CSE15 Abstracts

SIAM Conference on Computational Science and Engineering

Saturday, March 14 through
Wednesday, March 18, 2015
The Calvin L. Rampton
Salt Palace Convention Center
Salt Lake City, Utah, USA

Abstracts are printed as submitted by the authors.

Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Telephone: +1-215-382-9800 Fax: +1-215-386-7999
Conference E-mail: *meetings@siam.org*
Conference web: *www.siam.org/meetings/*
Membership and Customer Service:
800-447-7426 (US & Canada) or +1-215-382-9800 (worldwide)

IP1

Graph Data Analytics at Scale: Opportunities and Challenges

The four Vs of Big Data necessitate fundamentally different data analytics. A promising strategy toward understanding of a complex system's dynamics and function aims to extract features and relationships between them and to analyze how their evolution causes different functional system responses. Discovery and forecasting of patterns in such feature graphs can provide insights about the vulnerability of our nations energy infrastructure to disturbances, the spread of a cyber-security attack, or the anomalies in internode communication in high performance systems. This talk will present some opportunities and challenges in using this strategy for computational science and engineering applications.

Nagiza Samatova
North Carolina State University
Oak Ridge National Laboratory
samatova@csc.ncsu.edu

IP2

Model Reduction - Trouble with Scales?

Scientific and technological advances call for more and more complex models as well as systematic ways of complementing them by observational data. Despite the ever increasing computing capacity, ironically, the need for quantifiable model reduction concepts is also gaining increasing importance in numerous application contexts. Examples are large scale design or online optimization tasks, uncertainty quantification or inversion problems some of which may only become feasible through employing reduced models. Starting from a flow scenario with microscales this talk highlights several aspects of related model reduction strategies with particular focus on accuracy and stability guarantees, presence of small scales, singular perturbations, and high dimensionality. We address some of the key ingredients, revolving around error-residual relations, rate-optimality as a benchmark notion, adaptive or greedy methods, separation of variables. The discussion is illustrated by numerical examples.

Wolfgang Dahmen
RWTH Aachen
IGPM
dahmen@igpm.rwth-aachen.de

IP3

Petascale Finite Element Simulation of Real Worlds Complex Structure with Billions DOFs Model

Leading supercomputers offer the computing power of petascale, and exascale systems are expected to be available by the end of this decade. Supercomputers with more than tens of thousands of computing nodes, each of which has many cores cause serious problems in practical finite element software. We have been developing an open source parallel finite element software known as ADVENTURE, which enables very precise analyses of practical structures and machines using over 100 million to billions DOFs mesh. The basic parallel solution algorithms employed are the hierarchical domain decomposition method with balancing domain decomposition as preconditioner. In this talk, I explain several key technologies and one practical application, i.e. seismic response of nuclear power plant subjected to a strong earthquake.

Shinobu Yoshimura
The University of Tokyo
yoshi@sys.t.u-tokyo.ac.jp

IP4

Extreme-scale Multigrid in Space and Time

Multigrid methods are important techniques for efficiently solving huge linear systems and they have already been shown to scale effectively on millions of cores. Future exascale architectures will require solvers to exhibit even higher levels of concurrency (1B cores), minimize data movement, exploit machine heterogeneity, and demonstrate resilience to faults. While considerable research and development remains to be done, multigrid approaches are ideal for addressing these challenges. In this talk, we will discuss efforts to develop extreme-scale multigrid, including a new parallel time integration approach that has the potential for significant speedups over standard time stepping.

Robert Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

IP5

Statistical and Computational Challenges of Constraining Greenhouse Gas Budgets

Predicting future changes to the global carbon cycle (and therefore climate) and quantifying anthropogenic emissions of greenhouse gases (GHGs) both require an understanding of net GHGs emissions and uptake across a variety of spatial and temporal scales. This talk will explore some of the core scientific questions related to understanding GHG budgets through the lens of the statistical and computational challenges that arise. The focus will be on the use of atmospheric observations, and applications will include the natural and anthropogenic components of the methane and carbon dioxide budgets. The discussion will include issues related to the solution of spatiotemporal inverse problems, uncertainty quantification, data fusion, gap filling, and issues of "big data" arising from the use of satellite observations.

Anna Michalak
Carnegie Institution for Science
Stanford, CA
michalak@stanford.edu

IP6

Scaling Open Systems for Future Computational Challenges

Computational models are changing rapidly, partially in response to growing data size and advances in high-performance computing. Open approaches are well suited to this dynamic environment as they provide agile responses to complex, evolving code, and support the greater goal of ensuring reproducible science. This presentation introduces some open initiatives addressing big data and HPC and the role that software architectures and processes plays in advancing scientific computation. Also discussed are emerging trends including competitive challenges and active publications that will likely play an important role in the creation, development and deployment of computa-

tional software.

Will Schroeder
Kitware, Inc.
will.schroeder@kitware.com

IP7

A Calculus for the Optimal Quantification of Uncertainties

The past century has seen a steady increase in the need of estimating and predicting complex systems and making (possibly critical) decisions with limited information. With this purpose, this talk will describe the development of a form of calculus allowing for the (computational) manipulation of infinite dimensional information structures and its application to the optimal quantification of uncertainties in complex systems and the scientific computation of optimal statistical estimators/models. Specific examples will be discussed to illustrate how this form of calculus could also be used to facilitate/guide the process of scientific discovery.

Houman Owhadi
Applied Mathematics
Caltech
owhadi@caltech.edu

IP8

The Power of Matrix and Tensor Decompositions in Smart Patient Monitoring

Accurate and automated extraction of clinically relevant information from patient recordings requires an ingenious combination of adequate pretreatment of the data (e.g. artefact removal), feature selection, pattern recognition, decision support, up to their embedding into user-friendly user interfaces. The underlying computational problems can be solved by making use of matrix and tensor decompositions as building blocks of higher-level signal processing algorithms. A major challenge here is how to make the mathematical decompositions "interpretable' such that they reveal the underlying medically relevant information and improve medical diagnosis. The application of these decompositions and their benefits will be illustrated in a variety of case studies, including epileptic seizure onset localisation using adult and neonatal scalp EEG and Event-related potential analysis during simultaneous EEG-fMRI acquisition.

Sabine Van Huffel
ESAT-SCD(SISTA) Department of Electrical Engineering
Katholieke Universiteit Leuven
sabine.vanhuffel@esat.kuleuven.ac.be

IP9

Implications of Numerical and Data Intensive Technology Trends on Scientific Visualization and Analysis

Technology trends in numerically and data intensive computing have the potential to reshape and significantly advance how we visualize and analyze the results of scientific simulations. However, next generation numerically intensive supercomputers are bound by power and storage constraints. These require us to transition from standard post-processing visualization and analysis approaches to intelligent, automated in-situ ones. In addition, data intensive technology trends that support accessing and understanding our data using intuitive, web-based and query-driven interfaces are now the norm. In this talk, I will discuss these trends and several freely available, open-source approaches that leverage them.

James Ahrens
Los Alamos National Laboratory
ahrens@lanl.gov

SP1

Celebrating 15 Years of SIAM CSE

There can be no doubt that SIAM CSE has been a big success! We examine the growth of CSE in SIAM, and more broadly as a discipline, and look toward some of the challenges and opportunities for the future.

Linda R. Petzold
University of California, Santa Barbara
petzold@engineering.ucsb.edu

CP1

Computational Molecular Engineering: An Emerging Technology in Process Engineering

Molecular modeling and simulation has become a powerful tool which can be applied to many physical processes and properties of fluids on the molecular level. A shift in the accessible length and time scales due to massively parallel high-performance computing has greatly increased its potential. The novel molecular dynamics code ls1 mardyn, which scales excellently on up to 146 000 cores, is presented, highlighting the emergence of computational molecular engineering as a discipline.

Martin T. Horsch, Stephan Werth
University of Kaiserslautern, Germany
Laboratory of Engineering Thermodynamics
martin.horsch@mv.uni-kl.de, stephan.werth@mv.uni-kl.de

Christoph Niethammer, Colin Glass
High Performance Computing Center Stuttgart, Germany
niethammer@hlrs.de, glass@hlrs.de

Wolfgang Eckhardt, Philipp Neumann
Technische Universität München, Germany
Scientific Computing in Computer Science
eckhardw@in.tum.de, neumanph@in.tum.de

Hans-Joachim Bungartz
Technische Universität München, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

Jadran Vrabec
University of Paderborn, Germany
Thermodynamics and Energy Technology
jadran.vrabec@upb.de

Hans Hasse
University of Kaiserslautern, Germany
Laboratory of Engineering Thermodynamics

hans.hasse@mv.uni-kl.de

CP1

A Numerical and Computational Framework for Hierarchical Multi-Scale/multi-Physics Simulations

Multi-scale modeling (MSM) has become a dominant paradigm in materials modeling. The practical impact of MSM depends on its ability to utilize modern computing platforms. However, since there are no general numerical and computational frameworks for MSM, the vast majority of multi-scale material models or simulations are developed on a case-by-case basis. We present a formulation of an adaptive numerical and computational framework for MSM and analyze its performance for a number of challenging problems.

Jaroslaw Knap, Oleg Borodin, Carrie E. Spear, KENNETH W. Leiter, DAVID A. Powell, RICHARD C. Becker
U.S. Army Research Laboratory
jaroslaw.knap@arl.army.mil, oleg.a.borodin.civ@mail.mil, carrie.e.spear.ctr@mail.mil, kenneth.w.leiter.civ@mail.mil, david.powell@survice.com, richard.c.becker.civ@mail.mil

CP1

Amr Strategies for Scft Algorithm

We introduce an adaptive mesh refinement technique for solving the self consistent field theory mean-field optimization for an AB diblock copolymer. We use a 3D octree data structure and a level set based refinement method to solve a diffusion reaction equation. It reduces the required number of points for fine polymeric 3D structures and thus the data required to store is compressed consequently (1/3) without loss of accuracy on the physical observables.

Gaddiel Ouaknin
UCSB
gaddyouaknin18@gmail.com

Frederic G. Gibou
UC Santa Barbara
fgibou@engineering.ucsb.edu

CP2

Robust Multigrid Methods for Magnetohydrodynamics

Magnetohydrodynamic models are used for a wide range of plasma physics applications. The system of PDEs that characterizes these models is nonlinear, with strongly coupled fluid and electromagnetic interactions. The linear systems that result from linearization and discretization are typically difficult to solve. We consider multigrid-preconditioned GMRES to achieve efficient solution, and compare results for two potential relaxation methods, both motivated by well-known relaxation techniques for incompressible fluid dynamics.

Thomas Benson
Department of Mathematics
Tufts University
thomas.benson@tufts.edu

CP2

Dependance of the Convergence of Multigrid Methods on the Used Discretization

A lot of effort has been put into analyzing the convergence rate of multigrid methods depending on the used smoother and coarse grid operator. Effectively, the convergence rate of the solver is already being influenced when the discretization of the PDE is chosen. In order to solve PDEs as computationally efficient as possible, the choice of the discretization matters. We present analyses of the convergence rate and computational cost of multigrid methods for different discretizations.

Matthias Bolten
University of Wuppertal
Department of Mathematics
bolten@math.uni-wuppertal.de

CP2

Support Graph Smoothing

Large scale-free data mining tasks often require an efficient linear solver. Often standard iterative methods struggle do to the large size and complexity of the graphs. Computation of an optimal preconditioner for CG is challenging for a general scale-free graph. Support graph preconditioners have been a popular subject of study, but deserve more in depth study within a multilevel setting. We employ a support graph technique that serves as a relaxation method for AMG.

Alyson Fox
University of Colorado Boulder
fox.alyson@gmail.com

CP2

On Teh Efficiency of Nonlinear Multigrid Methods

Nonlinear multigrid methods such as the Full Approximation Scheme (FAS) and Newton-multigrid (Newton-MG) are widely-used as fast solvers for nonlinear PDEs of elliptic and parabolic type. This talk will consider Newton-MG and FAS iterations in a general setting to derive a theoretical approximation of the execution time of the algorithms. Specific examples will then be used to demonstrate the sharpness of these estimates for a range of nonlinear eliiptic and parabolic problems (the latter with implicit time-stepping). Our conclusion is that the Newton-MG approach is almost always the superior choice.

Peter K. Jimack
School of Computing
University of Leeds
p.k.jimack@leeds.ac.uk

CP2

Large-Scale Sparse Inverse Covariance Estimation

The Sparse Inverse Covariance Estimation problem arises in many statistical applications in Machine Learning. In this problem we estimate a sparse inverse of a covariance matrix of a multivariate normal distribution, by solving an l1 regularized optimization problem. Because of memory limitations, most algorithms are unable to handle large

scale instances of this problem. We present a block coordinate descent approach and a multilevel acceleration for solving the problem for such large-scale data sets. We further show that an additional debiasing phase improves the estimated matrix.

Eran Treister
Technion - Israel Institute of Technology, Israel
eran@cs.technion.ac.il

Javier Turek
Computer Science Department
Technion
javiert@cs.technion.ac.il

Irad Yavneh
Computer Science Department, Technion
irad@cs.technion.ac.il

CP3

CSE Education at JSC

Fostering a sound education of students and young researchers at bachelor, master and PhD level in high-performance computing, mathematics and computational science is an essential task of the Jülich Supercomputing Centre (JSC). This talk will give an overview of the educational activities and structures at JSC and informs about the guest student programme, the summer/winter schools for PhD students, joint bachelor and master courses with universities and the German Research School for Simulation Sciences (GRS).

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

CP3

Automatic Parallel Programming for Scientific Simulation

HiPro is an automatic parallel programming IDE designed for developing scientific simulation based on JASMIN, a domain-specific computational framework. It supports parallel programming through GUI and source code generation.The unique parallel part and all interfaces of the application are generated and implementation of sequential subroutines is the only part of the code left to be written manually for a programmer. It combines numerical mathematics with component-based programming to create ontological models for parallel simulations.

Li Liao
IAPCM
liliao@iapcm.ac.cn

Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
zhang_aiqing@iapcm.ac.cn

Zeyao Mo
Laboratory of Computational Physics, IAPCM
P.O. Box 8009, Beijing 100088, P.R. China
zeyao_mo@iapcm.ac.cn

CP3

Symbolic Representation and Automated Code Generation for Discontinuous Galerkin Finite Element Methods

We consider the application of discontinuous Galerkin finite element methods for the discretization of general fluid flow and multiphysics problems. By exploiting tools from symbolic differentiation, we present a simple programming environment which automatically generates the necessary code segments, leading to rapid development and testing of discontinuous Galerkin methods for a wide variety of problems. The proposed computational framework is demonstrated on a variety of problems arising in both incompressible and compressible fluid flows.

Nathan Sime
School of Mathematical Sciences
University of Nottingham
pmxns2@nottingham.ac.uk

Paul Houston
School of Mathematical Sciences
University of Nottingham, UK
Paul.Houston@nottingham.ac.uk

CP3

PoKiTT: an Efficient, Platform Agnostic Package for Thermodynamics, Kinetics, and Transport Properties within PDE Solvers

We introduce PoKiTT, a portable, lightweight library for performing data parallel calculations of thermochemical quantities commonly encountered in turbulent reactive flow simulations. PoKiTT uses Nebo, a domain specific language, to provide a platform agnostic implementation which will be ready for future exascale architectures. We demonstrate the performance benefits of using PoKiTT over Cantera in the context of a PDE solver for both CPU and GPU executions.

Nathan Yonkee
The University of Utah
nathan.yonkee@chemeng.utah.edu

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

CP4

A Framework for Parallel Fast Matrix Multiplication

We explore the performance of novel fast (Strassen-like) matrix multiplication algorithms in sequential and shared-memory parallel environments. We use a code generation framework to automatically implement over twenty fast matrix multiplication algorithms and to rapidly test algorithmic variations for performance tuning. Our implementations outperform Intel MKL for modest problem sizes. Furthermore, we find that Strassens algorithm is not always optimal in practice; in particular, other algorithms perform better for the multiplication of rectangular matrices.

Austin Benson
Stanford University
arbenson@stanford.edu

Grey Ballard

Sandia National Laboratories
gmballa@sandia.gov

CP4

Dynamic Runtime Scheduling for Dense Out-of-Core Matrix Computation on the Intel Xeon Phi

The talk describes the implementation of dense out-of-core matrix computations, such as Cholesky factorization, on the Intel Xeon Phi. The out-of-core algorithm is formulated as computation over submatrix tiles where the dependency is expressed as a directed acyclic graph. A multi-threaded runtime system launches asynchronous offload computations on the Intel Xeon Phi. Examples and results are presented.

Eduardo F. D'Azevedo
Oak Ridge National Laboratory
Mathematical Sciences Section
dazevedoef@ornl.gov

Ben Chan, Terrence Chong
Chinese University of Hong Kong
ss.sing@hotmail.com, tc92321@hotmail.com

Allan Morales
The George Washington University
arrm93@gwmail.gwu.edu

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

CP4

Optimization of Singular Vectors Computation

New SVD routines based on a two-step reduction of a general matrix to bidiagonal form were developed for Intel® MKL. In this talk we present some details of implemented optimizations: dynamic parallelization of singular vectors computations, speculative computations in QR and LQ factorizations, dynamic parallelization of the reduction of banded matrix to bidiagonal form and how these techniques are combined for achieving high performance.

Sergey V Kuznetsov, Nadezhda Mozartova
Intel Corporation
sergey.v.kuznetsov@intel.com,
nadezhda.mozartova@intel.com

CP4

Performance Study of a Randomized Dense Low-Rank Matrix Approximation Using Multiple Gpus

A standard method to compute low-rank approximations for dense matrices is truncated QR factorizations with pivoting, such as LAPACK DGEQP3, an important kernel in many scientific applications. In this talk, we study the performance of an algorithm based on randomized sampling to compute such factorizations of dense matrices for hybrid CPU/GPU architectures and show it can have comparable accuracy, better performance and reduced communications than DGEQP3.

Theo Mary
Universite de Toulouse, INPT ENSEEIHT
theo.mary@etu.enseeiht.fr

Ichitaro Yamazaki, Jakub Kurzak
University of Tennessee, Knoxville
iyamazak@icl.utk.edu, kurzak@icl.utk.edu

Piotr Luszczek
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
luszczek@eecs.utk.edu

Stanimire Tomov
University of Tennessee, Knoxville
tomov@icl.utk.edu

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@cs.utk.edu

CP5

Variational Bayesian Formulations for High-Dimensional Inverse Problems

The present paper advocates a Variational Bayesian (VB) approach for approximating the posterior density in stochastic inverse problems. In contrast to sampling techniques (e.g. MCMC, SMC), VB requires much fewer forward evaluations. Furthermore it enables learning of a suitable lower-dimensional subspace where most of the posterior probability lies, and reducing dramatically the number of unknowns. We demonstrate the accuracy and efficiency of the proposed strategy in nonlinear problems and non-Gaussian posteriors in view of biomedical applications.

Isabell Franck
Technische Universität München, Germany
franck@tum.de

Phaedon S. Koutsourelakis
Technische Universitat Muenchen
p.s.koutsourelakis@tum.de

CP5

Surrogate-Based Bayesian Model Ranking of Atomistic Models Incorporating the Fidelity of Surrogates

Approximate modeling of atomistic systems is crucial in designing new generation of material since ab initio simulations are prohibitively costly. We present how a Bayesian framework can rank a number of competing approximate models and also provide a basis concurrently exploit multiple models for the same system. We make use of Polynomial Chaos surrogates to accelerate the calculation and also account for the numerical error that is thus induced.

Hadi Meidani
University of Southern California
meidani@illinois.edu

Mike Kirby
University of Utah
School of Computing
kirby@cs.utah.edu

Dmitry Bedrov
University of Utah

d.bedrov@utah.edu

CP5

Minimal Set of Mechanisms Controlling Type I Interferon Differential Signaling

Type I interferon ligands differentially trigger a wide range of cellular responses through a common heterodimeric receptor. We seek to identify minimal configurations of cellular mechanisms in interferon signaling that could reproduce this observed behavior. We developed and applied a Bayesian model selection method tailored for linear steady state threshold models. The best models emphasize the importance of the rate of endocytosis and receptor binding dynamics within the endosome for interferon signaling.

Pencho Yordanov
ETH Zurich
pencho.yordanov@bsse.ethz.ch

Irene Otero-Muras, Joerg Stelling
Department of Biosystems Science and Engineering, ETH Zurich
irene.otero@bsse.ethz.ch, joerg.stelling@bsse.ethz.ch

CP5

Uncertainty Propagation Using Infinite Mixture of Gaussian Processes and Variational Bayesian Inference

Uncertainty propagation (UP) is a very challenging mathematical and computational problem. Among other things, UPs difficulty is due to the limited number of model evaluations, the curse of dimensionality, discontinuities, and multivariate responses with non-trivial correlations. In order to deal with all these problems simultaneously, we develop an infinite mixture of multi-output Gaussian process model. We train the model using variational Bayesian inference and we obtain highly competing results in porous flow problems.

Peng Chen, Nicholas Zabaras
Cornell University
pc423@cornell.edu, nzabaras@gmail.com

Ilias Bilionis
Purdue University
ebilionis@gmail.com

CP6

Data Based Regularization Methods

Often one is looking for preconditioners with special behavior on certain subspaces. So in Multigrid methods the preconditioner (smoother) should be defined to remove high frequency error components while in regularization problems the preconditioner (seminorm) should not recover the noise. In this talk we will present different methods to obtain such preconditioners especially in connection with ill-posed inverse problems by finding data dependent seminorms.

Thomas K. Huckle
Institut fuer Informatik
Technische Universitaet Muenchen
huckle@in.tum.de

CP6

Preconditioner Scaling for Finite Element Models of Turbulent Air/Water Flow in Coastal and Hydraulic Applications

While three-dimensional models of turbulent air/water flow are seeing application in coastal and hydraulic engineering, the computing resources required for simulating field problems is a major barrier to widespread adoption of the overall approach. One route to scalable solvers is exploiting the block structure of operators arising from discretizations of multi-phase Navier-Stokes equations. We present a study of recently developed approximate Schur complement factorizations to a stabilized finite element code for multi-phase flow.

Chris Kees
U.S. Army Engineer Research and Development Center
Coastal and Hydraulics Laboratory
christopher.e.kees@usace.army.mil

Aron Ahmadia
US Army Engineer Research and Development Center
aron@ahmadia.net

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory and CU Boulder
jed@jedbrown.org

Matthew Farthing
US Army Engineer Research and Development Center
matthew.w.farthing@usace.army.mil

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

CP6

A Scalable Newton-Krylov-Schwarz Method for Coupled Fluid-Structure Interaction Problems

Fluid-structure interaction is a challenging multi-physics problem. The difficulties are due to the high nonlinearity derived from the convective term of fluid problems, the constitutive relationship of solid materials, the dependency of the solution on the displacement of the moving fluid mesh, and unbalanced message passings between different processors resulting from the partition of unstructured meshes into a large number of parts. To overcome the difficulties, we study an inexact Newton-Krylov algorithm with a Schwarz preconditioner to solve the monolithically coupled fluid-structure system. We show that the proposed algorithm is scalable in terms of the iteration count and the total compute time on a supercomputer with a large number of processors.

Fande Kong, Xiao-Chuan Cai
Department of Computer Science
University of Colorado Boulder
fande.kong@colorado.edu, cai@colorado.edu

CP6

A Parallel Linear Solver Exploiting the Physical

Properties of the Underlying Mechanical Problem

The iterative solution of large systems of equations may benefit from parallel processing. However, using a straightforward domain decomposition in 'layered' geomechanical finite element models with significantly different stiffnesses may lead to slow or non-converging solutions. Physics-based domain decomposition is the answer to such problems, as explained in this paper and demonstrated on the basis of a few examples. Together with a two-level preconditioner comprising an additive Schwarz preconditioner that operates on the sub-domain level, an algebraic coarse grid preconditioner that operates on the global level, and additional load balancing measures, the described solver provides an efficient and robust solution of large systems of equations. Although the solver has been developed primarily for geomechanical problems, the ideas are applicable to the solution of other physical problems involving large differences in material properties.

Kees Vuik
Delft University of Technology
c.vuik@tudelft.nl

CP7

Parallel Graph Coloring for Scientific Computing

Graph coloring has many applications in scientific computing, such as parallel scheduling, sparse matrix reordering, and automatic differentiation. Often the graph coloring itself must be computed in parallel. We describe new software for parallel graph coloring on multicore and manycore architectures. Our implementation is based on the Kokkos library, which makes the code portable to many platforms, including Intel MIC and GPUs. We discuss algorithmic issues and show some preliminary results.

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

CP7

Solving Sparse Linear Systems on GPUs Based on the Biell Storage Format

Nowadays, the GPU has evolved into a highly parallel coprocessor which is suited to compute-intensive, highly parallel computation. Solving sparse linear systems is one of basic task for scientific and engineering computing. Achieving high performance of this task on GPUs is relatively challenging, especially when the matrix has no specific structure. For the general sparse matrices, we have proposed a new data structure based on the BiELL (bisection ELLPACK) format, which is designed to realize the load balance better and thus improve the performance of the SpMV (sparse matrix-vector multiplication) on GPUs. Now, we use this new format in iterative methods and preconditioning techniques for solving sparse linear systems. Numerical results on various matrices show that GMRES, CG and some preconditioning techniques using this new format have higher performance than that of using other formats on GPUs and their CPU counterparts.

Tongxiang Gu
Computer Network Information Center,
Chinese Academy of Science
gu_tx@sina.com

Cong Zheng
Graduate School of Chinese Academy of Engineering Physics
zhengcong55@126.com

Shou Gu
University of Electronic Science and Technology of China
294258779@qq.com

Xingping Liu
Institute of Applied Physics and Computational Mathematics
lxp@iapcm.ac.cn

CP7

Spectral Methods in PDE Solving: a Multi-GPU Framework

Spectral methods offer unconditionally stable time-stepping schemes for solving PDEs numerically. Although a parallel spectral scheme necessitates the parallelization of the FFT, all other operations can be executed asynchronously. We present a framework utilizing the benefits of multi-GPU hardwares and demonstrate the advantages and disadvantages of the method on the results obtained for surfactant assisted liquid phase separation in the water/CO2/hydrocarbon system, which is of high industrial importance nowadays on the field of improved recovery.

Gyula I. Toth
Wignes Research Centre for Physics
Budapest, Hungary
Gyula.Toth@ift.uib.no

Tatjana Kuztensova, Bjorn Kvamme
Department of Physics and Technology
University of Bergen
tatyana.kuznetsova@ift.uib.no, bjorn.kvamme@ift.uib.no

CP7

Direct Hierarchical Schur Method for Nested Dissection Reordered Linear Systems on Multi-GPUs

We propose a direct hierarchical Schur complement method to solve sparse symmetric positive-definite linear systems on multinode GPU clusters. By exploring the structures of the reordered coefficient matrices and developing a scheme to distribute data and schedule tasks, we can solve the problem efficiently. Our method solves the diagonal block and Schur subproblems by GPU-based Cholesky factorization to gain accelerated performance. Numerical results suggest the method is scalable for GPU-clusters with different node numbers.

Cheming Chu
Department of Mathematics, National Taiwan University
r00221036@ntu.edu.tw

Pochuan Wang
Institute of Applied Mathematical Sciences, National Taiwan
University
r02246011@ntu.edu.tw

Weichung Wang
Institute of Applied Mathematical Sciences
National Taiwan University

wwang@ntu.edu.tw

CP7

Optimizing Structured Grid Numerical Simulations for Numa-Multicore Systems

NUMA-multicore systems are now ubiquitous. Optimal data placement and data move are crucial to fully utilize these systems. In this talk, we share our experience in optimizing JASMIN framework for these systems. Firstly, we developed an efficient NUMA-aware heap manager to enforce data locality. Secondly, we designed a scalable NUMA-aware algorithm for structured grid data communications. These technologies are then integrated into JASMIN framework and numerical results shows that our approach improves significantly the performance of real-word applications on typical NUMA-multicore systems with 10K CPU cores.

Zhang Yang
Institute of Applied Physics and Computational Mathematics
China Academy of Engineering Physics
yang_zhang@iapcm.ac.cn

Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
zhang_aiqing@iapcm.ac.cn

Zeyao Mo
Laboratory of Computational Physics, IAPCM
P.O. Box 8009, Beijing 100088, P.R. China
zeyao_mo@iapcm.ac.cn

CP8

Scalable Alternative to Domain Decomposition

We present a novel algorithm based on decomposing the range of a PDE instead of the error. The method allows efficient use of W and full multigrid cycles in a parallel setting. A performance model predicts extreme weak scalability of the algorithm and numerical tests up to 1000 cores confirm the predicted scalability.

David A. Appelhans
University of Colorado, Boulder
dappelha@gmail.com

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

Steve McCormick
Department of Applied Math
CU-Boulder
stevem@colorado.edu

John Ruge
University of Colorado
jruge@colorado.edu

CP8

Reducing Communication Costs for Sparse Matrix Multiplication within Algebraic Multigrid

We consider the sequence of sparse matrix-matrix multiplications performed during the setup phase of algebraic multigrid. In particular, we show that the most commonly used parallel algorithm is often not the most communication-efficient one for all of the matrix-matrix multiplications involved. By using an alternative algorithm, we show that the communication costs are reduced (in theory and practice), and we demonstrate the performance benefit for both model and real problems on large-scale distributed-memory parallel systems.

Grey Ballard
Sandia National Laboratories
gmballa@sandia.gov

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

Christopher Siefert
Sandia National Laboratories
csiefer@sandia.gov

CP8

Avoiding Communication and Synchronization in Krylov Eigensolvers

Solving sparse eigenproblems is an important task in many engineering problems; Krylov subspace methods, e.g. IRLAN, TRLAN, are frequent solvers of choice. Krylov subspace eigensolvers have well-known difficulties scaling beyond 10,000 processors. This bottleneck is due to communication from dot products at each iteration rather than matrix-vector products. When matrix-vector products scale well, judiciously-applied preconditioners reduce total iterations; they shift work from globally-communicating dot products to locally-communicating matrix-vector products, and we witness improvements in scalability.

Alexander Breuer, Claire Eisner, Jaroslaw Knap, Kenneth Leiter
U.S. Army Research Laboratory
alexander.m.breuer.civ@mail.mil,
claire.g.eisner.ctr@mail.mil, jaroslaw.knap.civ@mail.mil,
kenneth.w.leiter2.civ@mail.mil

CP8

αSetup-Amg: An Adaptive Setup Based Amg Solver for Large-Scale Simulations with Long Time-Stepping

Coarse-level visiting is the main reason that causes loss of scalability for AMG solver on massive parallel computer. In our presented adaptive setup strategy, coarsening is performed based on the smoothing behavior on each level, instead of constructing via an independent setup phase in traditional procedure. As a results, doing relaxations on finer-levels as much as possible, while the required coarse-levels as less as possible. Realistic simulations on $O(10^4)$ cores show the improved scalability.

Xiaowen Xu
Institute of Applied Physics and Computational Mathematics
xwxu@iapcm.ac.cn

CP9

A Multiscale Finite Volume with Oversampling

Method to Simulate Low-Frequency Electromagnetic Geophysical Responses

In Geophysics, simulation of low-frequency Electromagnetic (EM) fields in highly heterogeneous, anisotropic media is computationally expensive. One reason being the multiple length scales that coexist in a given realistic setting. Discrete models require very fine meshes leading to solve large linear systems of equations. Here, we develop a multiscale finite-volume method with oversampling to reduce the size of the system of equations to be solved while retaining a good level of accuracy in the solution.

Luz Angelica A. Caudillo Mata
Earth, Ocean and Atmospherical Sciences Department
University of British Columbia
lcaudill@eos.ubc.ca

Eldad Haber
Department of Mathematics
The University of British Columbia
haber@math.ubc.ca

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

Christoph Schwarzbach
University of British Columbia
cschwarz@eos.ubc.ca

CP9

A Computational Shock-Tube for Reproducible Computational Experiments in Traumatic Brain Injury

We developed a computational shock tube with conservative finite volume methods and interface-shock interaction to complement experiments in traumatic brain injury(TBI). The 3D model was implemented using compressible Euler equations coupled with a Tammann-EOS. An experimental setup was simulated and yielded insights not available by experimental means, emphasizing the importance of geometry and yielding cavitation as a possible damage mechanism. The code is open-source to promote reproducible research.

Mauricio J. Del Razo
University of Washington
mauricio_delrazo@yahoo.com

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

David Cook
VA Hospital
dgcook@u.washington.edu

CP9

Universal Meshes for Problems with Moving Boundaries

We develop a framework for the design of high-order finite element methods for moving-boundary problems using a universal mesh: a stationary background mesh containing the domain of interest for all times that adapts to the geometry of the immersed domain by adjusting a small number of mesh elements in the neighborhood of the moving boundary. We illustrate the approach and compare it with conventional arbitrary Lagrangian Eulerian (ALE) schemes via numerical examples involving fluid-structure interaction.

Evan S. Gawlik
Stanford University
Institute for Computational and Mathematical Engineering
egawlik@stanford.edu

Adrian Lew
Stanford University
Mechanical Engineering
lewa@stanford.edu

CP9

Numerical Simulations of Biological Invasions

Biological invasions occur when there is a road on which an epidemic propagates faster than in the outlying fields adjacent to the roads. These types of invasions can be modeled using reaction-diffusion equations with varying parameters on the roads and in the outlying areas with coupling between the two. We will present a numerical method to study this problem. Comparisons with previous analytical work with a straight road will be presented, as well as numerical simulations of more complex road shapes.

Shilpa Khatri
Department of Mathematics
University of North Carolina at Chapel Hill
skhatri3@ucmerced.edu

Anna-Karin Tornberg
KTH
akto@kth.se

CP9

Asymptotic-Preserving Space-Time Discontinuous Galerkin Schemes for a Class of Relaxation Systems

We consider in this work a class of singularly perturbed hyperbolic balance laws that admit a diffusive limit. Such systems arise naturally in radiative transport applications if one starts with a Boltzmann description and expands the distribution function in spherical harmonics (i.e., the P_n approximation). One key difficulty in solving such systems is that standard numerical schemes have maximum time-step restrictions that vanish in the singular limit. Several approaches have been proposed in the literature to overcome this difficulty, many of which are based on splitting the equation into stiff and non-stiff pieces and using appropriate semi-implicit time-stepping methods. In this work we employ a different strategy in order to achieve asymptotic-preservation. We develop a scheme using a space-time discontinuous Galerkin approach. Several numerical test cases are used to validate the proposed scheme.

Anna Lischke
Iowa State University
alischke@iastate.edu

James A. Rossmanith
Iowa State University

Deparment of Mathematics
rossmani@iastate.edu

CP10

Finite-Difference Frequency-Domain Analysis of Photonic Devices with Periodic Structures Based on Domain Decomposition

We present an efficient algorithm and implementation of 3D finite-difference frequency-domain simulation of photonic devices. By proposing a new matrix-reordering scheme in the domain decomposition framework, our method exploits the homogeneous and periodic structures in photonic devices with Yee's mesh and achieves computing resources saving in memory usage and total runtime. The linear system solver is capable of solving the ill-conditioned problems and suitable for parallel computation via high-performance computers with GPU or many-core accelerators.

Cheng-Han Du
Department of Mathematics, National Taiwan University
b92006@csie.ntu.edu.tw

Pochuan Wang
Institute of Applied Mathematical Sciences, National Taiwan
University
r02246011@ntu.edu.tw

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

CP10

High Order Schemes Based on Operator Splitting and Deferred Corrections for Stiff Time Dependent Pdes

We consider quadrature formulas of high order in time based on Radau–type, L–stable implicit Runge–Kutta schemes to solve time dependent stiff PDEs. Instead of solving a large nonlinear system of equations, we develop a method that performs iterative deferred corrections to compute the solution at the collocation nodes. The numerical stability is guaranteed by a dedicated operator splitting technique that efficiently handles the stiffness of the PDEs and provides initial and intermediate solutions to the iterative scheme.

Max Duarte
Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory
MDGonzalez@lbl.gov

Matthew Emmett
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov

CP10

Monolithic Multi-Time-Step Coupling Methods for Transient Systems

We shall present new multi-time-step coupling methods for first- and second-order transient systems. The proposed methods are based on dual Schur domain decomposition, and employ techniques from differential/algebraic equations. We shall also provide a systematic theoretical study of these methods. Several numerical examples will be shown to illustrate the theoretical findings and the salient features of the proposed methods.

Saeid Karimi
University of Houston
skarimi@central.uh.edu

Kalyana Nakshatrala
University Of Houston - Main Campus
knakshatrala@uh.edu

CP11

A Parallel Fast Sweeping Method for Quadtrees and Octrees

We present a hybrid shared memory and message passing algorithm for the fast sweeping method on tree based adaptive grids. We utilize graph theory to decompose the tree into clusters of nodes that can be updated simultaneously via a shared memory parallelization model. Large scale parallelization is accomplished by domain decomposition with a message passing model. We present scaling results on a number of time-independent Hamilton Jacobi equations.

Miles L. Detrixhe
University of California Santa Barbara
mdetrixhe@engineering.ucsb.edu

CP11

A Communication Algorithm for the Patch-Based Multiblock Structured Mesh Applications

Multiblock structured mesh allows to handle complex configurations which are widely existed in computational physics applications. A Patch-based data structure is always used in applications with multiblock structured mesh to get satisfying parallel performance. However, such Patch-based data structure seriously challenges the block to block data communications. This talk presents an algorithm for such communication and introduces its integration to JASMIN infrastructure to support the peta-scale simulations while tens of thousands of processors are used. Performance results show its robustness.

Hong Guo
Institute of Applied Physics and Computational Mathematics
guo_hong@iapcm.ac.cn

CP11

Scalable Parallel Assembly for High-Performance Computing with Isogeometric and Higher-Order Finite Elements

Isogeometric and higher-order finite element methods lead to system matrices whose parallel assembly requires extensive communication. We introduce a distributed-memory domain decomposition strategy, which uniquely assigns each degree of freedom to one processor, therefore eliminating communication completely. We show that although contributions from interface elements need to be computed several times on different processors, our approach leads to a scalable and efficient parallel assembly. Algorithmic details and performance measurements for different examples

are presented.

Vasco Varduhn
Technische Universität München
vvarduhn@umn.edu

Dominik Schillinger
University of Minnesota
dominik@umn.edu

CP11

Exploring Communication Options with Adaptive Mesh Refinement

Finite difference and volume based codes comprise a significant portion of the workload on modern high performance computers. Many of these codes use Adaptive Mesh Refinement (AMR) as a computational strategy. We have developed miniAMR, a miniapp in the Mantevo suite, to explore AMR communication issues. We compare miniAMR to CTH, a shock hydrodynamics code, and use miniAMR to explore some communication strategies that may be necessary as we move to future architectures.

Courtenay T. Vaughan, Richard Barrett
Sandia National Laboratories
ctvaugh@sandia.gov, rfbarre@sandia.gov

CP11

A Communication Staging Technique for a Hierarchical Ocean Model

We study a communications staging technique applied to an algorithmically accelerated, free-surface, z-level ocean model. The ocean model is a hierarchical high-order / low-order model, which allows mapping to heterogeneous architectures and exploitation of advanced communication algorithms. We compare the benefit of communication staging between a current numerical method and a research method under development, within this high-order / low-order framework. We provide numerical examples to support our study and compare to traditional implementations.

Geoff Womeldorff, Chris Newman, Dana Knoll, Luis Chacón
Los Alamos National Laboratory
womeld@lanl.gov, cnewman@lanl.gov, nol@lanl.gov, chacon@lanl.gov

CP12

Asymptotic-Preserving Scheme for the Fokker-Planck-Landau-Maxwell System in the Quasi-Neutral Regime

This work deals with the numerical resolution of the Fokker-Planck-Maxwell system in the quasi-neutral regime. In this regime the stiffness of the stability constraints of classic schemes causes huge calculation times. That is why, we introduce a new stable numerical scheme consistent with the transitional and limit models. Such schemes are called Asymptotic-Preserving schemes in literature. This new scheme is able to handle the quasi-neutrality limit regime without any restrictions on time and space steps. Next, this approach can be easily applied to angular moment models by using a moments extraction. Finally, the efficientcy of the scheme is validated on the Batishev test case.

