference of Discrete Differential Operators

Partial differential equations (PDEs) are extensively used across scientific disciplines for modeling and describing various processes under consideration. Finite element, finite difference or similar methods allow to numerically solve them by discretizing the variables under consideration, e.g. space and time. In this discretized domain, differential operators can be approximated by matrices with nonzero coefficients for positions in the neighborhood of the considered point the so called stencil. We consider the inference of stencil weights for differential operators from linear, one-dimensional and two-dimensional PDEs using comprehensive regression techniques. Starting with the 1D case, we show that linear regression using an ordinary-least-squares (OLS) approach is able to recover mathematically meaningful stencils given a full-rank matrix of predictor variables. We discuss, how regularization techniques can allow the inference of the correct stencils even for rank-deficient matrices. The performance of the OLS-approach is compared for different orders of expansion of the differential operator. We discuss the impact of noise on the data in both predictor and predicted variables for various noise levels and compare different errors-in-variables approaches to mitigate the inherent consequences of noisy predictors. The presented techniques will be extended to the two dimensional case.

Yannis Schumann  
Helmut-Schmidt-University Hamburg  
Chair for High Performance Computing

yannis.schumann@hsu-hh.de

## MS82

### Learning Solution Adaptive Discretization Stencils from Data

Numerical methods for approximately solving partial differential equations (PDEs) are at the core of scientific computing. Often, this requires high-resolution or solution-adaptive discretizations in order to capture spatio-temporal features in the PDE solution. Systematically deriving solution-adaptive operator stencils is a current challenge and moreover requires one to know the PDE. In this talk, we present the STENCIL-NET, an artificial neural network for data-driven learning of problem- and resolution-specific local discretizations of unknown nonlinear PDEs. STENCIL-NET achieves numerically stable discretization of an unknown nonlinear PDE by spatially and temporally adaptive parametric pooling on regular Cartesian grids, and by incorporating knowledge about discrete time integration. Knowing the actual PDE is not necessary, as solution data is sufficient to train the network to learn the stencils. A once-trained STENCIL-NET can be used to predict solutions of the PDE on larger spatial domains and for longer times than it was trained for, addressing the problem of PDE-constrained extrapolation from data. We present numerical experiments on long-term forecasting of chaotic PDE solutions on coarse grids. We also present preliminary results on using differentiable convex layers to impose discrete stencil moment constraints, guaranteeing numerical consistency of the learned stencils.

Suryanarayana Maddu  
Technische Universität Dresden, Faculty of Computer Science  
Max Planck Institute of Molecular Cell Biology and Genetics  
maddukon@mpi-cbg.de

Dominik Sturm  
Center for Advanced Systems Understanding (CASUS)  
d.sturm@hzdr.de

Bevan Cheeseman  
ONI, Inc., Linacre House, Banbury Road, Oxford, OX2 8TA, UK  
bevan.cheeseman@gmail.com

Christian Müller  
Ludwig Maximilian University of Munich, Department of Statis  
Center for Computational Mathematics, Flatiron Institute  
christian.mueller@stat.uni-muenchen.de

Ivo F. Sbalzarini  
Max Planck Institute of Molecular Cell Biology and Genetics  
Technische Universität Dresden, Faculty of Computer Science  
ivos@mpi-cbg.de

## MS83

### A Comparison of Parallel Profiling Tools for Programs Utilizing the FFT

The overhead of the performance monitoring tools Craypat, FPMP, mpiP, Scalasca and TAU, are measured using default configurations likely to be chosen by a novice user

and shown to be small when profiling Fast Fourier Transform based solvers for the Klein Gordon equation based on 2decomp&FFT and on FFTE. Performance measurements help explain that despite FFTE having a more efficient parallel algorithm, it is not always faster than 2decomp&FFT because the compiled single core FFT is not as fast as that in FFTW which is used in 2decompFFT. This study was conducted on a Cray XC40 Shaheen II and will be implemented on Fugaku supercomputer as well and results will be displayed in this talk with the behaviour comparison of both platforms.

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

Benson Muite  
Kichakato Kizito  
benson\_muite@emailplus.org

## MS83

### Updates on Sequential and Parallel FFTX

We present an update on the design of both the sequential and parallel API of FFTX. The software package is developed as part of the DOE ExaScale effort by LBL, Carnegie Mellon University, and SpiralGen, Inc. We aim at translating the LAPACK/BLAS approach from the numerical linear algebra world to the spectral algorithm domain. FFTX is extending and updating FFTW for the exascale era and beyond while providing backwards compatibility. Patterns of FFTX call sequences capture higher level spectral algorithms and their variants, including convolutions, Poisson solvers, correlations, and numerical differentiation approaches that translate to FFT calls. The SPIRAL code generation and autotuning system—now available as open source under a BSD/Apache license—underpins the effort to provide performance portability.