Stephane Brull
Institut de Mathématiques de Bordeaux UMR 5251
Université Bordeaux
stephane.brull@math.u-bordeaux1.fr

Bruno Dubroca, d'Humière Emmanuel, Guisset Sébastien
CELIA
dubroca@math.u-bordeaux1.fr, dhumieres@celia.u-bordeaux1.fr, guisset@celia.u-bordeaux1.fr

CP12

Coarse Multiscale Timestepping for Problems in Plasma Physics with Equation-Free Projective Integration

Multiscale plasma problems are hard to simulate because the physics of micro and macro-scales are strongly linked. We propose a coarse-grained numerical scheme, based on equation-free projective integration, for a kinetic plasma system modelled by the Vlasov-Poisson equations, following the idea of [Shay, Drake, Dorland 2006]. A particle-in-cell (PIC) code is used to simulate the micro scale dynamics. As a first test case, we simulate the propagation and steepening of a nonlinear ion acoustic wave.

Paul Cazeaux, Jan Hesthaven
EPFL
paulcazeaux@gmail.com, jan.hesthaven@epfl.ch

CP12

Semi-Lagrangian Discontinuous Galerkin Schemes for the Relativistic Vlasov-Maxwell System

The Vlasov-Maxwell system describes the evolution of a collisionless plasma, represented through a probability density function (PDF) that self-interacts via the electromagnetic force. One of the main difficulties in numerically solving this system is the severe time-step restriction that arises from parts of the PDF associated with moderate-to-large velocities. The dominant approach in the plasma physics community is the so-called particle-in-cell method. The focus of the current work is on semi-Lagrangian methods. In particular, we develop a method based on high-order discontinuous Galerkin (DG) scheme in phase space, and an operator split, semi-Lagrangian method in time. The method is designed to be (1) high-order accurate, (2) mass conservative, and (3) positivity-preserving. The resulting scheme is applied to laser-plasma acceleration problems.

Pierson Guthrey
Iowa State University
Department of Mathematics
pguthrey@iastate.edu

James A. Rossmanith
Iowa State University
Deparment of Mathematics
rossmani@iastate.edu

CP12

Discontinuous Galerkin Deterministic Solvers of Boltzmann-Poisson Models of Hot Electronic Transport Using Empirical Pseudopotential Methods

We develop Discontinous Galerkin deterministic solvers of

Boltzmann-Poisson models of electronic transport, incorporating numerical full energy bands obtained from Empirical Pseudopotential Methods (EPM), to improve the semiconductor physical modeling related to energy bandstructure, charge carrier group velocity and scatterings. We will present the DG schemes related to electronic transport in both a conduction band and multi-band system. Simulations related to nano-devices such as $n^+ - n - n^+$ diodes and double gated MOSFETS will be presented.

Jose A. Morales Escalante
ICES, The University of Texas at Austin
jmorales@ices.utexas.edu

Irene M. Gamba
Department of Mathematics and ICES
University of Texas
gamba@math.utexas.edu

Yingda Cheng
Department of Mathematics
Michigan State University
ycheng@math.msu.edu

Armando Majorana
Dipartimento di Matematica e Informatica
University of Catania - Italy
majorana@dmi.unict.it

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

James R. Chelikowsky
Institute for Computational Engineering and Sciences
University of Texas at AUstin
jrc@utexas.edu

CP12

Vlasov-Poisson Simulations of Magnetized Plasmas Using High-Order Continuum Methods

The Vlasov-Poisson equation system, which describes collisionless plasma dynamics, can be solved in conservation-law form using continuum methods. A fourth-order accurate finite volume method has been implemented using the Chombo library to solve the governing equations in two spatial and two velocity dimensions. A new benchmark based on the Dory-Guest-Harris instability has been developed for validating magnetized plasma simulations in higher dimensional phase space. Extension of the method to cylindrical coordinates is described.

Genia Vogman
University of California - Berkeley
geniav@berkeley.edu

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

Uri Shumlak
Aerospace and Energetics Research Program
University of Washington
shumlak@uw.edu

CP13

Parallel Methods for Accelerated Multilevel Monte Carlo for Partial Differential Equations with Random Input

An improvement on previous work where MLMC for PDEs using finite element iterative solvers was accelerated by improving the initial guess during sampling. Using information gathered at previous samples, the number of iterations per Monte Carlo sample were greatly reduced. However, these methods did not allow for any parallel computation. In this work we propose alternate methods which can benefit from parallel processing.

Zane Colgin
Middle Tennessee State University
zac2g@mtmail.mtsu.edu

CP13

Topology Optimization under Manufacturing Uncertainties

The focus of this work is on incorporating manufacturing uncertainties in the topology optimization of micro and nano devices. The considered microfabrication process is photolithography, which transfers a mask pattern onto a substrate. Deviations between the print and the design occur due to light diffraction and process variations and these can change severely the design performance. Robust solutions can be obtained by including uncertainties in the design process. The modification increases significantly the computational cost and different strategies for its reduction are discussed.

Boyan S. Lazarov
Department of Mechanical Engineering
Technical University of Denmark
bsl@mek.dtu.dk

CP13

Reducing Dimensionality Through Active Subspaces, and the Effect of Gradient Approximations on the Associated Eigenpairs

Uncertainty quantification studies struggle in high dimensions; active subspaces are new tools for dimension reduction. Finding active subspaces requires gradient information which in practice is often unknown. To work around this problem, we can approximate gradients by fitting models to the data and computing the gradients associated with the models. We illustrate this principle using Local Linear Regressions, and investigate how approximation errors in the gradients affect the estimates of the eigenpairs. Examples are shown for selected test functions

Uno B. Vaaland
Norwegian University of Science and Technology, Norway
ubvaaland@gmail.com

Paul Constantine
Colorado School of Mines
Applied Mathematics and Statistics

pconstan@mines.edu

CP13

A Multi-Model Approach for Uncertainty Propagation and Model Calibration in CFD Applications

Monte Carlo-based uncertainty propagations are computationally expensive or even impractical for complex systems such as turbulent flows. Efforts to reduce sampling errors and modeling errors often compete for limited computational resources. Here we propose a multi-model Monte Carlo method that combines models of multiple fidelities to propagate uncertainties. A Gaussian process is used to construct the model discrepancy between high- and low-fidelity models to improve the results.

Jianxun Wang, Heng Xiao
Dept. of Aerospace and Ocean Engineering, Virginia Tech
vtwjx@vt.edu, hengxiao@vt.edu

CP13

Detecting Discontinuities and Localized Features Using Gaussian Processes

Constructing surrogates of physical models can be extremely difficult, especially when they exhibit sharp, or even discontinuous, variations. Iteratively-built tree-decompositions, or two-stage classification-regression approaches, have been able to partially deal with this problem. However, virtually all such techniques rely on intuitive ideas. Here, we develop a two-level-deep, potentially infinite mixture of Gaussian processes that can automatically detect local features with no ad hoc assumptions.

Ilias Bilionis
Purdue University
ebilionis@gmail.com

Nicholas Zabaras
Cornell University
nzabaras@gmail.com

CP14

Optimal Control of Level Sets

We present two level set approaches for numerical simulation of evolving interfaces, each based on PDE-constrained optimization problems. In the first one the optimal control procedure constrains the level set function so as to satisfy a conservation law and thus produces a mass conservative numerical solution. The second approach is designed to preserve the signed distance function property of the level set function by incorporating the residual of the Eikonal equation into the cost functional and prescribing a bilinear state equation. Both approaches are evaluated numerically.

Christopher Basting, Dmitri Kuzmin
Dortmund University of Technology
christopher.basting@math.tu-dortmund.de,
kuzmin@math.uni- dortmund.de

CP14

Fractional Powers of Finite Element Approximation for An Parabolic Optimal Control Problems

In this paper, a numerical theory based on fractional finite element approximations for an optimal control problem with pointwise control constraints is presented and analyzed. The state and co-state variables are approximated by the piecewise linear functions and the control is approximated by piecewise constant functions. We derive, a priori error estimates for both the control variable and the state variables. We illustrate with a numerical example to confirm our theoretical results.

Manickam Kandasamy
Periyar University, Salem 636011, INDIA
kkmmanickam@gmail.com

Periasamy Prakash
Department of Mathematics,Periyar University
Salem, INDIIA
kkmmanickam@gmail.com

CP14

Optimal Order Multigrid Preconditioners for Linear Systems Arising in the Semismooth Newton Method Solution Process of a Class of Control-Constrained Problems

In this work we discuss multigrid preconditioners for control-constrained optimal control problems constrained by semilinear elliptic PDEs. Building upon the existing work from linear-quadratic case, we study preconditoners for the submatrices of the reduced Hessian arising in the semismooth Newton solution process. The control is discretized using piecewise constant finite elements. It is observed that the resulting preconditioner is of optimal order with respect to the discretization. Analytical and numerical results are presented.

Jyoti Saraswat
Thomas More College
saraswj@thomasmore.edu

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

CP14

Multigrid Preconditioners for Stochastic Optimal Control Problems with Elliptic Spde Constraints

We consider an optimal control problem constrained by an elliptic SPDE, with a stochastic cost functional of tracking type. We use a sparse grid stochastic collocation approach to discretize in the probability space and finite elements to discretize in the physical space. To accelerate the solution process, we propose a deterministic multigrid preconditioner for the stochastic reduced KKT system, similar to the preconditioners introduced by Draganescu and Dupont for the deterministic PDE constrained problem.

Ana Maria Soane
Towson University
Department of Mathematics
asoane@umbc.edu

CP14

Numerical Realization of the Open Pit Mine Planning Problem

By reformulating a model for open pit mine planning due to [F. Alvarez et al., 2011], we obtain an optimization problem subject to viscosity solutions of an underlying Hamilton-

Jacobi PDE. We apply a monotone discretization scheme, which yields an optimal control problem of a system of ODEs. We present the algorithmic treatment and numerical results of this problem under consideration of the effort constraint, which is of non-incremetal type.

Nikolai Strogies
Humbold Universität zu Berlin
Matheon Mathematics for key technologies
strogies@mathematik.hu-berlin.de

Andreas Griewank
HU Berlin, Germany
griewank@math.hu-berlin.de

CP15

Fast Supercomputing Algorithms for Power System Operation and Control

We have developed new supercomputing algorithms for multicore architectures for the state estimation and power flow problems that present a step towards the real time power grid optimization and control. We have also developed allocation algorithms for Phasor Measurement Units (PMUs) for estimating the state of a nonlinear power system in real time. These algorithms involve partitioning of large power systems into several sub-systems, and multithreaded solving with PMU data constraints.

Eugene A. Feinberg
Stony Brook University
eugene.feinberg@stonybrook.edu

Bruce Fardanesh
New York Power Authority
bruce.fardanesh@nypa.gov

Muqi Li, Roman Samulyak
Stony Brook University
muqi.li@stonybrook.edu,
roman.samulyak@stonybrook.edu

George Stefopoulos
New York Power Authority
george.stefopoulos@nypa.gov

Gaurish Telang
Stony Brook University
gtelang@ams.sunysb.edu

CP15

Simulation-based Current Estimation in Magnetohydrodynamic Generators

Direct power generation via magnetohydrodynamic principals offers an increase in efficiency over traditional turbomachinery systems but commercialization is impeded by high lifecycle costs. The generators electrodes are damaged by the formation of high current density arcs. However, these arcs induce magnetic fields which are measurable nearby. The development of sensors to detect these arcs is critical in controlling the phenomenon. We produce simulations using the Mimetic Finite Differences and perform inversion by simulation-based parameter estimation.

Duncan A. Mcgregor
Oregon State University
Department of Mathematics
mcgregod@math.oregonstate.edu

Vrushali Bokil, Nathan L. Gibson
Oregon State University
bokilv@math.oregonstate.du,
gibsonn@math.oregonstate.edu

Charles Woodside
National Energy Technology Laboratory
rigel.woodside@netl.doe.gov

CP16

Data-Driven Uncertainty Quantification with Adaptive Sparse Grids in Subsurface Flow Simulations

We present a novel data-driven approach to propagate uncertainty through an expensive subsurface flow simulation. We remove the gap between the subjective approximation of the input's uncertainty and the unknown real distribution by applying sparse grid density estimation. We link the estimation to the adaptive sparse grid collocation method to propagate the uncertainty and obtain new refinement criteria. Our approach excels by speed, flexibility and thus can be applied in many fields from environmental to financial sciences.

Fabian Franzelin
Universität Stuttgart
fabian.franzelin@ipvs.uni-stuttgart.de

Sergey Oldayshkin
University of Stuttgart
sergey.oladyshkin@iws.uni-stuttgart.de

Benjamin Peherstorfer
ACDL, Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
pehersto@mit.edu

Dirk Pflüger
University of Stuttgart
Dirk.Pflueger@ipvs.uni-stuttgart.de

CP16

HPC and Model Reduction Algorithms for Large-Scale Simulation of Stochastic Wave Propagation Models

We discuss a new class of efficient iterative high-order high performance computing (HPC) model reduction algorithm for simulating wave propagation exterior to stochastic configurations containing very large numbers of particles. Such simulations are crucial in important medical applications such as in vivo and in vitro blood measurements using light scattering of blood cells, and topical climate science applications such as light scattering and absorption by atmospheric aerosols. Even simulation for a single deterministic configuration with a large number of particles is a large-scale computational challenge. The stochastic nature of the configuration leads to a larger dimensional model involving three spatial and several stochastic variables. Our approach provides a practically feasible HPC framework to compute highly accurate statistical moments to quantify uncertainties in stochastic configurations.

Mahadevan Ganesh
Colorado School of Mines

mganesh@mines.edu

CP16

Quantification of Structural Uncertainty in a Land Surface Model

This study identifies and quantifies model structural uncertainty in the Community Land Model, by fully exploring the high-dimensional model parameter space with efficient sampling and then evaluating the discrepancies between the corresponding numerical simulations and observations using wavelet decomposition and other spatiotemporal analysis approaches. The power spectra, dominant temporal scales and energy, and characteristic phase shift, are summarized to help quantify model structural uncertainty and identify the major processes contributing to such uncertainty.

Zhangshuan Hou, Maoyi Huang
Pacific Northwest National Lab
zhangshuan.hou@pnnl.gov, maoyi.huang@pnnl.gov

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Laura Swiler
Sandia National Laboratories
Albuquerque, New Mexico 87185
lpswile@sandia.gov

CP16

PDE-Constrained Optimization Applied to Core Flooding from Reservoir Engineering

This talk explores the impact of PDE-constrained optimization on the core flooding problem to determine rock-fluid parameters essential to effective reservoir simulation. Current core flooding technologies are time-consuming and involve manual inversion for parameters of interest on simplified physical models of fluid flow. We demonstrate, on simple model problems, that PDE-constrained optimization can automate the process and do so in a robust way that incorporates more realistic physical fluid flow modeling.

Caleb C. Magruder
Rice University, Computational and Applied Mathematics
cm47@rice.edu

Jeremy Brandman, Shivakumar Kameswaran
ExxonMobil Corporate Strategic Research
jeremy.s.brandman@exxonmobil.com, shivakumar.kameswaran@exxonmobil.com

CP16

Integration of Geophysical Fluid Dynamics and Fully 3D Fluid Dynamics to Simulate Multiphysics Coastal Ocean Flows

An integration of geophysical fluid dynamics and fully 3D fluid dynamics models is proposed to simulate multiphysics coastal ocean flows. This integration is the first of its kind and able to capture flow phenomena at spatial scales O (1) m – O (10,000) km. The approachs methodology and software development are discussed, and its unprecedented capabilities will be demonstrated in its applications to crucially important problems that are beyond the reach of other existing models.

Hansong Tang, Ke Qu
Dept. of Civil Eng., City College of NewYork, CUNY
htang@ccny.cuny.edu, kqu00@citymail.cuny.edu

CP17

An Efficient, Pressure Projection Method for Reacting Low-Mach Flow Simulations

A new explicit variable-density pressure projection method is proposed with a focus on transient low-Mach-number reacting flows. This method introduces a new form of the pressure Poisson equation suitable for use with an explicit algorithm. The density, assumed to be a function of an arbitrary set of transported scalars, is determined by solving a non-linear system of equations at each point in the domain. The proposed method is evaluated using several time-varying, variable-density test cases as well as an annular jet flow.

Amir Biglari
The Institute for Clean and Secure Energy
Department of Chemical Engineering, University of Utah
amir.biglari@utah.edu

Tony Saad
Institute for Clean and Secure Energy
University of Utah
saadtony@gmail.com

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

CP17

Energy-Stable Open Boundary Conditions for Two-Phase Flows

We present an effective open boundary condition, and an associated numerical algorithm, within the phase field framework for dealing with two-phase outflows or open boundaries. Two-phase outflows refer to situations where the interface between two immiscible incompressible fluids passes through open portions of the domain boundary. The proposed open boundary conditions ensure the energy stability of the two-phase system, even in situations where strong backflows or vortices occur at the two-phase outflow boundaries. Numerical examples involving two-phase inflows/outflows will be presented to demonstrate the effectiveness of the method when large density ratios and large viscosity ratios are involved and when strong backflows occur at the two-phase outflow boundaries.

Suchuan Dong
Purdue University
sdong@math.purdue.edu

CP17

A Low Mach Number Model for Moist Atmospheric Flows

We introduce a low Mach number model for moist atmospheric flows that accurately incorporates reversible moist processes in flows whose features of interest occur on advective rather than acoustic time scales. We numerically assess the validity of the more computationally efficient low

Mach number approximation for moist atmospheric flows by contrasting the low Mach number solution to reference solutions computed with a fully compressible formulation for a variety of test problems.

Max Duarte
Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory
MDGonzalez@lbl.gov

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

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

CP17

A Stable Projection Method for the Incompressible Navier-Stokes Equations on Arbitrary Geometries and Adaptive Quad/oc-Trees

We present a novel stable projection method for the incompressible Navier-Stokes equations on non-graded adaptive quad/oc-trees with arbitrary geometries. The viscosity is treated implicitly through a finite volume approach based on Voronoi partitions and the convective term is discretized with a semi-Lagrangian scheme, thus relaxing the time step restrictions.

Arthur Guittet
UCSB
arthur.guittet@gmail.com

CP18

Dual-Mixed Finite Element Methods for the Brinkman Problem

We develop a dual-mixed finite element method for the Brinkman problem of viscous flow in porous media. The primary unknowns are the fluid velocity, the fluid stress, and the deviatoric part of the velocity gradient. The method is stable and accurate for a wide range of problem parameters, including both the Stokes and Darcy limiting cases.

Jason Howell
College of Charleston
howelljs@cofc.edu

Noel J. Walkington
Department of Mathematical Sciences
Carnegie Mellon University
noelw@andrew.cmu.edu

CP18

Numerical Modeling of Non-Associated Flow Model by Successive Convex Optimization: Application in Incompressible Porous Media

In this study, we propose a numerical method to solve the PDEs of incompressible saturated porous media under dynamic condition and material nonlinearity. Two parallel schemes are implemented to handle the incompressibility condition and the non associated flow material model. The approach couples the Raviart-Thomas mixed and Galerkin finite elements to reach to a stable space based on the inf-sup condition. In the meanwhile, we cast the computation for material state update as a successive convex optimization problem.

Zahra S. Lotfian, Mettupalayam Sivaselvan
SUNY at Buffalo
zahrasad@buffalo.edu, mvs@buffalo.edu

CP18

Physically Motivated and Certified Approximation of Large Elastic Structures in Real-Time

We introduce a physically motivated, certified model reduction approach for the simulation of large component-based elastic structures as bridges. We build the system from a library of interoperable, parametrized components and apply a domain decomposition method. We compute the displacement field in real-time at a high FEA by using a reduced basis approximation within the component and physical modes on the interfaces/ports. The approximation is certified by error bounds based on local, component-wise error indicators.

Kathrin Smetana
Department of Mechanical Engineering
Massachusetts Institute of Technology
ksmetana@mit.edu

Phuong Huynh, David Knezevic
Akselos
phuong.huynh@akselos.com, david.knezevic@akselos.com

Anthony T. Patera
Massachusetts Institute of Technology
Department of Mechanical Engineering
patera@mit.edu

CP18

A Variational Multi-Scale Approach Using Linear Simplicial Finite Elements for Transient Viscoelastic Solid Mechanics

We present a variational multi-scale (VMS) approach for linear and nonlinear transient viscoelastic solid mechanics. The method is stable and second-order accurate on linear simplicial finite elements (including in the incompressible limit), with only one additional equation for pressure. Using a VMS decomposition, we model the fine-scale variables by residual-consistent formulations to maintain the accuracy. We assess the performance of the method by using manufactured solutions, and test it on 3D problems with complex geometries.

Xianyi Zeng, Guglielmo Scovazzi
Duke University
xy.zeng@duke.edu, guglielmo.scovazzi@duke.edu

CP18

High-Order Mixed Finite Elements for a Pressure Poisson Equation Reformulation of the Navier-Stokes Equations with Electric Boundary Conditions

Pressure Poisson equation (PPE) reformulations for the Navier-Stokes equations represent a class of methods that replace the incompressibility constraint by a Poisson equation for the pressure, with a suitable choice of the bound-

ary conditions so that the incompressibility is maintained. In this talk we present a mixed finite element methods for the Shirokoff-Rosales PPE reformulation, and demonstrate that this approach allows for arbitrary order of accuracy both in space and in time.

Dong Zhou
Temple University
dong.zhou@temple.edu

David Shirokoff
New Jersey Institute of Technology
david.g.shirokoff@njit.edu

Benjamin Seibold
Temple University
seibold@temple.edu

Rodolfo R. Rosales
Massachusetts Inst of Tech
Department of Mathematics
rrr@math.mit.edu

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

CP19

Hierarchical Hpk-Adaptivity for Isogeometric Analysis

Input your abstract, including TeX commands, here. Hierarchical h-adaptivity has since been implemented for tensor product B-splines, NURBS, and T-splines in the context of isogeometric analysis. We have implemented hpk-adaptivity in the hierarchical unstructured regime. We present several benchmark examples, which highlight the advantages of local hierarchical refinement, the necessity of multiple types of refinement, and the power of mixed refinement types. We present convergence results for hierarchical hpk-refinements, as well as efficient algorithms for handling hierarchical hpk-adaptivity.

Emily Evans
Mathematics Department
Brigham Young University
ejevans@math.byu.edu

Kevin Tew
Brigham Young University
kevin_tew@byu.edu

CP19

A Stencil Based Finite Element Method

A Stencil based FEM approach will be introduced which employs tensor-product grids with FEM discretizations. This approach eliminates costly FEM assembly, allows for stencil based computations which further reduces the need for sparse matrix storage and indirect memory access, and moreover is highly suitable for efficient geometric multigrid solvers. Complex boundaries and immersed (and moving) interfaces are treated locally with a new macro grid alignment technique. Examples and computational results will be presented to show the computational effectiveness of the approach.

Johan S. Hysing
Tokto Inst. of Technology
Aoki Laboratory
johan@sim.gsic.titech.ac.jp

CP19

An Anchored Analysis of Variance Petrov-Galerkin Projection Scheme for a Class of High Dimensional Elliptic Partial Differential Equations

High dimensional operator equations from state space estimation, molecular dynamics, and mathematical finance present an interesting numerical challenge because their dimensionality necessitates schemes beyond classical mesh based methods. In this work, we propose a novel *anchored separation of variables* function decomposition in conjuction with a Petrov-Galerkin projection scheme to overcome this "curse of dimensionality'. A class of model high dimensional elliptic partial differential equations are considered and numerical results using this nonlinear approximation are presented.

Matthew T. Li
University of Toronto Institute for Aerospace Studies
matt.li@mail.utoronto.ca

Christophe Audouze
University of Toronto
Institute for Aerospace Studies
c.audouze@utoronto.ca

Prasanth B. Nair
University of Toronto
pbn@utias.utoronto.ca

CP20

Ground States and Dynamics of Spin-Orbit-Coupled Bose-Einstein Condensates

We study analytically and asymptotically as well as numerically ground states and dynamics of two-component spin-orbit-coupled Bose-Einstein condensates (BECs) modeled by the coupled Gross-Pitaevskii equations (CGPEs). In fact, due to the appearance of the spin-orbit (SO) coupling in the two-component BEC with a Raman coupling, the ground state structures and dynamical properties become very rich and complicated.

Yongyong Cai
University of Wisconsin-Madison
yongyong.cai@gmail.com

Weizhu Bao
Dept. of Mathematics and Center for Computational Science an
matbaowz@nus.edu.sg

CP20

Fast Ewald Summation for Mixed Periodic Boundary Conditions Based on the Nonuniform Fft

We introduce a generalization of the NFFT based fast Ewald summation to mixed periodic boundary conditions. In our approach, we combine the corresponding Ewald formulas with the NFFT based fast summation. The new algorithms can be tuned to high accuracy and the perfor-

mance can be compared to those of well established methods for the fully periodic case (P^3M, P^2NFFT). In our talk we will present the main ideas and show numerical results.

Franziska Nestler, Michael Pippig
Chemnitz University of Technology
Department of Mathematics
franziska.nestler@mathematik.tu-chemnitz.de,
michael.pippig@mathematik.tu-chemnitz.de

CP20

Parallel Replica Dynamics with Spatial Parallelization for a Driven System

We explore the parallel efficiency and speedup for parallel replica dynamics (PRD) with spatial parallelization. In traditional PRD, each replica is assigned to an individual processor, which extends the time domain for the simulation. By adding spatial parallelization of replicas, we are able to simulate larger systems over the same extended time domain. Numerical results on a driven system indicate that this approach can lead to more efficient solutions for realistic physical systems.

Michael T. Stobb, Juan Meza
University of California, Merced
mstobb@ucmerced.edu, jcmeza@ucmerced.edu

Ashlie Martini
University of California
Merced
amartini@ucmerced.edu

CP21

Simulating Non-Dilute Transport in Porous Media Using a Tcat-Based Model

Predicting the transport of non-dilute species in fluids of variable density in porous media is a challenging problem. We use a thermodynamically constrained averaging theory (TCAT)-based model, which consists of a flow equation, a species transport equation, and closure relations. We rewrite the model as a system of two partial differential-algebraic equations. We use a stiff temporal integrator to perform 1D simulations. The model is nonlinear and non-smooth. We will discuss results and numerical difficulties.

Deena H. Giffen
North Carolina State University
dhannou@ncsu.edu

CP21

Hierarchical Model Reduction of the Navier-Stokes Equations for Incompressible Flows in Pipes

Hierarchical Model Reduction is a novel technique designed for reducing computational costs when solving incompressible flows in networks of pipes. It consists of a separate discretization of the axial and the transversal components of the flow. The former is approximated by finite elements, the latter by spectral methods. The local accuracy of the transversal discretization can be adaptively modulated. In this talk we present analysis and numerical results, having hemodynamics as reference application.

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

Simona Perotto
MOX - Modeling and Scientific Computing
Dipartimento di Matematica
simona.perotto@polimi.it

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

CP21

Interaction Between Toroidal Swimmers in Stokes Flow

The focus of this research has been devoted to study the interaction between two or more self-propelled toroidal swimmers in Stokes flow by applying the method of regularized Stokeslets. In the study of the interaction between two or more toroidal swimmers, we interpret these as three-dimensional, zero Reynolds number analogues of finite vortex dipoles in an ideal fluid. Then, we examine the stability of relative equilibria that can form for these swimmers when they are initially placed in tandem or abreast.

Jianjun Huang
Worcester Polytechnic Institute
jhuang@wpi.edu

Lisa J. Fauci
Tulane University
Department of Mathematics
fauci@tulane.edu

CP21

Brownian and Hydrodynamic Motion of Complex Shaped Particles in Straight and Branching Blood Vessels

We develop a computational methodology for the study of the motion of complex shaped nanoparticles within straight and branching vessels subject to thermal fluctuations. A framework based on Markovian fluctuating hydrodynamics of the fluid together with a non-Markovian Langevin dynamics perturbing motion of the particle is adopted. An important application of the method is the transport of nanocarriers within a blood vessel for targeted drug delivery. Supported by NIH through grant U01-B016027.

Yaohong Wang
Department of Mathematics
UCSB
yaohong@seas.upenn.edu

David Eckmann
Department of Anesthesiology and Critical Care
University of Pennsylvania
david.eckmann@uphs.upenn.edu

Ravi Radhakrish
Department of Bioengineering
University of Pennsylvania
rradhak@seas.upenn.edu

Helena Vitoshki, Portonovo Ayyaswamy
Department of Mechanical Engineering and Applied Mechanics
University of Pennsylvania

elenavit@seas.upenn.edu, ayya@seas.upenn.edu

CP21

An ALE-Phase-Field Method for Dynamic Wetting of Moving Particles

A hybrid method that uses an arbitrary Lagrangian-Eulerian technique to track solid particles and a phase-field method to capture fluid interfaces as well as moving contact lines is developed. The Navier-Stokes and the Cahn-Hilliard equations are are solved by a mixed finite element method on an adaptive triangular moving mesh. Numerical results on the interactions between floating particles, also known as the cheerio effect, and the effect of dynamic wetting in water-entry problems will be presented.

Pengtao Yue
Virginia Polytechnic Institute and State University
ptyue@math.vt.edu

CP22

Anchored ANOVA Petrov-Galerkin (AAPG) Projection Schemes for Parabolic Stochastic Partial Differential Equations

We present an intrusive method based on the combination of Hoeffding functional ANOVA decomposition with stochastic Galerkin projection for solving a class of high-dimensional parabolic stochastic PDEs. Enforcing the component functions of the approximate solution to be orthogonal with respect to an appropriate measure and using adapted test functions, the stochastic weak formulation is decoupled into independent low-dimensional subproblems. An a priori error analysis and numerical studies for stochastic diffusion models are provided.

Christophe Audouze
University of Toronto
Institute for Aerospace Studies
c.audouze@utoronto.ca

Prasanth B. Nair
University of Toronto
pbn@utias.utoronto.ca

CP22

Stochastic Low-Dimensional Modeling of Natural Convection Using Dynamically Orthogonal Decomposition

An efficient numerical method for studying the effect of stochastic parameters on natural convection is presented. In this methodology, the solution is approximated by a *generalized* Karhunen-Loeve expansion. The elements of the basis remain orthogonal for all times and they evolve according to the system dynamics to capture the energetically dominant stochastic subspace. The stochasticity can be introduced at the boundary conditions and source terms, a problem setup that includes a wide range of engineering problems.

Hessameddin Babaee
MIT
babaee@mit.edu

CP22

Local Polynomial Chaos Expansion for Linear Differential Equations with High Dimensional Random Inputs

We present a localized PC expansion for PDEs with random inputs, where most existing methods incur prohibitively high simulation cost. The local polynomial chaos method employs a domain decomposition technique to approximate the stochastic solution locally in a much lower dimensional random space. Our method applies the coupling conditions at the interfaces of the subdomains along with accurate samples to ensure both accuracy and high efficiency. We present the general mathematical framework of our methodology.

Yi Chen
Purdue University
chen411@purdue.edu

Dongbin Xiu
University of Utah
dongbin.xiu@utah.edu

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

claude gittelson
ETH Zurich
claude.gittelson@sam.math.ethz.ch

CP22

Variance Reduction in the Simulation of Stochastic Differential Equations

Variance reduction techniques are commonly used to enhance the efficiency of Monte Carlo simulations. This talk focuses on variance reduction for single and coupled systems of stochastic ordinary differential equations. Variance reduction techniques such as antithetic variates and control variates will be described and results presented.

David J. Horntrop
Dept of Mathematical Sciences, Center for Applied Math
New Jersey Institute of Technology
david.horntrop@njit.edu

CP22

Fully Implicit Runge-Kutta Methods for Multi-Channel Stiff Stochastic Differential Systems with Jumps

We discuss systems of ordinary SDEs with non-commutative multi-channel noise including jump-diffusion processes. Such systems arise in biochemical networks that involve reactions at different time scales. They are inherently stiff, both in deterministic and stochastic components, and change their stiffness with uncertainty. To resolve this issue we consider fully implicit split-step stochastic balanced Runge-Kutta methods and investigate their convergence, stability and positivity preserving properties. Numerical examples are provided to show the effectiveness of these methods.

Viktor Reshniak, Abdul Khaliq
Middle Tennessee State University
vr2m@mtmail.mtsu.edu, abdul.khaliq@mtsu.edu

Guannan Zhang
Oak Ridge National Laboratory

zhangg@ornl.gov

David A. Voss
Western Illinois University
d-voss1@wiu.edu

CP23

Space-Time Adaptive Multiresolution Simulations of the Compressible Euler Equations

Fully space-time adaptive multiresolution simulations of the compressible Euler equations applied to 2d and 3d Riemann problems will be presented. A new higher order local time stepping strategy is also proposed. The computational efficieny in terms of CPU time and memory compression will be assessed. The accuracy of the adaptive simulations with respect to fine grid computations will be studied.

Margarete O. Domingues
Instituto Nacional de Pesquisas Espaciais (INPE)
margarete.oliveira.domingues@gmail.com

Muller Lopes, Odim Mendes
INPE, Sao Jose dos Campos, Brazil
mullermslopes@gmail.com, odim.mendes@inpe.br

Kai Schneider
Universite de Provence, Aix-Marseille
Centre de Mathematiques et Informatique
kschneid@cmi.univ-mrs.fr

CP23

Cubic B-Spline Quasi-Interpolation Based Numerical Scheme for Hyperbolic Conservation Laws

The present work analyzes the Cubic B-Spline Quasi-Interpolation (CBSQI) based explicit numerical schemes for hyperbolic conservation laws. To improve the stability of the proposed numerical scheme, we modify the CBSQI scheme by adding an artificial diffusion to the numerical scheme through a diffusion parameter b. We derive a relation between the CFL condition and b, which ensures the monotonicity of the proposed numerical scheme. Further, the L^∞-error and the TVD property of the modified BSQI scheme are established.

Rakesh Kumar
IIT Bombay
rakeshk@math.iitb.ac.in

Sambandam Baskar
Indian Institute of Technology Bombay
baskar@math.iitb.ac.in

CP23

A Space-Time Finite Volume Differencing Method for Robust Higher Order Schemes for Transport Equations

A space-time discretization method is applied to construct robust higher-order schemes for transport equations. A unified space-time error for integral formulations are constructed using general weighted quadratures for flux integrals. Efficient quadrature approximations of sources are then sought to account for local space-time fluxes through a constrained minimization of error. Residual errors are then utilized to guide right time-steps to ensure error diminishing discretizations. Results for 1D heat and hyperbolic conservation laws will be demonstrated.

Yaw Kyei
North Carolina A&T State University
Greensboro, NC 27401
ykyei@ncat.edu

CP23

A Runge-Kutta Discontinuous Galerkin Method for Modeling Storm-Water Flow in Networks of Drainage Channels

A unique hybrid 1D/2D approach is presented to model water flow in networks of channels that naturally exist in coastal areas. The governing 1D and 2D shallow water equations are discretized using an RKDG method. Flows in individual channel branches are treated as 1D flows. These branches are conservatively coupled together at junctions using 2D elements. The accuracy of the model is established by comparing the simulation results to experimental data and other numerical models.

Prapti Neupane
Institute for Computation, Engineering and Science
University of Texas at Austin
prapti@ices.utexas.edu

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

CP23

Lagrangian Particle Method for Complex Flows

A new Lagrangian particle method improving the smooth particle hydrodynamics (SPH) has been developed for equations of compressible flows. The method eliminates two major deficiencies of SPH: the dependence on parameter called the smoothening length and the presence of large linear errors of SPH differential operators. Particle-based stable, high order upwinding schemes have been developed using moving weighted least squares. Rigorous verification tests and applications to complex free surface flows will be discussed.

Roman Samulyak, Hsin-Chiang Chen, Wei Li
Stony Brook University
rosamu@ams.sunysb.edu, hcchen@ams.sunysb.edu, wli@ams.sunysb.edu

CP24

Quantifying Scale Coupling and Energy Pathways in the Ocean

The oceans display energetic dynamics across a wide range of spatial scales, and researchers have long worked to better understand the energy coupling between these various scales. While there have been previous attempts to understand energy pathways, assumptions of homogeneity and isotropy have presented a limitation upon the applicability of the analyses. Here we present a more general technique, unrestricted by the usual assumptions of homogeneity or isotropy, which allows one to simultaneously probe the dynamics in both space and time. We makes use of a novel coarse-graining framework, which accounts for the spheri-

cal geometry of the problem, to directly analyze the coupling between scales. We apply this technique to strongly eddying high-resolution simulations using LANLs Parallel Ocean Program. We examine the extent to which the traditional paradigm for such pathways is valid at various locations such as in western boundary currents, near the equator, and in the deep ocean.

Hussein Aluie, Matthew Hecht
Los Alamos National Laboratory
hussein@rochester.edu, mhecht@lanl.gov

Geoffrey Vallis
University of Exeter
g.vallis@exeter.ac.uk

CP24

A Numerical Simulation of the Sediment Dynamics in a Three Dimensional Fluid Flow

In this study we use NaSt3D as fluid solver for incompressible two-phase flow problems in three dimensions. We apply this fluid solver to the problem of sediment transport processes. The main parts in sediment transport are bed load and suspension load. Both parts are calculated from the fluid velocities and are used to compute the transport of sediment masses. Single phase examples like dunes and ripples as well as two-phase phenomena like scouring at an obstacle can be reproduced by this model.

Markus Burkow
Institute for Numerical Simulation
University Bonn
burkow@ins.uni-bonn.de

Michael Griebel
Universitat Bonn
Inst fur Angewandte Mathematik
griebel@ins.uni-bonn.de

CP24

Openfoam Implementation of a New Subgrid-Scale Model for Large Eddy Simulation

We have recently proposed a novel subgrid scale model based on random vortex structures for large eddy simulation of the ocean. It is developed for homogeneous and incompressible flows. In this study, we validate the model by numerical simulations and compare with well-known subgrid scale models and direct numerical simulation. Turbulent channel flow is solved at different Reynolds numbers using OpenFOAM software. The results indicate the strengths of our model and the directions for improvement.

Rukiye Kara
Department of Mathematics
Mimar Sinan Fine Arts University
rukiye.kara@msgsu.edu.tr

Mine Caglar
Department of Mathematics
Koc University
mcaglar@ku.edu.tr

CP24

Free Surface Waves on a Horizontal Shear Flow

Free surface waves on a non-uniform mean flow are considered. The mean flow $U(y)$ varies with the transverse coordinate y but not the vertical. The domain is bounded on one side by a flat rigid vertical wall and unbounded on the other side. The mean flows considered are nonzero near the vertical wall and approach zero far from the wall, e.g. $U = e^y$. For large y where the mean flow is near-zero the waves are merely irrotational Stokes' waves. Near the wall the mean flow and the waves are rotational but still inviscid. Solutions are obtained using a nonuniform coordinate transformation that converts the free surface boundary condition into a modified Bessel equation. The solution for linear waves that are periodic along the wall is an expansion in Bessel of imaginary argument and imaginary order. Eigenvalues are found numerically.

John P. Mchugh
University of New Hampshire
john.mchugh@unh.edu

Gary Lapham
Maine Maritime Academy
gary.lapham@mma.edu

CP24

Three-Dimensional Wavelet-Based Adaptive Mesh Refinement Algorithm for Numerical Simulation of Atmospheric Global Chemical Transport

Accurate numerical modeling of multi-scale Global Chemical Transport Models (GCTMs) is a challenging task. Here we present Wavelet-based Adaptive Mesh Refinement (WAMR) method that allows efficient numerical simulation of the GCTMs by permitting two-three orders of magnitude finer local resolution than static-grid GCTMs for the same number of grid points. Therefore, WAMR provides a realistic opportunity to model efficiently challenging multi-scale GCTMs on existing computers by producing accurate results at a relatively low computational cost. Supported by NSF grant HRD-1036563

Artem N. Semakin, Yevgenii Rastigejev
North Carolina A&T State University
arte-semaki@yandex.ru, ye_rast@yahoo.com

CP25

Tensor Rank Prediction via Cross Validation

The use of higher order tensors in machine learning has become increasingly popular in recent years. Traditional techniques like clustering and principle component analysis can be extended to N-way data through the Canonical Polyadic (CP) tensor decomposition, but only if the rank is known ahead of time. Although computing the rank is NP-hard, we approximate it using cross validation, a statistical technique used to choose model complexity. We present results for dense, normally distributed data.