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

## MS83

### Beyond 2D Parallelization of Multi-Dimensional FFTs

Any multi-dimensional Fourier transform is typically decomposed as multiple 1D transforms applied in the different dimensions of the data set. Hence, most Fourier packages that implement parallel multi-dimensional Fourier transforms, focus on the parallelism across the 1D transforms. However, most of the existing frameworks do not offer flexible solutions to exploit the parallelism within the 1D FFT and also do not allow for changes of the data layout to exploit the parallelism. In the era of exascale, with systems with thousand of nodes and intricate network topologies, flexibility and parallel efficiency are key aspects all multi-dimensional DFT frameworks need to have in order to map and scale the computation appropriately. In this work, we present a framework that enables the development of a family of parallel FFT algorithms by 1) using different data layouts to distribute the data across the network, 2) exploiting both parallelism schemes, and 3) unifying the two parallelization schemes within a single framework. We show that the flexibility of choosing between different parallel multi-dimensional DFT algorithms allows for almost linear strong scaling results for problem sizes of  $1024^3$  on two supercomputers, namely the RIKEN's

K-Computer and Oakridge's Summit.

Doru Thom Popovici  
Lawrence Berkeley National Laboratory  
Lawrence Berkeley National Laboratory  
dtpopovici@lbl.gov

Martin D. Schatz  
Department of Computer Sciences  
The University of Texas at Austin  
martin.schatz@utexas.edu

Franz Franchetti, Tze Meng Low  
Carnegie Mellon University  
franzf@andrew.cmu.edu, lowt@andrew.cmu.edu

### MS83

#### Parallel Implementation of FFT in a Finite Field

In this talk, we propose a parallel implementation of fast Fourier transform (FFT) in a finite field. The butterfly operation of the FFT in a finite field includes modular multiplication, which can be computed faster by using Montgomery multiplication. We vectorized the butterfly operations using SIMD instructions and parallelized them using OpenMP. Performance results of parallel FFT in a finite field on manycore processors are reported.

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

### MS84

#### Narrow Band-Based Dynamic Load Balancing for Multiphase Flow Simulations

This talk presents a dynamic load balancing scheme for compressible two-phase flow simulations with the level-set ghost-fluid method. The proposed load balancing scheme is designed generally for dynamic workload variations on unstructured meshes using space-filling curves (SFCs). In case of the level-set ghost-fluid method, the workload imbalance arises from applying the costly interface-capturing algorithm only to the grid cells near the phase interface. To this end, we introduce an element masking (narrow band), that also enables an efficient code instrumentation through accurate, element-local wall time measurements. We compare state-of-the-art partitioning algorithms for SFCs and investigate the strong scaling behavior for up to  $O(1e6)$  cores. The results demonstrate near optimal parallel efficiency and an average performance gain of factor ten compared to previous, unbalanced simulations. The load balancing scheme is applied to simulate complex droplet dynamics, including the well-studied interaction of a shock wave with a singular droplet.

Daniel Appel  
University of Stuttgart  
daniel.appel@iag.uni-stuttgart.de

### MS84

#### Simulating Acoustic and Gravity Waves in ForestClaw

We present recent progress in the development the software platform ForestClaw, a parallel, dynamically adaptive library for solving PDEs on a nested hierarchy of logically

Cartesian grids. In this talk, we discuss software engineering challenges related to extending ForestClaw with the MAGIC (Model for Acoustic and Gravity wave Interactions and Coupling, J. Snively, Embry-Riddle, FL) and GEMINI (Geospace Environment Model Ion-Neutral Interactions, M. Zettergren, Embry-Riddle, FL) codes. A key challenge in this effort, funded as part of the DARPA AtmoSense program, is to factor existing MAGIC and GEMINI codes so they can take advantage of the adaptive capabilities in ForestClaw. A second challenge is to facilitate runtime exchange of solution data between GEMINI and MAGIC. We will show preliminary results from the AirWaves project, as well as code development with the core ForestClaw platform needed to support AirWaves.

Donna Calhoun  
Boise State University  
donna.calhoun@boisestate.edu

### MS84

#### High-Order Accurate Embedded-Boundary Discontinuous Galerkin Methods for High-Performance Computing Applications

This talk presents a numerical framework for high-performance computing applications based on high-order accurate embedded-boundary methods. The framework is based on discontinuous Galerkin formulations and suitable quadrature rules for embedded geometries, whose combined use allows resolving the curved geometries as well as solving the governing equations and enforcing boundary and interface conditions with high-order accuracy. The method uses block-structured Cartesian grids where the geometries are represented implicitly via a level set function. The intersection between the grids and the level set function generates the implicitly defined mesh, which consists of a collection of regular (hyper)rectangular cells plus a relatively small number of irregular cells cut by the embedded boundaries. High-order accuracy is achieved by using high-order quadrature rules for implicitly-defined domains and boundaries, while a cell-merging strategy addresses the small-cell problem. The space-discretized equations are then advanced in time using explicit high-order Runge-Kutta algorithms. The framework is implemented within AMReX, a software library that provides functionalities to write massively parallel applications using adaptive mesh refinement. Numerical test cases are presented in the areas of wave propagation and computational fluid dynamics and show the performance of the framework using various parallelization paradigms including general purpose GPU computing.