Woody N. Austin
University of Texas - Austin
austinwn@cs.utexas.edu

Tamara G. Kolda, Todd Plantenga

Sandia National Laboratories
tgkolda@sandia.gov, tplante@sandia.gov

CP25

Lu and Partial Orthogonalization Preconditioning for Conjugate Gradient Solution of Overdetermined Sparse Least Squares Problems

Let $PLU = AQ$ be a decomposition of a sparse rank n matrix A with m rows and n columns, $m > n$. Conjugate gradient iteration on the system $L^T L$ can be used to find x minimizing $||Ax - b||_2$. We compare LU and partial orthogonalization preconditionings and also a hybrid scheme using both techniques. Automated conversion to C++ from Matlab and Octave scripts is explored.

Gary W. Howell
North Carolina State University
gary_howell@ncsu.edu

Marc Baboulin
University of Paris-Sud/INRIA
marc.baboulin@inria.fr

CP25

On a priori and a posteriori Eigenvalue/eigenvector Error Estimates for Nonlinear Eigenvalue Problems

In this talk we present the recent results in finite element approximations for nonlinear eigenvalue problems, with nonlinearity in the spectral parameter. New a priori and a posteriori eigenvalue/eigenvector error estimates are introduced and verified using the Residual Inverse Iteration Method (RINVIT) for nonlinear eigenvalue problems. Various numerical examples arising from applications in structural mechanics and electromagnetics are discussed to display the performance of our approach. This is a joint work with Daniel Kressner (EPF Lausanne).

Agnieszka Miedlar
TU Berlin
miedlar@math.tu-berlin.de

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

CP25

Efficient Low-Rank Solutions of Generalized Lyapunov Equations

An iterative method for the low-rank approximate solution of a class of generalized Lyapunov equations is studied. At each iteration, a standard Lyapunov is solved using Galerkin projection with an extended Krylov subspace method. This Lyapunov equation is solved inexactly, thus producing a nonstationary iteration. The inexactness criteria for convergence is provided by a new theorem. These tools, together with others presented, comprise an efficient algorithm. Numerical experiments indicate that this method is competitive vis-à-vis the current state-of-the-art methods, both in terms of computational times and storage needs.

Daniel B. Szyld
Temple University
Department of Mathematics
szyld@temple.edu

Setephen Shank
Massachussets Instittute of Technology
sshank@mit.edu

Valeria Simoncini
Universita' di Bologna
valeria.simoncini@unibo.it

CP25

An Implementation and Analysis of the Refined Projection Method For (Jacobi-)Davidson Type Methods

The computation of interior eigenvalues of large sparse matrices remains a challenging problem. Compared to the Rayleigh-Ritz projection, the refined projection method is a more effective way to extract Ritz pairs and achieve monotonic convergence but with a much higher computational cost. We analyze four different implementations of refined projection and present a new efficient approach to compute interior eigenvalues accurately for (Jacobi-)Davidson type methods. Numerical experiments demonstrate the effectiveness and accuracy of the presented method.

Lingfei Wu
Department of Computer Science
College of William & Mary
lfwu@cs.wm.edu

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

CP26

Updating and Downdating Techniques for Networks

The total communicability of a network is the sum of the entries in the exponential of its adjacency matrix. This quantity offers a good measure of connectivity and can be useful in the design of networks having certain desirable properties. I will discuss algorithms that can be used to construct networks that are sparse and have a large total communicability. Computational results will be provided.

Francesca Arrigo
Università degli Studi dell'Insubria
farrigo87@gmail.com

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

CP26

Dynamic Causal Modelling of Brain-Behaviour Relationships

Dynamic Causal Modelling (DCM) of neuroimaging data has become a standard tool for identifying the structure and plasticity of brain networks that respond to the experimental manipulation (e.g., sensory stimuli or task de-

mands). DCM, however, does not explain how distributed brain responses are causally involved in the production of behaviour. Here, we propose a generic extension of DCM that captures how experimental manipulations are transformed, through large-scale brain networks, into behaviour (e.g. choices, reaction times).

Jean Daunizeau, Lionel Rigoux
ICM
jean.daunizeau@gmail.com, lionel.rigoux@gmail.com

CP26

Using Space Filling Curves to Find An Element That Contains a Given Point

Many techniques used in computational science and engineering use a 2D triangulation or 3D tesselation of given region, for example finite element analysis or computer graphics. Often one needs to locate a triangle or tetrahedron that contains a given point. For example, to evalute a finite element solution at an arbitrary point, a containing element must be located. We present a new fast algorithm for this operation which is based on space filling curves.

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

CP26

Topology Backs Holistic Medicine

The holistic concept in alternative medical practice upholds that "all of people's needs should be taken into account." In other words, the body is seen as a whole. The holistic point of view can be scientifically validated by determining whether different bodily variables are related to one another. In this study we provide strong evidence supporting this claim. We establish the existence of at least one fully non-linear relationship involving two bodily variables.

Fernando Schwartz
Department of Mathematics
University of Tennessee
fernando@math.utk.edu

Louis Xiang
The Chinese University of Hong Kong
louisxiang1994@gmail.com

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

MS1

Local Methods in Network Science

Local methods for methods for network analysis return a property of the network without looking even at the entire network. Thus, their runtimes are typically sublinear in the size of the input network. I'll discuss recent work on using local methods to compute centrality vectors such as the PageRank vector and the heat kernel vector. These involve approximately solving linear systems and approximate matrix exponentials, respectively. We'll also see how these primitives give highly scalable algorithms for community detection.

David F. Gleich
Purdue University
dgleich@purdue.edu

MS1

Identifying the Largest Entries in Matrix Multiplication

Consider matrices A and B of size $m \times p$ and $p \times n$. We wish to determine the indices of the largest entries in the $m \times n$ product matrix $C = AB$ without explicitly calculating C. For instance, if A represents an adjacency network of social connections, then $C = AA'$ (i.e., $B = A'$) represents the number of common neighbors for any pair of nodes. The largest entries correspond to those pairs with the most common neighbors. Alternatively, it may be the case that A and B are latent variables in a prediction task, and the largest entries in $C = AB$ correspond to the most likely pairings. The matrices A and B may be dense with $p \ll m, n$ or general sparse matrices (in which case C may or may not be dense). We propose a sampling-based method to efficiently identify the largest entries in C. Each sample produces a pair (i, j), and the pairs that are sampled most frequently correspond to the largest entries in C in expectation. Specifically, the probability of choosing pair (i, j) is proportional to c_{ij}^2. The number of samples for a given confidence level is independent of the size of the matrices, and in practice $s \ll mn$. The cost of the method is $O(\text{nnz}(A) + \text{nnz}(B))$ for preprocessing and $O(\log(\text{nnz}(A)) + \log(m) + \log(n))$ per sample.

Grey Ballard, Tamara G. Kolda
Sandia National Laboratories
gmballa@sandia.gov, tgkolda@sandia.gov

Ali Pinar
Sandia National Labs
apinar@sandia.gov

C. Seshadhri
Univ. California Santa Cruz
csesha@gmail.com

MS1

Mining Uncertain Networks

Abstract not available at time of publication.

Evamaria Terzi
Boston University
evimaria@cs.bu.edu

MS1

Network Science of Brain Networks

Abstract not available at time of publication.

Zoltan Toroczkai
University of Notre-Dame
toro@nd.edu

MS2

Convex Biclustering

In the biclustering problem, we seek to simultaneously group observations and features. We present a convex

formulation of the biclustering problem that possesses a unique global minimizer and an iterative algorithm, COBRA, that is guaranteed to identify it. The key contributions of our work are its simplicity, interpretability, and algorithmic guarantees. We demonstrate the advantages of our approach, which includes stably and reproducibly identifying biclusterings, on simulated and real microarray data.

Eric Chi
Rice University
echi@rice.edu

Richard G. Baraniuk
Rice University
Electrical and Computer Engineering Department
richb@rice.edu

Genevera Allen
Rice University
tba@rice.edu

Ronald Coifman
Yale University
Department of Computer Science
coifman@math.yale.edu

MS2

Cluster-based Reduced-order Modelling: From Shear Flows to Engine Tumble Motion

We propose a cluster-based ROM strategy to distil nonlinear mechanisms in an unsupervised manner. This strategy uses cluster analysis to partition snapshot data into a small number of representative states in the state space. The transitions between the states are dynamically modelled as Markov process. CROM has potential applications for the systematic identification of physical mechanisms of complex dynamics, for the identification of precursors to desirable and undesirable events, and for flow control design exploiting nonlinearities.

Eurika Kaiser
Institute PPRIME, CNRS
eurika.kaiser@univ-poitiers.fr

Bernd R. Noack
Institut PPRIME, CNRS
bernd.noack@univ-poitiers.fr

Laurent Cordier, Laurent Cordier
Institute PPRIME, CNRS
laurent.cordier@univ-poitiers.fr,
laurent.cordier@univ-poitiers.fr

Andreas Spohn
Institute PPRIME, ENSMA
andreas.spohn@ensma.fr

Marc Segond, Markus W Abel
Ambrosys GmbH
marc.segond@ambrosys.de, markus.abel@ambrosys.de

Guillaume Daviller
CERFACS, France
daviller@cerfacs.fr

Jan sth, Sinisa Krajnovic
Chalmers University of Technology
jan.osth@chalmers.se, sinisa@chalmers.se

Yujun Cao
Institute PPRIME, ENSMA
RENAULT, France
yujun.cao@ensma.fr

Jacques Boree
Institute PPRIME, ENSMA
jacques.boree@ensma.fr

Robert K. Niven
UNSW/ADFA, Australia
r.niven@iinet.net.au

Louis N. Cattafesta
FCAAP, Florida State University, USA
cattafes@gmail.com

MS2

Self-Tuning Complex Systems

Abstract not available at time of publication.

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

MS2

The Impact of L1 optimization in Nonlinear PDE

At this time almost everyone interested in finding sparse solutions to discrete equations is aware that l1 optimization plays a key role. However that fact that L1 optimization, using the techniques developed for l1, is a very powerful tool in nonlinear PDE and numerical analysis is less widely known. H. Brezis, in 1974, showed that adding an L1 type penalty to calculus of variations problems which lead to elliptic equations plus a signum type terms gives solutions with compact support. I will discuss this, numerical aspects, applications to Schrodinger equations, obstacle problems, numerical homogenization and certain high dimensional PDE's.

Stanley J. Osher
University of California
Department of Mathematics
sjo@math.ucla.edu

MS3

Modulus of Families of Walks on Graphs

The modulus of a family of paths in a continuum provides a quantitative assessment of the "richness' of the family: large families of short paths have larger modulus than small families of long paths. In the discrete setting, the concept of modulus can be linked to several graph-theoretic quantities including shortest path, minimum cut, and effective resistance. This talk will cover connections among these concepts, some applications, and a numerical algorithm for computing the modulus.

Nathan Albin
Kansas State University

albin@math.ksu.edu

MS3

Graph Directed Topic Modeling

Recent advancements in dictionary learning or topic modeling algorithms find data representations using sparse coding techniques such as L1 or non-parametric Bayesian regularizations. Concurrently, the graph community has developed algorithms to segment data into clusters using similarity measures. In this work, we integrate the graphical models and dictionary learning concepts to infer representations directed by a graphical structure, enforcing prior correlations among documents. The utility of this technique is demonstrated on text and image data.

Arjuna Flenner
Naval Air Weapons Station
aflenner@mac.com

Cristina Garcia-Cardona
Los Alamos National Laboratory
cgarciac@lanl.gov

MS3

Building Graphs to Analyze Big Data

There has been increasing demand to understand the data around us. The flood of social media requires new mathematics, methodologies and procedures to extract knowledge from massive datasets. Spectral methods are graph based techniques that uses eigenfunctions of a graph to extract the underlying global structure of a dataset. The construction of these, application dependent, graphs require new mathematical ideas that extend data representation, distance, topic modeling and sparsity. The product is often massive matrices that push the limits of matrix computation. This talk looks at applications to analyzing text, images, Twitter microblogs and content based search.

Blake Hunter
UCLA Mathematics Department
bhunter@cmc.edu

MS3

An Incremental Reseeding Strategy for Clustering

In this work we propose a simple and easily parallelizable algorithm for multiway graph partitioning. The algorithm alternates between three basic components: diffusing seed vertices over the graph, thresholding the diffused seeds, and then randomly reseeding the thresholded clusters. We demonstrate experimentally that the proper combination of these ingredients leads to an algorithm that achieves state-of-the-art performance in terms of cluster purity on standard benchmarks datasets. Moreover, the algorithm runs an order of magnitude faster than the other algorithms that achieve comparable results in terms of accuracy. This a joint work with X. Bresson, H. Hu, A. Szlam and J. von Brecht.

Thomas Laurent
Loyola Marymount University
thomas.laurent@lmu.edu

MS4

Adaptive h-refinement for Reduced-order Models via Basis Splitting

This talk discusses a method to adaptively refine reduced-order models *a posteriori* without requiring additional full-order-model solves. The technique is analogous to mesh-adaptive h-refinement: it enriches the reduced-basis space online by 'splitting' a given basis vector into several vectors with disjoint support. The splitting scheme is defined by a tree structure constructed offline via recursive k-means clustering of the state variables using snapshot data. The method identifies the vectors to split online using a dual-weighted-residual approach that aims to reduce error in an output quantity of interest. The resulting method generates a hierarchy of subspaces online without requiring large-scale operations or full-order-model solves. Further, it enables the reduced-order model to satisfy *any prescribed error tolerance* regardless of its original fidelity, as a completely refined reduced-order model is mathematically equivalent to the original full-order model.

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

MS4

Online-Adaptive Reduced Bases for Parametric Problems

We address the topic of projection-based model order reduction for parametric systems. If the solution manifold under varying parameters or time is complex, single reduced projection spaces are not sufficient to achieve global accurate approximation. Locality with respect to parameter or time for generating submodels in the offline phase is a solution but may be misleading or yield redundant models. Therefore, we propose to generate reduced models by online-greedy procedures from dictionaries of basis elements.

Bernard Haasdonk
University of Stuttgart
haasdonk@mathematik.uni-stuttgart.de

MS4

A Nonlinear Trust Region Framework for PDE-Constrained Optimization Using Progressively-Constructed Reduced-Order Models

The large computational cost associated with high-fidelity simulations has limited their use in many-query scenarios (optimization and UQ). A nonlinear trust-region framework using Reduced-Order Models (ROMs) is introduced as a means to accelerate PDE-constrained optimization. A progressive approach is employed to construct a ROM during the optimization procedure. On problems from aerodynamic shape optimization, the framework reduces the number of queries to the high-dimensional model by a factor of 4 with no loss in accuracy.

Matthew J. Zahr
Stanford University
University of California, Berkeley
mzahr@stanford.edu

Charbel Farhat
Stanford University

CFarhat@stanford.edu

MS5

A Bound-Plus-Equality Constrained Quadratic Minimization Algorithm for Inverse Problems

Solutions of inverse problems are often obtained by solving a quadratic minimization problem, e.g., the least squares solution in linear cases. Prior information can be incorporated into such problems by adding constraints to the quadratic minimization. Moreover, L^1 regularization methods, such as total variation and wavelet-based regularization, can be implemented by solving a constrained quadratic minimization problem. In this talk, we present an algorithm for solving bound-plus-equality constrained quadratic minimization problems, with applications in imaging.

Johnathan M. Bardsley
University of Montana
bardsleyj@mso.umt.edu

Marylesa Howard
National Security Technologies, LLC
howardmm@nv.doe.gov

MS5

Statistically Motivated Preconditioners and Stopping Criteria for Biomedical Inverse Problems

The solution of large-scale, ill-posed, possibly underdetermined linear systems of equation arises in many application to linear and nonlinear biomedical inverse problems. In this talk we show how statistical a-priori expectation about the solution and statistical description of the noise in the right-hand side can be exploited to improve the resolution of the linear system and design solid stopping rules for Krylov-type iterative methods.

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS5

Numerical Implementation of a New Class of Forward-Backward-Forward Diffusion Equations for Image Restoration

In this talk, we present the implementation and numerical experiments demonstrating new forward-backward-forward nonlinear diffusion equations for noise reduction and deblurring, developed in collaboration with Patrick Guidotti and Yunho Kim. The new models preserve and enhance the most desirable aspects of the closely-related Perona-Malik equation without allowing staircasing. By using a Krylov subspace spectral (KSS) method for time-stepping, the properties of the new models are preserved without sacrificing efficiency.

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

Patrick Guidotti
University of California at Irvine
gpatrick@math.uci.edu

Yunho Kim
Ulsan National Institute of Science and Technology
Department of Mathematical Sciences
yunhokim@unist.ac.kr

MS5

Recycling Krylov Subspaces for Parametric Linear Systems Arising from Hyperspectral Diffuse Optical Tomography

The imaging of chromophore concentrations using Diffuse Optical Tomography (DOT) data can be mathematically described as an ill-posed and non-linear inverse problem. The reconstruction algorithm for hyperspectral data even using a linearized Born model is prohibitively expensive, both in terms of computation and memory. We discuss novel computational strategies for reducing the computational cost based on a recycling Krylov subspace approach for a class of parameteric linear systems. We will demonstrate the resulting computational gains and the validity of our approach by comparison with synthetic experiments.

Arvind Saibaba
Department of Electrical and Computer Engineering
Tufts University
arvind.saibaba@tufts.edu

Misha E. Kilmer, Eric Miller
Tufts University
misha.kilmer@tufts.edu, elmiller@ece.tufts.edu

MS6

The Dirichlet-Neumann Iteration and Unsteady Thermal Fluid Structure Interaction

We consider unsteady thermal fluid structure interaction to model industrial gas quenching in steel forging, where steel is cooled using high pressured gas. The models are the compressible Navier-Stokes equations and the nonlinear heat equation. In time, a previously developed efficient adaptive higher order time integration scheme with linear extrapolation within a partitioned Dirichlet-Neumann framework for the coupling is considered. The iteration is surprisingly fast and we present some analysis that explains this phenomenon.

Philipp Birken
Department 10, Mathematics and Natural Sciences
University of Kassel, Germany
philipp.birken@na.lu.se

Azahar Monge Sanchez
Lund University
Centre for the Mathematical Sciences
azahar.monge@na.lu.se

MS6

Multi-level Acceleration of Strongly Coupled Fluid-structure Interaction with Manifold Mapping

Strongly coupled partitioned fluid-structure interaction problems require multiple coupling iterations per time step. In order to reduce the number of coupling iterations, the

manifold mapping algorithm is applied, which originates from multi-fidelity optimization. Preliminary computations showed a gain of factor four in terms of the number of coupling iterations with two grid levels. To gain a larger speedup, the use of more than two levels is investigated for a three-dimensional engineering case, namely a hydrofoil.

David Blom
Delft University of Technology
d.s.blom@tudelft.nl

Alexander H. van Zuijlen, Hester Bijl
Faculty Aerospace Engineering
Delft University of Technology, NL
a.h.vanzuijlen@tudelft.nl, h.bijl@tudelft.nl

MS6

Multirate GARK Schemes

Multirate GARK schemes define a multirate extension of GARK schemes, generalized additive Runge-Kutta schemes. These allow for exploiting multirate behaviour in both the right-hand sides and in the components in a rather general setting, and are thus especially useful for coupled problems in a multiphysics setting. We discuss two types of MGARK schemes: IMEX methods, which makes fully use of the different dynamics and stability properties of the coupled system; and fully implicit schemes, which inherit the stability properties from both underlying schemes without any coupling constraint.

Michael Guenther
Bergische Universitaet Wuppertal
guenther@math.uni-wuppertal.de

Adrian Sandu
Virginia Polytechnic Institute and
State University
sandu@cs.vt.edu

MS6

Partitioned Fluid-Structure Interaction on Massively Parallel Systems

Multi-physics applictions such as fluid-structure interaction (FSI) need massively parallel computations to resolve all relevant spatial and temporal scales. Yet, many classical coupling approaches limit the scalability of the overall simulation. A partitioned approach allows to reuse highly scalable single-physics codes. We discuss different milestones towards massively parallel FSI, including inter-solver parallelism, parallel communication, and parallel data mapping. All methods are implemented in the coupling library preCICE. Test cases with different solvers are presented.

Florian Lindner, Miriam Mehl
Universität Stuttgart
florian.lindner@ipvs.uni-stuttgart.de,
miriam.mehl@ipvs.uni-stuttgart.de

Benjamin Uekermann
TU Munich
uekerman@in.tum.de

MS7

Envelopes: Subspace Methods for Efficient Estimation in Multivariate Statistics

An envelope is a nascent construct for increasing efficiency in multivariate statistics. Improvements in efficiency are made possible by recognizing that the data may contain extraneous variation that is immaterial to the purpose of the analysis. This leads to an active subspace – an envelope – for enveloping the material information and thereby reducing variation. Beginning with the multivariate linear model, we will discuss various types envelopes and how they act to improve efficiency.

Dennis Cook
University of Minnesota
dennis@stat.umn.edu

MS7

Active Subspaces in Theory and Practice

We review the theory of the active subspace method from a dimensionality reduction perspective. We describe how a space with lower dimension than the input space can be constructed using gradient information for the quantities of interest. Bounds are constructed that relate the error in reduced order models of the quantities of interest to the SVD of the Covariance matrix of the input-output Jacobian. Applications from aerospace engineering and oil reservoir simulation will be presented.

Eric Dow
Massachusetts Institute of Technology
ericdow@mit.edu

MS7

Mathematical Foundations of Subspace Selections

In this talk we present and analyze algorithms which allow us to detect the relevant subspaces characterizing a certain nonlinear phenomenon, exclusively from random collections of relatively few data. We discuss the theoretical guarantees given by these algorithms both for samples drawn at random according to specific clustering distributions and for samples drawn without any specific clustering. We show applications in determining regularization parameters in image denoising, without need of knowing the noise level a priori.

Massimo Fornasier
Technical University Munich
massimo.fornasier@ma.tum.de

Valeriya Naumova
Simula Research Laboratory
valeriya@simula.no

MS7

Order Determination for Dimension Reduction Using An Alternating Pattern of Spectral Variability

Similar eigenvalues in random matrices leads to highly variable eigenvectors; otherwise, eigenvector variability tends to be small. We exploit this phenomenon to estimate the rank of a fixed, unknown matrix. The proposed method combines eigenvalue drops and eigenvector variability to pinpoint the rank. Under general conditions, we establish the consistency of the new estimator. We also compare the proposed method with other order-determination methods

by simulations and in an applied setting.

Bing Li
Department of Statistics
Penn State
bing@stat.psu.edu

Wei Luo
Penn State University
wul132@psu.edu

MS8

Filtering Unstable Quadratic Dissipative Systems

Data assimilation refers to the combination of noisy measurements of a physical system with a model of the system in order to infer the state and/or parameters. In the context of numerical weather prediction the underlying model is typically an unstable dynamical system. Unstable dynamical systems can be stabilized, and hence an estimate of the solution recovered from noisy data, provided two conditions hold. First, observe enough of the system: in particular, the unstable modes. Second, weight the observed data sufficiently over the model. This talk will illustrate this for the 3DVAR filter applied to three unstable quadratic dissipative dynamical systems of increasing dimension: the Lorenz 1963 model, the Lorenz 1996 model, and the 2D Navier-Stokes equation.

Kody Law
SRI UQ Center, CEMSE, KAUST
kodylaw@gmail.com

Andrew Stuart
Mathematics Institute,
University of Warwick
a.m.stuart@warwick.ac.uk

Daniel Sanz-Alonso
University of Warwick
d.sanz-alonso@warwick.ac.uk

Abhishek Shukla
Warwick
a.shukla@warwick.ac.uk

MS8

Conditions for Successful Data Assimilation in High Dimensions

We show that numerical data assimilation can be successful only if an effective dimension of the problem is not excessive. This effective dimension depends on the noise in the model and the data, and can be moderate even when the number of variables is huge. We analyze several data assimilation algorithms, including particle filters, and show that well-designed particle filters can solve most of those data assimilation problems that can be solved in principle.

Matthias Morzfeld
Department of Mathematics
Lawrence Berkeley National Laboratory
mmo@math.lbl.gov

Alexandre Chorin
Department of Mathematics
University of California at Berkeley
chorin@math.berkeley.edu

MS8

High Dimensional Non-Gaussian Bayesian Inference with Transport Maps

Characterizing high dimensional posterior distributions in the context of nonlinear and non-Gaussian Bayesian inverse problems is a well-known challenging task. A recent approach to this problem seeks a deterministic transport map from a reference distribution to the posterior. Thus posterior samples can easily be obtained by pushing forward reference samples through the map. In this talk, we address the computation of the transport map in high dimensions. In particular, we propose a scalable adaptive algorithm that exploits recent ideas in dimensionality reduction for Bayesian inverse problems.

Alessio Spantini, Youssef M. Marzouk
Massachusetts Institute of Technology
spantini@mit.edu, ymarz@mit.edu

MS8

Ensemble Methods for Large-Scale PDE-Constrained Bayesian Inverse Problems

Sampling techniques are important for large-scale high dimensional Bayesian inferences. However, general-purpose technique such as Markov chain Monte Carlo is intractable. We present an ensemble transform algorithm that is rooted from the optimal transportation theory. The method transforms the prior ensemble to posterior one via a sparse optimization. We develop methods to accelerate the computation of the transformation. Numerical results for large-scale Bayesian inverse problems governed by PDEs will be presented.

Kainan Wang
Texas A&M University
kainan@ices.utexas.edu

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

MS9

Matrix-Free Interior-Point Method for Large Scale Machine Learning Problems

We present a general interior point method for piecewise linear quadratic (PLQ) penalties. These penalties are ubiquitous in signal processing, inverse problems, and machine learning applications; examples include the L2, L1, Huber, Vapnik, hinge loss, elastic net, and many others. We exploit a conjugate representation of these penalties to design an interior point method for the entire class. The representation also gives rise to a calculus that makes it possible to handle compositions, addition, and smoothing of PLQ penalties, as well as exploit specific problem structure. The method is available in an open source package called IPsolve. Future work focuses on matrix free implementations of IPsolve. Joint work with James Burke, Gianluigi Pillonetto, and Dominique Orban.

Aleksandr Aravkin
IBM T.J. Watson Research Center

saravkin@us.ibm.com

MS9

Matrix Free Methods for Large-Scale Nonlinear Constrained Optimization

We present two methods for solving exact penalty subproblems for nonlinear constrained optimization problems on product sets that arise when solving large-scale optimization problems. The first is a novel iterative re-weighting algorithm (IRWA) that iteratively minimizes quadratic models of relaxed subproblems while automatically updating a relaxation vector. The second approach is based on alternating direction augmented Lagrangian (ADAL) technology applied to our setting. The main computational costs of each algorithm are the repeated minimizations of convex quadratic functions which can be performed matrix-free. Both algorithms are globally convergent under loose assumptions, and each requires at most $O(1/\epsilon^2)$ iterations to reach epsilon-optimality of the objective function. Experiments exhibit the ability of both algorithms to efficiently find inexact solutions. However, in certain cases, these experiments indicate that IRWA can be significantly more efficient than ADAL.

James V. Burke
University of Washington
Department of Mathematics
burke@math.washington.edu

MS9

Gauges, Duality, and Phase Retrieval

Gauge functions significantly generalize the notion of a norm, and gauge optimization is the class of problems for finding the element of a convex set that is minimal with respect to a gauge. These conceptually simple problems appear in a remarkable array of applications. I illustrate these ideas in the context of the phase retrieval problem, and show how they lead to new algorithmic approaches for large problems.

Michael Friedlander
Mathematics Department
University of California, Davis
mpf@math.ucdavis.edu

MS9

Anatomy of a Matrix-Free Interior-Point Solver for Convex Optimization

Recent updates to PDCO (Saunders et al.) implement multiple matrix-free and semi-matrix-free options, including CG with constraint preconditioner, a limited-memory LDL factorization preconditioner, and several formulations of the Newton equations as a system or least-squares problem. Most take advantage of regularization. The new object-oriented design lets users select a variant appropriate for their problem, and implement new system solvers. We illustrate the new solvers on large-scale problems.

Dominique Orban
GERAD and Dept. Mathematics and Industrial Engineering
Ecole Polytechnique de Montreal
dominique.orban@gerad.ca

MS10

Difference Potentials Method for Parabolic Models in Irregular Domains

The Difference Potentials Method (DPM) was originally designed as a computationally efficient framework for the numerical approximation of the solutions to elliptic problems in irregular domains in 2D and 3D. Additionally, DPM can handle general boundary conditions with equal ease. Recently DPM was extended with high-order accuracy to parabolic models with variable coefficients and interfaces. I will discuss this extension and illustrate the performance of the approach with several 1D and 2D examples.

Jason Albright
University of Utah
Math Department
albright@math.utah.edu

Yekaterina Epshteyn
Department of mathematics
University of Utah
epshteyn@math.utah.edu

MS10

Multidimensional Embedded Finite Difference Methods which Satisfies Energy Estimates

Abstract not available at time of publication.

Adi Ditkowsky
Department of Applied Mathematics
Tel Aviv University
adid@post.tau.ac.il

MS10

High Order Cut Finite Elements Methods

We will present recent work on high order finite element methods for cut and composite meshes. The common theme is to use Nitsche's method to enforce conditions on boundaries or interfaces, or to solve PDEs on surfaces. The instabilities that may occur can be resolved in various ways for example using a mesh-based element-wise association or by adding stabilization terms to the variational form. Both theoretical and numerical results will be shown.

August Johansson
Simula School of Research and Innovation Simula Research Lab
Norway
august@simula.no

Mats G. Larson
Department of Mathematics
Umea University
mats.larson@math.umu.se

MS10

High-Order Accurate Difference Potentials Methods for the Stokes–Darcy Problem

The Stokes–Darcy problem is a multiphysics model coupling Stokes flow with flow in porous media. In my talk I will present an efficient, high-order accurate numerical

method for this problem. This method, based on the Difference Potentials Method, uses uniform Cartesian grids which do not conform to boundaries or interfaces, and is uniformly high-order accurate. I will present numerical results to test the new method. This talk is based on joint work with Y. Epshteyn.

Yekaterina Epshteyn
Department of mathematics
University of Utah
epshteyn@math.utah.edu

Kyle R. Steffen
University of Utah
Department of Mathematics
steffen@math.utah.edu

MS11

Finite Element Multigrid Framework for Mimetic Finite Difference Discretizations

We are interested in the efficient multigrid method of linear systems of equations discretized from the mimetic finite difference (MFD) schemes which work on general unstructured and irregular grids and result in discrete operators that satisfy the exact sequence connecting grad, div and curl operators on the continuous level. We derive such MFD schemes from the standard finite element spaces. Using the finite element framework, we are able to analyze the convergence of the MFD discretizations and construct efficient multigrid methods for the MFD discretizations of elliptic partial differential equations based on the local Fourier analysis. Finally, we present several numerical tests to demonstrate the robustness and efficiency of the proposed multigrid methods.

Francisco José Gaspar
University of Zaragoza
fjgaspar@unizar.es

Xiaozhe Hu
The Pennsylvania State University
xiaozhe.hu@tufts.edu

Carmen Rodrigo
University of Zaragoza
carmenr@unizar.es

Ludmil Zikatanov
Department of Mathematics
The Pennsylvania State University
ltz@math.psu.edu

MS11

New Multigrid Methods for Saddle Point Problems

We have developed new smoothers for the Stokes and linear elasticity problems. Using the multigrid Poisson solve, we precondition the indefinite system from the finite element discretization of these saddle point problems. We prove the resulting multigrid algorithms are contractions with the contraction number depending on the regularity of the solution but independent of the mesh level.

Susanne Brenner
Department of Mathematics
Louisiana State University
brenner@math.lsu.edu

Hengguang Li
Department of Mathematics
Wayne State University
hli@math.wayne.edu

Li-yeng Sung
Department of Mathematics and CCT
Louisiana State University
sung@math.lsu.edu

MS11

Adaptive Regularization Strategies for Nonlinear PDE

We will discuss an adaptive regularization strategy for stabilizing Newton-like iterations on a coarse mesh, developed in the context of adaptive finite element methods for nonlinear PDE. This methods allows the adaptive algorithm to start on a coarse mesh where the problem data is badly resolved and the linearizations feature indefinite and ill-conditioned Jacobians. Stable configurations of the regularized coarse-mesh iterates are used to refine the mesh, leading to sufficient resolution of the data to accurately approximate the PDE solution. We will discuss the use of a positive semidefinite penalty term which is adapted with both mesh refinements and iterations of the nonlinear solver. Numerical examples demonstrate the effectiveness of the method and illustrate the distinct phases of the solution process.

Sara Pollock
Texas A&M University
snpolloc@math.tamu.edu

MS11

Robust Multilevel Preconditioners for Elliptic Problems with Discontinuous Coefficients

We develop robust and efficient multilevel solvers for the large scale linear systems arising from finite element discretizations of general second order elliptic equations in heterogeneous materials. In this talk, we will discuss the influence of the discontinuous (and possibly anisotropic) coefficients to the overall performance of the multilevel preconditioners.

Yunrong Zhu
Department of Mathematics
University of California at San Diego
zhuyunr@isu.edu

MS12

Efficient Eigensolver Algorithm on Accelerator-Based Architecture

The enormous gap between the high-performance capabilities of GPUs and the slow interconnect between them has made the development of numerical software that is scalable across multiple GPUs extremely challenging. We describe a successful methodology on how to address the challenges –starting from our algorithm design, kernel optimization and tuning, to our programming model– in the development of a scalable high-performance symmetric eigenvalue and singular value solver.

Azzam Haidar, Piotr Luszczek
Department of Electrical Engineering and Computer Science

University of Tennessee, Knoxville
haidar@utk.edu, luszczek@eecs.utk.edu

Stanimire Tomov
Innovative Computing Laboratory, Computer Science Dept
University of Tennessee, Knoxville
tomov@eecs.utk.edu

MS12

GPGPU Acceleration of the Ams Eigensolver Using Magma

The automatic multilevel substructuring (AMS) technology is a state-of-the-art eigensolver for large-scale eigenvalue problems and has been frequently used to speed up the eigensolution process for finite element models, especially for noise and vibration (N&V) analysis in the automotive industry. To further accelerate the eigensolution process using the AMS technology, it is essential to take advantage of high performance GPGPU accelerators. MAGMA project by University of Tennessee provides various numerical linear algebra subroutines on GPGPU and those subroutines have been facilitating the adoption of GPGPU acceleration to this technology. Recently, MAGMA team developed a new high-performance algorithm for symmetric dense eigenproblems, which has been a performance bottleneck for more than a decade on multicore architectures. By adopting this new high-performance algorithm, the performance of the AMS technology can be significantly accelerated. In this paper/talk, we will introduce the new high-performance MAGMA algorithm and demonstrate the AMS performance benefits using those eigensolution routines.

Mintae Kim
SIMULIA Dassault Systemes
mintae.kim@3ds.com

Luis Crivelli, Michael Wood, Cristian Ianculescu, Vladimir Belsky
SIMULIA-Dassaults systemes
luis.crivelli@3ds.com, michael.wood@3ds.com, cristian.ianculescu@3ds.com, vladimir.belsky@3ds.com

MS12

High-Performance Computation of Pseudospectra

This talk introduces several high-performance variations of Van Loan's triangularization followed by inverse iteration algorithm which involve parallel reduction to (quasi-)triangular or Hessenberg form followed by interleaved Implicitly Restarted Arnoldi iterations driven by multi-shift (quasi-)triangular or Hessenberg solves with many right-hand sides. Since multi-shift (quasi-)triangular solves can achieve a very high percentage of peak performance on both sequential and parallel architectures, such an approach both improves the eigenvaluesciency of sequential pseudospectral computations and provides a high-performance distributed-memory scheme. Results from recent implementations within Elemental (P. et al.) will be presented for a variety of large dense matrices and practical convergence-monitoring schemes will be discussed

Jack L. Poulson
Computational Science and Engineering
Georgia Institute of Technology
jack.poulson@gmail.com

MS12

Towards Materials Design with Extreme-Scale Quantum Simulations

With ever improving accuracy of ab initio electronic structure methods, quantum simulations have now become a predictive tool that can be used to search for new materials with desired properties. For simulations tools to be relevant for such searchers, calculations for individual compounds have to be very reliable and optimized to minimize the machine and energy footprint, as they have to be repeated automatically tens or hundreds thousands of times. Today, this can be accomplished with clusters that have hybrid CPU-GPU nodes. We will discuss the algorithmic developments that were necessary to run modern electronic structure codes on such systems.

Thomas C. Schulthess
Swiss National Supercomputing Center
schulthess@cscs.ch

Anton Kozhevnikov
ETH Zürich
Institut f. Theoretische Physik
antonk@ethz.ch

Solcà Raffaele
ETH
Eidgenössische Technische Hochschule Zürich
rasolca@student.ethz.ch

MS14

Opportunities and Challenges in First Principles Models of Materials

First principles models based on quantum mechanics have provided a significant insight into materials and molecules in recent years, and hold the promise of an ability to design new materials properties one atom at a time. However, these models also raise a number of important challenges. First, many involve rather ad hoc approximations to the Schrodinger's equation. For example, the mathematical basis of widely used exchange and correlation functionals of density functional theory is unclear. Second, many of these models are extremely computationally demanding, thereby limiting their applicability to relatively small computational cells. However, interesting properties of materials require calculations that are orders of magnitude beyond what is currently possible. This talk will give a brief overview of state of the art first principle methods, and provide opinions on opportunities and challenges where mathematics and computational sciences can lead to significant impact.

Kaushik Bhattacharya
Howell N. Tyson Sr. Professor of Mechanics
California Institute of Technology
bhatta@caltech.edu

MS14

New Liquid-Crystal Based Models and Technologies

Research on liquid-crystal based suspensions is rapidly advancing motivated by applications in materials science as well as in biological systems. From proposals for new

display technologies and nanofluidic devices to more fundamental questions about the mechanisms of clustering and de-clustering in systems of particles, new experimental findings call for new modeling and analysis efforts. Efficient engineering of these systems requires advances to current understanding of isotropic fluid colloids, in that the existence of structure in the liquid matrix affords new opportunities for flow control, processing, and suspension stability. One of the newly observed phenomena is the quadratic dependence of the drift velocity of particles moving in an ionic liquid crystal matrix, on the direction transverse to the applied electric field. This property may allow for significant technological applications such as the development of AC- electrophoresis. I will present a survey of the current research status on liquid crystal colloids pointing to challenging mathematical and computational issues, such as in the case of particles of typical size above 50nm. Another line of research involving liquid crystal elastomers is the modeling and development of microdevices, such as pumps and valves. A basic mechanism is the ability of the anisotropic elastomer to control the evolution of liquid crystal defects. Devices made of liquid crystal elastomers, in addition to exhibiting excellent shape memory, hardly require any micromachining processes. I will outline some of the main mathematical issues and their research status.

Carme Calderer
University of Minnesota
calde014@umn.edu

MS14

Opportunities in Computational Science: Genomes, Mesoscale and Closing the Loop

Strategies for advancing computational science will be drawn from three concepts now in their infancy: the Materials Genome(1) mesoscale science(2), and closing the loop among computation, synthesis, characterization and targeted materials outcomes(3). Each of these strategies addresses the fundamental challenge of materials science: capturing and controlling increasingly complex and functional materials behavior. Examples will be drawn from battery science, composite materials and biology. 1. http://www.sciencedirect.com/science/article/pii/S1359028614000060
2. http://science.energy.gov/~/media/bes/pdf/reports/files/OFMS_rpt.pdf
3. http://www.nsf.gov/mps/advisory/mpsac_other_reports/materials_instrumentation-final_from_subcommittee.pdf

George Crabtree
Argonne National Laboratory
crabtree@anl.gov

MS14

Problems in Pattern Formation, Geometry and Design of Materials

We propose to study some mathematical problems, combining geometry and analysis, that arise from practical questions of materials design. Manipulating the microstructure of a material can radically change its mechanical responses. For example, the basic engineering design goal of actuation of nematic glass sheets is to determine imprinted director distributions in the flat sheet, in order to achieve particular actuated shapes on exposure to appropriate stimuli. Another rich source of related analytical questions is provided by the heterogeneous incompatibilities of strains, present in bulk and in thin structures, associated with growth, swelling or shrinkage, plasticity, etc. In this vein, one class of mathematical problems we will discuss is related to wrinkling of thin films and the symmetry breaking in the prestress-to-wrinkle-to-crumple transition, where some insight may be gained by simplifying a complex system in terms of a reduced set of coordinates. The second class of problems concerns the actuation of thin nematic glass sheets. Yet another class of problems relates to the interaction of nonlinear pdes and mechanics of materials in the prediction of static and dynamic microstructure patterns.