Vincenzo Gulizzi  
LBNL  
vgulizzi@lbl.gov

### MS84

#### Hierarchical Mesh Based Coupling of Solution Methods for the Simulation of Multiphase Flows

The simulation of multiphase flows based on Eulerian-Eulerian or Eulerian-Lagrangian methods on high performance computers requires an efficient coupling and a simultaneous execution of the numerical solvers for the fluid and particles or gas phase. Especially, communication overhead and load imbalances can lead to a low computational efficiency. In this presentation an efficient coupling strategy based on hierarchical Cartesian meshes is presented that enables large-scale multiphase simulations. A method that

has previously been successfully used for a multiphysics application, i.e., a computational fluid dynamics method coupled with a computational aeroacoustics method, is extended to the modeling of gas-liquid bubbly flows. The accurate description of such multiphase flows is essential for process engineering or the machinery industry. Simulations of a bubbly turbulent channel flow, which mimics the transport phenomena in an electrochemical machining (ECM) process, demonstrate the applicability of the method.

Daniel Lauwers, Julian Vorspohl  
RWTH Aachen University  
d.lauwers@aia.rwth-aachen.de,  
j.vorspohl@aia.rwth-aachen.de

Matthias Meinke  
Institute of Aerodynamics  
RWTH Aachen University  
m.meinke@aia.rwth-aachen.de

Wolfgang Schröder  
RWTH Aachen University  
office@aia.rwth-aachen.de

#### MS85

##### **Architectures View Point of Hardware-Software Co-Design for Heterogeneous Computing**

ISA has been an excellent abstract to decouple microarchitecture and applications. However, rapid progress in new circuit technologies and expanding applications make the boundaries become a blur. Furthermore, new technologies such as processing near memory demand a new set of ISAs to reflect the computation changes. In this talk, Ill present a few examples that show the trade-off between software implementation vs. hardware support. First, Ill focus on processing near memory. When is good to offload computation on PNM can be studied with both analytical models and also cycle-level simulation. Second, Ill discuss the design trade-off on heterogeneous computing platforms especially CPU+GPU architectures. Ill also discuss how to abstract the applications and architecture to estimate the benefits of new ISA support.

Hyesoon Kim  
Georgia Institute of Technology  
hyesoon@cc.gatech.edu

#### MS85

##### **EPEEC: The European Joint Effort Towards Highly Productive Programming for Heterogeneous HPC**

The European joint Effort toward a Highly Productive Programming Environment for Heterogeneous Exascale Computing (EPEEC) is a Horizon 2020 project from the European Commission involving 10 partners, which include research centres, universities, and SMEs (<https://epeec-project.eu>). Started 2018 and already on its last mile, the partners have been advancing European technology in programming environments specifically addressed to turn upcoming overwhelmingly complex supercomputers into manageable platforms for domain application developers. EPEECs components include tools, programming models, and applications. After having released intermediate software prototypes, EPEEC will soon release final software

components.

Antonio J. Pena  
Barcelona Supercomputing Center  
antonio.pena@bsc.es

#### MS85

##### **Magma: Towards High Performance and Cross Platform Algorithms for Numerical Linear Algebra**

MAGMA is an open source library for dense and sparse linear algebra algorithms on heterogeneous systems with GPUs. It provides many important algorithms such as optimized BLAS and LAPACK for GPUs, including dense matrix factorizations, linear system solvers, least square problems, eigenvalue problems, batch linear algebra, as well as sparse iterative solvers. With an emphasis on performance, MAGMA provides many routines that often outperform those provided by the vendor's numerical software. This talk provides an overview of the latest heterogeneous computing algorithms in MAGMA, and highlights the recent efforts for adding support and optimizations for AMD GPUs. The talk also presents preliminary performance results of new developments that are lined up for future releases.

Stanimire Tomov  
University of Tennessee  
tomov@icl.utk.edu

Ahmad Abdelfattah  
University of Tennessee, Knoxville  
ahmad@icl.utk.edu

#### MS86

##### **Task-Based Parallel Programming for Scalable Algorithms : Application to Matrix Multiplication**

Task-based programming models have succeeded in gaining attraction in the numerical linear algebra community thanks to how they have relieved part of the burden of building portable distributed-memory parallel algorithms. In increasingly larger, more heterogeneous clusters of computers, these models appear as a way to maintain and enhance more complex algorithms. However, task-based programming models lack flexible, elegant expression for scalable algorithms such as Communication Avoiding algorithms. We show that the Sequential Task Flow paradigm can be extended to write a compact yet efficient scalable General Matrix Multiplication. This extension required fewer modifications to the StarPU runtime system. The final implementation is competitive with state-of-the-art libraries as it outperforms them in application-specific scenarios or on irregular logical computing grids, using homogeneous multicore clusters.

Emmanuel Agullo  
INRIA  
emmanuel.agullo@inria.fr

Alfredo Buttari  
CNRS-IRIT-Université de Toulouse, France  
alfredo.buttari@irit.fr

Abdou Guermouche  
LaBRI-INRIA futurs  
abdou.guermouche@labri.fr

Julien Herrmann

IRIT, Université de Toulouse, CNRS  
julien.herrmann@irit.fr

Antoine Jegou  
IRIT, INPT  
antoine.jego@toulouse-inp.fr

### MS86

#### Meeting the Computational Challenges of Giant Optical Telescopes' Real-time Control

We present a distributed-memory implementation of a real-time controller for giant optical telescopes. Based on task-based programming models and powered by dynamic runtime systems, we leverage the computational power of the underlying massively parallel hardware resources and scale the computational astronomy application to an unprecedented level. The newly developed numerical algorithm increases the performance of the adaptive optics control loop by continuously outsmarting the atmospheric turbulence in real-time. This hardware-software solution is essential to obtain sharp images with current and future giant telescopes under construction. It may eventually help to better detect planets around other stars by pushing the limits of the largest astronomical telescopes on Earth.