Marta Lewicka
Department of Mathematics
University of Pittsburgh
lewicka@pitt.edu

MS15

High Order Semi-Lagrangian Discontinuous Galerkin Schemes for First and Second Order PDEs

High order semi-lagrangian DG schemes are proposed for some linear first and second order multi-dimensional time-dependant PDEs, based on splitting techniques and one-dimensional gaussian quadrature formula. Stability and convergence is proved for the splitting and for variable coefficients. Extention to some non-linear PDEs will be also discussed.

Olivier Bokanowski
Université Paris-Diderot
boka@math.jussieu.fr

MS15

High-order Discontinuous Galerkin Methods for Some Kinetic Models

In this talk, I will present our recent progress in developing high-order discontinuous Galerkin methods for simulating some kinetic models, such as Vlasov-Maxwell equations, some discrete-velocity models in the diffusive limit, and the BGK model. Theoretical results will be discussed in terms of conservation, stability, accuracy, and asymptotic preserving property, together with some numerical examples.

Fengyan Li
Rensselaer Polytechnic Institute
lif@rpi.edu

MS15

Convergence of Semi-Discrete Stationary Wigner Equation with Inflow Boundary Conditions

Making use of the Whittaker-Shannon interpolation formula with shifted sampling points, we propose in this paper a well-posed semi-discretization of the stationary Wigner equation with inflow boundary conditions. The convergence of the solutions of the discrete problem to the continuous problem is then analyzed, providing certain regularity of the solution of the continuous problem.

Tiao Lu
Peking University
tlu@math.pku.edu.cn

Ruo Li
School of Mathematical Science
Peking University
rli@math.pku.edu.cn

Zhangpeng Sun
Peking University
sunzhangpeng@pku.edu.cn

MS15

Self-Organized Hydrodynamics in An Annular Domain: Modal Analysis and Nonlinear Effects

The self-organized hydrodynamics model of collective behavior is studied on an annular domain. A modal analysis of the linearized model around a perfectly polarized steady-state is conducted. It shows that the model has only pure imaginary modes in countable number and is hence stable. Numerical computations of the low-order modes are provided. The fully non-linear model is numerically solved and nonlinear mode-coupling is then analyzed. Finally, the efficiency of the modal decomposition to analyze the complex features of the nonlinear model is demonstrated.

Hui Yu
Université de Toulouse; UPS, INSA, UT1, UTM ;
h.yu@imperial.ac.uk

MS16

The Role of Intraclot Transport in the Dynamics of Platelet Deposition and Coagulation Under Flow

We present a spatial-temporal continuum model of blood clot formation in response to vascular injury that incorporates platelet activation, adhesion, and cohesion, treats coagulation reactions comprehensively, and includes the influence of flow and thrombus growth on one another. We focus on transport of platelets and proteins within the clot itself and how it influences the progression of the coagulation reactions and thus has profound effects on the growth and ultimate structure of the thrombus.

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

Karin Leiderman
Applied Mathematics
UC Merced
kleiderman@ucmerced.edu

MS16

Modeling Cardiac Electro-Fluid-Mechanical Interaction

Integrative models of the heart that account for the coupling between cardiac mecahnics, fluid dynamics, and electrophysiology promise to serve as powerful platforms for understanding cardiac diseases and for treatment planning and optimization. This talk will describe progress towards the development of such electro-fluid-mechanical models of the heart using the framework of the immersed boundary (IB) method, with a focus on new image-based model of the heart and work to develop validated high-resolution models of cardiac fluid dynamics.

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

MS16

An Integrative Model of Lamprey Locomotion Using the Immersed Boundary Method

The lamprey is considered a model organism for neurophysiology and locomotion studies. Here we present a 2D, integrative, multi-scale model of the lamprey's anguilliform (eel-like) swimming driven by neural activation and muscle kinematics coupled to body interactions with fluid surroundings and implemented using the immersed boundary method. Effects on swimming speed and cost (metabolic work) at each scale as well as the role of feedback on performance are presented.

Christina Hamlet
Tulane University
Department of Mathematics
chamlet@tulane.edu

Eric Tytell
Tufts University
Department of Biology
eric-tytell@tufts.edu

Lisa J. Fauci
Tulane University
Department of Mathematics
fauci@tulane.edu

MS16

Modeling Escherichia Coli Chemotaxis in a Fluid

The hydrodynamics of Escherichia coli is modeled by coupling the chemotaxis equations of a simplified phosphorylation cascade with the method of regularized Stokeslets of the fluid motion. For a slow enough diffusion rate of the attractant gradient, simulations have consistently resulted in a biased random walk of the majority of cells towards the highest concentration of attractant, chemotactic behavior. The results demonstrate how the phosphorylation affects the run and tumble mechanism of swimming bacteria.

Hoa Nguyen
Trinity University
hnguyen5@trinity.edu

MS17

On First Experiments for Nuclear Engineering Applications on Intel Xeon Phi

Next generation of supercomputers integrate many core computational units and will involve multilevel of parallelism such as Intel Xeon Phi. For nuclear engineering, numerical simulation is based on verified and validated large simulation codes. Moving to these new computing architectures, will imply deep modifications and all existing codes will not be adapted easily. The subject of this talk is to present first experiments on porting existing CEA applications for nuclear engineering on Xeon Phi architecture.

Christophe Calvin
CEA-Saclay/DEN/DANS/DM2S
christophe.calvin@cea.fr

MS17

Opportunities and Challenges in Developing and Using Scientific Libraries on Emerging Architec-

tures

Scalable manycore and accelerator based systems are becoming the norm, and will soon be the dominant platforms for very high end systems. In the development of new algorithms and libraries for these systems, common themes have emerged that, if appreciated and adopted by application and library developers, will improve portability, performance and resilience in the future. In this presentation, we give an overview of some of these themes, and discuss strategies for application development efforts today that will have lasting value in the future.

Michael Heroux
Sandia National Laboratories
maherou@sandia.gov

MS17
Managing Portability for ASC Applications

Large production-quality multi-physics software contains tens of thousands of inner-loop kernels that must run effectively on multiple processor architectures and a variety of memory subsystems. We have identified a set of four key encapsulation idioms that allow us to manage portability in legacy applications while imposing minimal restrictions on the way developers write software. We will discuss our experience integrating these idioms into large software projects, and present performance behavior for select applications.

Jeff Keasler
Lawrence Livermore National Laboratory
keasler1@llnl.gov

Richard Hornung
Lawrence Livermore National Lab
hornung@llnl.gov

MS17
ppOpen-APPL/HEXA: A Framework for Development of Parallel FEM/FVM Applications on Intel Xeon Phi

ppOpen-HPC is an open source infrastructure for development and execution of large-scale scientific applications on post-peta-scale supercomputers with automatic tuning (AT). ppOpen-APPL/HEXA is a part of ppOpen-HPC and a special framework for development of parallel FEM/FVM codes with voxel-type hexahedral meshes. In this talk, outline of development, optimization and utilization of ppOpen-APPL/HEXA is described. Moreover, details of optimization of preconditioned iterative solvers, and matrix assembling procedures for Intel Xeon Phi processors are presented.

Kengo Nakajima
The University of Tokyo
Information Technology Center
nakajima@cc.u-tokyo.ac.jp

MS18
Fast Solvers for Wave Propagation and Scattering by General Structures

We present fast spectral solvers for Partial Differential Equations that address some of the main difficulties associated with simulation of realistic problems of propagation and scattering in the frequency domain. Based on fast high-order methods for evaluation of integral operators these algorithms can solve, with high-order accuracy, problems of electromagnetic and acoustic scattering for large and complex three-dimensional geometries. A variety of applications will be presented which demonstrate the significant improvements the new algorithms can provide over the accuracy and speed resulting from other approaches.

Oscar P. Bruno
California Institute of Technology
obruno@caltech.edu

MS18
Scalable Algorithms for Density Matrix Calculations of Cavity Quantum Electrodynamic Systems

If a coupled quantum systems shows some level of entanglement, that system can not simply be described by a single state. They exist in mixed states. Furthermore, real quantum systems exhibit dephasing and decoherence, which requires a statistical description of the system. This is done through the density matrix formulation. Density matrices describe the system as a statistical ensemble of several pure quantum states.The time dynamics of the density matrix are governed by the Lindblad master equation, which has involves operator multiplication from both the right and left of the density matrix. We study the time dynamics of a system consisting of quantum dots coupled to a single plasmon mode. The size of the density matrix grows quickly; a physically reasonable system of 16 quantum dots requires a matrix dimension of $50 * 2^{16}$. At this size, extreme scale computing must be utilized. Aside from the novel physics applications, we are currently studying how best to treat this system, through different time stepping schemes (RK and exponential time integrator) and different parallelization algorithms.

Matthew Otten
Department of Physics
Cornell University
mjo98@cornell.edu

MS18
Electromagnetic Power Absorption and Plasmon Resonances on Rough Conducting Surfaces

In this talk we present high-order integral equation methods for the evaluation of electromagnetic wave scattering and absorption by dielectric/conducting bumps and cavities on penetrable half-planes. The numerical far- and near-fields exhibit excellent convergence as discretizations are refined –even at and around points where singular fields and infinite currents exist. The methods presented in this talk are applied to study the absorption of electromagnetic power that results from the presence of defects on flat conducting surfaces. The performance of the integral equation solvers herein discussed allows for the accurate evaluation of electromagnetic fields at and around the surface of the conducting material where relevant physical phenomena, such as skin-depth effects, occur. Finally we discuss the application of the high-order integral equation methods to the evaluation electromagnetic fields resulting from excitation of plamon resonances on rough conducting surfaces.

Carlos A. Perez-Arancibia
Computing and Mathematical Sciences
California Institute of Technology
cperezar@caltech.edu

Oscar P. Bruno
California Institute of Technology
obruno@caltech.edu

MS18

Generalized Combined Sources Integral Equations for Helmholtz Transmission Problems

We present a new class of well conditioned integral equations for the solution of two and three dimensional scattering problems by homogeneous penetrable scatterers. Our novel boundary integral equations result from representations of the fields inside and outside the scatterer as combinations of single and double layer potentials acting on suitably defined regularizing operators. The regularizing operators are constructed to be approximations of the admittance operators that map the transmission boundary conditions to the exterior and respectively interior Cauchy data on the interface between the media. The latter operators can be expressed in terms of Dirichlet-to-Neumann operators. We refer to these regularized boundary integral equations as Generalized Combined Source Integral Equations (GCSIE). The ensuing GCSIE are shown to be integral equations of the second kind in the case when the interface of material discontinuity is a smooth curve in two dimensions and a smooth surface in three dimensions.

Catalin Turc
Department of Mathematical Sciences
New Jersey Institute of Technology
catalin.c.turc@njit.edu

MS19

Generalized Radiative Transfer: Accounting Accurately for Unresolved Variabilities at No Computational Cost, Yet Without Homogenization

Generalized Radiative Transfer (GRT) is formalized as the integral 1D radiative transfer equation in a uniform medium, but where the exponential functions in the propagation kernel are replaced with two closely related power laws. An adapted Markov chain numerical solution was developed. GRT theory applies to challenging transport problems with highly variable optical properties over a broad range of scales. Another application is to the efficient and yet accurate computation of broadband spectral responses.

Anthony B. Davis, Feng Xu
Jet Propulsion Laboratory
California Institute of Technology
anthony.b.davis@jpl.nasa.gov, feng.zu@jpl.nasa.gov

MS19

Quadrature-Based Moment Methods for Radiation Transport

Radiation transport can be described by a kinetic equation with free transport and scattering operators. The high dimensionality of the density function makes the numerical solution very challenging. Available strategies range from Monte-Carlo approaches, to collocation such as discrete ordinate methods, to moment methods. An important aspect of radiation solvers is related to how they represent the phase-space dependence of the density function, which can lead to so-called ray effects where radiation travels along fixed directions in physical space. To overcome these effects, moment methods can use a delta-function representation with adaptive quadrature nodes, or a continuous representation. In general, the positivity of the reconstructed density function is not guaranteed with traditional moment methods, and positivity constraints or filtering are required. Here we explore quadrature-based moment methods that reconstruct the density function on the unit sphere. The transported moment set is composed of the trigonometric moments of the spherical angles. Compared to quadrature-based reconstruction on the unit square, the additional constraint of periodicity on the unit sphere requires special considerations. Here we employ Szegö quadrature and an extension of CQMOM on the unit circle. The density function is reconstructed using the extended quadrature method of moments (EQMOM) with a von Mises kernel density function. Special attention is paid to the hyperbolicity of the moment transport equations, and to the numerical algorithms to handle realizable spatial fluxes, scattering and the diffusion limit.

Rodney O. Fox
Iowa State University
Department of Chemical and Biological Engr.
rofox@iastate.edu

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

Ming Tse P. Laiu
Department of Electrical and Computer Engineering
University of Maryland - College Park
mtlaiu@umd.edu

Frederique Laurent
Laboratoire EM2C - CNRS UPR 288
Ecole Centrale Paris
frederique.laurent@em2c.ecp.fr

Marc Massot
Laboratoire EM2C - UPR CNRS 288
Ecole Centrale Paris
marc.massot@em2c.ecp.fr

MS19

Stability of P_N Approximations for the Radiative Transfer Equation in the Free Streaming Limit

Based on a mixed variational framework we will investigate solvability of the radiative transfer equation in the free streaming limit. Using this framework we will discuss the classical P_N approximations and related stability issues. As an alternative we present an equivalent stabilized variational method.

Matthias Schlottbom
Fachbereich Mathematik
AG Numerik und Wissenschaftliches Rechnen
schlottbom@uni-muenster.de

Herbert Egger
Numerical Analysis and Scientific Computing
Darmstadt University of Technology
egger@mathematik.tu-darmstadt.de

MS19

On Combining Moment Methods and Discrete-Velocity-Schemes for Solving the Boltzmann Equa-

tion

Varius methods exit to solve the Boltzmann Equation deterministically. Moment methods and discrete-velocity-schemes are two possibilities both using very different approaches and serving very different purposes. Moment methods focus on physical variables and come with build-in conservation properties, Galilean invariance, easy coupling to CFD, etc. Discrete velocity schemes can approximate very complex distribution functions and are relatively easy to implement. This talk will discuss the similarities of both approaches and possible combinations, like the use of physically-adaptive velocity grids.

Manuel Torrilhon
Schinkelstrasse 2
RWTH Aachen University
torrilhon@rwth-aachen.de

MS20

ODTLES: A Multiscale Approach for Highly Turbulent Flows

ODTLES combines the ability of ODT to resolve molecular effects with the ability of Large Eddy Simulations (LES) to describe 3D large scale flows. ODTLES is based on an extended LES scale separation ansatz wherein molecular effects and the turbulent cascade are simulated by ODT while a standard 3D numerical approach time advances the 3D large scale flow. Simulation results for highly turbulent flows reveal the model and numerical properties of ODTLES.

Christoph Glawe
BTU Cottbus
glawechr@tu-cottbus.de

Heiko Schmidt
BTU Cottbus
Mechanical Engineering
schmidth@tu-cottbus.de

Alan Kerstein
Consultant
alan.kerstein@gmail.com

MS20

Particle-Scalar Field Interactions in One-Dimensional Turbulence

The development of Lagrangian particle models in the context of one-dimensional turbulence (ODT) allows the study of the interactions of particles with various scalar fields in a novel sense. We describe finite Stokes-number effects of the relative dispersion of scalars and particles. We also show the wide range of interaction time scales associated with particle slip across scalar gradients and with scalar-field diffusion possible with full ODT resolution of scalar fields.

John C. Hewson
Sandia National Laboratories
jchewso@sandia.gov

Guangyuan Sun, David O. Lignell
Brigham Young University
gysungrad@gmail.com, davidlignell@byu.edu

MS20

One-dimensional Turbulence Simulation: Overview and Application to Soot Formation in Non-premixed Flames

An overview of one-dimensional turbulence (ODT) simulation of turbulent flows is presented along with recent advances and application to soot formation in nonpremixed flames. ODT is a computationally affordable stochastic model that resolves a full range of scales and can be applied to flow regimes not available to direct numerical simulation. Soot formation, radiation, and flame emission is an important and challenging multiphysics problem to which we apply ODT for fundamental insight and model validation.

David O. Lignell, Victoria Lansinger
Brigham Young University
davidlignell@byu.edu, vlansinger@gmail.com

John C. Hewson
Sandia National Laboratories
jchewso@sandia.gov

MS20

Multiphase Turbulent Reacting Flow Simulations Using ODT

Abstract Results for application of the One-Dimensional Turbulence model to multiphase turbulent combustion of coal is considered. Results are compared to a pilot-scale reactor and the ability of the model to capture ignition delay in this situation is evaluated. Because of its low cost relative to DNS and LES, ODT can be used to explore high-fidelity thermochemistry models. We explore several models for heterogeneous and homogeneous reactions and evaluate their ability to reproduce the observed experimental data.

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

Babak Goshayeshi
University of Utah
b.goshayeshi@gmail.com

MS21

The most current list of participating companies is available at www.siam.org/meetings/cse15/career.php.

For the most recent list of participating companies visit http://www.siam.org/meetings/cse15/career.php

Bill Kolata
SIAM
kolata@siam.org

MS22

Cilia Beating Patterns are not Hydrodynamically Optimal

We examine the hydrodynamic performance of two cilia

beating patterns reconstructed from experimental data. In their respective natural systems, the two beating patterns correspond to: (A) pumping-specialized cilia, and (B) swimming-specialized cilia. We compare the performance of these two cilia beating patterns as a function of the metachronal coordination in the context of two model systems: the swimming of a ciliated cylinder and the fluid pumping by a ciliated carpet. Three performance measures are used for this comparison: (i) average swimming speed/pumping flow rate; (ii) maximum internal moments generated by the cilia; and (iii) swimming/pumping efficiencies. We found that, in both models, pattern (B) outperforms pattern (A) in almost all three measures, including hydrodynamic efficiency. These results challenge the notion that hydrodynamic efficiency dictates the cilia beating kinematics, and suggest that other biological functions and constraints play a role in explaining the wide variety of cilia beating patterns observed in biological systems.

Hanliang Guo
Aerospace and Mechanical Engineering
University of Southern California
guohanliang@gmail.com

Janna C. Nawroth
Harvard School of Engineering
jnawroth@gmail.com

Yang Ding
Beijing Computational Science Research Center
dingyang@csrc.ac.cn

Eva Kanso
University of Southern California
kanso@usc.edu

MS22

A Numerical Method for Doubly-Periodic Stokes Flow Near a Wall

We present a numerical method for computing doubly-periodic Stokes flow in the presence of a wall. By finding fundamental solutions to Stokes equation in Fourier space exactly, in practice only an inverse FFT has to be computed. Our algorithm can be used to effectively model arrays of pulmonary or nodal cilia. We match theoretical, numerical and experimental results.

Franz M. Hoffmann
Mathematics Department
Tulane University
fhoffma@tulane.edu

MS22

Computation of the Regularized Image Systems for Doubly-Periodic Brinkman Flow in the Presence of a Wall

A fast summation method of Ewald type is developed for Brinkman flow due to doubly-periodic arrays of regularized forces near a plane wall. Method of images is applied to find the Brinkman flow due to one regularized force; then Poisson summation formula is applied to arrive at the final formula for the doubly-periodic flow. Initial simulation results for fluid-cilia interaction problems are also presented.

Hoang-Ngan Nguyen
University of California, Merced
zhoangngan@gmail.com

Karin Leiderman
Applied Mathematics
UC Merced
kleiderman@ucmerced.edu

MS22

Accelerated Boundary Integral Simulations for Fluid-Structure Interactions in Periodic Stokes Flow

Boundary integral formulations for Stokes flow involves the periodic summation of Green's functions. The Ewald summation formula for the Stokeslet under triply period boundary conditions is well known (Hasimoto, 1959). Based on this formula, we have developed a spectrally accurate FFT based method to speed up the computations. Such methods are also needed for the Stresslet, and for systems that are only quasi-periodic. Ewald summation formulas and corresponding FFT based Spectral Ewald methods for such cases will be discussed. Boundary integral formulations and simulation results for fluid-structure problems employing Spectral Ewald methods will also be presented.

Anna-Karin Tornberg
KTH
akto@kth.se

MS23

Symmetry-Preserving Conservative Lagrangian Scheme for Compressible Euler Equations in Two-Dimensional Cylindrical Coordinates

In applications such as astrophysics and inertial confinement fusion, there are many three-dimensional cylindrical-symmetric multi-material problems which are usually simulated by Lagrangian schemes in the two-dimensional cylindrical coordinates. For this type of simulation, a critical issue for the schemes is to keep spherical symmetry in the cylindrical coordinate system if the original physical problem has this symmetry. In the past decades, several Lagrangian schemes with such symmetry property have been developed, but all of them are only first order accurate. In this talk, we develop a second order cell-centered Lagrangian scheme for solving compressible Euler equations in cylindrical coordinates, based on the control volume discretizations, which is designed to have uniformly second order accuracy and capability to preserve one-dimensional spherical symmetry in a two-dimensional cylindrical geometry when computed on an equal-angle-zoned initial grid. The scheme maintains several good properties such as conservation for mass, momentum and total energy, and the geometric conservation law. Several two-dimensional numerical examples in cylindrical coordinates are presented to demonstrate the good performance of the scheme in terms of accuracy, symmetry, non-oscillation and robustness. This is a joint work with Chi-Wang Shu.

Juan Cheng
Institute of Applied Physics and Computational Mathematics
Beijing, China
cheng_juan@iapcm.ac.cn

MS23

Superconvergent HDG Methods for Third-Order

Equations in One-Space Dimension

We design and analyze *the first* hybridizable discontinuous Galerkin methods for third-order linear equations in one-space dimension. The methods are defined as discrete versions of characterizations of the exact solution in terms of local problems and transmission conditions. They provide approximations to the exact solution u and its derivatives $q := u'$ and $p := u'$ which are piecewise-polynomials of degree k_u, k_q and k_p, respectively. We consider the methods for which the difference between these polynomial degrees is at most two. We prove that all these methods have superconvergence properties which allows us to prove that their numerical traces converge at the nodes of the partition with order at least $2\,k+1$, where k is the minimum of k_u, k_q, k_p. This allows us to use an element-by-element post-processing to obtain new approximations for u, q and p converging with order at least $2k+1$ uniformly.

Bo Dong
University of Massachusetts Dartmouth
bdong@umassd.edu

Yanlai Chen
Department of Mathematics
University of Massachusetts Dartmouth
yanlai.chen@umassd.edu

Bernardo Cockburn
School of Mathematics
University of Minnesota
cockburn@math.umn.edu

MS23

A Local Discontinuous Galerkin Scheme for the Patlak-Keller-Segel Chemotaxis Model

In this talk, a new local discontinuous Galerkin method is designed for solving the Patlak-Keller-Segel (PKS) equation. We give the error analysis and also prove the positive-preserving property for the scheme on 1D and 2D structured meshes. Numerical simulations will be presented which confirm the predicted convergence rate.

Yang Yang
Michigan Technological University
yyang7@mtu.edu

Xingjie Li
Brown University
xingjie_li@brown.edu

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

MS23

Optimal Error Estimates for Discontinuous Galerkin Methods Based on Upwind Biased Fluxes for Linear Hyperbolic Equations

We analyze discontinuous Galerkin methods using upwind-biased numerical fluxes for time-dependent linear conservation laws. In one dimension, optimal a priori error estimates of order $k+1$ are obtained when piecewise polynomials of degree at most k ($k \geq 0$) are used. We extend the analysis to the multidimensional case on Cartesian meshes when piecewise tensor product polynomials are used. Numerical experiments are shown to demonstrate the theoretical results.

Xiong Meng
harbin institute of technology
University of East Anglia
xiongmeng@hit.edu.cn

MS24

A Novel Elliptic Solver Based on RBF-Finite Differences for Understanding the Earth's Electric System

The Global Electric Circuit (GEC) represents the Earth's electric link between solar, upper and lower atmosphere processes, and cloud system dynamics. It is modeled by an elliptic PDE in 3D spherical-like geometry with highly variable coefficients. Due to the several computational challenges that it presents, a novel radial basis function-generated finite differences (RBF-FD) solver has been developed. Several features of RBF-FD have proven to be very beneficial. Among all, the ability to easily handle irregular boundaries such as the Earths topography stands out. In this talk, we present this novel RBF-FD elliptic solver and the new techniques for handling irregular geometries that was motivated by it.

Victor Bayona
National Center for Atmospheric Research (NCAR)
vbayona@ucar.edu

Natasha Flyer
National Center for Atmospheric Research
Institute for Mathematics Applied to Geosciences
flyer@ucar.edu

MS24

Guidelines to Modeling the Navier-Stokes and Euler Equations with RBF-FD

In applications of radial basis functions (RBFs) for fluid modeling, infinitely smooth RBFs have traditionally been used due to the spectral convergence properties. However, fluid flows in nature can exhibit complex features such that spectral accuracy cannot be realized on resolutions that are observable or practical. A novel approach for modeling with RBF-generated finite differences (RBF-FD) is presented by using polyharmonic spline RBFs (PHS RBF) together with higher-order polynomials. The approach is tested on nonhydrostatic compressible atmospheric flows in limited area domains. Test cases include flows exhibiting Kelvin-Helmholtz instabilities, turbulence, and bubble updrafts, simulating cloud entrainment. General guidelines are given as to how to choose the parameters involved, such as the degree of the PHS RBF, the order of the polynomials, stencil size, node layout, etc.

Natasha Flyer
National Center for Atmospheric Research
Institute for Mathematics Applied to Geosciences
flyer@ucar.edu

Louis Wicker
NOAA, NSSL
louis.wicker@noaa.gov

Gregory Barnett
Department of Applied Mathematics
University of Colorado, Boulder

gregory.barnett@colorado.edu

MS24

A High-Order RBF-Based Leray Projection Method for the Incompressible Stokes and Navier-Stokes Equations

We present a novel pressure-free Leray projection method for the solution of the incompressible Stokes and Navier-Stokes equations in two dimensions. The discrete Leray projector is constructed via generalized interpolation with divergence-free Radial Basis Functions (RBFs). This RBF-based Leray projection method does not require that one specify boundary conditions for the pressure. Results demonstrate that the RBF-based Leray projection method demonstrates high orders of convergence in both time and space.

Edward Fuselier
High Point University
Department of Mathematics and Computer Science
efuselie@highpoint.edu

Varun Shankar
School of Computing
University of Utah
vshankar@math.utah.edu

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

MS25

Distributed Contraction of Tensors

Unlike distributed matrix multiplication, which has been extensively studied, limited work has been done in characterizing distributed tensor contractions. In this talk, a characterization is presented for a family of communication-optimal distributed tensor contraction algorithms on torus networks. The framework uses three fundamental communication operators to generate communication-efficient algorithms for arbitrary tensor contractions. Trade-off involving use of additional memory for reduced inter-processor communication is also addressed. Experimental results on a BlueGene/Q system demonstrate good scalability on 256K cores.

P Sadayappan, Samyam Rajbhandari
Ohio State University
psaday@gmail.com, samyamrb@gmail.com

Sriram Krishnamoorthy
Pacific Northwest National Laboratory
sriram@pnnl.gov

Pai-Wei Lai
Ohio State University
laip@cse.ohio-state.edu

MS25

A Framework for Distributed Tensor Computations

Branches of scientific computing, express data as a tensor which can be viewed as a multi-dimensional analog of a matrix. For instance, chemical methods heavily rely on the tensor contraction operation, which can be viewed as a generalization of matrix-matrix multiplication. The algebra associated with tensors in these applications is referred to as multi-linear algebra, the multi-dimensional analog of linear algebra. Problems in this area of research frequently require the use of distributed-memory computing architectures to compute the desired method or operation. One common approach to computing multi-linear operations acting on tensors on such architectures is to cast the operation in terms of linear operations acting on matrices and utilize high-performance linear algebra libraries to compute the result. Unfortunately, the underlying linear algebra library may incur higher network communication costs than is necessary as some structure of the multi-linear objects is lost when cast to linear objects. In this work, we introduce a notation for describing distributions of tensor data on processing grids, relate redistributions of data to well-studied collective communications, and describe a method to systematically derive and analyze algorithms for the tensor contraction operation.

Martin D. Schatz
Department of Computer Sciences
The University of Texas at Austin
martin.schatz@utexas.edu

Tamara G. Kolda
Sandia National Laboratories
tgkolda@sandia.gov

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

MS25

Tensor Computation for Chemistry and Material Science

Tensors are ubiquitous in many-body quantum physics that supports chemistry and materials science. We will discuss the types of tensor structures that arise in numerical simulation of many-body physics and the resulting computational challenges. We will also discuss TiledArray, a modern C++ library for block-sparse tensor computation and how it addresses some of these challenges. (see github.com/ValeevGroup/tiledarray for more information).

Justus Calvin
Virginia Tech
justusc@vt.edu

Edward F. Valeev
Virginia Polytechnic Institute and State University
efv@vt.edu

MS25

Exploiting Multiple Tensor Symmetries though Block Diagonalization

Suppose an order-6 tensor $\mathcal{A}$ has the property that the value in entry $\mathcal{A}(i_1, i_2, i_3, j_1, j_2, j_3)$ does not change if the i-indices are permuted, or if the j-indices are permuted, or if the i-indices as a group are swapped with the j-indices as a group. Such a tensor has multiple symmetries and if it is smartly unfolded into a matrix A, then A itself has interesting structure above and beyond ordinary symmetry. In the case of the given example, there are permutation

matrices Γ_1 and Γ_2 (both involving Kronecker products and perfect shuffles) such that both $\Gamma_1 A \Gamma_1^T$ and $\Gamma_2 A \Gamma_2^T$ equal A. We show how to compute a structure-preserving, low-rank approximation to A using LDL^T with diagonal pivoting together with a very cheap block diagonalization that is performed at the start. The full exploitation of structure has ramifications for efficiency and applications.

Charles Van Loan
Cornell University
Department of Computer Science
cv@cs.cornell.edu

Joseph Vokt
Cornell University
joseph.vokt@gatech.edu

MS26

The Role of Microdomains and Ephaptic Coupling in Cardiac Action Potential Propagation

Much of our theoretical understanding of cardiac propagation is built on the premise that gap junctional coupling between cells is the primary means of electrical coupling. In this talk, I will describe some of the newly discovered features of propagation that result from consideration of microdomains and spatially inhomogeneous extracellular potential and in doing so, provide new, indirect evidence for the importance of ephaptic, or field effect, coupling.

James P. Keener
University of Utah
keener@math.utah.edu

MS26

Strongly Scalable Numerical Approaches for Modeling Cardiac Electromechanics at High Spatiotemporal Resolution

This study explores the feasibility of high resolution models of bidirectionally coupled cardiac electro-mechanics which resolve cardiac anatomy at a para-cellular resolution. A novel algebraic multigrid method is presented, custom-tailored for non-linear mechanics, which is shown to be strongly scalable up to 4k cores when using a human whole heart model. Benchmark results demonstrate that a single heart beat can be simulated in about 15 minutes at full anatomical and biophysical detail.

Christoph Augustin
Medical University of Graz
Institute of Biophysics
christoph.augustin@medunigraz.at

Aurel Neic
University of Graz, Institute for Mathematics
Graz, Austria
aurel.neic@uni-graz.at

Manfred Liebmann, Gundolf Haase
University of Graz
manfred.liebmann@uni-graz.at, gundolf.haase@uni-graz.at

Gernot Plank
Institute of Biophysics
Medical University of Graz, Austria
gernot.plank@medunigraz.at

MS26

Towards Multidomain Modeling of Cardiac Electrophysiology

Mathematical models are established tools for studies of cardiac tissue physiology and disease. Previously, we introduced an approach for multidomain mathematical modeling of tissue electrophysiology (Sachse et al, Ann Biomed Eng. 2009;37(5):874-89), which allows us to describe tissues as a mixture of cells with variable electrophysiological properties and intercellular electrical coupling. Here, we provide insights into our methodology to establish a microstructural basis for this and other types of models. The approach is based on high-resolution three-dimensional confocal microscopy and image quantification.

Frank B. Sachse
Cardiovascular Research and Training
Institute and Bioengineering Department
frank.sachse@utah.edu

Thomas Seidel
Cardiovascular Research and Training Institute
University of Utah
thomas.seidel@utah.edu

Jan Christoph Edelmann
Cardiovascular Research and Training Institute
University of
University of Utah
edelmann@cvrti.utah.edu

Gunnar Seemann
Karlsruhe Institute of Technology
Germany
gunnar.seemann@kit.edu

MS26

Multi-Scale Modeling of the Failing Heart: from Molecule to Patient

Heart failure results in remodeling of the heart at the molecular, cellular, tissue, organ and organ system scales. Here we describe a new cellular model of the cardiac myocyte and the changes that occur in genetically engineered mouse models of heart failure. We then extend the analysis to multi-scale models of arrhythmia and finally apply these approaches to patient-specific models of atrial fibrillation and heart failure.

Andrew D. McCulloch
Department of Bioengineering
University of California San Diego
amcculloch@ucsd.edu

Christopher Villongco, Adarsh Krishnamurthy
UC San Diego
ctvillongco@gmail.com, adarsh@ucsd.edu

David Krummen
UC San Diego
VA San Diego Medical Center
dkrummen@ucsd.edu

Kevin Vincent
UCSD

vincent@ucsd.edu

MS27

Overview of the Field and the Community of Fast Multipole Methods

At the last SIAM CSE conference, at least four mini symposia were devoted to fast algorithms for integral equations, and I am sure there will be as much interest in the topic at this conference, or more. It's been almost 30 years since the fast multipole method (FMM) was invented, but the activity in the field is getting hot. The landscape of computer hardware has contributed to this, with multi-core and then many-core architectures placing severe constraints on communication bandwidth. FMM has a uniquely favorable communications pattern, making it more attractive as an alternative for various applications. This talk will overview the state of the art and of the community, as the algorithm matures.

Lorena A. Barba
Department of Mechanical and Aerospace Engineering
George Washington University
labarba@gwu.edu

MS27

N-body Methods in Computational Science and Engineering

N-body methods find numerous applications in science and engineering. In this talk, I will outline the evolution of the methods since the early methods and the adoption of their use in fields ranging from gravitational simulations to machine learning.

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS27

Computer Science Aspects of Fast Multipole Methods

Abstract not available at time of publication.

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

MS27

The Geometry of the Fast Multipole Methods

An overview of the geometric aspect of the different kinds of fast multipole methods.

Lexing Ying
Stanford
lexing@math.stanford.edu

MS28

Low-Complexity Stochastic Modeling of Turbulent Flows

Abstract not available at time of publication.

Mihailo R. Jovanovic
Electrical and Computer Engineering
University of Minnesota
mihailo@umn.edu

MS28

Data Mining and Coarse Graining for Network Evolution Problems

In dynamic problems where what evolves is the structure/connectivity of a complex network, coarse graining requires the identification of a small set of macroscopic observables (variables) that hopefully suffice to characterize the evolutionary dynamics. Data mining tools (and, in particular, diffusion maps) can be used to provide such a parametrization of the dynamics, using graph metrics to provide the pairwise similarities between "nearby" graphs. We illustrate this approach in data sets from different (static as well as dynamically evolving) families of graphs and use the resulting variables to enhance dynamic modeling.

Yannis Kevrikidis
Princeton University
yannis@princeton.edu

MS28

Data-Driven Modeling of Complex Systems with Control

Analysis of large-scale datasets generated from complex systems with external control is becoming increasingly important for the engineering, applied, and biological sciences. Here, the development of a new method, which extends Dynamic Mode Decomposition (DMD), incorporates the effect of control to extract low-order models from high-dimensional, complex systems. The method, called Dynamic Mode Decomposition with control (DMDc), was developed for the purpose of studying infectious disease spread and designing intervention campaigns.

Joshua L. Proctor
Intellectual Ventures
JoshLProctor@gmail.com

MS28

A Deim Induced Cur Factorization

Abstract not available at time of publication.

Danny C. Sorensen
Rice University
sorensen@rice.edu

MS29

Geometric Methods in Image Processing, Networks, and Machine Learning

We present new methods for segmentation of large datasets with graph based structure. The method combines ideas from classical nonlinear PDE-based image segmentation with fast and accessible linear algebra methods for computing information about the spectrum of the graph Laplacian. The goal of the algorithms is to solve semi-supervised and unsupervised graph cut optimization problems. I will present results for image processing applications such as image labeling and hyperspectral video segmentation, and results from machine learning and community detection in social networks, including modularity optimization posed

as a graph total variation minimization problem.

Andrea L. Bertozzi
UCLA Department of Mathematics
bertozzi@math.ucla.edu

MS29

Sampling of Dynamic Graphs and Recovery of the Spectral Properties

Massive networks have become ubiquitous today. To monitor such very large networks, one must resort to sampling these networks: a succinct subset of nodes and associated edges are extracted from the original network. In this paper, we present an algorithm to recover the dominant eigenvectors of a very large graph using different realistic sampling strategies. We evaluate the performance of the algorithm on realistic synthetic datasets and real networks.

Nathan D. Monnig
Department Applied Mathematics
University of Colorado Boulder
nathan.monnig@colorado.edu

Conrad Hougen
University of Colorado at Boulder
conrad.hougen@gmail.com

Francois G. Meyer
University of Colorado at Boulder
USA
fmeyer@colorado.edu

MS29

Consistency of Variational Partitioning of Point Clouds

I will discuss variational problems arising in machine learning and their consistency as the number of data points goes to infinity. Consider point clouds obtained as random samples of a measure on a Euclidean domain. Graph representing the point cloud is obtained by assigning weights to edges based on the distance between the points. Many machine learning algorithms are based on minimizing a functional on this graph. Among them are balanced cuts and spectral methods for clustering. We will discuss under what conditions do the minimizers of graph based functionals converge to a well defined limit.

Nicolas Garcia Trillos, Dejan Slepcev
Carnegie Mellon University
ngarciat@andrew.cmu.edu, slepcev@math.cmu.edu

James von Brecht
University of California, Los Angeles
jub@math.ucla.edu

Thomas Laurent
Loyola Marymount University
thomas.laurent@lmu.edu

Xavier Bresson
Swiss Federal Institute of Technology (EPFL)
Institute of Electrical Engineering
xavier.bresson@epfl.ch

MS29

A Panoply of Graph-ported PDEs and Processes

In fields such as image processing, data analysis and community detection, the data sets are often modeled as a graph in which the nodes represent the data points and the edges encode some relationship between the nodes, relevant to the task at hand. In recent years, people have studied classical continuum partial differential equations PDE models from image analysis, but formulated on graphs to be applicable to data analysis problems. These studies show interesting connections between continuum results and the analogous problems on graphs. In this talk we will explore some of these PDE type problems formulated on graphs and their connections with both continuum results and notions from graph theory. Examples include threshold dynamics, modularity optimization, Ohta-Kawasaki minimization, and their relation to graph cuts, mean curvature flow, and bootstrap percolation.

Yves van Gennip
Nottingham University
y.vangennip@nottingham.ac.uk

MS30

Symplectic Model Reduction for Hamiltonian Systems

Proper symplectic decomposition (PSD) is introduced as a symplectic model reduction technique using the symplectic Galerkin projection. Our aim is two-folded. First, to achieve computational savings for large-scale Hamiltonian systems. Second, to preserve the symplectic structure of the original system. As an analogy to the classical POD-Galerkin approach, the PSD is designed to build a symplectic subspace to fit empirical data, while the symplectic Galerkin projection constructs a low-order Hamiltonian system on the symplectic subspace. The proposed technique can preserve system energy, volume of flow, and stability. Therefore, it can be better suited for model reduction of hyperbolic PDEs as compared to the classical POD-Galerkin approach. The stability, accuracy, and efficiency of the proposed technique are illustrated through numerical simulations of several linear and nonlinear wave equations.

Kamran Mohseni
University of Florida, Gainesville, FL
mohseni@ufl.edu

Liqian Peng
University of Florida
liqianpeng@ufl.edu

MS30

An Adaptive and Efficient Greedy Procedure for the Optimal Training of Parametric Reduced-Order Models

An adaptive and efficient approach for constructing parametrically robust reduced-order models is developed. The approach is based on a greedy sampling of the underlying high-dimensional model along with a fast and efficient surrogate based procedure for exploring the parametric space to identify parameters for which the error is likely to be high. The procedure is illustrated on several cases, in-

cluding the realistic prediction of the response of a V-hull vehicle to underbody blasts.

Arthur Paul-Dubois-Taine
Stanford
artpdt@stanford.edu

David Amsallem
Stanford University
amsallem@stanford.edu

MS30

Error Estimation for Hyper-Reduced Elastoviscoplastic Models

We propose an a posteriori error indicator related to hyper-reduced predictions of simulation outputs. We restrict our attention to elastoviscoplastic materials having an incremental variational formulation of the constitutive equations, when the outputs are extracted by a continuous function of the displacements. We obtain an upper bound of the approximation error due to the hyper-reduced formulation. We show numerical results on oligocyclic fatigue. Computational speed-ups are preserved while errors are estimated.

David Ryckelynck
MINES Paristech
David.Ryckelynck@mines-paristech.fr

MS30

Geometric Methods in Adaptive Model Order Reduction

The first part of this presentation addresses an original *manifold learning* approach to nonlinear flow problems, where geometric information is used to choose adaptively local prediction neighborhoods for each ROM approximation. The method provides a sensible tool for detecting gaps in the design of experiment and thus for adaptive refinement. The second part deals with adaptive projection-based ROMs. We propose a method for adjusting a projection subspace to a given parameter condition by optimizing a suitable goal function on the *Grassmann manifold*.

Ralf Zimmermann
DLR Braunschweig, Germany
ralf.zimmermann@tu-bs.de

Thomas Franz
Institute of Aerodynamics and Flow Technology
DLR, German Aerospace Center
thomas.franz@dlr.de

MS31

Flexible Krylov Subspace Methods for Shifted Systems with Multiple Right Hand Sides

Several PDE-based inverse problems require repeated solution to large-scale shifted systems of the form $(K + z_j M)x_{j,k} = b_k$ for $j = \ldots, n_z$ and $k = 1, \ldots, n_s$. We propose a flexible Krylov subspace method that uses multiple preconditioners of the form $K+\tau M$ and discuss extensions to multiple right hand sides using both recycling and block methods. Numerical experiments using synthetic examples from a model inverse problem, Hydraulic Tomography, will demonstrate the performance of our algorithms.

Tania Bakhos
Institute for Computational and Mathematical Engineering
Stanford University
taniab@stanford.edu

Arvind Saibaba
Department of Electrical and Computer Engineering
Tufts University
arvind.saibaba@tufts.edu

Peter K. Kitanidis
Dept. of Civil and Environmental Engineering
Stanford University
peterk@stanford.edu

MS31

An Iterative Algorithm for Large-Scale Tikhonov Regularization

In this talk, we describe a hybrid iterative approach for computing solutions to large scale inverse problems via Tikhonov regularization. We consider a hybrid LSMR approach, where Tikhonov regularization is used to solve the subproblem of the LSMR approach. One of the benefits of the hybrid approach is that semiconvergence behavior can be avoided. In addition, since the regularization parameter can be estimated during the iterative process, the regularization parameter does not need to be estimated a priori, making this approach attractive for large scale problems. We consider various methods for selecting regularization parameters in the hybrid LSMR framework and discuss stopping criteria for the iterative method. Numerical examples from image processing illustrate the benefits and potential of the new approach.

Julianne Chung
Virginia Tech
jmchung@vt.edu

Katrina Palmer
Department of Mathematical Sciences
Appalachian State University
palmerk@appstate.edu

MS31

The Arnoldi-Tikhonov Framework: Choice of Regularization Parameters and Matrices

Iterative Krylov subspace methods play a central role in the regularization of large-scale inverse problems. We describe some new regularization parameter choice techniques for an Arnoldi-Tikhonov method. In addition, we describe how to adaptively choose a regularization matrix by exploiting the previous approximations, and we introduce two new strategies to approximate regularization terms weighted in a generic norm. The first strategy is exploits adaptive preconditioning, and the second strategy is based on a restarting procedure.

Silvia Gazzola
University of Padova
gazzola@math.unipd.it

James G. Nagy
Emory University
Department of Math and Computer Science

nagy@mathcs.emory.edu

Paolo Novatti
University of Padova
novati@math.unipd.it

MS31

Unbiased Predictive Risk and Discrepancy Principles Applied for LSQR Solutions of Ill-posed Least Squares

Numerous methods exist for finding regularization parameters when solving ill-posed linear inverse problems using full expansion methods such as the singular value decomposition (generalized SVD). Interpreting the propagation of noise through the Krylov subspace has limited development of robust parameter estimators for iterative solutions. We demonstrate the use of unbiased predictive risk estimators and discrepancy principles applied at the subspace level. The work is illustrated for a large scale three dimensional inversion of gravity data.

Rosemary A. Renaut
Arizona State University
School of Mathematical and Statistical Sciences
renaut@asu.edu

MS32

Half-Imex Time Integrators for Large Scale Simulations of Turbulent Incompressible Flows

We develop high order time integrators for the incompressible Navier-Stokes equations. This system is an index 2 differential-algebraic problem. For the velocity, we consider an explicit treatment of the nonsymmetric convective term, keeping the stiff viscous term implicit. At every stage, a pressure Poisson problem has to be solved. We combine the time integration with scalable domain decomposition solvers, and show the efficiency of this approach for large core counts (orders of hundreds of thousands).

Santiago Badia
International Center for Numerical Methods in Engineering
Universitat Politècnica de Catalunya, Barcelona, Spain
sbadia@cimne.upc.edu

Oriol Colomes
International Center for Numerical Methods in Engineering
ocolomes@cimne.upc.edu

MS32

Overview of Added-Mass Partitioned Algorithms for FSI Simulations

In recent work we have developed some partitioned algorithms for fluid-structure interaction (FSI) problems that couple various combinations of compressible and incompressible fluids with compressible bulk solids, rigid solids and structural beams. These algorithms overcome the added-mass instability for light solids that has plagued the traditional partitioned approach. This talk will give an overview of these new added-mass-partitioned (AMP) algorithms along with their implementation using deforming overlapping grids.

William Henshaw
Rensselaer Polytechnic Institute
henshw@rpi.edu

MS32

Overcoming the Added Mass Instability for Coupling Incompressible Flows and Elastic Beams

A new partitioned algorithm for coupling incompressible flows with elastic beams is described that overcomes the added-mass instability for light solids. The algorithm requires no sub-iterations and is fully second-order accurate. The new scheme is shown to be stable, even for very light beams, through the analysis of a model problem. The approach is then applied to the simulation of FSI problems involving beams undergoing large deformations using deforming composite grids.

Longfei Li
Department of Mathematical Sciences
University of Delaware
longfei@udel.edu

MS32

Partitioned Algorithms for FSI Problems Involving Elastic Solids Coupled to Compressible and Incompressible Fluids

New partitioned algorithms for the simulation of elastic solids coupled to compressible and incompressible fluids are described. For the compressible case, the coupling uses a interface project based on the solution of a fluid-solid Riemann problem, while Robin-Robin coupling conditions are derived for the incompressible case. Both coupling methods result in stable schemes, without sub-iterations, even for FSI problems with large added-mass effects. Numerical results are presented which verify the accuracy of these new added-mass partitioned (AMP) algorithms.

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

MS33

An Approach to Big Data in Inverse Problems

We develop innovative approaches to address the big data challenge in large-scale inverse problems and UQ governed by expensive PDEs. Various numerical results will be presented to demonstrate the effectiveness of our approaches.

Ellen B. Le
UT Austin
ellenle@ices.utexas.edu

Aaron Myers
UT Austin, ICES
aaron@ices.utexas.edu

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

MS33

Active Subspaces for the Design of Supersonic Low-

Boom Aircraft

We apply the active subspace method to a realistic simulation-based design problem: the Lockheed N+2 low-boom supersonic passenger jet. The method discovered a low-dimensional linear subspace of inputs that explained a majority of the variability in drag, lift, and sonic-boom. We exploited this subspace to find an optimal design at reduced computational cost, as well as automatically generate deformation modes that yield intuitive interpretations such as twist and camber.

Trent W. Lukaczyk
Stanford University
twl26@stanford.edu

Juan J. Alonso
Department of Aeronautics and Astronautics
Stanford University
jjalonso@stanford.edu

Francisco Palacios
Stanford University
fpalacios@stanford.edu

MS33

Parameter Selection Techniques for Disease Models

We discuss parameter selection techniques for nonlinearly parameterized, dynamic, disease models. The objective is to develop techniques to determine the sets of identifiable or influential parameters for these models. We compare the performance of Morris screening, Sobol analysis employing analysis of variance, Bayesian analysis, and active subspace techniques utilizing singular value decompositions (SVD) and QR factorizations for models with both single and multiple responses. We also detail the necessity of employing such parameter selection techniques before Bayesian model calibration and uncertainty propagation for models having a large number of inputs.

Jared Cook
Asbury University
jared.cook@asbury.edu

Nicholas Myers
University of Wisconsin-Milwaukee
njmyers@uwm.edu

Nina Ning
George Washington University
ning9386@gwmail.gwu.edu

Mami Wentworth
NC State
mtonoe@ncsu.edu

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

MS33

Likelihood-Informed Dimension Reduction for Bayesian Inverse Problems

The intrinsic dimensionality of an inverse problem is affected by prior information, the observations, and the properties of the forward operator. In many problems, changes from the prior to the posterior can be confined to a relatively low-dimensional subspace. We define and identify such a subspace, called the likelihood-informed subspace (LIS). We show that significant computational savings can be achieved by using this subspace to approximate the posterior distribution.

Tiangang Cui
Massachusetts Institute of Technology
tcui@mit.edu

James R. Martin
University of Texas at Austin
Institute for Computational Engineering and Sciences
jmartin@ices.utexas.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

Antti Solonen
Lappeenranta University of Technology
Department of Mathematics and Physics
antti.solonen@gmail.com

Alessio Spantini
Massachusetts Institute of Technology
spantini@mit.edu

Luis Tenorio
Colorado School of Mines
ltenorio@mines.edu

MS34

Sparse Grid and Reduced Basis Approximation of Bayesian Inverse Problems

We present a computational reduction framework for efficient and accurate solution of Bayesian inverse problems on high- or infinite-dimensional parameter spaces that commonly face the curse of dimensionality and large-scale computation. For the approximation of high or infinite dimensional integration, we take advantage of sparsity in the parametric solution maps in novel dimension-adaptive sparse grid interpolation and quadrature algorithms. For large scale problems, we also exploit intrinsic sparsity in the solution map and the high-fidelity approximation and propose a novel, goal-oriented reduced basis method.

Peng Chen
ETH Zurich
peng.chen@sam.math.ethz.ch

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

MS34

Goal-Oriented Model Adaptivity for Inference

We present a goal-oriented adaptive method for inverse problems based on the use of multiple models. Previous work has addressed the use of goal-oriented grid adaptation in inverse problems, but has been restricted to the single model paradigm. Our method develops estimates for the error in the prediction quantity of interest due to the use

of a lower fidelity model in the inference process. We also present numerical experiments illustrating the efficiency of the method.

Vikram Garg
ICES, U Texas, Austin
vikram.v.garg@gmail.com

Harriet Li
MIT
kameeko@mit.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@MIT.EDU

MS34

An Empirical Objective Bayes Method for Large Inverse Problems

The Geostatistical Approach is a general, flexible, and comprehensive formalism for solving underdetermined parameter estimation problems (inverse problems). The method is increasingly used for problems with many unknowns and observations. I will review some of the tools, all of which exploit the problem structure, to make this approach computationally tractable for large problems, including Hierarchical Matrices and Fast Multipole Method for fast matrix vector products, compression methods for dimensionality reduction, and iterative methods.

Peter K. Kitanidis
Dept. of Civil and Environmental Engineering
Stanford University
peterk@stanford.edu

MS34

Applying UQ Approaches to Random Ordinary Differential Equations

Random Ordinary Differential Equations (RODEs) can offer an alternative concept for Stochastic Differential Equations via path-wise solutions of ODEs, avoiding Ito calculus and frequently using coloured noise instead of white noise processes for the underlying random effects. We focus on the numerical solution of RODEs and their connection to UQ techniques to improve computational performance for applications such as earthquake-induced motion of wireframe buildings.

Tobias Neckel
Technische Universität München
neckel@in.tum.de

Hans-Joachim Bungartz
Technische Universität München, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

Alfredo Parra
Technische Universität München
alfredo.parra@tum.de

MS35

Matrix-Free Solvers for Robust PCA and Distance Matrix Completion

The problem of (stable) robust PCA is to separate a data matrix into low-rank, sparse and noise components. Due to the cost of standard algorithms, which require full or partial SVDs at every step, the problem is extremely computationally difficult for large-scale applications. We present a new method that incorporates more information than a traditional first-order method, but without increased cost per iteration, and show very favorable results compared to other recent solvers. We also present variants of this algorithm. Lastly, we discuss the specific challenges of a variant of distance matrix completion that is applicable for locating sites on chromosomes.

Stephen Becker
UC Boulder
stephen.becker@colorado.edu

Aleksandr Aravkin
IBM T.J. Watson Research Center
saravkin@us.ibm.com

Aurelie Lozano
IBM Research
T. J. Watson Research Center
aclozano@us.ibm.com

MS35

Matrix Free Quadratic-penalty Methods for PDE-constrained Optimization

The large scale of seismic waveform inversion makes matrix free implementations essential. We show how to exploit the quadratic penalty structure to construct matrix free reduced-space and full-space algorithms, which have some advantages over the commonly used Lagrangian based methods for PDE-constrained optimization. This includes the construction of effective and sparse Hessian approximations and reduced sensitivity to the initial guess. A computational bottleneck is the need to solve a large least squares problem with a PDE block. When direct solvers are not available, we propose a fast matrix free iterative approach with reasonable memory requirements. It takes advantage of the structure of the least squares problem with a combination of preconditioning, low rank decomposition and deflation.

Bas Peters
UBC
bpeters@eos.ubc.ca

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

MS35

Compressing Clustered Data using Sparse NMF

We propose a method for storing clustered data compactly, where each cluster is a collection of many similar datasets in matrix form (such as a set of related images). We first compute basis elements for each cluster using low-rank SNMF via PDCO. We then use the preprocessing approach of Gillis (2012) to sparsify the full set of (already sparse) basis elements, without significant loss of quality.

Michael A. Saunders

Systems Optimization Laboratory (SOL)
Dept of Management Sci and Eng, Stanford
saunders@stanford.edu

San Kim
ICME, Stanford University
sankim@stanford.edu

MS35

Dimensionality Reduction and Uncertainty Quantification for Inverse Problems

Many inverse problems in science and engineering involve multi-experiment data and thus require a large number of forward simulations. Dimensionality reduction techniques aim at reducing the number of forward solves by (randomly) subsampling the data. In the special case of non-linear least-squares estimation, we can interpret this compression of the data as a (low-rank) approximation of the noise covariance matrix. We shoe that this leads to different design criteria for the subsampling process. Furthermore, the resulting low-rank structure can be exploited when designing matrix-free methods for estimating (properties of) the posterior covariance matrix. Finally, we discuss the possibility of estimating the noise covariance matrix itself.

Tristan van Leeuwen
Mathematical Institute
Utrecht University
t.vanleeuwen@uu.nl

MS36

A Fourth Order Accurate Embedded Boundary Method for the Wave Equation in Second Order Form

A fourth-order accurate embedded boundary method for the scalar wave equation with Dirichlet or Neumann boundary conditions is described. The method is based on a compact Pade-type discretization of spatial derivatives together with a Taylor series method (modified equation) in time. A novel approach for enforcing boundary conditions is introduced which uses interior boundary points instead of exterior ghost points. This technique removes the small-cell stiffness problem for both Dirichlet and Neumann boundary conditions, is more accurate and robust than previous methods based on exterior ghost points, and guarantees that the solution is single-valued when slender bodies are treated. Numerical experiments are presented to illustrate the stability and accuracy of the method as well as its application to problems with complex geometries.

Daniel Appelo
University of New Mexico
appelo@unm.edu

MS36

High-Order Numerical Methods for Elliptic Interface Problems

In our talk we consider elliptic equations with piecewise smooth coefficients in irregular domains separated by arbitrary shaped interfaces. Our numerical approach for these problems is based on generalized Calderon's operators and the Difference Potentials Method. The developed algorithm easily handles curvilinear boundaries, variable coefficients and general boundary conditions. High-order accuracy is achieved efficiently since only simple Cartesian grids used regardless of the interface shape or boundary. The performance of the numerical method is illustrated in several numerical examples in 2D.

Michael Medvinsky
University of Utah
Department of Mathematics
medvinsky.michael@gmail.com

Yekaterina Epshteyn
Department of mathematics
University of Utah
epshteyn@math.utah.edu

Semyon V. Tsynkov
Department of Mathematics
North Carolina State University
tsynkov@math.ncsu.edu

Eli Turkel
Tel-Aviv University
eliturkel@gmail.com

MS36

A Nitsche Stabilized Fictitious Domain Finite Element Method for the Wave Equation

We give a weak formulation for solving the wave equation ($\ddot{u} = \nabla^2 u + f$) on a 2-dimensional fictitious domain. In the spatial finite element discretization, boundaries do not conform to element boundaries. Boundary conditions are enforced weakly by Nitsche's method. Additional penalty terms ensure a non-stiff temporal system. We give optimal a priori error estimates: second order for $u - u_h$ and $\dot{u} - \dot{u}_h$ and first order for $\nabla(u - u_h)$, for linear elements in L_2-norm. Numerical experiments verify the theory. Extensions to higher orders are discussed.

Simon Sticko
Uppsala University
Department of Information Technology
simon.sticko@it.uu.se

Gunilla Kreiss
Division of Scientific Computing
Uppsala University
gunilla.kreiss@it.uu.se

MS37

Cascadic Multilevel for Saddle Point Least-Squares Methods

We present a cascadic multilevel method for approximating variational formulations of symmetric saddle point systems. The discretization algorithm is based on the availability of families of stable finite element pairs and on fast and accurate solvers for symmetric positive definite systems. On each fixed level an efficient solver such as the gradient or the conjugate gradient algorithm for inverting a Schur complement is implemented. As a main application of our approach we define the "saddle point least-squares' method for solving first order systems of PDEs and relate it with the Bramble-Pasciak's least-squares approach. We apply our *saddle point least-squares* method to discretizing the time harmonic Maxwell equations. The variable of interest for the Maxwell equations, the magnetic and electric vector fields, become the constrained variable in the *saddle*

point least-squares formulation, and bases or stiffness matrices associated with these variables are not needed. The cascadic saddle point least-squares iterative discretization we build does not involve edge elements or spaces of bubble functions, and is suitable for preconditioning.

Constantin Bacuta
University of Delaware
Department of Mathematics
bacuta@math.udel.edu

MS37

Multigrid Method for Linear Elasticity with Weakly Imposed Symmetry

We present a new abstract framework for the convergence analysis of the subspace correction methods applied to the system of linear algebraic equations associated to the differential operator arising from the finite element-based discrete systems of the linear elasticity with weakly imposed symmetry. The method is proven to be convergent uniformly with respect to the mesh size and parameters that appear in the equations. A special case of our theory can provide a transparent and improved analysis for the multigrid methods developed for the pseudo-stress formulations of Stokes equation. Some sample numerical experiments for two dimensional cases are provided to illustrate the theory.

Young Ju Lee
Department of Mathematics
Rutgers University
yjlee@txstate.edu

MS37

Solver for Structure-Preserving Discretization of Incompressible MHD Equations

A uniform solver for a stable structure-preserving finite element discretization of incompressible MHD equation will be introduced. Numerical results are provided to verify the structure-preserving property and demonstrate the effectiveness of the solver.

Yicong Ma
Department of Mathematics
The Pennsylvania State University
ma_y@math.psu.edu

Xiaozhe Hu
The Pennsylvania State University
xiaozhe.hu@tufts.edu

Jinchao Xu
Pennsylvania State University
xu@math.psu.edu

MS37

Modeling and Numerical Studies for Fluid-Structure Interactions

In this talk, I will present our recent study on a dynamic fluid-structure interaction (FSI) problem involving with a rotational and elastic structure by using the arbitrary Lagrangian-Eulerian (ALE) method for fluid model in Eulerian description and the St.Venant-Kirchhoff structural model in Lagrangian description, and design its monolithic mixed finite element approximation. In addition, our studies on fictitious domain method and full Eulerian phase field method for FSI problems will be also introduced. In particular, in terms of ALE method, we developed a nonlinear rotational and deformable structural model for FSI for the first time. The technique of master-slave relations is employed to realize the interfacial kinematic condition on the interface of fluid and structure. Velocity is adopted as the principle unknown to reformulate the structural model. Our algorithm can also handle a large and irregular fluid flow channel in which the rotational structure is immersed with ALE method. We use Newtons method to linearize, and Galerkin/least-square (GLS) and streamline-upwind/Petrov-Galerkin (SUPG) schemes to stabilize the mixed finite element discretization of fluid equations with ALE approach. Numerical experiments are successfully carried out for a simplified turbine in 2D and a realistic turbine in 3D which is deforming as well as spinning around its rotation of axis cases due to the fluid impact.

Pengtao Sun
University of Nevada Las Vegas, Math. Dept.
pengtao.sun@unlv.edu

MS38

Algorithms for Hessenberg-Triangular Reduction in Parallel

A recent improvement of the parallel QZ algorithm has made the initial reduction to Hessenberg-Triangular form a new bottleneck in the solution of generalized non-symmetric matrix eigenvalue problems. We propose a new distributed algorithm for Hessenberg-Triangular reduction based on a novel static scheduling technique. Experiments demonstrate that the new algorithm improves on the current state-of-the-art and also identifies the factors that limit scalability. The developed code is ScaLAPACK compatible.

Björn Adlerborn
Umeå University
Computing Science and HPC2N
adler@cs.umu.se

Lars Karlsson
Umeå University, Dept. of Computing Science
larsk@cs.umu.se

Bo T. Kågström
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se

MS38

Avoiding Communication in Distributed-Memory Tridiagonalization

We discuss theoretical and practical improvements to parallel symmetric tridiagonalization algorithms, augmenting existing algorithms with more communication-efficient parallel QR factorization kernels and multiple steps of band reduction. Our theoretical analysis suggests that our approach can reduce interprocessor communication and memory bandwidth cost by $O(p^{1/6})$ on p processors, using some additional memory. We validate our analysis by implementing a three-step approach, which uses our key ideas to reduce interprocessor communication, albeit by a constant factor rather than $O(p^{1/6})$.

Grey Ballard

Sandia National Laboratories
gmballa@sandia.gov

James W. Demmel
University of California
Division of Computer Science
demmel@berkeley.edu

Nicholas Knight
UC Berkeley
knight@cs.berkeley.edu

Edgar Solomonik
University of California at Berkeley
solomon@eecs.berkeley.edu

MS38

A Parallel Multishift QZ Algorithm with Aggressive Early Deflation for Distributed Memory HPC Systems

Appearing frequently in applications, generalized eigenvalue problems represent one of the core NLA problems. We propose a parallelization of the QZ algorithm by Moler and Stewart that incorporates all modern ingredients of dense eigensolvers, such as multishift and aggressive early deflation techniques. To deal with (possibly many) infinite eigenvalue, a new parallel deflation strategy is developed. Numerical experiments for random and application examples demonstrate the effectiveness of our algorithm on distributed memory HPC systems.

Björn Adlerborn, Bo T. Kågström
Umeå University
Computing Science and HPC2N
adler@cs.umu.se, bokg@cs.umu.se

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

MS38

Performance Evaluation of Sparse Matrix-Vector Multiplication Using GPU/MIC Cluster

Sparse matrix-vector multiplication (SpMV) is an important computational kernel for many applications. Recently, the number of computing systems equipped with NVIDIA's GPU and Intel's Xeon Phi coprocessor based on the MIC architecture has been increasing. In this talk, we present a performance evaluation of parallel SpMV using GPU/MIC cluster. As shown by the results, the implementation for the GPU/MIC cluster attained higher performance than the implementation for the CPU cluster in some matrices.

Hiroshi Maeda
Graduate School of Systems and Information Engineering
University of Tsukuba
maeda@hpcs.cs.tsukuba.ac.jp

Daisuke Takahashi
Faculty of Engineering, Information and Systems
University of Tsukuba
daisuke@cs.tsukuba.ac.jp

MS40

Integrating Mathematical Modeling and Computer Simulation with Experimental Synthesis and Characterization of Materials

To design materials through computational modeling and simulation requires close collaborations among applied mathematicians, materials theorists, and experimentalists. The author will discuss his experiences working with experimental groups in synthesis and characterization of materials and with applied math groups in numerical computation. Examples to be discussed include designing domain structures and properties of oxide thin films. The author will share his views on the ideal collaboration mechanisms among applied mathematicians, materials theorists, and experimentalists.

Long-qing Chen
Department of Materials Science and Engineering
Pennsylvania State University
lqc3@psu.edu

MS40

Mathematical Challenges in Nonequilibrium Approaches to Amorphous Solids: Quantifying Disorder, Predicting Plasticity, Accelerating Simulation

Amorphous solids, materials that exhibit an absence of long range order, exist in every class of materials. Despite their ubiquity few methodologies exist to quantify their structure, predict their response or achieve the simulation time scales necessary to understand their complex behavior. I will review some of the underlying gaps in our mathematical knowledge that make the study of amorphous solids so demanding. Some of these issues limit progress not only in amorphous solids, but in atomic simulation generally. In particular I will discuss the difficulty of quantifying disorder in a mathematically rigorous way through generalizations of the Frank-Kasper criterion. I will also consider the challenges that face modeling the plastic flow of materials that bridge the gap between elastic solids and liquids. Finally I will consider the limitations that atomistic simulations face when studying materials that have inherent time scales that may span a broad range from nanoseconds to seconds.

Michael Falk
Materials Science and Engineering
Johns Hopkins University
mfalk@jhu.edu

MS40

Structural Optimization and 3D Printing

We have known since the 1980's that some structural optimization problems call for designs with a high degree of spatial complexity. However such results were mainly viewed as theoretical benchmarks, since the manufacture of spatially complex designs has been difficult. The development of 3D printing (also known as additive manufacturing) calls for re-examination of this area. With 3D printing, spatially complex structures are relatively easy to manufacture. It is natural to ask how this capability can be used to good effect. The optimization of 3D printed structures

raises many challenging questions, most of them open.

Robert V. Kohn
Courant Institute of Mathematical Sciences
New York University
kohn@cims.nyu.edu

MS40

DNA-Functionalized Nanoparticle Assembly and Crystallization

The selectivity of DNA recognition inspires an elegant protocol for designing versatile nanoparticle (NP) assemblies. We use molecular dynamics simulations to analyze static and dynamic aspects of the assembly process and identify key ingredients for the assembly of NP superlattices through DNA hybridization. A scale-accurate coarse-grained model captures the relevant contributions to the kinetics of the DNA hybridization process and is able to recover all experimentally reported to date binary superlattices of DNA functionalized nanoparticles. Using multiscale modeling we show that through very slow cooling, DNA functionalized nanoparticles assemble into superlattices with a specific crystal habit. We reproduce polyhedra growth in silico, and confirm the Wulff shape for the BCC. Due to defects including twinning and stacking faults in the lattice, the FCC system does not show any uniform shape. Simulated crystal habits of both the BCC and FCC system are consistent with experiments.

Monica Olvera De La Cruz
Northwestern University
m-olvera@northwestern.edu

MS41

A New Penalization Method for the Shallow Water Equations with Applications to Global Ocean Flow

We propose a new mass and energy conserving Brinkman boundary condition penalization for the rotating shallow water equations. This penalization does not lead to higher wave speeds in the solid region. The error estimates for the penalization are derived analytically and verified numerically for linearized one dimensional equations. The penalization is implemented in a conservative dynamically adaptive wavelet method for the rotating shallow water equations on the sphere with bathymetry and coastline data from NOAA's ETOPO1 database. This code could form the dynamical core for a future global ocean model. The potential of the dynamically adaptive ocean model is illustrated by using it to simulate the 2004 Indonesian tsunami and wind-driven global ocean circulation.

Nicholas Kevlahan
Department of Mathematics & Statistics
McMaster University
kevlahan@mcmaster.ca

MS41

Imposing Dirichlet and Neumann Conditions in Fourier Pseudospectral Methods Using Volume Penalization

The volume penalization method for computing confined fluid and plasma flows in complex geometries while using Fourier pseudo-spectral discretizations is reviewed. The mathematical principle of this technique is described and simple examples for imposing Dirichlet and Neumann boundary conditions are given. Applications for fluid and plasma turbulence in two and three space dimensions illustrate the efficiency of the method.

Kai Schneider
Universite de Provence, Aix-Marseille
Centre de Mathematiques et Informatique
kschneid@cmi.univ-mrs.fr

MS41

High-Order Fourier-Penalty Methods for PDEs on Irregular Domains

Penalty methods offer an attractive approach for solving partial differential equations (PDEs) on domains with curved or moving boundaries. The new penalized PDE is then attractive to solve since one no longer needs to actively enforce the boundary conditions. Despite the simplicity, these methods have suffered from poor convergence rates. I will show how to systematically construct a new class of penalization terms which improve the order of convergence. I will then demonstrate that the new approach allows one to devise higher order, stable, Fourier spectral methods to solve the penalized PDE.

David Shirokoff
New Jersey Institute of Technology
david.g.shirokoff@njit.edu

MS42

Designing Self-Propelling Microgel Swimmer

Using dissipative particle dynamics, we design a hydrogel micro-swimmer that is actuated by a periodically applied stimulus. The gel has an X-shaped geometry and two bonded layers, one of which is responsive to the stimulus and the other is passive. When the stimulus is turned on, the responsive layer swells and causes the swimmer to deform. We demonstrate that when such stimulus-induced deformations occur periodically, the gel swimmer effectively propels forward through the viscous fluid.

Alexander Alexeev
George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
alexander.alexeev@me.gatech.edu

Svetoslav Nikolov, Peter Yeh
Georgia Institute of Technology
snikolov3@gatech.edu, pyeh3@gatech.edu

MS42

The Effect of Curved or Flat Edges Microchannels on Vortex Entrapment of Particles as seen in Lattice-Boltzmann Simulations

Numerical simulations were conducted to investigate if and how curved or flat channel wall surfaces, as part of the surface topography in micro-flow channel walls, would impact vortex development and vortex entrapments of uniform-size particles in pulsating and non-pulsating flows. This research also demonstrated the sensitivity of the location of particle-trapping vortices to particle concentrations.

John Carrola
Southwest Research Institute
john.carrola@swri.org

Hakan Basagaoglu

Southwest Research Institute
Mechanical Engineering Division
hakan.basagaoglu@swri.org

MS42

Scaffold-free Three-dimensional Hepatocyte Assembly for Liver Tissue Engineering

We demonstrate a scaffold-free hepatocyte assembly method, which enables generation of repeating and symmetric cellular structures with high cell-packing density and viability. Standing waves established at the airliquid interface are used to pack cells in 3D architecture without scaffolding materials. The pattern of 3D architecture is determined by the topography of standing waves and can be dynamically controlled by vibrational frequency and amplitude of standing waves.

Utkan Demirci
Stanford University
School of Medicine
utkan@stanford.edu

Pu Chen
Stanford Medicine
Department of Radiology
puchen@stanford.edu

MS42

Brownian Motion of Arbitrarily Shaped Particles Confined in Two-Dimensions

I will present our studies on Brownian motion of microfabricated boomerang particles in two-dimensions. We show that due to translation-rotation coupling, mean squared displacements of single particles exhibit a nonlinear crossover from short time faster to long time slower diffusion, and the mean displacements for fixed initial orientation are biased. I will also show an analytical model based on Langevin theory to explain how the symmetry of particle shape affect Brownian motion.

Qi-Huo Wei
Kent State University
Liquid Crystal Institute
qwei@kent.edu

MS43

A Kokkos Implementation of Albany: A Performance Portable Multiphysics Simulation Code

Albany is a C++ object-oriented, parallel, unstructured-grid, implicit finite element code for solving partial differential equations (PDEs) in various fields of engineering applications and multiphysics simulations in particular. Such complex simulations require to be run on modern Supercomputers and performance portability across variety of HPC architectures has become a critical issue. We have developed a performance portable implementation of the Albany code, based on Kokkos programming model from Trilinos. A new Albany-Kokkos implementation is a single code base which runs and is performant on diverse HPC architectures, and is expected to be performant on future architectures that are supported by the Kokkos library.

Irina Demeshko, H. Carter Edwards, Michael Heroux,
Roger P. Pawlowski
Sandia National Laboratories
ipdemes@sandia.gov, hcedwar@sandia.gov,
maherou@sandia.gov, rppawlo@sandia.gov

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

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

MS43

Evaluations of Directive Based Programming Model for GPUs and Extensions for Performance Portability

Recently, the number of heterogeneous supercomputers using accelerators such as GPUs is increasing. These systems have been showing remarkable performance on real-world applications such as Computational Fluid Dynamics (CFD). However, porting of legacy CPU-based applications to systems with accelerators is still big challenge. One of the reasons of this problem is low-level programming model such as CUDA, which is the most widely used for porting of these applications. To solve this problem, existing approach includes high-level programming model such as OpenACC but it still has a problem about performance portability between devices which have different performance characteristics about data layout. To improve the performance portability, we suggest an abstraction of data layout. We also implement and evaluate a directive-based source-to-source (extended OpenACC to OpenACC) translator that automatically transforms data layout to suitable form.

Tetsuya Hoshino
Tokyo Institute of Technology
hoshino@matsulab.is.titech.ac.jp

Naoya Maruyama
RIKEN Advanced Institute for Computational Science
nmaruyama@riken.jp

Satoshi Matsuoka
Tokyo Insitute of Technology
matsu@is.titech.ac.jp

MS43

Optimization of Preconditioned Iterative Linear Solvers Using Openmp/openacc on Gpu and Mic

OpenMP is widely used for accelerating programs on CPU and MIC. Also, some people use OpenACC on mainly GPU. These parallel programming environment can make parallel programs easy with similar programming fashion. However, in order to obtain good performance, users may have to make differ programs in view of target architectures. In this talk, we show the implementation and performance of ICCG method on OpenMP and OpenACC. We compare the obtained performances with each other.

Satoshi Ohshima
The University of Tokyo
Information Technology Center
ohshima@cc.u-tokyo.ac.jp

Masaharu Matsumoto
Information Technology Center
The University of Tokyo
matsumoto@cc.u-tokyo.ac.jp

Takahiro Katagiri
The University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

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

Kengo Nakajima
The University of Tokyo
Information Technology Center
nakajima@cc.u-tokyo.ac.jp

MS43

OCCA: An Extensible Portability Layer for Many-Core Programming

There are a number of relatively popular APIs for programming many-core CPUs, GPUs, and accelerators including OpenMP, pThreads, CUDA, OpenCL, and OpenACC. This manufacturer driven fragmentation echoes the scattershot approach to distributed computing before the Message Passing Interface standard became the de facto standard. We have developed OCCA as a manufacturer independent runtime library and abstraction layer. OCCA enables a programmer who is comfortable with the notions of parallel loops and barriers to write parallel threaded kernel code that is parsable at runtime as either OpenMP, pThreads, CUDA, or OpenCL. The OCCA library is easily extensible and has native interfaces that can be called from C++, C, F90, Python, Java, or Julia. I will show examples drawn from numerical PDEs that indicate some application compute kernels can achieve good performance across several platforms while other examples show that more than one kernel implementation may be required to achieve best possible performance across platforms.

Tim Warburton
Rice University
Department of Computational and Applied
timwar@rice.edu

David Medina
Rice University
dsm5@rice.edu

Amik St-Cyr
Computation and Modeling,
Shell International E&P, Inc.
amik.st-cyr@shell.com

MS44

An Efficient Spectral Element Helmholtz Solver with an Accurate Treatment for Transparent Boundary Condition for Periodic Lossy Media

We present a high-order spectral element method for solving layered media scattering problems featuring an operator that transparently enforces the far-field boundary condition. The incorporation of this Dirichlet-to-Neumann (DtN) map into the spectral element framework is a novel aspect of this work, and the resulting method can accommodate plane-wave radiation of arbitrary angle of incidence. In order to achieve this, the governing Helmholtz equations subject to quasi-periodic boundary conditions are rewritten in terms of periodic unknowns. We construct a spectral element operator to approximate the DtN map, thus ensuring nonreflecting outgoing waves on the artificial boundaries introduced to truncate the computational domain. We present an explicit formula that accurately computes the Fourier data involved in the DtN map on the spectral element discretization space. Our solutions are represented by the tensor product basis of one-dimensional Legendre-Lagrange interpolation polynomials based on the Gauss-Lobatto-Legendre grids. We will demonstrate the scattered field in singly and doubly layered media with smooth and nonsmooth interfaces.

Ying He
UC Davis
yinghe@math.ucdavis.edu

MS44

Electromagnetic Field Enhancement for Metallic Nano-gaps

There has been increasing interests in electromagnetic field enhancement and extraordinary optical transmission effect through subwavelength apertures in recent years, due to its significant potential applications in biological and chemical sensing, spectroscopy, terahertz semiconductor devices, etc. In this talk, I will present a quantitative analysis for the field enhancement when an electromagnetic wave passes through small metallic gaps. We focus on a particular configuration when there is extreme scale difference between the wavelength of the incident wave, the thickness of metal films, and the size of gap aperture. Based upon a rigorous study of the perfect electrical conductor model, we show that enormous electric field enhancement occurs inside the gap. Furthermore, the enhancement strength is proportional to ratio between the wavelength of the incident wave and the thickness of the metal film, which could exceed 10000 due to the scale difference between the two. On other hand, there is no significant magnetic field enhancement inside the gap. The ongoing work along this research direction will also be discussed.

Junshan Lin
Auburn University
jzl0097@auburn.edu

MS44

Highly tuned hybrid MPI/OpenACC implementation with GPUDirect communication for electromagnetic solvers based on spectral element discretization

I will present a collaborative work on highly tuned hybrid MPI/OpenACC implementation with GPUDirect communcation for solving electromagnetic systems based on high-order spectral element discretizations. The OpenACC implementation covers the full solution routines including element-by-element operator evaluation as well as tuned MPI communication kernels to effect the near-neighbor flux exchanges. I will demonstrate performance and analysis on up to 16,384 GPUs on the Cray XK7 supercomputer showing more than 3x speedup, compared to the MPI-only CPU performance on the same number of nodes (262,144 CPUs) for the problem sizes of up to 6.9 billion data points.

MiSun Min

Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

MS44

A High Order Perturbation of Surfaces Method for Simulating Surface Plasmons on Periodic Gratings

In this talk we describe a High Order Perturbation of Surfaces (HOPS) method for simulating the scattering of electromagnetic waves by periodic grating structures. The method amounts to Nystrom's method applied to a class of Integral Equations for the Helmholtz equation on periodic domains inspired by the recent work of Fokas and collaborators on novel solution formulas for boundary value problems. These Integral Equations have a number of advantages over standard alternatives including: (i.) ease of implementation (high-order spectral accuracy is realized without sophisticated quadrature rules), (ii.) seamless enforcement of the quasiperiodic boundary conditions (no periodization of the fundamental solution, e.g. via Ewald summation, is required), and (iii.) reduced regularity requirements on the interface proles (derivatives of the deformations do not appear explicitly in the formulation). We show how these can be efficiently discretized and utilized in the simulation of surface plasmon excitations on periodic metal gratings in configurations of engineering interest.

David P. Nicholls
University of Illinois at Chicago
nicholls@math.uic.edu

MS45

Performance of Parallel Algorithms for Particle Transport on Massively Parallel Architectures

We describe algorithms that execute multi-octant discrete-ordinate transport sweeps in the minimum possible stage count for a given partitioning across processors and given work-unit aggregation. We describe automated selection of the best partitioning and aggregation for a given problem on a given machine. Spatial grids can be locally unstructured within structured blocks. We present results to $O(10^6)$ parallel processes, including results from nested parallel algorithms in which each MPI process uses multiple threads.