Jesse Cranney  
The Australian National University  
jesse.cranney@anu.edu.au

Damien Gratadour  
Observatoire de Paris  
damien.gratadour@obspm.fr

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

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

Hatem Ltaief  
King Abdullah University of Science & Technology (KAUST)  
hatem.Ltaief@kaust.edu.sa

### MS86

#### Floating Point Overflow and the Parallel Solution of Triangular Equations

This talk is relevant to anyone interested in the development of libraries for dense numerical linear algebra. Standard linear systems, nonsymmetric standard and generalized eigenvalue problems and Sylvester matrix equations can all be solved by reducing the matrices to triangular form. Many algorithms for solving triangular equations are based on substitution. Unfortunately, substitution is very vulnerable to floating point overflow and more so when the representational range is reduced compared with IEEE double precision floating point arithmetic. LAPACK realizes a family of algorithms for solving triangular equations without suffering from floating point overflow. The subroutines are all sequential scalar codes that can only achieve a very small fraction of the peak flop rate. Recently, we have developed algorithms that are parallel and blocked. In this talk we exhibit well-conditioned triangular equa-

tions for which the solution exceeds the representational range. We explain how to prevent floating point overflow without sacrificing performance. The underlying principles have been integrated into the task-parallel library StarNEig for solving dense nonsymmetric standard and generalized eigenvalue problems.

Carl Christian Kjelgaard Mikkelsen  
Umeå University  
Department of Computing Science  
spock@cs.umu.se

### MS86

#### Template Task Graph (TTG): Flow Graph Programming for Fine-Grained Distributed Memory Parallelism

TTG is a general-purpose programming model for shared and distributed memory parallel computers. It is particularly suited for expressing irregular data-driven computation patterns (numerical calculus on adaptively refined meshes, block sparse tensor algebra, and others) that appear in many domains of science and engineering. TTG represents a program as a graph of template tasks (hence the name); flow of key,value pairs over the graph generates tasks that are executed by a low-level task runtime. By abstracting the low-level task runtime TTG aims to offer high-level declarative composition with portable execution on distributed memory and heterogeneous platforms. This talk will highlight the initial implementation of TTG (as a C++17 library, abstracting PaRSEC and MADNESS distributed-memory tasking runtimes) and its use for implementing prototypical algorithms in linear algebra and numerical calculus.

Edward Valeev  
Virginia Tech  
Department of chemistry  
evaleev@vt.edu

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

George Bosilca  
University of Tennessee Knoxville  
bosilca@icl.utk.edu

Thomas Herault  
Univ. of Tennessee - Knoxville  
herault@icl.utk.edu

Joseph Schuchart  
ICL  
University of Tennessee Knoxville  
schuchart@icl.utk.edu

Poornima Nookala  
Stony Brook University  
poornimavinaya.nookala@stonybrook.edu

### MS87

#### Particle Resampling and Mesh Quasi-Interpolation in Delta-F Methods

In  $\delta f$  PIC plasma simulations, particle weight growth may lead to significant error, while marker particle motion can

lead to both inadequate phase space resolution and poor load balancing in large-scale simulations. We introduce a class of particle resampling techniques that address these issues, detailing the effective selection of new particle phase space position and culling low-impact particles. In addition, we relate the resampling procedure to the general problem of particle-mesh projection, and examine how different selections impact weight positivity, projection accuracy, phase space locality, and moment preservation. We discuss modifications to cater to code coupling and simulations using a total- $f$  approach. Finally, we illustrate the method at scale, using XGC, a global 5D total- $f$  gyrokinetic PIC code.

Varis Carey  
Theodon Consulting  
varis.carey@theodonconsulting.com

Albert Mollén, Robert Hager  
Princeton Plasma Physics Laboratory  
amollen@pppl.gov, rhager@pppl.gov

### MS87

#### Conservation and Monotonicity in Plasma Kinetic Simulations

The Landau collision operator provides an accurate model for small-angle dominated Coulomb collisions in hot plasmas. Here we review a new discrete gradient dependent, marker-particle scheme for simulating the standard Landau operator. The inclusion of discrete gradients, a first integral preserving integrator historically used for Hamiltonian systems, ensures conservation of mass, momentum, and kinetic energy, as well as preservation of the positive semi-definite production of entropy in long-time simulations. This new scheme was tested on standard test problems, such as Landau Damping and Two-Stream Instability, using particle-in-cell (PIC) simulations written using the PETSc library.

Daniel Finn  
University at Buffalo  
dsfenn@buffalo.edu

### MS87

#### Ultracold Neutral Plasmas: An Ideal Testbed for Numerical Plasma Models

Ultracold neutral plasmas (UCNPs), created by photoionization of a cold gas, are an excellent platform for studying neutral plasmas in far more complex environments such as plasma in the Sun's atmosphere, white dwarf stars, and inertial-confinement fusion devices. UCNPs have clean and controllable initial conditions, with ion temperatures below 1 K, tunable electron temperatures from 1 – 1000 K, and densities ranging from  $10^6 - 10^{12} \text{ cm}^{-3}$ , and precise diagnostics and have been used to benchmark theories and numerical techniques for molecular-dynamic, kinetic, and hydrodynamic descriptions of the plasma, including those that capture strongly coupled physics. Magnetized UCNPs are of current interest because of the interplay of magnetization and strong coupling, connection to plasma confinement, and modification of recombination dynamics in strong fields. We recently demonstrated the magnetic confinement of UCNPs at the null of a biconic cusp, or quadrupole magnetic field. In this talk, I will discuss what makes UCNPs amenable to a wide variety of computational modeling techniques and will overview our recent progress developing a two-fluid magnetohydrody-

dynamic code to model UCNP evolution in the presence of quadrupole magnetic fields. Research supported by NSF/DOE Partnership in Basic Plasma Science and Engineering through the NSF Award Number 2107709 and the NSF Graduate Research Fellowship Program under Grant No. 1842494.

Grant Gorman, Gregory Szytko, MacKenzie Warrens,  
Stephen Bradshaw, Thomas Killian  
Rice University  
gg25@rice.edu, gms12@rice.edu, mkw9@rice.edu,  
stephen.bradshaw@rice.edu, killian@rice.edu

### MS87

#### Modeling the Solar Corona: Magnetohydrodynamics and Beyond

Understanding the Sun's corona, its outer layer consisting of high-temperature, low-density magnetized plasma, is crucial to understanding the Sun's complex interactions with the rest of the Solar System. Magnetohydrodynamics (MHD), where the plasma is treated as a conducting fluid, is a common framework for modeling behavior in the corona. However, additional mechanisms of note in coronal plasmas, such as electron-mediated thermal conduction and magnetic reconnection, often necessitate treatment of temporal and spatial scales outside of the strict MHD regime. This necessity of multi-scale treatment poses a challenge from a computational standpoint. In this presentation, we discuss our current efforts at numerically modeling the solar atmosphere in an MHD framework, as well as the potential for integrating kinetic-scale physics to better capture the behavior of the corona.

Gregory Szytko, Stephen Bradshaw, Grant Gorman  
Rice University  
gms12@rice.edu, stephen.bradshaw@rice.edu,  
gg25@rice.edu

### MS88

#### A Communication-Centric Approach for Designing Flexible Accelerators

Deep Neural Networks (DNN) have demonstrated highly promising results across computer vision and speech recognition, and are becoming foundational for ubiquitous AI. The computational complexity of these algorithms and a need for high energy-efficiency has led to a surge in research on hardware accelerators. To reduce the latency and energy costs of accessing DRAM, most DNN accelerators are spatial in nature, with hundreds of processing elements operating in parallel and communicating with each other directly. DNNs are evolving at a rapid rate - leading to myriad layer/operator types of varying shapes. Given a DNN there can be myriad computationally efficient implementations (e.g., via pruning) - leading to structured and unstructured sparsity. Finally, a given DNN can be tiled and partitioned in myriad ways to exploit data reuse. All of the above can lead to irregular dataflow patterns within the accelerator substrate. Getting high mapping efficiency for all these cases is highly challenging in accelerators today that are often tightly coupled 2D grids with rigid near-neighbor connectivity. In this talk, we will demonstrate a systematic methodology for understanding data reuse opportunities within the algorithm, determine the cost vs benefit for efficiently exploiting them in hardware and a communication-centric methodology for accelerator design, that can provide  $\sim 100\%$  mapping efficiency

for arbitrary DNNs shapes, sparsity ratios and mappings.

Tushar Krishna

Georgia Institute of Technology  
tushar@ece.gatech.edu

## MS88

### Exploiting Parallelism in Large Scale Deep Learning Model Training: From Chips to Systems to Algorithms

We live in a world where hyperscale systems for machine intelligence are increasingly being used to solve complex problems ranging from natural language processing to computer vision to molecular modeling, drug discovery and recommendation systems etc. A convergence of breakthrough research in machine learning models and algorithms, increased accessibility of hardware systems at cloud scale for research and thriving software ecosystems are paving the way for an exponential increase in model sizes. Effective parallel processing techniques and large clusters of accelerators will be required to train these models economically. What are some of the scale challenges we face, given the enormous cost and time to train tera-scale models? What are the compute, memory and networking requirements to implement these models? How can a novel approach with fine-grained parallelism at the hardware level help? What kind of model decomposition techniques are required to exploit the innate parallelism in deep learning models? Attend this session to learn about how Graphcore aims to address these challenges. Get to know our Intelligent Processing Unit (IPU) - a purpose-built hardware accelerator with a unique MIMD architecture - designed to address the most demanding compute and memory bandwidth needs of modern ML models. Our network disaggregated architecture uniquely positions us to build highly scalable systems (IPU-PODs) with thousands of accelerators aimed at exploiting parallelism.