Marvin L. Adams, Michael Adams, W. Daryl Hawkins, Timmie Smith
Texas A&M University
mladams@tamu.edu, i.am.michael.adams@gmail.com, dhawkins@tamu.edu, timmie@tamu.edu

Lawrence Rauchwerger
Computer Science Department
Texas A&M University
rwerger@cs.tamu.edu

Nancy Amato
Texas A & M University
amato@tamu.edu

Teresa S. Bailey
LLNL
bailey42@llnl.gov

Robert Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

MS45

Particle-Particle, Particle-Mesh Methods for Electromagnetic Problems

In this talk we revisit the topic of subcell methods inside of traditional particle-in-cell methods. The subcell method proposed in this work is based on using the boundary integral tree code as a subcell model inside of a multidimensional PIC code. 1D introduces some some greatly reduces issues associated with numerical heating and provides greater accuracy in the calculation. This work investigates issues of numerical seating and accuracy in a multi-D setting. Further potential roadmap is presented for producing heating time domain electromagnetic particle-particle particle-mesh method that is well-suited to the simulation of non-relativistic plasmas.

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

Eric Wolf
Michigan State University
wolferi1@msu.edu

MS45

Stochastic Galerkin Method for Hamilton-Jacobi Equations with Uncertainty

Abstract not available at time of publication.

Jingwei Hu
The University of Texas at Austin
hu342@purdue.edu

Shi Jin
Shanghai Jiao Tong University, China and the University of Wisconsin-Madison
jin@math.wisc.edu

Dongbin Xiu
University of Utah
dongbin.xiu@utah.edu

MS45

Uncertainty Quantification in Kinetic Theory

In this talk we will study generalized polynomial chaos (gPC) approach to transport equation with uncertain coefficients/inputs and show that they can be made asymptotic-preserving, in the sense that in the diffusion limit the gPC scheme for the transport equation approaches to the gPC scheme for the diffusion equation with random diffusion coefficient. This allows the implemention of the gPC method without numerically resolving (by space, time, and gPC modes) the small mean free path for transport equation in the diffusive regime. We will also discuss more general hyperbolic equations with (random) geometric source terms and the relevant stochastic well-balanced property.

Shi Jin
Shanghai Jiao Tong University, China and the University of Wisconsin-Madison

jin@math.wisc.edu

Dongbin Xiu, Xueyu Zhu
University of Utah
dongbin.xu@utah.edu, xzhu@sci.utah.edu

MS46

Testing the Multilayer Shallow Shelf Approximation Against Higher-order Models

In this talk, I will introduce a new hybrid ice flow model (called MSSA) which generalises the Shallow Shelf Approximation (SSA) by a Multilayer approach. Advantageously, the MSSA keeps intact the elliptic structure of the SSA and is mechanically exhaustive (it contains both: membrane and vertical stresses) while being of 2D complexity. Comparative results will be presented to assess the mechanical and the computational performances of the MSSA with respect to classical higher-order models.

Guillaume Jouvet
Freie Universität Berlin
Institut für Mathematik
jouvet@vaw.baug.ethz.ch

MS46

A Finite Element Three-Dimensional Stokes Ice Sheet Dynamics Model with Enhanced Local Mass Conservation

In this paper, we present and discuss a new finite element Stokes ice sheet dynamics model that enforces local element-wise mass conservation by enriching the pressure finite element space by adding the discontinuous piecewise constant pressure space to the Taylor-Hood pressure space. Through various numerical tests based on manufactured solutions, benchmark test problems, and the Greenland ice-sheet, we demonstrate that, for ice-sheet modeling, the enriched Taylor-Hood finite element model remains highly accurate and efficient, and is physically more reliable and robust compared to the classic Taylor-Hood finite element model.

Wei Leng
Chinese Academy of Sciences
State Key Laboratory of Scientific and Engineering Computing
wleng@lsec.cc.ac.cn

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Max Gunzburger
Florida State University
School for Computational Sciences
mgunzburger@fsu.edu

MS46

Improving Grounding Line Discretization using an Embedded-Boundary Approach in BISICLES

Correctly representing grounding line dynamics is of fundamental importance in modeling marine ice sheets. We have developed a grounding-line discretization based on the Chombo embedded-boundary cut-cell framework. This promises better representation of grounding lines vs. a traditional stair-step discretization on Cartesian meshes like those used in the block-structured AMR BISICLES code. Also, the fundamental discontinuous nature of flow across the grounding line is respected.

Daniel Martin
Lawrence Berkeley National Laboratory
dfmartin@lbl.gov

Peter O. Schwartz
Lawrence Berkeley Laboratory
poschwartz@lbl.gov

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

MS46

On the Development and Performance of a First Order Stokes Finite Element Ice Sheet Dynamical Core Built Using Trilinos Software Components

This talk describes the new Albany/FELIX parallel, scalable and robust First-Order (FO) Stokes finite element ice sheet code developed using Trilinos libraries. Focus will be on the computational aspects of the code: verification; multilevel preconditioning to achieve good parallel scalability for FO systems; homotopy continuation techniques for robust nonlinear solves; many-core performance portability. Coupling of Albany/FELIX to other land ice dycores (CISM, MPAS) for dynamic simulations and global climate runs will also be discussed.

Irina K. Tezaur
Sandia National Laboratories
ikalash@sandia.gov

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

Mauro Perego
CSRI Sandia National Laboratories
mperego@fsu.edu

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

MS47

The most current list of participating companies is available at www.siam.org/meetings/cse15/career.php.

For the most recent list of participating companies visit http://www.siam.org/meetings/cse15/career.php

Bill Kolata
SIAM
kolata@siam.org

MS48

Application of a Speed Up Fast Direct Solver for the Solution of the Lippmann-Schwinger Equation

In this work, we applied the HODLR fast solver of Sivaram

and Darve on the Lippmann-Schwinger integral equation for the solution of the Helmholtz problem. We split and discretize the domain by using an adaptive level-restricted tree structure based on its contrast function. To speed up the solver, we use the series representation of the Hankel functions to pre-compute tables that can be used with the tree structure to obtain the far field and local iterations between the discretizartion points.

Carlos C. Borges
Worcester Polytechnic Institute
borges@cims.nyu.edu

Lise-Marie Imbert-Gerard
CIMS, New York University
imbertgerard@cims.nyu.edu

Sivaram Ambikasaran
Department of Mathematics
Courant Institute of Mathematical Sciences
sivaram@cims.nyu.edu

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

MS48

Yee Scheme Coupled with Linear Current in Magnetic Plasmas with Varying Coefficients

We analyze the stability of the Yee scheme for non stationary Maxwell's equations coupled with a linear current model, with highly varying coefficients. Indeed the usual procedure is unstable for physical situations that correspond to strongly magnetized plasmas in X-mode (TE) polarization. We propose to use first order clustered discretization of the vectorial product that gives back a stable coupling. Validation is performed with physically based problems. The equivalent time harmonic problem (elliptic equation) yields a similar behavior.

Bruno Despres
University Paris VI
Laboratory LJLL
despres@ann.jussieu.fr

Martin Campos Pinto
CNRS-LJLL UPMC
campos@ann.jussieu.fr

Stéphane Heuraux
Institut Jean Lamour, UMR 7198 CNRS-University Lorraine
stephane.heuraux@univ-lorraine.fr

Filipe da Silva
Instituto de Plasmas e Fusao Nuclear, Instituto Superior
tanatos@ipfn.ist.utl.pt

MS48

Approximation of Degenerate Elliptic Equations with Muckenhoupt Coefficients: a priori and a posteriori Analyses and Efficient Solvers

In many problems of interest it is unavoidable to consider elliptic equations with coefficients that either degenerate or become singular. An example of this is the α-harmonic extension for the localization of fractional powers of uniformly elliptic operators. In this talk we will develop an approximation theory for weighted Sobolev spaces when the weight belongs to the Muckenhoupt class A_p, which is a class of weights that may degenerate or blow up. This will provide the tools necessary for an optimal a priori error analysis of discretizations of such equations. In addition, we will discuss multilevel methods for the solution of such problems and prove their nearly uniform convergence. Finally, we will show how to exploit the structure of the weight to develop a posteriori error estimators.

Abner J. Salgado
Department of Mathematics
University of Tennessee
asalgad1@utk.edu

MS48

Discontinuous Enrichment Method for Problems with Variable Coefficients

Abstract not available at time of publication.

Radek Tezaur, Charbel Farhat
Stanford University
rtezaur@stanford.edu, CFarhat@stanford.edu

Irina K. Tezaur
Sandia National Laboratories
ikalash@sandia.gov

MS49

Superconvergence of Discontinuous Galerkin Methods for Hyperbolic Equations in Two Space Dimensions

We study the superconvergence of discontinuous Galerkin methods for hyperbolic equations in two space dimensions. We first consider periodic boundary condition, and will prove that with piecewise k-th degree polynomial, the error between the numerical solution and the exact solution at the downwind-biased Radau points are k+2-th order accuracy. Moreover, at downwind point, the error is 2k+1-th order accurate. For Dirichlet boundary condition, we have to make a special projection of the exact solution at the inflow boundary to obtain similar results. Numerical experiments will be given to verify our numerical analysis.

Waixiang Cao
Beijing Computational Science Research Center
wxcao@csrc.ac.cn

Yang Yang
Michigan Technological University
yyang7@mtu.edu

Zhimin Zhang
Wayne State University
Department of Mathematics
ag7761@wayne.edu

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

MS49

Runge-Kutta Discontinuous Galerkin Methods for

the Relativistic Vlasov-Maxwell System

The relativistic Vlasov-Maxwell (RVM) system is a kinetic model of plasma to describe the phenomena when the charged particles move in the relativistiic regime, and the collisions of the particles are omitted. In this talk, we propose several discontinuous Galerkin (DG) methods with Runge-Kutta time discretizations to solve the RVM system. When the DG methods with the standard polynomial spaces are used, the mass conservation of the system is shown both theoretically and numerically. However, due to the special formulation of the total energy for the RVM system, the energy conservation cannot be obtained. Therefore, we proposed DG methods with non-polynomial spaces to preserve the total energy of the system. We discuss the conservation properties of both semi-discrete and fully discrete schemes, the error estimates, and validate the theoretical results numerically. Numerical experiments including streaming Weibel instability and wakefield acceleration are also presented.

He Yang
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
hyang2@stanford.edu

Fengyan Li
Rensselaer Polytechnic Institute
lif@rpi.edu

MS49

A Simple DG Scheme for Acoustic Wave Equations with Curved Interfaces and Boundaries

Consider solving acoustic wave equations with presence of curved boundary or interfaces. The conventional high-order discontinuous Galerkin scheme on straight-sided elements suffers from second order errors due to piece-wise segment approximation to the curve. We propose a simple flux correction to reduce the errors by projecting quadrature points for line integration onto curved interfaces, and evaluating numerical fluxes at projection points. For curved interfaces, numerical tests demonstrate that this simple modification may reduce interface error and non-physical diffractions. For curved boundary conditions, with the assumption that the exact solution can be smoothly extended, the local truncation error of the modified DG scheme is high order. Accuracy tests will be shown to demonstrate the effectiveness of this simple correction.

Xiangxiong Zhang
Massachusetts Institute of Technology
Department of Mathematics
zhan1966@purdue.edu

MS49

Runge-Kutta Discontinuous Galerkin Method with a Simple and Compact Hermit Weno Limiter

In this talk, we propose a new type of weighted essentially non-oscillatory (WENO) limiter, which belongs to the class of Hermite WENO (HWENO) limiters, for the Runge-Kutta discontinuous Galerkin (RKDG) methods solving hyperbolic conservation laws. This new HWENO limiter is a modification of the simple WENO limiter proposed recently by Zhong and Shu. Both limiters use information of the DG solutions only from the target cell and its immediate neighboring cells, thus maintaining the original compactness of the DG scheme. The goal of both limiters is to obtain high order accuracy and non-oscillatory properties simultaneously. The main novelty of the new HWENO limiter in this talk is to reconstruct the polynomial on the target cell in a least square fashion while the simple WENO limiter is to use the entire polynomial of the original DG solutions in the neighboring cells with an addition of a constant for conservation. This modification improves the robustness in the computation of problems with strong shocks or contact discontinuities, without changing the compact stencil of the DG scheme. Numerical results for both one and two dimensional equations including Euler equations of compressible gas dynamics are provided to illustrate the viability of this modified limiter.

Jun Zhu
Nanjing University of Aeronautics and Astronautics
zhujun@nuaa.edu.cn

Xinghui Zhong
Michigan State University
zhongxh@math.msu.edu

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

Jianxian Qiu
Xiamen University
jxqiu@xmu.edu.cn

MS50

Compact Scattered RBF-FD Stencils for PDEs on Surfaces

We present a novel high-order meshfree method for the solution of PDEs on smooth surfaces embedded in R^3, with applications in biology and chemistry. The approach is a combination of local Hermite RBF interpolation and iterated application of the projected gradient, and thus only requires a set of points and the corresponding normal vectors. A diagonally dominant sparse matrix is obtained for fourth order discretizations of the Laplace–Beltrami operator, ensuring computational efficiency.

Erik Lehto
Department of Mathematics
Royal Institute of Technology
elehto@kth.se

Varun Shankar
School of Computing
University of Utah
vshankar@math.utah.edu

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

MS50

Solving PDEs on the Sphere via Novel Galerkin Method using Highly Localized Kernel Bases

In this talk we will discuss a novel, kernel-based meshless Galerkin method for numerically solving partial differential equations on the sphere. In particular, we will apply this method to treat a general PDE describing stationary heat

conduction in an inhomogeneous, anisotropic medium on S^2. The Galerkin method used to do this employs spatially well-localized, "small footprint, " robust bases for the associated kernel space. The stiffness matrices arising in the problem have entries decaying exponentially fast away from the diagonal. Discretization is achieved by replacing the stiffness matrix with one whose entries are computed by a very efficient kernel quadrature formula for the sphere. The discretized stiffness matrix retains the exponential decay in its off diagonal entries. We will present error estimates for this problem and give a numerical example illustrating the results obtained.

Francis J. Narcowich
Department of Mathematics
Texas A&M University
fnarc@math.tamu.edu

Joseph Ward
Texas A&M University
Department of Mathematics
jward@math.tamu.edu

Stephen Rowe
Department of Mathematics
Texas A&M University
srowe@math.tamu.edu

MS50

A Least Squares-RBF Approach to Transport Problems on Surfaces

In this this talk, we aim to solve hyperbolic PDEs such as the wave equation on a given manifold S, using interpolation with Radial Basis Functions (RBFs). Our approach is to use the process of approximating the differential operator suggested and outlined originally by Wright. We seek to address challenges involved in solving hyperbolic problems on surfaces. One primary challenge is the need for added hyper-viscosity in solving hyperbolic problems. This talk discusses how hyper-viscosity can be implemented and the possibility of using an RBF-Least Squares method for avoiding hyper-viscosity altogether.

Daryl J. Springer
Department of Mathematics
Arizona State University
Daryl.Springer@asu.edu

MS50

Quadrature on Spheres and Other Manifolds Based on Kernels

Quadrature formulas for spheres, the rotation group, and other compact, homogeneous manifolds are important in a number of applications. The purpose of this talk is to present coordinate independent quadrature formulas associated with certain classes of positive definite and conditionally positive definite kernels that are invariant under the group action of the homogeneous manifold. In particular these formulas are accurate-optimally so in many cases-and stable under anincreasing number of nodes and in the presence of noise, provided the set X of quadrature nodes is quasi-uniform.

Joseph Ward
Texas A&M University
Department of Mathematics
jward@math.tamu.edu

MS51

High-Order Finite Element Methods for Cardiac Electrophysiology

Finite element and finite difference methods are widely used for solving problems in computational cardiac electrophysiology. However, their computational cost prohibits their direct use in online clinical practice during ablation therapy. In this talk we illustrate how high-order finite element methods and surface representations of the left atrial chamber can be used to achieve greater accuracy, reduced computation times and increased scalability to bring modelling closer to achieving clinical utility.

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

Sergey B. Yakovlev
University of Utah
School of Computing
yakovs@sci.utah.edu

Rheeda Ali
Department of Bioengineering
Imperial College London
rheeda.ali07@imperial.ac.uk

Nicholas Peters
National Heart and Lung Institute
Imperial College London
n.peters@imperial.ac.uk

Mike Kirby
University of Utah
School of Computing
kirby@cs.utah.edu

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

MS51

Three-Dimensional Modeling of Ca2+ Signaling in Healthy and Failing Cardiomyocytes

Calcium (Ca2+) is an important ion that drives contraction of muscle cells. In cardiac cells, Ca2+ dependent steps of the contractile cycle involve influx of the ion into the cytosol from the extracellular membrane (sarcolemma) and sarcoplasmic reticulum (SR), diffusion to the myofibril where it activates force generation, and uptake at the sarcolemma and SR. The presence of intracellular structures including the myofibrils, organelles (SR, mitochondria), and myriad protein crowders restrict the volume, through which Ca2+ diffuses and thus reduces the ions apparent diffusion rate. Theories that account for the arrangement and excluded volume of intracellular diffusion obstacles provide accurate estimates of the apparent diffusion rate. We extend one of these theories, homogenization, to explore the impact of diffusional obstacles and nonuniformly distributed Ca2+ fluxes or buffers on Ca2+ signaling using realistic cell geometries of from healthy and failing hearts.

Peter Kekenes-Huskey

University of Kentucky
pkekeneshuskey@ucsd.edu

MS51

Assessing the Credibility of Computational Models of Cardiac Electrophysiology

There is a much interest in clinical applications of computational models of cardiac electrophysiology, which will require the reliability of model predictions to be thoroughly investigated. Rigorous methods for assessing credibility of computational models have been developed within engineering and physical sciences ('verification, validation and uncertainty quantification'). However, applying such techniques to evaluate highly complex physiological models such as cardiac electrophysiology models is extremely challenging. In this talk we will discuss work bridging this gap.

Pras Pathmanathan
Department of Computer Science
University of Oxford
pras.pathmanathan@fda.hhs.gov

Richard Gray
U.S. Food and Drug Administration
richard.gray@fda.hhs.gov

MS51

Multi-Scale Modeling in Cardiac Electrophysiology: What Are the Challenges in Front of Us?

Mathematical models of cardiac electrophysiology across many scales, from single ion channels to the whole organ, have been developed. However, a multi-scale modeling framework, which is key for linking random molecular events to organ functions and diseases, has not been developed. This talk will focus on what progress we have made, what are the hurdles in front of us, and discuss how mathematics may help to overcome these challenges.

Zhilin Qu
University of California
Dept of Medicine
zqu@mednet.ucla.edu

MS52

Conservative and Accurate Geometric Transport Methods for Discontinuous Variables in Turbulent Multi-physics Two-phase Flows

Simulating multiphase flows presents significant challenges: flow variables exhibit discontinuities across the phase interface, complex microscale dynamics arise due to surface tension, and the interface develops highly complex corrugations. Fluid turbulence and multi-physics processes such as electro-hydrodynamics exacerbate these difficulties. We will discuss recently advanced geometric techniques capable of addressing these challenges conservatively and with second order accuracy. Applications ranging from interface-turbulence interaction in homogeneous flows to electro-hydrodynamic fuel atomization will be presented.

Olivier Desjardins
Cornell University
Department of Mechanical and Aerospace Engineering
olivier.desjardins@cornell.edu

MS52

Modeling and Simulation of Multimaterial Compressible Flows

Abstract not available at time of publication.

Marianne M. Francois
Los Alamos National Laboratory
mmfran@lanl.gov

MS52

Methods for Computing Turbulent Phase Interface Dynamics Across Multiple Scales

Dynamics of turbulent interfaces, e.g. the primary atomization of liquid jets, can span several orders of magnitude in scales ranging from centimeter scales of the injector down to the sub-micron scale of the smallest drops. Strategies and examples to efficiently couple these scales and enable predictive simulations will be discussed, including coupling Eulerian interface capturing methods to Lagrangian drop models and dual-scale approaches to enable LES-type simulations of the phase interface dynamics in turbulent flows.

Marcus Herrmann
School of Mechanical, Aerospace, Chemical and Materials Engi
Arizona State University
marcus.herrmann@asu.edu

MS52

Direct Numerical Simulations of Multiphase Flow: Now What?

After briefly reviewing the history and current status of direct numerical simulations of multiphase flows, two future challenges for such simulations are discussed. The first is how to use the vast quantity of data now available to improve modeling of the large scale or average flow and the other is how to efficiently include isolated processes, often due to additional physics, taking place on length and time scales much smaller than the dominant flow scale.

Gretar Tryggvason
University of Notre Dame
Department of Aerospace and Mechanical Engineering
gretar.tryggvason.1@nd.edu

MS53

Discovering Underlying Nonlinear Dynamics of Complex Systems from Data

Abstract not available at time of publication.

Steven Brunton
University of Washington
sbrunton@uw.edu

MS53

Common Manifold Learning Using Alternating Diffusion for Multimodal Signal Processing

One of the true challenges in signal processing is to distinguish between different sources of variability. In this

work, we consider the case of multiple multimodal sensors measuring the same physical phenomenon, such that the properties of the physical phenomenon are manifested as a hidden common source of variability (which we would like to extract), while each sensor has its own sensor-specific effects. We will address the problem from a manifold learning standpoint and show a method based on alternating products of diffusion operators and local kernels, which extracts the common source of variability from multimodal recordings. The generality of the addressed problem sets the stage for the application of the developed method to many real signal processing problems, where different types of devices are typically used to measure the same activity. In particular, we will show application to sleep stage assessment. We demonstrate that through alternating-diffusion, the sleep information hidden inside multimodal respiratory signals can be better captured compared to single-modal methods.

Ronald Coifman
Yale University
Department of Computer Science
coifman@math.yale.edu

Roy Lederman
Yale University
roy.lederman@yale.edu

Ronen Talmon
Technion
ronenta2@gmail.com

MS53

Title Not Available at Time of Publication

Abstract not available at time of publication.

Surya Ganguli
Stanford University
sganguli@stanford.edu

MS53

Data-Driven Model Reduction to Support Decision Making in Complex Systems

The next generation of complex engineered systems will be endowed with sensors and computing capabilities that enable new modes of decision-making. For example, new sensing capabilities on aircraft will be exploited to assimilate data on system state, make inferences about system health, and issue predictions on future vehicle behavior—with quantified uncertainties—to support critical operational decisions. Model reduction is one way to achieve this challenging task, in particular through data-driven reduced models that exploit the synergies of physics-based computational modeling and physical data.

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@MIT.EDU

MS54

PhySense: Social and Organizational Network Activity Simulation for Signature Generation and Extraction

We present *PhySense* an agent-based temporal event simulator that addresses signal propagation in online social and organizational networks. Probabilistic linear and nonlinear models for node activity and edge interactions characterize network activity. *PhySense* includes algorithms to post-process the resulting temporal data to generate metrics such as node-to-node influence, activity-based centrality and transfer entropy. The framework serves as a virtual laboratory for generation and extraction of signatures to understand phenomena in a class of social networks.

Vikram Jandhyala
University of Washington
vj@u.washington.edu

Arun Sathanur
Dept. of Electrical Engineering
University of Washington
arunsv@uw.edu

MS54

Computational Analysis of Ensemble Neural Data Recorded From An Insect Brain

Methods to simultaneously monitor neural activity across an ensemble of neurons with fine temporal precisions and across multiple trials/conditions have become standard practices in systems neuroscience. However, analysis of such datasets to reveal the underlying neural representation on a trial-by-trial basis has been a challenge. Since, these neural datasets are high dimensional: neurons vs. time vs. trials; we present methods that take advantage of this data-structure to extract meaningful response features in these neural datasets.

Debajit Saha, Chao Li
Washington University in St. Louis
sahad@seas.wustl.edu, li.chao@wustl.edu

Barani Raman
Biomedical Engineering
Washington University in St. Louis
barani@wustl.wdu

MS54

Computational Tools and Methods for Signature Discovery (Session Overview)

Scientists engaged in signature discovery have developed numerous analytic algorithms for data processing, feature extraction, machine learning, and classification. These algorithms are written in a wide variety of programming languages and utilize data stored in many types of databases. We will discuss analytic tools we developed for signature discovery and a supporting framework that allows researchers to easily use (and reuse) analytic codes, regardless of the language in which they were originally written.

Landon H. Sego
Pacific Northwest National Laboratory
landon.sego@pnnl.gov

MS54

Statistics, Learning, and Optimization for Data Analysis and Visualization

A variety of technologies in diverse areas such as medical imaging, industrial inspection, and oil and gas share common underlying challenges. Many of these problems lend themselves to statistical methods that entail estima-

tion, learning, or regression. We motivate these ideas from some traditional problems in image processing and develop the concepts of nonparametric modeling with applications to data analysis and visualization more generally. Applications from simulations, demographics, and industrial processes will be discussed.

Ross Whitaker
School Of Computing, Scientific Computing and Imaging Inst.
University of Utah
whitaker@cs.utah.edu

MS55

Real-Time Data-to-Decision Using Adaptive Surrogate Modeling Strategies

This talk presents an approach to achieve a dynamic data-to-decision process in real-time. We use a non-intrusive adaptive strategy that combines reduced-order modeling and localization techniques. The methodology is demonstrated for a self-aware autonomous vehicle that relies on the ability to dynamically process sensor data, assess system state and capabilities, and make decisions accordingly. We consider the specific case of real-time structural assessment in the presence of damage.

Laura Mainini
Politecnico di Torino
lmainini@mit.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@MIT.EDU

MS55

An Occam's Razor Strategy for Field Estimation from Wall-Mounted Sensors

We present a field estimation technique suitable for experimental closed-loop control. The method relies on an offline/online strategy where an approximation basis is learnt using information provided by very few, wall-mounted, sensors. Offline derivation of the basis uses sparsity promotion techniques and an informative sequence while online estimation is achieved by sparse recovery from the sensors information only. The method is illustrated with the flow around a cylinder and its performance is compared with POD.

Kevin Kasper
SATIE - ENS Cachan, France
kevin.kasper@ens-cachan.fr

Lionel Mathelin
LIMSI - CNRS
mathelin@limsi.fr

Mohamed Abbas-Turki, Hisham Abou-Kandil
SATIE - ENS Cachan, France
mohamed.abbas-turki@satie.ens-cachan.fr, hisham.abou-kandil@satie.ens-cachan.fr

MS55

Thermal Reduced Order Model Adaptation to Aero-Thermo-Structural Interactions

Predicting the behavior of hypersonic vehicles components requires fully coupled structural/aerodynamic/thermal analyses. This effort being computationally very demanding on full order models, reduced order modeling techniques have been developed for this multidisciplinary problem. Central to these approaches is the selection of appropriate bases to represent the solutions. This talk focuses on the determination both a-priori and during computations of bases to represent the temperature field which are representative of the multidisciplinary interactions.

Andrew Matney
Arizona State University
andrew.matney@asu.edu

Marc P. Mignolet
Arizona State
marc.mignolet@asu.edu

MS55

Numerical Study of Local Reduced Basis with Adaptive Training for Incompressible Navier-Stokes Flows

In this work, we combine a set of reduced-order representation techniques for simulating incompressible Navier-Stokes flows over a range of physical parameters and illustrate that an adaptive time-integration scheme can efficiently accelerate the generation of snapshots as well as the simulations with the reduced-order representations. The number of training configurations selected in snapshots is adaptively increased until the norms of residuals are reduced below a user-specified tolerance. The reduced-order representation is further accelerated by combining the ideas of local bases and hyper-reduction of nonlinear terms. The accuracy and efficiency of the proposed method is illustrated on numerical examples with parameter sweeps.

Yuqi Wu
Department of Applied Mathematics, University of Washington
yuqiwu@uw.edu

Ulrich Hetmaniuk
University of Washington
Department of Applied Mathematics
hetmaniu@uw.edu

MS56

Hierarchical Tensor Approximation of Parameter-Dependent PDEs

Parametric PDEs appear in a large number of applications, as e.g. in uncertainty quantification or optimisation. Typically, the amount of data to approximate and represent the solution scales exponentially in the parameter dimension. Therefore, a crucial task is to develop special numerical techniques that rely on data-sparsity in order to cope even with high parameter dimensions. In this talk, we will discuss low-rank tensor techniques that allow to reduce the complexity to a linear dependence on the parameter dimension. In particular, our aim is to adaptively construct an approximation of the solution in the hierarchical tensor format from a relatively small set of data samples. Once this approximation from an offline computation is available, the evaluation of quantities of interest becomes a cheap online task. Moreover, the explicit tensor representation can be used to compute stochastic properties of the solution in a straightforward way. The potential of this approach is

illustrated by numerical examples.

Jonas Ballani
EPF Lausanne
jonas.ballani@epfl.ch

Lars Grasedyck
Max-Planck-Institute
Math in Sciences
lgr@igpm.rwth-aachen.de

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

MS56

Tensor Format Representations and Optimal Model Reduction for Uncertainty Quantification

In this talk I will give an introduction to the basic principles of tensor format representations with a special focus on applications from uncertainty quantification.

Mike Espig
Max Planck Institute for Mathematics in the Sciences
mike.espig@alopax.de

MS56

High-Dimensional Tensor Sampling

The hierarchical tensor format allows for the low-parametric representation of tensors even in high dimensions d. On the one hand, this format provides a robust framework for approximate arithmetic operations with tensors based on rank truncations, which can be exploited in iterative algorithms. On the other hand, it can be used for the direct approximation of high-dimensional data stemming, e.g., from the discretisation of multivariate functions. In this talk, we discuss several strategies for an adaptive approximation of tensors in the hierarchical format by black box sampling techniques.

Lars Grasedyck
Max-Planck-Institute
Math in Sciences
lgr@igpm.rwth-aachen.de

Jonas Ballani
EPF Lausanne
jonas.ballani@epfl.ch

Melanie Kluge
RWTH Aachen
kluge@igpm.rwth-aachen.de

MS56

Novel Tensor-Product Representations for Uncertaintiy Quantification

Hierarchical Tucker tensor format (HT - Hackbusch tensors) and Tensor Trains (TT- Tyrtyshnikov tensors, I.Oseledets) have been introduced recently for low rank tensor product approximation. Hierarchical tensor decompositions are based on sub space approximation by extending the Tucker decomposition into a multi-level framework. Therefore they inherit favorable properties of Tucker tensors, e.g they offer a stable and robust approximation, but still enabling low order scaling with respect to the dimensions. For many high dimensional problems, hard to be handled so far, this approach may offer a novel strategy to circumvent the curse of dimensionality. For uncertainty quantification we cast the original boundary value problem, with uncertain coefficients problem into a high dimensional parametric boundary value problem, discretized by Galerkin method. The high dimensional problem is cast into an optimization problems, constraint by the restriction to tensors of prescribed ranks r. This problem could be solved by optimization on manifolds, or more simply by alternating least squares. Since the norm of the underlying energy-space is a cross norm, preconditioning is required only for the spatial part and e.g. performed by standard multi grid approaches, e.g BPX. This leads to a modification of the orthogonality of the used component tensors. A important task is a posteriori error control, with respect to the spatial discretization and also w.r.t. the tensor product approximation.

Reinhold Schneider
Technische Universitat Berlin
schneidr@math.tu-berlin.de

MS57

Fluid-composit Structure Interaction

We focus of the interaction between a multi-layered, composite structure, and the flow of an incompressible, viscous fluid, giving rise to a fully coupled, nonlinear moving boundary, fluid-multi-structure interaction problem. Examples include arterial walls interacting with blood flow, and oil platforms interacting with water. We present a novel stable, modular, loosely coupled scheme for the numerical simulation of the coupled problem, and prove that the scheme converges to a weak solution of the nonlinear, fully coupled problem. Our numerical and analytical results reveal new physical regularizing mechanism: the presence of a thin fluid-structure interface with mass regularizes the evolution of the entire FSI solution.

Suncica Canic
Department of Mathematics
University of Houston
canic@math.uh.edu

Martina Bukac
University of Notre Dame
martina@nd.edu

Boris Muha
University of Zagreb, Croatia
boris@math.hr

MS57

Second Order Embedded Boundary Methods for Fluid-Structure Interaction

This talk presents a second order accurate embedded boundary method for fluid-structure interaction. First, a spatially second order accurate embedded boundary method for first order hyperbolic systems is presented. Second, existing time discretizations for fluid-structure interaction are reviewed, with an emphasis on partitioned integrators, and then an extension for embedded boundary methods is described. Finally, the framework is verified on a set of fluid structure interaction test problems.

Alex Main, Charbel Farhat

Stanford University
alexmain@stanford.edu, CFarhat@stanford.edu

MS57

A Tetrahedral Method for Transient Nonlinear Dynamics Computations in Solids, Fluids and Coupled Fluid Structure Problems

A new mixed tetrahedral finite element formulation is presented, aimed at transient dynamic computations of fluids, solids, and fluid/structure interaction. It utilizes very simple approximation spaces: Piece-wise linear continuous functions for displacements and pressures, and discontinuous constants for the deviatoric part of the stress tensor (if present, as in solids.) A variational multiscale stabilization eliminates possible pressure checkerboard instabilities. Extensive simulations of compressible fluids, nonlinear elastic/visco-elastic/elastic-plastic solids, and fluid/structure interaction will conclude the presentation.

Guglielmo Scovazzi, Xianyi Zeng
Duke University
guglielmo.scovazzi@duke.edu, xy.zeng@duke.edu

Brian Carnes, David Hensinger
Sandia National Laboratories
bcarnes@sandia.gov, dmhensi@sandia.gov

MS57

Fractional Modeling of Brain Aneurysms

The arterial wall is typically described using integer-order PDEs. Recently, 1D simulations indicate that fractional-order models is a powerful alternative because they are less sensitive to the parameter estimation. We develop numerical methods for fractional-order models, and for the first time employ them in 3D fluid-structure interaction aneurysm simulations. Comparison studies indicate that although the 3D models are more sensitive to the fractional order compared to 1D cases, they are insensitive to the relaxation parameters.

Yue Yu
Brown University
yuy214@lehigh.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS58

Exploiting Active Subspaces to Quantify Uncertainty in the Numerical Simulation of the HyShot II Scramjet

We present the computational analysis of the reactive flow in a hypersonic scramjet engine with emphasis on effects of uncertainties in the operating conditions. The active subspaces methodology is used to characterize the effects of the input uncertainty on the scramjet performance. In addition to discussing the computational cost benefits associated with this dimension reduction technique, we focus on interpretation of the active variable and how it compares to expertise/intuition and sensitivity analysis

Michael A. Emory
Stanford University
memory@stanford.edu

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

Johan Larsson
University of Maryland
jola@umd.edu

MS58

Influence of Surface and Subsurface Parameter Uncertainty and Sensitivity on the Latent Heat Flux Using An Integrated Hydrologic Model

Integrated hydrologic models simulate surface, subsurface and land-surface water and energy fluxes and require parameters to characterize land cover, describe hydraulic properties and specify initial and boundary conditions. Uncertainty will be propagated through ParFlow, a nonlinear integrated hydrologic model, using a Monte Carlo simulation approach. Results will then be used to understand parameter sensitivity through an active subspace technique; identification of the active subspace will help explain which parameters result in the greatest model variability.

Jennifer Jefferson
Colorado School of Mines
jejeffer@mymail.mines.edu

Reed M. Maxwell
Department of Geology and Geologic Engineering
Colorado School of Mines
rmaxwell@mines.edu

MS58

Dimension Reduction in MCMC using Active Subspaces

Markov Chain Monte Carlo is a useful tool for sampling from the posterior distribution in Bayesian inverse problems, though it often struggles to converge in high dimensions. We discover and exploit active subspaces—an emerging set of tools for dimension reduction—in the misfit term of the likelihood to focus sampling along important directions in the parameter space, which improves mixing and convergence. We develop the framework and demonstrate its capabilities on a 100-dimensional model inverse problem.

Carson Kent
Colorado School of Mines
ckent@mines.edu

Paul Constantine
Colorado School of Mines
Applied Mathematics and Statistics
pconstan@mines.edu

MS58

Discovering An Active Subspace in a Single-Diode Solar Cell Model

Single-diode models for solar cells contain several parameters, so sensitivity analyses and UQ can be computationally expensive. We employ active subspaces for the solar cell's

the maximum power functiona function of the single-diode model parametersto discover a one-dimensional subspace that enables us to reduce the dimension for UQ parameter studies.

Mark Campanelli
National Renewable Energy Laboratory
mark.campanelli@nrel.gov

Brian Zaharatos
Colorado School of Mines
bzaharat@mymail.mines.edu

MS59

Operator Weighted MCMC on Function Spaces

Many inference problems require exploring the posterior distribution of high-dimensional parameters, which in principle can be described as functions. We introduce a family of operator-weighted MCMC samplers that can adapt to the intrinsically low rank and locally complex structure of the posterior distribution while remaining well defined on function space. Posterior sampling in a nonlinear inverse problem and a conditioned diffusion process are used to demonstrate the efficiency of these dimension-independent operator-weighted samplers.

Tiangang Cui
Massachusetts Institute of Technology
tcui@mit.edu

Kody Law
SRI UQ Center, CEMSE, KAUST
kodylaw@gmail.com

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS59

Estimation of Parameters of Chaotic Dynamic Systems

To estimate parameters of chaotic dynamical systems a measure to quantify the likelihood function of chaotic variability (the 'distance' between different trajectories) is needed. We review problems encountered by previously used methods and propose a method related to fractal dimension concepts. The methodology is illustrated using classical chaotic examples, the Lorenz 63 and Lorenz 95 systems, as well as higher dimensional fluid systems.

Heikki Haario
Lappeenranta University of Technology
Department of Mathematics and Physics
heikki.haario@lut.fi

MS59

Mapped Stochastic Newton Sampling

Different sampling algorithms for inverse UQ problems with high-dimensional parameters are reviewed and compared. A method that combines ideas from stochastic Newton MCMC sampling and implicit sampling is proposed and its efficiency is studied numerically.

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

MS60

RBF Response Surfaces with Inequality Constraints

In this talk, we describe how to construct RBF interpolants that have given function values at some sample points and satisfy upper and lower bound constraints at other points. Our approach is based on a constrained quadratic minimization problem that leads to a unique, parsimonious interpolant in which RBF centers appear only as they are needed to enforce an equality or inequality constraint. We also show that this method always improves the native space error over an interpolant subject to the equality constraints alone.

David Bindel
Cornell University
bindel@cs.cornell.edu

MS60

Miso: Mixed-Integer Surrogate Optimization for Computationally Expensive Black-Box Problems

We present an algorithm framework that uses surrogate models for solving mixed-integer global optimization problems whose objective function evaluation requires a computationally expensive computer simulation, and thus the analytical description and the derivatives are not available (black-box). Using a range of benchmark problems and two applications arising from groundwater remediation, we show that a combination of stochastic and deterministic sampling techniques leads to significantly improved solutions as opposed to using only a single sampling strategy.

Juliane Mueller
Lawrence Berkeley National Lab
juliane.mueller2901@gmail.com

MS60

Parallel Surrogate Global Optimization with Pareto Centers for Single Objective Expensive Functions

In single objective optimization with surrogates there are typically two criteria when searching on the surrogate surface for the next x for expensive function evaluation, a) that $s(x)$, the surrogate surface, is small and b) evaluating $f(x)$ will improve the overall quality of the accuracy of the surrogate approximation. In this new algorithm, a multi-objective optimization method is used to select the next points for multiple parallel processors. Solution accuracy is shown to improve.

Christine A. Shoemaker
Cornell University
cas12@cornell.edu

Tipaluck Krityakierne
Center for Applied Mathematics, Cornell University
tk338@cornell.edu

Taimoor Akhtar
Cornell University
taimoor.akhtar@gmail.com

MS60

Efficient Multi-Start for Global Optimization in Accelerator Design

We present a scalable, asynchronous algorithm for locating several/many high-quality local minimizers of expensive computer simulations for which derivatives are unavailable. Our method adjusts the amount of exploration/refinement performed depending on the observed function values and user desires. We highlight how our method scales as more computational resources become available. We lastly employ our algorithm on particle accelerator design problems, performing our computational studies on some of the largest supercomputers in the US.

Jeffrey M. Larson, Stefan Wild
Argonne National Laboratory
jmlarson@anl.gov, wild@mcs.anl.gov

MS61

Symmetry-adapted No-core Shell Model for First Principle Lage Scale Computations of Atomic Nuclei

We introduce Symmetry-adapted No-core Shell Model (SA-NCSM) approach—an emerging tool for first principle studies of light and medium mass nuclei. We review the pillars of the SA-NCSM framework and discuss techniques and algorithms that facilitate its implementation in form of a highly scalable computer code for modern petascale systems. We also outline future research directions and developments.