Saurabh Kulkarni

Graphcore  
saurabhk@graphcore.ai

## PP0

### A Parallel Finite Element Solver for Skeletal Muscle Deformation

Recent advances in understanding how skeletal muscle deforms have led to a new, more informative multiphysics model that views muscles as composite biomaterials, drawing a connection to the deformation of nonlinear solids and active materials. This model yields a three-dimensional, highly nonlinear system of partial differential equations (PDEs). On the other hand, the presence of important physiological entities across multiple length scales and highly heterogeneous material coefficients require finite element meshes to be particularly refined. Therefore, traditional computers may not be able to store all this information, or it may take too long to solve the resulting linear systems of equations. In this work, a parallel solver for skeletal muscle deformation is developed. The code is designed to run primarily on a distributed architecture and contains Message Passing Interface (MPI) capabilities. In particular, the decomposition of the mesh is achieved through the p4est library, while the linear algebra components are handled using the Trilinos library. The core of the code has been developed using deal.II, a C++ library for the finite element solution of PDEs. Scaling analyses show that this distributed implementation represents

a step above an existing shared-memory implementation in which the design pattern WorkStream is used to enable multithreading.

Javier Almonacid

Simon Fraser University  
javiera@sfu.ca

Nilima Nigam

Dept. of Mathematics  
Simon Fraser University  
nigam@math.sfu.ca

James Wakeling

Department of Biomedical Physiology and Kinesiology  
Simon Fraser University  
wakeling@sfu.ca

## PP0

### Sundials: Suite of Nonlinear and Differential/Algebraic Equation Solvers

SUNDIALS is a suite of robust and scalable integrators and solvers for systems of ordinary differential equations, differential-algebraic equations, and nonlinear equations. The suite consists of six independent packages, CVODE(S), ARKODE, IDA(S), and KINSOL, and is designed to be easily incorporated into existing simulation codes with minimal information from the user. In this poster, we overview the capabilities of the SUNDIALS suite, discuss recent results from applications codes ranging from additive manufacturing to combustion and fusion energy, and highlight efforts to add performance instrumentation and measurement capabilities as well as automated performance regression testing. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Cody J. Balos

Lawrence Livermore National Laboratory  
balos1@llnl.gov

## PP0

### Experiments Towards Parallel Adaptive In-Situ Visualization

The visualization of simulation data that arises from large-scale numerical simulations is a challenging task. Visualization as post-processing step has the disadvantage of writing large amounts of data to disk, which can be slow to the point of impracticality. Relying on a dedicated third-party library for visualization often incurs duplicating data or converting it to a prescribed external data structure, as well as a sizable increase in code, executable, and memory complexity. We develop algorithms to visualize the simulation data using the simulation data structure. Specifically, we work with the widely known distributed adaptive octree structure and extend it to support in-situ visualization. One current approach of ours is to revive the radiosity method, which requires only a re-projection of the radiosity results to visualize the scene from a different view point in the rendered scene. Radiosity works by computing the pairwise light transfer between surface patches in a scene, which we determine by enforcing strict local radiation balance. Our algorithm exploits the octree-based data structure by using recursive top-down search algorithms and recursively excluding non-visible surface patches. The AMR data is distributed using the Morton space-filling curve,

which we use to parallelize the radiosity system setup and solver. We present the current state of our research, a parallel and natively supported radiosity solver operating on a distributed octree data structure.

Tim Griesbach  
University of Bonn  
tim.griesbach@uni-bonn.de

Carsten Burstedde  
Universität Bonn  
burstedde@ins.uni-bonn.de

## PP0

### Scalable Fine-Grid Empirical Performance Modeling via Tensor Completion

Rising architectural complexity demands aggressive kernel parameterization to achieve performance portability across an increasingly diverse array of architectures. Latent patterns in kernel performance are often localized across configuration spaces and consequently difficult to model. We observe that many multivariate kernels feature execution times that exhibit low-rank structure. We therefore develop performance models that characterize a multivariate kernel's performance as a low-rank tensor represented explicitly as a summation of multi-dimensional outer products (i.e., CP Decomposition). We leverage tensor completion to learn these expansions from a sparse dataset, which generalizes parametric linear regression performance models by implicitly capturing interactions and (nonlinear) correlations among kernel parameters. Systematic increase in tensor rank improves prediction accuracy, which we show to be necessary to model nonlinear behavior present at fine-grained scales. We demonstrate that the enhanced scalability offered by low-rank representations enables performance prediction accuracy to exceed that achieved by state-of-the-art regression methods across a suite of kernels.

Edward Hutter  
University of Illinois at Urbana-Champaign  
Department of Computer Science  
hutter2@illinois.edu

Edgar Solomonik  
University of Illinois at Urbana-Champaign  
solomon2@illinois.edu

## PP0

### On Centrality Resilience of Complex Networks

The resilience of a complex network relates to how well certain of its properties are retained under attacks. In this paper, we study centrality resilience, that is, resilience to path-based centralities: closeness and betweenness centralities. Our goal is to develop attack models that can disrupt the rank of the top path-based centrality nodes. Attack models are useful in disrupting flow through a network with minimum disruption, such as in avoiding the spread of epidemics by vaccinating/quarantining super-spreaders. However, finding high centrality nodes in a large network, although polynomial time is extremely expensive for large-scale networks. We show that there are several expander-like subgraphs that are embedded in complex networks, and influential nodes with high centrality are located within these expander graphs. We present an efficient method for identifying these regions and thereby the influential nodes, using snowball sampling and the core-

periphery structure of the sampled graph. We further demonstrate that the resilience of the network is dependent on the number and size of the expander-like subgraphs. Based on these observations we develop attack models to disrupt the flow through the network.

Fariba Afrin Irany, Sanjukta Bhowmick  
University of North Texas  
faribaafrinirany@my.unt.edu, sanjukta.bhowmick@unt.edu

## PP0

### Take-Away Impartial Combinatorial Game on Different Geometric and Discrete Structures

Following from the winning strategy for a Take-Away Impartial Combinatorial Game on only Oddly Uniform or only Evenly Uniform Hypergraphs in the Ph.D. Dissertation of Dr. Kristen Barnard (an Assistant Professor of Mathematics at Berea College), T. H. Molena found the new winning strategy for a Take Away Game on neither Oddly nor Evenly Uniform Hypergraphs during her Undergraduate Independent Research opportunity. However, those neither Oddly nor Evenly Uniform Hypergraphs must meet the given special requirements. In a Take-Away Game on hypergraphs, two players take turns to remove the vertices and the hyperedges of a hypergraph. In each turn, a player must remove either only one vertex or only one hyperedge. When a player chooses to remove one vertex, all of the hyperedges that contain the chosen vertex are also removed. When a player chooses to remove one hyperedge, only that one chosen hyperedge is removed. Whoever removes the last vertex wins the game. All of the new theorems in this research paper are in agreement with the previous theorems in Dr. Kristen Barnard's Ph.D. Dissertation.

T. H. Molena  
North Carolina State University  
thmolena@gmail.com

## PP0

### A Multiresolution Discrete Element Method for Triangulated Objects with Implicit Timestepping

Simulations of many rigid bodies colliding with each other sometimes yield particularly interesting results if the colliding objects differ significantly in size and are non-spherical. The most expensive part within such a simulation code is the collision detection. We propose a family of novel multiscale collision detection algorithms that can be applied to triangulated objects within explicit and implicit time stepping methods. They are well-suited to handle objects that cannot be represented by analytical shapes or assemblies of analytical objects. Inspired by multigrid methods and adaptive mesh refinement, we determine collision points iteratively over a resolution hierarchy, and combine a functional minimisation plus penalty parameters with the actual comparison-based geometric distance calculation. Coarse surrogate geometry representations identify no collision scenarios early on and otherwise yield an educated guess which triangle subsets of the next finer level potentially yield collisions. They prune the search tree, and furthermore feed conservative contact force estimates into the iterative solve behind an implicit time stepping. Implicit time stepping and non-analytical shapes often yield prohibitive high compute cost for rigid body simulations. Our approach reduces these cost algorithmically by one to two orders of magnitude. It also exhibits high vectorisation

efficiency due to its iterative nature.

Peter J. Noble  
University of Durham  
peter.j.noble@durham.ac.uk

Tobias Weinzierl  
Durham University  
tobias.weinzierl@durham.ac.uk

## PP0

### Some HPC-Specific C++ Extensions in Llmv

Even though DSLs gain popularity, C++ remains a dominant language for scientific computing in HPC. Many selling points of DSLs can directly be explained from four shortcomings of C++, where the language design aligns inaptly with the hardware. (i) The C++ language yields classes (structs) with a large memory footprint. (ii) The C++ language naturally leads towards an AoS data design as a class is the primary modelling entity for programmers. (iii) The C++ language does not offer built-in support for distributed memory parallelisation (MPI). (iv) The C++ language lacks support for multifaceted (floating point) data precision. We present a pragma-based LLVM extension which allows computational scientists to annotate their programs with data layout, MPI and precision information. Our compiler extension then yields native MPI and C++ code. This greatly improves the productivity of numerical developers and hides realisation complexity.

Tobias Weinzierl  
Durham University  
tobias.weinzierl@durham.ac.uk

Pawel Radtke  
Computer Science, Durham University  
pawel.k.radtke@durham.ac.uk

## PP0

### HDBSCAN\* on GPUs: Can We Do It Better?

HDBSCAN\* is a recent clustering algorithm for a multi-dimensional data. It addresses several drawbacks of the original DBSCAN algorithm, such as requiring expert knowledge to tune the parameters or the inability to find clusters of different densities. Despite sharing a similar name, the underlying computation is entirely different: first constructing a minimum spanning tree; then forming a dendrogram, and finally producing a flat clustering. Parallelizing each part is a challenging task, more so on GPUs. We present novel algorithms for each phase, all running on GPUs. We are interested in applying HDBSCAN\* to detect clusters in a large cosmological data with tens of millions of data points (galaxies!). Therefore, we currently limit ourselves to the two and three-dimensional datasets for the experiments. We describe the bottlenecks we encountered, and our current solution for them. The question, however, remains: can we do any better?