Tomas Dytrych
Louisiana State University
tdytry1@lsu.edu

MS61

Add, Multiply, Divide and Conquer: On-the-fly Algorithms for Many-body Calculations

Quantum many-body systems can be cast as a large-dimension matrix eigenvalue problem. The dimensionality is not the only challenge; even with very sparse matrices, simply storing the non-zero matrix elements can require terabytes or even petabytes. I discuss how on-the-fly methods can dramatically reduce storage by factorizing the matrix exactly using additive and multiplicative quantum numbers. Any machine, whether a laptop or a supercomputer, can thus run much larger problems.

Calvin W. Johnson
San Diego State University
cjohnson@mail.sdsu.edu

MS61

Multi-Level LOBPCG Method in MFDn

Many Fermion Dynamics for nuclei (MFDn) is a state-of-the-art software package for nuclei configuration interaction calculation. The locally optimal block preconditioned conjugate gradient (LOBPCG) method is used in MFDn to compute a few lowest eigenpairs of the very large, sparse many-body nuclear Hamiltonian matrix. We present a multi-level strategy that exploits the structure of the Hamiltonian matrix and accelerate the convergence of the LOBPCG method. Some preliminary results of preconditioning are also presented.

Meiyue Shao
Lawrence Berkeley National Laboratory
myshao@lbl.gov

MS61

Derivative-free Optimization Techniques in ab initio Nuclear Structure Calculations

Many tasks in ab initio nuclear structure calculations may rely on search techniques, such as tuning of nucleon-nucleon and three-nucleon interactions to fit light nuclei, optimizing of basis functions, and reducing 3-body force effects by optimizing freedoms in chiral effective-field theory NN interactions. This field is still under-served, however, by the existing optimization software, mainly due to the complexity of the underlying function evaluations. In this talk, we will first show a flexible framework for interfacing the MFDn package for performing ab initio nuclear structure calculations with derivative-free optimization packages. We demonstrate a few examples of tuning the interaction parameters and compare results using POUNDeRs and QNSTOP optimization algorithms, and note on efficient function evaluations when several nuclei are fitted simultaneously.

Masha Sosonkina
Old Dominion University
masha@scl.ameslab.gov

MS62

Hdg Method for Linear Elasticity

This paper presents a new hybridizable discontinuous Galerkin (HDG) method for linear elasticity on general polyhedral meshes, based on a strong symmetric stress formulation. The key feature of this new HDG method is the use of a special form of the numerical trace of the stresses, which makes the error analysis different from the projection-based error analyzes used for most other HDG methods. For arbitrary polyhedral elements, we approximate the stress by using polynomials of degree k¿=1 and the displacement by using polynomials of degree k+1. In contrast, to approximate the numerical trace of the displacement on the faces, we use polynomials of degree k only. This allows for a very efficient implementation of the method, since the numerical trace of the displacement is the only globally-coupled unknown, but does not degrade the convergence properties of the method. Indeed, we prove optimal orders of convergence for both the stresses and displacements on the elements. These optimal results are possible thanks to a special superconvergence property of the numerical traces of the displacement, and thanks to the use of a crucial elementwise Korn's inequality.

Weifeng Qiu
City University of Hong Kong

qiuw78@gmail.com

MS62

HDG Methods for the p-Laplacian

We propose a hybridizable discontinuous Galerkin (HDG) method for the p-Laplacian equation. The numerical scheme inherits two distinctive features when the solution is sufficiently regular. First, when using polynomial of degree $k \geq 0$ for $_h$, $_h$ and u_h, our scheme exhibits optimal $k+1$ order of convergence for all of them in L^2, L^p and L^∞ norm. Second, when $k \geq 1$, by applying an element-by-element postprocessing technique, we can obtain a new approximation u_h^* that converges at order $k+2$ to u. In addition, we formulate two efficient nonlinear conjugate gradient algorithms to solve our HDG methods from the minimization standpoint. We carefully compare these two algorithms in terms of convergence speed, memory requirement and robustness. To improve the efficiency further, we construct a class of remarkable preconditioners based on the unifying hybridization framework. In addition, for the first time, we postprocess the numerical trace $\widehat{u}_h$ for constructing a quality initial guess as an implementation technique.

Bernardo Cockburn, Jiguang Shen
School of Mathematics
University of Minnesota
cockburn@math.umn.edu, shenx179@math.umn.edu

MS62

Parallel hp-multigrid for HDG

We present a parallel algorithm for hp-adaptive multigrid for the HDG method. We first coarsen in p, smoothing only on the skeletal system, until we obtain an HDG system with $p = 1$. We switch to a linear CG grid at this stage and continue to coarsen in h. This allows us to create a deep grid hierarchy and obtain optimal multigrid convergence. We compare the performance of different smoothers and with our recent work on high-order multigrid using a CG discretization. We also present early scalability results from a parallel implementation.

Tan Bui
University of Texas at Austin
tanbui@ices.utexas.edu

Hari Sundar
Institute for Computational and Engineering Sciences
University of Texas at Austin
hari@cs.utah.edu

MS63

Solving a Parameterized Eigenvalue Problem from Regularized Total Least Squares

The solution of an ill-conditioned total least squares (TLS) problem by a regularization approach of Golub et al. [SIAM J. Matrix Anal. Appl. 21(1):185-194,1999] is considered. That regularized TLS formulation can be viewed as a parameterized eigenvalue problem. The approach given here is a Newton method that iterates for the parameter and the associated parameterized eigenvalue. The two nonlinear equations are formulated from a standard regularization bound constraint and a spectral function from the parameterized eigenvalue problem. The Jacobian of the system can be computed inexpensively and can be proven to be nonsingular under certain reasonable assumptions. A basic Newton method is presented as well two approximate Newton methods based upon Krylov space and generalized Krylov space projections. The resulting algorithms are applied to a problem from high-resolution image reconstruction that is appropriately modeled by total least squares.

Jesse L. Barlow
Penn State University
Dept of Computer Science & Eng
barlow@cse.psu.edu

Geunseop Lee
The Pennsylvania State University
gxl174@cse.psu.edu

Haoying Fu
Penn State University
Computer Science & Engineering Department
haoyingfu@gmail.com

MS63

Performance Evaluation of EigenExa Dense Eigensolver on the Oakleaf-Fx Supercomputer System

We have developed a dense eigensolver named EigenExa, which employs the divide-and-conquer method for a banded matrix (currently penta-diagonal). To capture its current performance in detail, we conducted an evaluation by using the whole supercomputer system Oakleaf-FX, which consists of 4800 nodes and has the almost same architecture as the K computer. In this talk, we present the results of the evaluation, which shows advantages of our eigensolver and issues we face.

Takeshi Fukaya
RIKEN, Japan
Japan
takeshi.fukaya@riken.jp

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

MS63

Scaling Comparison of Dense Eigensolvers and Purification Techniques to Large Node Counts

This talk focuses on the case when large parallel resources are available for computing the spectral projector of a relatively small matrix. This is the case, for example, in quantum chemistry where large parallel resources are needed for computing the Fock matrix, and then all the resources are available to diagonalize the Fock matrix. We examine the scaling of communication and computation for various dense eigensolvers, and also for a Newton-Schulz approach for computing the spectral projector which is based on matrix-matrix multiplications.

Xing Liu, Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
xing.liu@gatech.edu, echow@cc.gatech.edu

MS63

Revisiting SVD(A) through EIG(T) for

Sca/LAPACK

It is well known that the SVD of a matrix A, $A = USV^T$, can be obtained from the eigenpairs of the matrix $C_V = A^TA$ or $C_U = AA^T$. Alternatively, the SVD can be obtained from the eigenpairs of the augmented matrix $C_{UV} = [0\ A, A^T\ 0]$. This presentation focuses on C_{UV} in the context of bidiagonal matrices and Sca/LAPACK's tridiagonal eigensolvers, and discusses accuracy and implementation issues.

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

MS64

Higher-Order Filtered Methods for Nonlinear Partial Differential Equations

The theory of viscosity solutions has been effective for representing and approximating weak solutions of nonlinear partial differential equations such as Hamilton-Jacobi and second-order elliptic equations. The approximation theory of Barles and Souganidis requires that numerical schemes be monotone. However, monotone schemes have limited accuracy. We introduce a framework for using monotone schemes as the foundation for filtered schemes, which are almost-monotone. This allows us to construct higher-order discretisations that provably converge to the viscosity solution of the underlying PDE.

Brittany Froese
University of Texas at Austin
bfroese@math.utexas.edu

Adam M. Oberman
McGill U.
adam.oberman@mcgill.ca

MS64

High Order Methods for Traveltime and Amplitude in Geometrical Optics

In geometrical optics approximation for the Helmholtz equation, accurate and efficient computation of traveltime and amplitude is important. We present an approach to resolving the source singularities and high order methods to compute the traveltime and amplitude efficiently. Numerical examples are presented to demonstrate the methods.

Songting Luo
Department of Mathematics,
Iowa State University
luos@iastate.edu

Jianliang Qian
Department of Mathematics
Michigan State University
qian@math.msu.edu

Robert Burridge
Department of Mathematics and Statistics
University of New Mexico
burridge137@gmail.com

MS64

High-Order Gas Evolution Model for Computational Fluid Dynamics

The foundation of compressible flow solver is based on the first order gas dynamic model, i.e., the Riemann solution of the Euler equations, where the spatial and temporal discretizations are decoupled. Here, we will present a high-order gas evolution model based on the Boltzmann equation, and develop the corresponding high-order schemes. In the high-order gas-kinetic schemes, the spatial and temporal discretization are fully coupled starting from a piecewise discontinuous high-order initial condition. The necessity to couple spatial and temporal evolution nonlinearly is important to construct high order compact schemes.

Kun Xu
University of Science & Technology
Hong Kong
makxu@ust.hk

MS64

An Efficient Spectral Method for the Euler-Lagrange Equations of Minimum Action Methods

Minimum Action Method is one of the popular approaches to study phase transitions of dynamical systems. According to the Freidlin-Wentzell theory, at zero temperature limit, transitions happen along the path which takes minimum action. Minimum Action Method converts a phase transition problem to an optimal control problem. Traditional method to solve the optimal control problem usually first discretize the problem, then use a general purpose optimization package to solve the discretized problem. In this talk, we will propose another approach, which first derive the Euler-Lagrange equation of the minimum action, then solve this high-order partial differential equation directly. By taking advantage of good properties of the Euler-Lagrange equation, we establish an efficient spectral solver. Numerical results for phase transitions of both ODEs and PDEs, gradient and non-gradient systems will be presented.

Haijun Yu
LSEC
Academy of Mathematics and Systems Science, China
hyu@lsec.cc.ac.cn

Xiaoliang Wan
Louisiana State University
Department of Mathematics
xlwan@math.lsu.edu

MS65

Title Not Available at Time of Publication

Abstract not available at time of publication.

Mary Galvin-Donoghue
National Science Foundation
Div. of Materials Research

mgalvind@nsf.gov

MS65

Materials from Mathematics

We present examples of new materials whose synthesis was guided by some essentially mathematical ideas. These materials undergo phase transformations from one crystal structure to another, without diffusion. The underlying mathematical theory was designed to identify alloys that show exceptional reversibility of the transformation. Some of these alloys convert heat to electricity (without a separate electrical generator), and provide an interesting possible route to recover the vast amounts of energy stored on earth at small temperature difference. The broader lessons from this research are: 1) mathematics can be used for unexpected discovery, but nonstandard approaches are helpful (what if?, rather than what? and why?), 2) an algorithmic approach (theorems $\Longrightarrow$ algorithms) seems to be useful, 3) a functioning feedback loop with experiment is desirable, 4) a useful development would be to integrate modern applied mathematics with the synthesis/characterization tools that play such an essential role in materials science.

Richard James
Department of Aerospace Engineering and Mechanics
University of Minnesota
james@umn.edu

MS65

Computational Materials Design: Challenges in Practical Applications

Abstract not available at time of publication.

Sadasivan Shankar
Harvard University
sshankar@seas.harvard.edu

MS65

Title Not Available at Time of Publication

Abstract not available at time of publication.

Michael S. Vogelius
Rutgers University, New Brunswick
vogelius@math.rutgers.edu

MS66

Penalty Methods for the Hyperbolic System Modeling the Wall-Plasma Interaction in a Tokamak

One main challenge for producing energy using magnetic confinement fusion in a tokamak is the control of wall-plasma interactions. Thus, numerical tools with efficient implementations of the boundary conditions are needed. This talk presents several penalty methods. A penalization method for a non linear hyperbolic problem is provided and analyzed theoretically and numerically: this method does not generate any boundary layer and the convergence rate, when the penalty parameter goes to 0, is optimal.

Philippe Angot, Thomas Auphan, Olivier Guès
Aix-Marseille Université
philippe.angot@univ-amu.fr, tauphan@cmi.univ-mrs.fr, olivier.gues@univ-amu.fr

MS66

A Dispersionless Fourier Method for the Maxwell Equations Using Volume Penalization

We develop numerical Fourier methods for solving the time dependent boundary value problem for Maxwell's equations in the vicinity of perfect electric conductors. The high order methods are obtained by analytically modifying the standard Maxwell equations with a volume penalization term. The absence of strong dispersive effects and the efficiency of the fast Fourier transform make this approach potentially well suited for high frequency applications. We demonstrate the approach with several simulations of scattering and wave guide problems (with the inclusion of perfectly matched layers where appropriate).

Ryan Galagusz
McGill University
Department of Electrical Engineering
ryan.galagusz@mail.mcgill.ca

MS66

Fourier based PDE Solution on Complex Domains

Fourier continuation/extension (FC) methods allow for highly-accurate Fourier representations of non-periodic functions. FFT speed algorithms, including the FC(Gram) algorithm, have been developed for these Fourier approximations and subsequently applied to the solution of partial differential equations. In particular, a new analysis approach has been utilized to ensure unconditional stability for some FC alternating direction solvers. A unique character of solutions generated from the Fourier approximations is the lack pollution error for wave propagation problems while avoiding many of the limitations of traditional spectral methods. Various methods for the solution of partial differential equations and other applications of FC algorithms will be discussed.

Mark Lyon
University of New Hampshire
mark.lyon@unh.edu

MS66

New Active Penalty Methods with Applications to Fluid Flow

In this talk I will present a new active penalty method. This method relies on a construction that systematically improves the amplitude of the boundary layer that results from the penalty term. As a result, convergence rate is improved. I will present a Fourier Spectral 4th order globally convergent version of this approach for the heat equation with internal non-mesh-conforming Dirichlet boundary conditions. I will then extend this to the solution of incompressible Navier-Stokes equations in the presence of solid non-mesh-conforming boundaries. I will finish by discussing convergence rate and stability of the method in general.

Jean-Christophe Nave
McGill University
jcnave@math.mcgill.ca

MS67

The Use of Residual-Based Compact Schemes for

Industrial Les

The use of suitable high accurate numerical techniques is of the utmost importance for the simulation of turbulent flows since they enable capturing flow structures from large to small scales at an acceptable computational cost. In this sense, compact schemes appears to be particularly attractive because of their spectral-like accuracy. Moreover, for industrial aerodynamics computations, robustness and shock capturing capabilities are also crucial issues. In this work, we discuss the capabilities of a family of high order residual-based compact (RBC) schemes, characterized by odd orders of accuracy and a genuinely multidimensional numerical dissipation, for the implicit large eddy simulation (ILES) of compressible turbulent flows in complex configurations.

Paola Cinnella
ENSAM, ParisTech
paola.cinnella@ensam.eu

Cédric Content
Laboratoire DynFluid
Arts et Métiers ParisTech
cedric.content@ensam.eu

Luca Sciacovelli
Laboratoire DynFluid, Arts et Métiers Paristech
and Politecnico di Bari
luca.sciacovelli@ensam.eu

MS67

The Application of High Order Dgm for Resolved and Wall-Modeled Les of Full Scale Turbomachinery Passages

Large Eddy Simulations of turbulent flows require highly accurate discretisations in order to minimise the impact of discretisation errors for inevitably underresolved computations. Although extremely accurate methods are used in academia, industry relies low order finite volume methods, mainly due to their geometric flexibility. The discontinuous Galerkin Method promises to be an enabler of high-resolution LES in an industrial context, as it provides high-accuracy on unstructured meshes and therefore complex geometry, but also provides excellent parallel scaling. It furthermore offers the perspective of adaptive computations. The current contribution describes the development of an industrial CFD code for LES, including validation, optimisation and extreme scalability, and concludes with practical examples.

Koen Hillewaert
Cenaero
Cenaero
koen.hillewaert@cenaero.be

Corentin Carton de Wiart
Cenaero
corentin.carton@cenaero.be

Guillaume Verheylewegen
Université catholique de Louvain
guillaume.verheylewegen@uclouvain.be

Ariane Frère
Cenaero
ariane.frere@cenaero.be

MS67

Validation of a High-Order Implicit Les Solver for the Simulation of a Low-Reynolds-Number Vertical-Axis Wind Turbine

We use a high-order DG scheme for the ILES simulation of a straight-bladed Vertical Axis Wind Turbine (VAWT). We validate the solver by simulating the flow field about a single static NACA0012 airfoil over a range of angles of attack. Using an ALE-scheme in both 2D and 3D, we simulate a real-world VAWT configuration at low chord Reynolds number (about 50,000) for which experimental data is available, and determine the accuracy for various flow conditions.

Samuel Kanner
University of California, Berkeley
kanners@berkeley.edu

Per-Olof Persson
University of California Berkeley
Dept. of Mathematics
persson@berkeley.edu

MS67

Applications of the Spectral/hp Element Method to Complex Flow Geometries

As industrial requirements evolve to require both transient and scale-resolving capabilities, the use of under-resolved DNS and implicit LES methods is needed to capture the essential flow features. In this talk, we give some examples of industrially-relevant simulations, and highlight some of the challenges that arise when performing simulations in complex three-dimensional geometries. Furthermore we demonstrate how the spectral/hp element method combined with appropriate stabilisation and discretisations can resolve flow features in these complex domains.

David Moxey
Imperial College London
d.moxey@imperial.ac.uk

Joaquim Peiro
Dept of Aeronautics
Imperial College London, UK
j.peiro@imperial.ac.uk

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

MS68

Unit and Conquer Algorithms for Large Eigenvalue Problems

The emerging new multi-level and heterogeneous computer architectures, gearing up the road to exascale computing, are requiring numerical algorithmic revisions and their corresponding implementations to achieve performance scalability. We present unite and conquer approach as a solution to achieve this goal. That consists to accelerate the rate of convergence of a restarted method by coupling either synchronously or asynchronously, several restarted methods called also co-methods. Some experiments validating

the approach will be presented.

Nahid Emad
PRiSM laboratory and Maison de la Simulation,
University of
Nahid.Emad@prism.uvsq.fr

MS68

Performance of Algebraic Multigrid Preconditioners for Large-Scale Finite Element Simulations

Finite element method (FEM) approaches are used for the large-scale, high fidelity simulation of many important physical phenomena, e.g. fluid flow or magnetohydrodynamics (MHD). Our solution approach for the resistive magnetohydrodynamics (MHD) equations employs an FEM discretization with an algebraic multigrid preconditioned Newton-Krylov method on unstructured meshes. We present scaling results for resistive MHD test cases, including results on 500,000 cores on an IBM Blue Gene/Q platform.

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

MS68

Intelligent Iterative Methods : the Future of Parallel and Distributed Runtime Tuned Linear Algebra?

We discuss some recent comparisons on clusters of accelerators between orthogonal, incompletely orthogonal and non-orthogonal Krylov Basis computing, focusing on communications and orthogonality accuracy. Then, we discuss some results for the ERAM method with respect to the restarting strategies. We survey some smart-tunning strategies we proposed and evaluated for some of the Krylov method parameters. As a conclusion, we propose auto-tuning strategies for future hybrid methods on future exascale hypercomputers, on the road to intelligent linear algebra methods mixing both distributed and parallel programming models.

Serge G. Petiton
University Lille 1, Science and Technologies - CNRS
serge.petiton@lifl.fr

MS68

Divide and Conquer Algorithms for Large Hermitian Eigenvalue Problems

A number of eigenvalue problems are considered from the point of view of highly parallel distributed environments. We first discuss a polynomial filtering technique for extracting extreme or interior eigenvalues of large sparse matrices. This general approach can be effective in the situation when a large number of eigenvalues is sought, as is the case in electronic structure calculations for example. The method presented relies on a combination of the Lanczos algorithm with partial reorthogonalization and polynomial filtering based on least-squares polynomials. We also discuss the problem of updating eigenvalues and eigenvectors of large dense matrices when a small rank perturbation is applied. Here several different practical scenarios are considered: partial or full spectrum to be computed, interior or extreme eigenvalues, etc.

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

Vasilis Kalantzis
CEID, School of Engineering
University of Patras, Greece
kalantzis@ceid.upatras.gr

MS69

Augmenting the One-Shot Framework by Additional Constraints

The multi-step Oneshot method for design optimization problems has been successfully implemented for various applications. To this end, a slowly convergent primal fixed point iteration of the state equation is augmented by an adjoint iteration and a corresponding preconditioned design update. Within this talk we present a modification of the method that allows for additional equality constraints besides the usual state equation. A retardation analysis and the local convergence of the method in terms of necessary and sufficient conditions are given, which depend on key characteristics of the underlying problem and the quality of the utilized preconditioners.

Torsten F. Bosse
Humboldt-Universität zu Berlin,
Institut für Mathematik
mail@torsten-bosse.de

Andreas Griewank
HU Berlin, Germany
griewank@math.hu-berlin.de

MS69

Fixed-Point Iterations for Simultaneous One-Shot Optimization of Unsteady Flows

The One-shot method has proven to be very efficient in optimization with steady PDEs which are solved by fixed-point iterations. We provide a framework that extends the method to unsteady problems that are solved by classical time-marching schemes. The One-shot method is applied to an optimal control problem with unsteady incompressible Navier-Stokes equations. The unsteady fixed-point iteration is further improved applying adaptive time scales. Opportunities and first results on integrating One-shot optimization into parallel-in-time simulation will be presented.

Stefanie Günther, Nicolas R. Gauger
RWTH Aachen University
stefanie.guenther@scicomp.uni-kl.de,
nicolas.gauger@rhrk.uni-kl.de

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

MS69

Towards Second Order One-Shot Methods in the

Context of Shape Calculus

One-shot methods aim at efficient implementation of gradient based optimization methods. Highest efficiency is reached, if second order methods can be employed. This poses a certain challenge for shape optimization methods, which employ shape calculus for efficiency reasons. The challenge lies in the fact that the space of shapes is not a linear space. This talk proposes the usage of Riemannian shape manifolds in order to carry over second order properties to shape optimization problems. Numerical results for PDE constrained shape optimization problems are provided.

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

MS69

On An Extension of the Augmented Lagrangian Approach for One-Shot Optimization

One-shot approaches for design optimization augment the solution of the state equation with an adjoint solver yielding approximate derivatives to change the design. The coordination of these iterative processes is well established when only the state equation serves as equality constraint. We propose a modified augmented Lagrangian function for the handling of additional equality constraints. We show that this augmented Lagrangian can be used in gradient-based optimization to solve the original design task.

Andrea Walther
Universität Paderborn
andrea.walther@uni-paderborn.de

Nicolas R. Gauger
RWTH Aachen University
nicolas.gauger@rhrk.uni-kl.de

MS70

Energy-Conserving Schemes for Vlasov-Type Systems

In this talk, we present the discontinuous Galerkin (DG) methods to solve the Vlasov-Maxwell system. The scheme employs DG discretizations for both the Vlasov and the Maxwells equations, resulting in a consistent description of the probability density function and electromagnetic fields. We prove that using this description the total particle numbers are conserved, and the total energy could be preserved upon a suitable choice of numerical flux for the Maxwells equations and the underlying polynomial spaces on the semi-discrete level, if boundary effects can be neglected. We further established error estimates based on several flux choices. We test the scheme on the Weibel instability and verify the order and conservation of the method.

Yingda Cheng
Department of Mathematics
Michigan State University
ycheng@math.msu.edu

MS70

High Order Asymptotic Preserving Projective Integration Methods

In order to introduce new asymptotic preserving schemes for kinetic equations in regimes leading to hyperbolic or diffusive systems of conservation laws appearing e. g. in some models of radiative transfer or fluid-particle interactions, we apply the projective integration method developed by Gear and Kevrekidis in the context of large multiscale differential systems appearing in Chemistry. The crucial point of this work is our obtaining high order in time asymptotic preserving schemes.

Pauline Lafitte
Project team SIMPAF, INRIA Lille
pauline.lafitte@ecp.fr

Annelies Lejon, Ward Melis
KU Leuven
annelies.lejon@cs.kuleuven.be, ward.melis@cs.kuleuven.be

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

MS70

Effective High-Order Diffusive Moment Closures with the StaRMAP Software

StaRMAP is a software to efficiently compute linear moment models of radiative transfer (e.g., PN, SPN, FPN) using staggered grid finite difference approximations. We present the extension of the method and software to diffusive closures that possess a "viscosity" in the highest resolved moment. We show how the new term can be incorporated into the numerical approach without incurring a reduction in convergence order or a requirement for undesirably small time steps. Then, using the StaRMAP code, the benefits and drawbacks of various diffusive closures are discussed. The StaRMAP code is available for download, and the results are easily reproducible.

Benjamin Seibold
Temple University
seibold@temple.edu

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

MS70

An Asymptotic-preserving Scheme for Linear Kinetic Equation with Fractional Diffusion Limit

We present an asymptotic-preserving scheme for the linear Boltzmann equation with fractional diffusion limit. The equilibrium is a fat tail function, which disables any truncation in the velocity space numerically. The stiffness adds more numerical difficulty. Our idea is based on a macro/micro/tail decomposition, where the macro/micro components support on a compact velocity space and decompose the equation following a reshuffled Hilbert expansion. The tail that collects all the rest information solves a limit equation.

Li Wang
UCLA Mathematics
liwang@math.ucla.edu

MS71

Multilevel Methods for Forward and Inverse Ice

Sheet Modeling

We present our work on recovering ice sheet modeling parameter fields (such as basal slipperiness) from observations, using both deterministic and Bayesian inversion. The scalability of our adjoint- and Hessian-based methods is determined by the scalability of two sub-problems: the solution of the state PDEs (Stokes equations of ice sheet dynamics), and the approximation and preconditioning of the parameter-to-observation Hessian. For the former problem, we compare the effectiveness of geometric and algebraic multigrid within the solution of the state PDEs; for the latter, we discuss the use of multilevel approximations to improve on Hessian approximation by low-rank updates. The scalability of our work is tested on full-scale models of the Antarctic ice sheet.

Toby Isaac
ICES
The University of Texas at Austin
tisaac@ices.utexas.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS71

Assessment of Finite Element Schemes for Accurate Modeling of the Grounding Line

Modeling grounding line dynamics is critical in order to improve the projections of the contribution of the ice sheets to sea level rise. The most advanced way of modeling the grounding line consists of formulating a contact criterion, where the normal stress at the base of the ice is compared to the ocean pressure in the vicinity of the grounding line. The grounding line is also the location of a sharp change in boundary condition, which generates a singularity in the pressure field. The contact condition is therefore applied at a location where the basal stress is most challenging to capture accurately. Here, we present a new approach for the validation of Stokes Free-Surface flow with friction boundary conditions based on analytical-numerical solutions. We then compare the performance of several full-Stokes finite element solvers and show that the solvers currently used in the glaciological community are not capable of capturing the pressure field when there is a non-penetration condition. These results show that the grounding line cannot be accurately modeled with the current solvers. Finally, we propose an algorithm that reaches a precision at the percent level in the pressure field for a typical ice flow regime. This work has vast implications for the modeling of grounding lines and the ice sheets in a warming climate.

Mathieu Morlighem
University of California - Irvine
mathieu.morlighem@uci.edu

Jerome Monnier
Mathematics Institute of Toulouse
jerome.monnier@insa-toulouse.fr

Helene Seroussi
Jet Propulsion Laboratory
helene.seroussi@jpl.nasa.gov

Nathan Martin
IMT CNRS/INSA/UPS, Toulouse, France.
nathan.mar7in@gmail.com

MS71

Advances on Ice-Sheet Model Initialization using the First Order Model

Ice sheet initialization is critical for performing reliable forward simulations. We propose an adjoint-based optimization algorithm for the ice-sheet initialization, where we invert for basal topography and basal friction fields simultaneously, minimizing the mismatch between: 1. observed and computed surface velocity data, and 2. surface mass balance forcing and modeled flux divergence. We show the effectiveness of the initialization on the Greenland ice sheet. We mention ongoing work on quantifying uncertainties in the optimized parameters.

Mauro Perego
CSRI Sandia National Laboratories
mperego@fsu.edu

Stephen Price
Los Alamos National Laboratory
sprice@lanl.gov

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

Charles Jackson
UTIG
University of Texas at Austin
charles@ig.utexas.edu

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Irina K. Tezaur
Sandia National Laboratories
ikalash@sandia.gov

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

MS71

Uncertainty Quantification for Large-Scale Bayesian Inverse Problems with Application to Ice Sheet Models

We consider the estimation of the uncertainty in the solution of large-scale ice sheet inverse problems within the framework of Bayesian inference. The ice flow is modeled as a three-dimensional, creeping, viscous, incompressible, non-Newtonian fluid via the nonlinear Stokes equations.

The observational data come from InSAR satellite measurements of surface ice flow velocity, and the uncertain parameter field to be inferred is the basal friction parameter. We show that the work required for applying our framework–measured in number of forward (and adjoint) ice sheet model solves–is independent of the state and parameter dimensions.

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Toby Isaac
ICES
The University of Texas at Austin
tisaac@ices.utexas.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS72

The most current list of participating companies is available at www.siam.org/meetings/cse15/career.php.

For the most recent list of participating companies visit http://www.siam.org/meetings/cse15/career.php

Bill Kolata
SIAM
kolata@siam.org

MS73

Fast Direct Solver based on the Cyclic Reduction Algorithm and Hierarchical Matrix Arithmetic for the Solution of Variable-coefficient Elliptic PDEs

We present a fast direct solver for variable-coefficient elliptic partial differential equations on a Cartesian product mesh. We combine the Cyclic Reduction algorithm with Hierarchical Matrix arithmetic to reduce the fill-in blocks, resulting in an $O(Nk^2 \log^2 N)$ flop quasi-linear complexity solver. We compare against three techniques exploiting hierarchical low-rankness in the context of matrix bisection, nested dissection, and multifrontal ordering. Comparisons highlight algorithmic differences and implementation details with particular consideration to memory locality.

Gustavo Chavez
KAUST
gustavo.chavezchavez@kaust.edu.sa

George M. Turkiyyah
American University of Beirut
gt02@aub.edu.lb

Rio Yokota
King Abdullah University of Science and Technology
rio.yokota@kaust.edu.sa

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

MS73

A Polarized-Trace Preconditioner for 2D Helmholtz and Frequency Domain Full-Waveform Inversion

Full-waveform inversion, a method for recovering Earth's physical parameters by matching seismic observations with simulations, can be treated in the frequency domain. A scalable solver for Helmholtz's equation is necessary to make feasible imaging in high resolution and in 3D; however, this remains an open problem. We present recent developments on a domain-decomposition solver for the acoustic Helmholtz equation, based on the notion of polarization, that demonstrates substantial scalability and performance improvements over the state-of-the-art.

Russell Hewett
Total
russell.hewett@total.com

Leonardo Zepeda-Nunez
MIT
lzepeda@math.mit.edu

Laurent Demanet
Professor of Mathematics, MIT
laurent@math.mit.edu

MS73

Fast Multipole Method as a Preconditioner for Finite Discretizations of Elliptic Boundary Value Problems

We employ the Fast Multipole Method to precondition sparse iterative solvers. FMM solves certain elliptic PDEs with O(N) complexity via kernels of high thread uniformity, high arithmetic intensity, and relaxed synchronization compared with factorization-based or multilevel approaches. Combined, these features make FMM an interesting preconditioner on future architectures, even on problems where FMM is not an "exact" solver. We compare FMM-preconditioned Krylov solvers with multilevel and sparse direct solvers on a variety of elliptic problems.

Huda Ibeid, Rio Yokota
King Abdullah University of Science and Technology
huda.ibeid@kaust.edu.sa, rio.yokota@kaust.edu.sa

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

MS73

Generalized Plane Waves Adapted to Varying Coefficients

This talk will focus on new shape functions adapted to the scalar wave equation with smooth coefficients. It follows the idea of approximating the solution by basis functions that have the appropriate oscillatory behavior. The local design procedure of these generalized plane waves is developped to fit the varying coefficients. High order approximation is achieved, provided that a sufficient number of basis functions is used. Both theoretical and numerical

aspects will be investigated.

Lise-Marie Imbert-Gérard
Courant Institute of Mathematical sciences, NYU
imbertgerard@cims.nyu.edu

MS74

Accelerating MCMC with Parallel Local Approximations

In many inference problems, the cost of MCMC analysis is dominated by repeated evaluations of expensive forward models. We have previously shown that when the model is well behaved, locally-constructed surrogates can significantly reduce the number of evaluations required by MCMC while preserving the asymptotic exactness of the resulting inference. We extend those results with parallel computing resources and adjoint techniques, demonstrating strong improvements in run-time on challenging example inference problems.

Patrick R. Conrad, Youssef M. Marzouk
Massachusetts Institute of Technology
p.conrad@warwick.ac.uk, ymarz@mit.edu

Natesh Pillai
Statistics
Harvard
pillai@fas.harvard.edu

Aaron Smith
University of Ottawa
asmith3@math.stanford.edu

MS74

Convex Relaxations of Polynomial Imaging Problems

Recent work on phase retrieval suggests that quadratic optimization problems can sometimes be solved by lifting into the space of PSD matrices. This approach can be seen as a special case of sum-of-squares moment relaxation, which is known in theory to convexify much more general polynomial optimization problems. We show that this set of ideas offers a fresh point of view toward resolving the hard nonlinearities of inverse wave scattering. Joint work with Augustin Cosse.

Laurent Demanet
Professor of Mathematics, MIT
laurent@math.mit.edu

Augustin Cosse
MIT
acosse@math.mit.edu

MS74

Fast Algorithms for Linear Inverse Problems with Gaussian Priors

A common theme in Gaussian process methods for inverse problems is the need to work with dense covariance matrices. Standard approaches based on the Cholesky decomposition have cubic complexity. In this talk, we describe a method to construct a generalized Cholesky decomposition of many commonly used covariance matrices (e.g., squared exponential, Matérn, rational quadratic) in linear time. The algorithm is based on hierarchical matrix approximation and borrows heavily from fast multipole-type ideas for compressing structured linear operators. The factorization allows us to efficiently apply the matrix and its inverse (inference), apply the matrix square root (sampling), and compute the log-determinant (likelihood calculations), among other related capabilities. We anticipate that such fast techniques will have a significant impact on large-scale inversion. This is joint work with Lexing Ying.

Kenneth L. Ho
Department of Mathematics
Stanford University
klho@stanford.edu

Lexing Ying
Stanford
lexing@math.stanford.edu

MS74

Stochastic inadequacy operators with applications to chemical kinetics

We investigate model form uncertainty for a reaction mechanism model in hydrocarbon combustion. In a typical reaction, the complete mechanism is either not well-understood, or too complex to effectively use as part of a larger combustion problem, necessitating a reduced model. To account for the discrepancy between the full model and its reduced version, we propose an additive, linear, probabilistic formulation. This representation is encoded in a random matrix, whose entries are calibrated using a hierarchical Bayesian scheme. In particular, this formulation is designed to respect certain physical constraints, but also be flexible enough to apply to multiple reactions.

Rebecca Morrison
UT Austin
rebeccam@ices.utexas.edu

Robert D. Moser
University of Texas at Austin
rmoser@ices.utexas.edu

MS75

Multilevel Collocation with Radial Basis Functions

In this talk, we discuss multilevel radial basis function collocation methods for solving elliptic partial differential equations on bounded domains. The approximate solution is constructed in a multilevel fashion, each level using compactly supported radial basis functions with smaller support radius on an increasingly fine mesh. A convergence theory is given, which builds on recent theoretical advances for multilevel interpolation using compactly supported radial basis functions. If time permits, we also discuss the condition numbers of the arising systems as well as the effect of a simple, diagonal preconditioner.

Patricio Farrell
Weierstrass Institute for Applied Analysis and Stochastics
farrell@maths.ox.ac.uk

MS75

A Reduced Radial Basis Function Method for Partial Differential Equations on Irregular Domains

We propose and test the first Reduced Radial Basis Function Method (R2BFM) for solving parametric partial dif-

ferential equations on irregular domains. The two major ingredients are RBF-FD solver and a collocation-based model reduction approach that systematically generates a reduced-order approximation whose dimension is orders of magnitude smaller than the total number of RBF centers. The resulting algorithm is demonstrated through two- and three-dimensional test problems.

Alfa Heryudono
University of Massachusetts Dartmouth
Department of Mathematics
aheryudono@umassd.edu

Yanlai Chen
Department of Mathematics
University of Massachusetts Dartmouth
yanlai.chen@umassd.edu

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

Akil Narayan
UMass Dartmouth
anarayan@umassd.edu

MS75

Kernel-based Image Reconstruction from Scattered Radon Data

A novel kernel-based algebraic reconstruction method for image reconstruction from scattered Radon data is proposed. The reconstruction relies on generalized Hermite-Birkhoff interpolation by positive definite kernel functions in combination with a well-adapted regularization of the Radon transform. This leads to a very flexible image reconstruction method, which works for arbitrary distributions of Radon lines, unlike in classical Fourier-based methods relying on the filtered back projection formula. The good performance of the proposed kernel-based image reconstruction method is supported by numerical examples and comparisons.

Armin Iske
University of Hamburg
Department of Mathematics
armin.iske@uni-hamburg.de

MS75

RBF-Based Partition of Unity Collocation Methods for the Numerical Solution of PDEs

Numerical methods based on radial basis function (RBF) approximation are attractive because they can achieve high-order accuracy and can handle non-trivial geometries. Global RBF methods can be computationally expensive. Therefore, we propose a localized approach where RBFs are used in a partition of unity framework. We have shown theoretically that the resulting method converges spectrally in the node distance and algebraically in the patch size. We present numerical experiments for elliptic and parabolic PDEs.

Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

Alfa Heryudono
University of Massachusetts Dartmouth
Department of Mathematics
aheryudono@umassd.edu

MS76

Human Fetal Growth Model of Hypoplastic Left Heart Syndrome

Abstract not available at time of publication.

Adarsh Krishnamurthy
UC San Diego
adarsh@ucsd.edu

MS76

A Computational Model of Reverse Cardiac Growth in Response to Mechanical Stimulus

Abstract not available at time of publication.

Lik Chuan Lee
College of Engineering
Michigan State University
lclee@egr.msu.edu

MS76

Modeling Growth and Remodeling in Heart Muscle Tissue

Abstract not available at time of publication.

Joakim Sundnes
Simula Research Laboratory
sundnes@simula.no

MS76

Finite Element Models of Growth and Remodelling in the Infarct Injured Left Ventricle

Abstract not available at time of publication.

Samuel Wall
Simula Research Laboratory
Center for Biomedical Computing
samwall@simula.no

MS77

Tribits: Tribal Build, Integrate, and Test System

The Tribal Build, Integrate, and Test System (TriBITS) is a framework built on top of the open-source CMake tools which is designed to handle large software development projects involving multiple independent development teams and multiple source repositories. TriBITS also defines a complete software development, testing, and deployment system consistent with modern agile software development best practices. TriBITS is used by Trilinos, the Advanced Simulation of Light Water Reactors (CASL) VERA codes, and other projects.

Roscoe Bartlett
Oak Ridge National Laboratory

bartlettra@ornl.gov

MS77

Computational Model Builder and ParaView Catalyst: Empowering HPC Workflows

A key metric for computational scientists is how quickly they gain insight into their problem from simulations. Barriers include developing complex input models and analyzing results. We discuss how Computational Model Builder (CMB) can create simulation input models starting from the problem geometry. We show multiple CMB configurations targeting various simulators. Additionally, we discuss how to perform in situ analysis and visualization with ParaView Catalyst to reduce the time spent post-processing simulation results.