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

## PP0

### Distributed Memory Tensor Completion For Gen-

## eralized Loss Functions

Tensor completion is a useful tool applied in many fields such as data science, machine learning, signal processing on large scale data. Recent works have shown that using a specific loss function based on insights about the data distribution of the tensor may lead to a better model for completion. We provide novel algorithms and systems infrastructure which enable efficient parallel implementation of algorithms for tensor completion with generalized loss functions. Specifically, we consider alternating minimization, coordinate minimization, and a quasi-Newton (generalized Gauss-Newton) method. We compare these methods to stochastic gradient descent method and show that using second order information leads to more accurate factors in relatively less time in a parallel setting. To make possible tensor completion for very sparse tensors, we introduce new multi-tensor primitives, for which we provide specialized parallel implementations and compare these to the state of the art. We provide microbenchmarking results on the Stampede2 supercomputer to demonstrate the efficiency of the new primitives.

Navjot Singh  
University of Illinois at Urbana Champaign  
navjot2@illinois.edu

Edgar Solomonik  
University of Illinois at Urbana-Champaign  
solomon2@illinois.edu

## PP0

### Location Estimation from the Ground Up: A New Approach to Teaching Least-Squares Algorithms

High-performance computing uses algorithms to resolve mathematical models of a wide range of phenomena. Courses and textbooks that focus on the algorithms and performance-engineering techniques often provide little motivation because explaining the models requires domain knowledge that many students lack. The poster will present a course and a new SIAM book that present a wide range of least-squares minimization problems and algorithms using models that require almost no background to understand, namely models that estimate location from observations. The problems and algorithms that are covered include linear least squares problem (both full rank and rank deficient), the QR algorithm and the SVD, non-linear least-squares, space-state (Kalman) problems, and mixed-integer least squares. The location-estimation problems that motivate the algorithms include essentially all the techniques used in GPS, as well as many advanced techniques used in other location-estimation systems. Location-estimation models discussed in the book can thus motivate many types of least-square minimization problems. The book can obviously also be used in courses that expose Mathematics and Computer Science undergraduates to statistical estimation techniques and algorithms without requiring any background in electrical engineering or statistics. The poster will spur lively discussions on teaching numerical algorithms, a topic that many SIAM PP attendees are highly interested in.

Sivan A. Toledo  
Tel Aviv University  
stoledo@tau.ac.il

## PP0

### Intelligently Offloading Tasks to Intelligent Hard-

**ware**

Many HPC applications suffer from bandwidth and latency constraints as well as lacking message progression. These problems become particularly severe for sophisticated scientific codes with tasks, multiscale methods, adaptive meshes, multiphysics, multi-numeric, and so forth, as they induce non-deterministic, non-homogeneous data exchange. Smart network devices can serve as a man-in-the-middle between the CPU and the network and thus could champion message transferral. We present an architecture where compute ranks offload tasks to Mellanox BlueFields which either compute those tasks themselves or, where appropriate, issue those tasks to underworked ranks elsewhere in the cluster in a way that is sensitive to both workload imbalances and bottlenecks such as network congestion. Novel hardware in the form of programmable network cards offer the opportunity to offload data transfer responsibilities and thus increase code reactivity with regards to unexpected messages and allow the compute ranks to focus on compute throughput. Our work leads into a generic library for nonpersistent load balancing that is orchestrated in the fabric of cluster supercomputers: it has a minimal API and can easily be integrated into simulation codes that phrase their computations in tasks.

Mark Turner, Tobias Weinzierl  
Durham University  
mark.turner@durham.ac.uk,  
tobias.weinzierl@durham.ac.uk

Holger Schulz  
University of Durham  
holger.schulz@durham.ac.uk

Philipp Samfass  
Technical University of Munich  
samfass@in.tum.de

**PP0****Interoperability of Resilience Strategies for Simplicity and Performance**

Resilience at high scales necessitates minimizing and localizing the impact of failures and their recovery. Current techniques for resilience in high performance computing (HPC) applications use highly efficient, well tested libraries, but rely on a global tear-down and re-launch which places an unnecessary burden on the recovery process. However, these resilience strategies are simple for developers to implement and are already present in many current applications. Work on fault tolerant MPI specifications - such as User Level Fault Mitigation (ULFM) - has enabled more localized recovery, but at the cost of high development costs. We present the integration of existing, high performance recovery libraries with Fenix, a library built on top of ULFM to simplify recovery and emulate a more typical global tear-down process without paying the costs. Using Fenix to bridge the gap between local and global recovery processes lets the users step-by-step localize recovery with a series of minor changes to their application code. This integration leads to better performance in failure recovery and serves as a simplified bridge to the types of software resilience that may be needed to enable performant and energy-efficient future hardware. [SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.]

Matthew Whitlock  
College of Computing

Georgina Institute of Technology  
mwhitlo@sandia.gov

Nic Morales  
Sandia National Labs  
nmmoral@sandia.gov

Hemanth Kolla, Jackson Mayo  
Sandia National Laboratories  
hnkolla@sandia.gov, jmayo@sandia.gov

Bogdan Nicolae  
MCS Division  
Argonne National Laboratory  
bnicolae@anl.gov

George Bosilca  
University of Tennessee Knoxville  
bosilca@icl.utk.edu

Keita Teranishi  
Sandia National Laboratories  
Livermore, California  
knteran@sandia.gov