Andrew Bauer
Kitware Inc.
andy.bauer@kitware.com

Patrick O'Leary
Kitware, Inc., USA
patrick.oleary@kitware.com

Robert O'Bara, Berk Geveci
Kitware
bob.obara@kitware.com, berk.geveci@kitware.com

MS77

Developing Hydra-TH: A Vertical, VERA-integrated Application based on the Hydra Toolkit

Hydra is an extensible C++ toolkit that provides a rich suite of software components for rapid scientific application development. Hydra supports multiple discretization techniques and many different physics. This talk describes the development of Hydra-TH, a vertical application for thermal-hydraulics specifically targeted for direct integration into VERA, the Virtual Environment for Reactor Applications. Here, attention is given to the unique aspects of Hydra that enable Hydra-TH (single/multiphase flow physics) to be automatically integrated into VERA.

Mark Christon
Los Alamos National Laboratories
christon@lanl.gov

Jozsef Bakosi
Computational Physics Group
Los Alamos National Laboratory
jbakosi@lanl.gov

Markus Berndt
Los Alamos National Laboratory
berndt@lanl.gov

Andrew Bauer
Kitware Inc.
andy.bauer@kitware.com

Alan Stagg
Oak Ridge National Laboratory
staggak@ornl.gov

Balasubramanya Nadiga
Los Alamos
balu@lanl.gov

Patrick O'Leary
Kitware, Inc., USA
patrick.oleary@kitware.com

MS77

Software Quality with the Open Source Tools CMake, CDash, CTest

CMake has been used on several large projects such as KDE, ITK, Hydra, VTK, ParaView, VXL, Trilinos and CMake itself. In addition to building software, CMake provides a testing client (CTest) that integrates with the web-based CDash testing server. This talk will cover how these tools can be leveraged in the context of an integrated development environment to properly manage the software process. This speeds up development while maintaining a high quality code-base.

Bill Hoffman
Kitware
bill.hoffman@kitware.com

MS78

Teaching Computing to Engineers

The term "computational thinking" became quite popular some years ago with a viewpoint piece at the CACM by Wang (2006), but the term goes back to Papert 10 years before. Wang explained it as the ability to think algorithmically and apply problem-solving computation in other fields. We've been talking about teaching computational thinking ever since. With an interest in reforming how we teach computing to engineering students, I stumbled onto an extension of this idea that defines the pedagogical value of computation. With modern tools for interactive computing (e.g. IPython Notebooks), it becomes possible to learn by computing, to actually create knowledgesimilar to what we do in scientific computing to make discoveries, but in education. I will describe how this thinking has been inspiring several educational initiatives aiming at making computing a core pedagogical instrument.

Lorena A. Barba
Department of Mechanical and Aerospace Engineering
George Washington University
labarba@gwu.edu

MS78

Teaching Data Science from a Computer Science Perspective: Experience from a First Mooc

In Spring 2013 and Summer 2014, we ran a first massively open online course in Data Science, attracting over 160,000 students. In this talk, Ill describe our motivation for designing this course, some unique aspects of the curriculum, our results from the first two offerings, and finally some recommendations around how training students to participate in data-intensive science has become essentialy to prepare them for both research and industry roles.

Bill Howe
University of Washington
billhowe@cs.washington.edu

MS78

Opportunities and Experiences with Teaching Computational Science from the Very Start of Uni-

versity Studies

Over the last 10 years, newcomers to science programs at the Univ. of Oslo have been exposed to a tandem of analytical and numerical methods together with computer programming. The talk will provide examples on how such an approach enhances the understanding of mathematical concepts and how it enables a new pedagogical approach, more relevant working styles, and a tighter coupling to research in traditional science subjects, with physics as primary example.

Hans Petter Langtangen
Center for Biomedical Computing
Simula Research Laboratory and University of Oslo
hpl@simula.no

MS78

Teaching Statistical Computing to Undergraduates

We frame the minisymposium in the context of our experience teaching computing with data to undergraduates. We discuss our recent efforts to automate grading student computational work in Python and R. Automating grading can improve pedagogy by making it easier to assign students more work, reducing latency in providing feedback, and enabling TAs to spend more time working directly with students by freeing them from the "busy work" of grading.

Kenneth J. Millman
University of California, Berkeley
millman@berkeley.edu

Philip B. Stark
Statistics
University of California, Berkeley
stark@stat.berkeley.edu

MS79

Parallel Scalable FETI Methods for Nonlinear Problems

The solution of nonlinear problems requires fast and highly scalable parallel solvers. FETI-DP domain decomposition (DD) methods are parallel solution methods for implicit problems. A common iterative DD approach for nonlinear problems is a Newton-Krylov-DD strategy where the nonlinear problem is linearized using a Newton method. Then, the linear system associated with the tangent stiffness matrix is solved with a preconditioned Krylov space method. The preconditioner is an efficient and parallel scalable domain decomposition method where local subdomain problems and a sufficiently small global problem have to be solved. The local problems are inherently parallel, the global problem is needed to obtain numerical and parallel scalability. FETI-DP domain decomposition methods have been shown to be scalable for linearized elasticity problems on up to 65 536 cores of a BlueGene/P supercomputer (JUGENE, JSC, Germany) in 2009. Recently, nonlinear versions of the well-known FETI-DP methods for linear problems have been introduced. Here, the nonlinear problem is decomposed directly before linearization. The new approaches have the potential to reduce communication and to show a significantly improved performance, especially for problems with localized nonlinearities, compared to a standard Newton-Krylov-FETI-DP approach. Another new approach can be viewed as a strategy to further localize computational work and to extend the scalability of FETI-DP methods towards extreme-scale supercomputers. Here, a recent nonlinear FETI-DP method is combined with an approach that allows an inexact solution of the FETI-DP coarse problem. We combine the nonlinear FETI-DP domain decomposition method with an algebraic Multigrid method and obtain a hybrid nonlinear domain decomposition/multigrid method. Parallel scalability results for up to 262 144 cores on the MIRA BlueGene/Q supercomputer (Argonne National Laboratory, USA) are shown for our new implementation.

Axel Klawonn
Universität Duisburg-Essen, Germany
axel.klawonn@uni-koeln.de

Martin Lanser
Mathematical Institute
Universit
mlanser@math.uni-koeln.de

Oliver Rheinbach
Technische Universität Bergakademie Freiberg
oliver.rheinbach@math.tu-freiberg.de

MS79

Fluid-Structure-Interaction in Computational Hemodynamcis using Nonlinear Hyperelastic Arterial Wall Models

We consider the fluid-structure-interaction problem in a blood vessel. We follow a monolithic coupling approach (P. Crosetto, S. Deparis, G. Fourestey, A. Quarteroni, Parallel algorithms for fluid structure-interaction problems in haemodynamics. SISC, 33(4), 1598-1622, 2011), applying a Convective Explicit approach for the fluid. To obtain an accurate prediction of transmural stresses we make use of sophisticated nonlinear material models for the vessel wall. Fortunately, such models have been developed in the past and their parameters have been adapted to experimental data. Here, we use an anisotropic, polyconvex hyperelastic material model for the structure (D. Balzani, P. Neff, J. Schröder, G.A. Holzapfel, A polyconvex framework for soft biological tissies. Adjustment to experimental data. IJSS, 43(20), p. 6052–6070, 2006). The coupled simulations build on the LifeV software library (LifeV Software Library, www.lifev.org) and FEAP (R.L. Taylor, Finite Element Analysis Program,http://www.ce.berkeley.edu/projects/feap/). Absorbing boundary conditions on the outflow are imposed to reduce reflections.

Daniel Balzani
Institute of Mechanics
Universit
daniel.balzani@uni-due.de

Simone Deparis
Ecole Politechnique Federale de Lausanne
simone.deparis@epfl.ch

Simon Fausten
Institute of Mechanics
Universit
simon.fausten@uni-due.de

Davide Forti
Chair of Modeling and Scientific Computing,
Ecole Politechnique Federale de Lausanne
davide.forti@epfl.ch

Alexander Heinlein
Mathematical Institute
Universit
alexander.heinlein@uni-koeln.de

Axel Klawonn
Universität Duisburg-Essen, Germany
axel.klawonn@uni-koeln.de

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

Oliver Rheinbach
Technische Universität Bergakademie Freiberg
oliver.rheinbach@math.tu-freiberg.de

Jörg Schröder
Institute of Mechanics
Universit
j.schroeder@uni-due.de

MS79

A Scalable Monolithic Solver for the Coupling of a Finite Element and a Finite Volume Method for Fluid-Structure-Interaction

We use two different discretization methods for the fluid and the structure sub-problems in a monolithic approach, i.e both sub-problems as well as the coupling conditions are assembled into one large algebraic system. Our solver bases on the application of Newtons method using iterative solvers with different multi-level methods. In this talk we discuss both, the coupling approach of the two different discretizations as well as the efficient solution of the arising large nonlinear system.

Johannes Steiner
University of Lugano
johannes.steiner@usi.ch

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

MS79

Adaptive Spectral Deferred Correction Methods for Cardiac Simulation

The electrical excitation of the heart as described by the monodomain equations exhibits a wide range of temporal and spatial scales. In this talk we will explore the possibilities of combining spectral deferred correction (SDC) methods for time stepping with spatial and temporal adaptivity, in particular using optimized DIRK sweeps, interleaving mesh refinement with SDC iteration, inexact linear solves, and local time stepping. The efficiency of corresponding adaptive algorithms are illustrated at some numerical examples.

Martin Weiser
Zuse Institute (ZIB)
Berlin, Germany
weiser@zib.de

MS80

A High-order Finite Element Method for Moving Boundary Problems Using Universal Meshes

Abstract not available at time of publication.

Evan Gawlik
Stanford University
tba@siam.org

MS80

Towards a Hybrid Parallelization of Chebyshev Filtered Subspace Iteration

Abstract not available at time of publication.

Maria Kranjcevic
University of Zagreb
tba@siam.org

MS80

Computational Reduction Strategies for Bifurcation and Stability Problems

Abstract not available at time of publication.

Giuseppe Pitton
SISSA, Trieste, Italy
tba@siam.org

MS80

Nonlinear Frequency Response Analysis of Mechanical Vibrations based on Isogeometric Discretization and Model Order Reduction

Abstract not available at time of publication.

Oliver Weeger
TU Kaiserslautern
tba@siam.org

MS81

Low-Rank Adaptive Tensor Approximation

A strategy to mitigate the curse of dimensionality", which roughly means that the computational work, needed to approximate a given function within a desired target accuracy increases exponentially in the spatial dimension, is to seek problem dependent dictionaries with respect to which the function pos- sesses sparse approximations. In this talk we highlight some recent developments centering on the adaptive solution of high dimensional elliptic operator equations in terms of linear combinations of particularly adapted rank-one tensors. An adaptive iterative coarse-to-fine algorithm is presented that produces near-minimal rank approximation in stable hierarchical tensor formats where the adaptive identification of corresponding subspaces and sparse approximations of the low-dimensional tensor factors are intertwined. We highlight the main conceptual ingredients as well as some essential obstruction due to the fact that the underlying energy spaces are not endowed with cross norms. The theoretical results are illustrated by some numerical tests.

Wolfgang Dahmen

RWTH Aachen
IGPM
dahmen@igpm.rwth-aachen.de

MS81

A Theory for Model Verification

Model verification tries to capture the solution u to a parametric PDE when the parameters that determine u are not known but other information about u is present. The form of this other information is typically (i) a posteriori measurements in the form of linear functionals applied to u and (ii) knowledge that u is well approximated by elements of a known finite dimensional space V. We describe theoretically, the best we can approximate u from this knowledge and then discuss numerical algorithms which perform near these best bounds.

Ronald DeVore
Department of Mathematics
Texas A&M University
rdevore@math.tamu.edu

MS81

High-Order Digital Nets for Parametric and Stochastic Operator Equations

We present sparsity theory for direct and inverse problems for PDEs with infinite-dimensional uncertain input parameters, stemming from parametrization of distributed uncertainty.It reduces the direct and inverse problem to integration problem over infinite-dimensional parameter spaces. Based on a holomorphy condition on the parameteric dependence in [1], we present regularity estimates for the parametric integrand functions and for uniform prior measure on the parameter uncertainty in classes of weighted RKHS introduced in [2]. Related recent results (joint with J. Dick, F. Kuo, T. LeGhia and D. Nuyens) [3] on dimension independent convergence rates of the deterministic, higher order QMC quadrature for integrand functions in weighted function spaces will be presented. The density of the posterior measure in Bayesian inverse problems as considered in [5, 4] is shown to belong to the class of admissible integrand functions with a hybrid of product and sPOD weights. Research supported in part under ERC AdG 247277.

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

MS81

Multivariate Decomposition Methods ∞-Variate Problems

We present the *Multivariate Decomposition Method* for approximation of functions of infinitely many variables It works for functions that admit a decomposition $f = \sum_u f_u$, where the sum is with respect to finite subsets u of positive integers, and for each $u = (i_1, \ldots, i_k)$ the function f_u depends only on $x_{i_1}, \ldots, x_{i_k}$. For a number of weighted spaces of such functions, the complexity of the corresponding problem is small.

Grzegorz W. Wasilkowski
University of Kentucky
Department of Computer Science
greg@cs.uky.edu

MS82

Title Not Available at Time of Publication

Abstract not available at time of publication.

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

MS82

Portable Programming and Runtime Support for Application-Controlled Resilience in Large-Scale Scientific Applications

The Global View Resilience (GVR) system supports flexible, scalable, application-controlled resilience with a portable abstraction versioned, distributed arrays. We have evaluated GVRs utility on both a number of mini-apps (miniMD, miniFE), and larger applications (a PCG solver, OpenMC, ddcMD, and Chombo). Our results show programming effort (code change) is small ($< 1\%$ code) and localized. Studying the same applications, we find that GVR version-based resilience can be achieved with low overhead for all of them.

Andrew A. Chien
The University of Chicago
Argonne National Laborator
achien@cs.uchicago.edu

Hajime Fujita
Department of Computer Science
University of Chicago
hfujita@cs.uchicago.edu

Zachary Rubenstein
Department of Computer Science, The University of Chicago
zar1@uchicago.edu

Nan Dun, Aiman Fang
Department of Computer Science
University of Chicago
dun@cs.uchicago.edu, aimanf@cs.uchicago.edu

Ziming Zheng
Department of Computer Science, The University of Chicago
zimingzheng@uchicago.edu

MS82

Resilient Programming Models

Today's leadership computing systems and scalable clusters already exhibit significant failure rates that impact overall results rates and scientific productivity. Current issues are expected to persist, and new issues to emerge. In this presentation we discuss several abstract programming models that permit application and algorithm developers to reason about and develop resilience capabilities. We present the basic models and then discuss specific strategies for how to design and implement applications to be resilient to system failures such as local process loss, silent data corruption, and the performance variability that is

inherent in aggressive failure detection and correction.

Michael Heroux
Sandia National Laboratories
maherou@sandia.gov

MS82

MPI Fault Tolerance: The Good, The Bad, The Ugly

The MPI forum is currently investigating the inclusion of fault tolerance as a feature in the MPI specification. This issue has raised and continuous to raise some controversy about what MPI implementations can reasonably be expected to provide and what is useful for application developers. In this talk I will present the main proposals that are currently on the table, their advantages and the main concerns against them. The goal of this talk is to expand the discussion and to gather feedback that will help the MPI forum to come to a solution that is helpful for the larger HPC community.

Martin Schulz
Lawrence Livermore National Laboratory
schulzm@llnl.gov

MS83

Preconditioning for Various Cahn-Hilliard Systems

The Cahn-Hilliard equation models the motion of interfaces between several phases and has many applications including materials science and image inpainting. Besides their smooth formulations we study the nonsmooth ones which result in variational inequalities. The focus is on the efficient solution of the arising large and sparse linear systems. Further, we study fractional-in-space versions of Cahn-Hilliard systems. Numerical results illustrate the efficiency of the approaches.

Jessica Bosch
MPI Magdeburg
bosch@mpi-magdeburg.mpg.de

Martin Stoll
Max Planck Institute, Magdeburg
stollm@mpi-magdeburg.mpg.de

MS83

Null-Space Based Preconditioners for Saddle-Point Systems

We derive a formula for the inverse of a nonsingular saddle-point matrix whose leading block is maximally rank-deficient, which is based on the null-space of that block. We then use the formula to develop a class of indefinite block preconditioners. When a sparse form of the null space is approximately available, a preconditioned Krylov iterative solver converges rapidly. We give a couple of examples, discuss spectral properties, and present some numerical results that validate the analysis.

Ron Estrin
The University of British Columbia
ronestrin756@gmail.com

Chen Greif
Department of Computer Science
The University of British Columbia
greif@cs.ubc.ca

MS83

Unreduced Symmetric KKT Systems Arising from Interior Point Methods

We consider symmetrized KKT systems arising in the solution of quadratic programming problems by Interior Point methods. Two strictly related and well established formulations for such systems are studied, with particular emphasis on their spectral properties and how these are affected by preconditioning. Constraint and augmented preconditioners are considered here. Both a theoretical and experimental analysis is conducted in order to assess which of the two formulations should be preferred when solving large scale problems.

Mattia Tani, Valeria Simoncini
Universita' di Bologna
mattia.tani2@unibo.it, valeria.simoncini@unibo.it

Benedetta Morini
Dipartimento di Ingegneria Industriale
Universita' di Firenze
benedetta.morini@unifi.it

MS83

Convergence of Stationary Iteration with Indefinite Preconditioner

The relationship of diagonal dominance ideas to the convergence of stationary iterations is well known. There are a multitude of situations where such considerations are used to guarantee convergence when the splitting matrix (the preconditioner) is positive definite. In this talk we will describe and prove sufficient conditions for convergence of a stationary iteration based on a splitting with an indefinite preconditioner. Simple examples covered by this theory from Optimization and Economics will be described.

Andy Wathen
Oxford University, UK
wathen@maths.ox.ac.uk

MS84

Quantifying Errors in a Probabilistic Solution to Stochastic Inverse Problems for Physics-Based Models

We define a measure-theoretic framework to formulate and solve a stochastic inverse problem for deterministic physics-based models. Computational algorithms are presented and analyzed to solve problems within this framework involving high dimensional input parameter and output data spaces. Sources of statistical and deterministic errors are identified and full a priori and a posteriori error analyses on computed probabilities of events are presented including numerical examples.

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

MS84

Region of Influence Sensitivty Analysis for Time Dependant Problems

We review existing strategies for implementing adjoint-

based sensitivity analysis on HPC architectures. We focus on strategies based on space-time local adjoint problems, posed in the "Region of Influence"(RoI), motivated by the application to uncertainty quantification of DNS-scale turbulent combustion. In addition, we discuss ramifications and potential benefits from implementing this approach on potential exascale systems.

Varis Carey, Robert D. Moser
University of Texas at Austin
variscarry@googlemail.com, rmoser@ices.utexas.edu

MS84

Adaptive Measure-Theoretic Inverse Techniques for High Dimensional Parameter Domains and Complex Multi-Scale Models

The application of uncertainty quantification methodologies to computationally complex models often involves the exploration of a high-dimensional parameter space. It is well known that this endeavor is plagued by the "curse of dimensionality." We explore the scalability of adaptive algorithms for solving the stochastic inverse problem within the context of a measure-theoretic solution framework to with the goal of reducing error in approximating the probability of implicitly defined "rare" events.

Lindley C. Graham
University of Austin at Texas
Institute for Computational Engineering and Sciences (ICES)
lgraham@ices.utexas.edu

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

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

MS84

Optimizing Quantities of Interest in High Dimensions to Improve Solutions to Inverse Problems

The predictive capabilities of physics-based models are improved by reliably decreasing the size of the sets defining the uncertain input parameters. These sets are often inferred by solution to an inverse problem. We explore techniques for identifying the optimal quantities of interest within a high dimensional output data set for use in the inverse problem to improve the predictive capabilities of a model. Numerical results on physically relevant models are provided.

Scott Walsh
University of Colorado Denver
scott.walsh@ucdenver.edu

MS85

Modeling Magneto-Mechanical Interactions in Deformable Solids

A new class of "smart" materials, magneto-sensitive materials (MSMs), is at the frontier of research in material design. The mechanical properties of these materials change dramatically in the presence of an applied magnetic field. Such materials undergo finite deformations and the magneto-mechanical coupling requires nonlinear models. Incompressible MSMs are then modeled by heuristic equations that minimize the elastic and electromagnetic energy in the form of a coupled free energy. In this talk, we focus on modeling and developing finite-element methods for incompressible magneto-elasticity, and we present theoretical results and numerical experiments for such problems. The exterior of the material body is also taken into consideration due to the magnetic field, resulting in a much bigger discrete system. Even within uncoupled incompressible elasticity, many open questions remain about optimal finite-element approximation and these, of course, impact our choices for the more complicated coupled problems considered here.

James H. Adler, Luis Dorfmann
Tufts University
james.adler@tufts.edu, luis.dorfmann@tufts.edu

Dong Han
Tufts University
Department of Mathematics
dong.han@tufts.edu

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

Chris Paetsch
Tufts University
chris.paetsch@tufts.edu

MS85

First-Order System Least Squares for Isotropic and Anisotropic Materials in Hyperelasticity

We present least squares finite element methods based on the conservation of linear momentum and nonlinear constitutive equations for hyperelastic materials. Our approach is motivated by a well-studied least squares formulation for linear elasticity. This idea was already generalized to nonlinear elasticity using a special Neo-Hooke model. We recall essential theorems and illustrate the performance using adaptive refinement strategies. Additionally we extend this approach to transversely isotropic materials adapting the isotropic model and give examples.

Benjamin Müller
University of Duisburg-Essen
benjamin.mueller@uni-due.de

Gerhard Starke
University of Duisburg-Essen
Fakultät für Mathematik
gerhard.starke@uni-due.de

Jörg Schröder
Universität Duisburg-Essen
Institut für Mechanik, Abteilung Bauwissenschaften
j.schroeder@uni-due.de

MS85

Advanced Finite Element Methods for Chemo-Electromechanical Skeletal Muscle Mechanics

Simulating the mechanical behavior of (soft) biological ma-

terials is challenging as one does not only need to consider a highly nonlinear material behavior within a complex domain, but also often needs to consider different scales, e.g. the cellular scales, to accurately model functional aspects, e.g. for simulating the chemo-electromechanical behavior of skeletal muscles. Furthermore, due to high uncertainties within material parameters, one needs fast, efficient, and accurate solution strategies.

Oliver Rohrle, Thomas Heidlauf, Mylena Mordhorst, Daniel Wirtz
Institute of Applied Mechanics
University of Stuttgart, Germany
roehrle@simtech.uni-stuttgart.de, heidlauf@mechbau.uni-stuttgart-de, mylena.mordhorst@mechbau.uni-stuttgart.de, daniel.wirtz@simtech.uni-stuttgart.de

MS85

Momentum Balance Accuracy in Finite Element Methods for Elastoplasticity

First-order system least squares formulations involving stresses and displacement as process variables are investigated for elastoplasticity models. Optimal order convergence is shown for the stress approximation with respect to the H(div) norm using Raviart-Thomas finite elements in the context of von Mises flow with isotropic hardening. In particular, the implications on the accuracy of momentum balance and surface forces is studied. Computational results for a realistic benchmark problem illustrate the effectiveness of this approach.

Gerhard Starke
University of Duisburg-Essen
Fakultät für Mathematik
gerhard.starke@uni-due.de

MS86

A Coupled CFD CAA Adjoint Method for Aeroacoustic Optimization and Error Estimation

Adjoint techniques have played a major role in gradient based steady state aerodynamic optimization, especially in 3D, since the full sensitivity vector of a single objective function with respect to any number of design variables can be computed with a single adjoint solution, at a cost roughly equivalent to a single flow solution. On the contrary adjoint methods for unsteady problems have received less attention, their development having been hindered by the inherent computational cost and the complexity of the associated flow physics. Recently unsteady adjoint techniques have been proposed for both two-dimensional and three-dimensional problems. In this work we apply the unsteady adjoint method to a two dimensional blade vortex interaction noise problem. An Euler near field flow solver is coupled to an FW-H aeroacoustic code to propagate the noise to a far field observer. The discrete adjoint solvers for both the flow and the acoustic codes are derived by exact linearization and transposition of each subroutine in the analysis code and finally coupled together. We apply the newly developed adjoint solver to blade vortex interaction noise in the context of gradient based optimization, to investigate optimal passive noise minimization technique, and a posteriori error estimation to improve the accuracy of the coupled aeroacoustic analysis.

Enrico Fabiano
University of Wyoming
efabiano@uwyo.edu

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

Jay Sitaraman
University of Wyoming
jsitaraman@uwyo.edu

MS86

Efficient Approaches for Optimal Active Flow Control

For efficient optimal active control of unsteady flows, the use of consistent and robust adjoint approaches is a first essential ingredient. For the generation of discrete adjoint solvers, we discuss the use of Automatic Differentiation (AD) and its combination with checkpointing techniques. Furthermore, we discuss so-called one-shot methods. Here, one achieves simultaneously convergence of the primal state equation, the adjoint state equation as well as the design equation.

Nicolas R. Gauger
RWTH Aachen University
nicolas.gauger@rhrk.uni-kl.de

Anil Nemili, Emre zkaya
TU Kaiserslautern
anil.nemili@rhrk.uni-kl.de, emre.oezkaya@rhrk.uni-kl.de

Stefanie Günther
RWTH Aachen University
stefanie.guenther@scicomp.uni-kl.de

MS86

Aerodynamic Design for Unsteady Flows Using An Adjoint Approach

A discrete adjoint-based design methodology for unsteady turbulent flows on three-dimensional dynamic overset unstructured grids is described. The methodology supports both compressible and incompressible flows and is amenable to massively parallel computing environments. The approach provides a general framework for performing highly efficient and discretely consistent sensitivity analysis. Meshes consisting of mixed-element topologies and overset component grids are supported, where grids may be static, dynamic, or deforming, including any combination thereof. An overview of a broad range of aerospace applications for which the implementation has been demonstrated will be shown.

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

Boris Diskin
National Institute of Aerospace, Hampton, VA
bdiskin@nianet.org

MS86

Optimal Wall-Forcing for Compressible Wall-Bounded Flows Using Adjoint Techniques

Adjoint operators are vastly used in many applications, ranging from linear stability analysis and optimization to flow control. Although highly valuable, the derivation of

these operators in the case of compressible and higher-order solvers is often a tedious task. This is particularly due to the complexity of the governing equations and of the higher-order numerical schemes. In this study, we have adopted a novel technique for the evaluation of the direct and adjoint operators directly from the flow solvers (Fosas et al., 2012) which requires minimal additional programming efforts, and automatically takes into account subsequent modifications in the governing equations and boundary conditions. This approach is applied to a compressible, staggered and curvilinear framework, allowing higher-order interpolation and derivation schemes. The original nonlinear solver has been used to perform large-scale turbulent flow simulations, and shown to scale well on up to 62K processors. The adjoint solver features similar performances. The developed methodology is first validated in the context of a three-dimensional compressible boundary layer, by extracting the optimal initial condition, and comparing the results to that of the linear stability analysis. Finally, this method is used to extract the optimal wall-forcing in a compressible wall-bounded flow.

Taraneh Sayadi, Peter Schmid
Imperial College London
sayadi@illinois.edu, peter.schmid@imperial.ac.uk

MS87

Stable and Robust Hybridized Discontinuous Galerkin Methods for High Reynolds Number Flow Problems

We present an output-based anisotropic hp-adaptive method for high-order hybrid discontinuous Galerkin (HDG) discretizations of the Navier-Stokes equations. The adaptive framework uses a discrete adjoint on refined spaces to compute error estimates. The effectiveness of refinement options is determined by solving local adjoint problems. We discuss the effects of different stabilization methods on solver robustness and compare the effectiveness of HDG and traditional discontinuous Galerkin (DG) methods.

Krzysztof Fidkowski, Johann Dahm
University of Michigan
kfid@umich.edu, jdahm@umich.edu

MS87

Multiscale Hybridizable Discontinuous Galerkin Methods

We present the recent development of multiscale hybridizable discontinuous Galerkin (MHDG) methods for computational electromagnetics and fluid dynamics. The essential ingredients are (i) a HDG discretization of the underlying PDEs at the subdomain level to parametrize the numerical solution in terms of a Lagrange multiplier; (ii) a judicious choice of the numerical flux to provide stability and consistency; and (iii) a global jump condition that enforces the continuity of the numerical flux across subdomain boundaries to arrive at a global weak formulation in terms of the Lagrange multiplier. The MHDG methods inherit all the properties of the HDG method and possess additional advantages. First, they reduce the globally coupled unknowns to the approximate trace of the solution on subdomain boundaries, thereby leading to a significant reduction in the degrees of freedom. Second, they make use of the similarity of the subdomain structures to efficiently accommodate very large-scale simulations. And third, the MHDG methods lend themselves efficient for parallelization. We will present several examples in nanophotonics and turbomachinery to demonstrate the performance of the MHDG method.

Cuong Nguyen
Massachusetts Institute of Technology
cuongng@mit.edu

Joel Saa-Seoane
M.I.T.
jsaa@mit.edu

David Moro, Francisco J. Roca, Jaime Peraire
Massachusetts Institute of Technology
dmoro@mit.edu, xeviroca@mit.edu, peraire@MIT.EDU

MS87

A Computational Framework for Target-Based HP-Adaptation in Compressible Flow Simulation Using HDG Methods

We present a conceptual overview of our computational framework which includes both standard and hybridized DG methods and offers both isotropic and anisotropic hp-adaptation based on adjoints. The framework is designed in such a way that both discretizations and physical models can be exchanged rapidly by making heavy use of object-orientation and C++ templates. We will show examples in a variety of flow regimes.

Georg May, Michael Woopen
AICES
RWTH Aachen
may@aices.rwth-aachen.de, woopen@aices.rwth-aachen.de

MS87

To CG or HDG: Updates on Our Comparative Study

Since the inception of discontinuous Galerkin (DG) methods for elliptic problems, there has existed a question of whether DG methods can be made more computationally efficient than continuous Galerkin (CG) methods. Fewer degrees of freedom, approximation properties for elliptic problems together with the number of optimization techniques, such as static condensation, available within the CG framework made it challenging for DG methods to be competitive until recently. However, with the introduction of a static-condensation-amenable DG method, the hybridizable discontinuous Galerkin (HDG) method, it has become possible to perform a realistic comparison of CG and HDG methods when applied to elliptic problems. In this talk, we extend upon an earlier 2D comparative study, providing numerical results and discussion of the CG and HDG method performance in three dimensions. The comparison categories covered include steady-state elliptic and time-dependent parabolic problems, various element types and serial and parallel performance.

Mike Kirby, Sergey B. Yakovlev
University of Utah
School of Computing
kirby@cs.utah.edu, yakovs@sci.utah.edu

MS88

Partitioned Low Rank fast and Efficient Compression of Absorbing Boundary Conditions for the

Helmholtz Equation

Absorbing layers are sometimes required to be impractically thick to offer an accurate Absorbing Boundary Condition (ABC) for the Helmholtz equation in heterogeneous media. In previous work [BR and Demanet, submitted, 2014], we used matrix probing to compress an ABC from a few exterior Helmholtz solves with random Dirichlet data. We now present an algorithm (nearly linear in the dimension of the matrix) for applying this compressed ABC using Partitioned Low Rank matrices.

Rosalie Belanger-Rioux
Massachusetts Institute of Technology
Department of Mathematics
robr@math.mit.edu

Laurent Demanet
Professor of Mathematics, MIT
laurent@math.mit.edu

MS88

Comparison Between DG and Finite Difference Methods for Acoustics with Smooth Coefficients

This work analyzes the computational efficiency of two types of numerical methods, finite difference (FD) and discontinuous Galerkin (DG) methods, in the context of 2D acoustic equations in pressure-velocity form with smooth coefficients. The acoustic equations, which model propagation of sound waves in elastic fluids, are used throughout seismic imaging with applications in oil prospecting. The ubiquity of smooth trends in real data, and thus in the acoustic coefficients, validates the importance of this novel study. Previous work, from the discontinuous coefficient case of a two-layered media, demonstrates the efficiency of DG over FD methods but does not provide insight for the smooth coefficient case. Running-times and floating point operations are compared, relative to a prescribed accuracy, for standard 2-2 and 2-4 staggered grid FD methods, and a standard DG methods

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

MS88

Analysis and Numerical Approximation for Adsorption Models

We focus on the structure of an adsorption model as systems of conservation laws (multicomponent case for adsorption), with equilibrium and non-equilibrium type nonlinearities, where the latter are associated with microscale diffusion. We also work with an unusual type isotherm called Ideal Adsorbate Solution, which is defined implicitly. For the IAS adsorption system, we show sufficient conditions that render the system hyperbolic. We also construct numerical approximations for equilibrium and nonequilibrium models.

Francis P. Medina
Oregon State University
medinaf@math.oregonstate.edu

Malgorzata Peszynska
Department of Mathematics
Oregon State University
mpesz@math.oregonstate.edu

MS88

Microseismic Event Location Via Full Waveform Inversion

The process of hydraulic fracturing involves injecting large volumes of water into impermeable rocks such as shale in an effort to create flow paths for fluids such as oil and gas. This process often generates tiny earthquakes (or microseismic events) which in turn can be used as passive seismic sources for imaging the subsurface. In this work we describe estimating the location of these microseismic events and the uncertainty inherent in this process.

Susan E. Minkoff
University of Texas at Dallas
Dept of Mathematical Sciences
sminkoff@utdallas.edu

MS89

Recovering Exponential Accuracy in Spectral Methods Involving Piecewise Smooth Functions with Unbounded Derivative Singularities

Techniques will be presented to overcome the Gibbs phenomenon for functions with unbounded derivative singularities, resulting in recovering exponential accuracy in the maximum norm from the knowledge of the first N Fourier coefficients or standard collocation point values of such functions. With these post-processing methods, we are able to obtain exponential accuracy of spectral methods applied to linear transport equations involving such function.

Zheng Chen
Iowa State University
zchen@iastate.edu

MS89

Efficient High-Order Algorithms for Solving Drift-Diffusion Systems

I will discuss about recent developments of spectral element method (SEM) for solving drift-diffusion equations, with applications in semi- conductor device simulation, biological ion channels problems, etc. The drift-diffusion system is a non-linear system, involving the coupling of two transport equations for the carrier concentrations with the Poisson equa- tion for the electric potential. I will present our SEM algorithms, focusing on stable, efficient, and accurate time-splitting schemes, properly designed for the high-order spectral element discretizations. I will demonstrate the computational results for the study of potassium channel in a biological membrane, provided with the validation.

Ying He
UC Davis
yinghe@math.ucdavis.edu

MS89

Estimating Residual Stresses in Arteries by an Inverse Spectral Technique

A mathematical model is studied to estimate residual stresses in the arterial wall using intravascular ultrasound (IVUS) techniques. A BVP is formulated for the nonlinear,

slightly compressible elastic wall, the boundary of which is subjected to a quasi-static blood pressure, and then an idealized model for IVUS is constructed by superimposing small amplitude time harmonic vibrations on large deformations. Using the classical theory of inverse Sturm-Liouville problems and op- timization techniques, an inverse spectral algorithm is developed to approximate the residual stresses, given the first few eigenfrequencies of several induced pressures.

Sunnie Joshi
Temple University
sjoshi@temple.edu

MS89

Force-based Blended Atomistic-to-continuum Coupling Method for Crystals: Theory and Computations

We formulate a multiscale method based on blending atomistic and continuum forces. We present a comprehensive error analysis whichis valid in two and three dimensions, for finite many-body interactions, and in the presence of defects. Based on a precise choice of blending mechanism, the error estimates are considered in terms of degrees of freedom. The numerical experiments confirm and extend the theoretical predictions, and demonstrate a superior accuracy of our method over other schemes.

Xingjie Li
Brown University
xingjie_li@brown.edu

MS90

On-Lattice and off-Lattice Hybrid Simulation Using the Smoldyn Software

The Smoldyn simulator is a tool for modeling biochemical spatial organization on nanometer to micron size scales. It represents proteins and other molecules of interest as individual point-like particles that diffuse, react, and interact with membranes, all in continuous space. Although effective, this is computationally expensive, so colleagues and I recently added on-lattice capabilities to Smoldyn as well. Simulated molecules are able to freely diffuse back and forth between the two regions of space.

Steven Andrews
Fred Hutchinson Cancer Research Center
steven.s.andrews@gmail.com

MS90

MCell/CellBlender: An Environment for Spatially Realistic Simulation of Cellular Microphysiology

MCell is a simulation kernel for biophysically realistic 3D simulations of reaction-diffusion processes occurring in cells. A major stumbling block for new and experienced users of MCell is the effort required to create 3D models. To remedy this situation we have created CellBlender, a complete integrated development environment for creating and visualizing MCell and SBML/Spatial models. An overview of the design of CellBlender will be presented, including the model creation/simulation/visualization pipeline.

Thomas M. Bartol
The Salk Institute
bartol@salk.edu

MS90

From Macroscopic to Microscopic Simulations Using Mesord

MesoRD simulates both the spatial and stochastic aspects of intracellular chemical reactions using a spatially discrete and temporally continuous framework. Spatially discrete stochastic simulators may give incorrect results for bi-molecular reactions. We have incorporated a solution to this problem into MesoRD. Using these new features of MesoRD we could simulate a number of simple examples where it is shown how strikingly important it is to choose a correct modelling-framework for the problem at hand.

David Fange
Uppsala University
Department of Molecular Biology
david.fange@icm.uu.se

MS90

E-Cell System Version 4.0: an Integrated Platform for Single-particle-level Simulations

Here, we present a novel software for cellular simulations, E-Cell System version 4, which provides an integrated platform with a rule-based modeling environment, bioimaging visualizations and a variety of single-particle level simulation algorithms including an exact event-driven particle algorithm, the enhanced Greens Function Reaction Dynamics method, and a microscopic lattice-based method, Spatiocyte. Moreover, we also introduce the parallelization techniques for these particle methods toward the whole-cell-scale simulation on high-performance computers.

Kazunari Kaizu, Kozo Nishida, Masaki Watabe, Arjunan Satya, Kazunari Iwamoto, Koichi Takahashi
RIKEN Quantitative Biology Center, Osaka
kaizu@riken.jp, knishida@riken.jp, masaki@riken.jp, satya@riken.jp, kiwamoto@riken.jp, ktakahashi@riken.jp

MS91

Multiscale Model Reduction for PDE-Constrained Optimization

In parameter estimation and data fitting problems, reduced order models are usually built from the forward map that depends on the unknown parameters. This implies that the reduced model is parameter dependent. Typically, this dependence is ignored and the reduced model is either not updated in the process of the data fitting process or, its dependency is not considered when computing gradients and Hessians. The consequences of this approach can lead to inefficient algorithms as well as unacceptable solution. In this talk we discuss a framework that allow us to the differentiation of the reduced model with respect to the parameter. We show that this framework can yield robust estimates for some model inverse problems.

Eldad Haber
Department of Mathematics
The University of British Columbia
haber@math.ubc.ca

Lars Ruthotto
Department of Mathematics and Computer Science
Emory University

lruthotto@emory.edu

MS91

Model Mis-Specification and Model Reduction - Connecting the Dots

In addressing large-scale inverse problems, great care needs to be devoted to attainment of appropriate balance of inexactness throughout the various stages of modeling and inversion. Disregard to such objective, either entails redundant computation or impairment of the overall fidelity of the inversion process. Model reduction is instrumental in trading-off fidelity for computation, yet, in some situations, it is essential to perform the opposite action, and enhance model fidelity and thereby inversion output. In this talk, we shall describe the interplay between model reduction and model mis-specification mitigation and provide a generic infrastructure for model re-specification based upon a hybrid first principles and data-driven approach.

Lior Horesh
Mathematical Sciences & Analytics
IBM TJ Watson Research Center
lhoresh@us.ibm.com

MS91

Inference for Prediction in Nonlinear Systems

When only sparse data are available for the calibration of a many-parameter model that will be used to make only a few decisions or predictions, not all parts of the posterior parameter space are either informed by the data or informative to the output. This talk presents an adaptive sampling method to preserve the posterior predictive distribution without fully exploring the large posterior parameter distribution.

Harriet Li
MIT
kameeko@mit.edu

MS91

Model Reduction for Some Inverse Problems in Finance

Implied volatility is a key value in financial mathematics. We discuss some of the pros and cons of the standard ways to compute this quantity, i.e. numerical inversion of the well-known Black-Scholes formula or asymptotic expansion approximations, and propose a new way to directly calculate the implied variance in a local volatility framework based on the solution of a quasilinear degenerate parabolic partial differential equation using POD and DEIM. Numerical results prove the quality of our approach compared to other methods.

Ekkehard W. Sachs, Marina Schneider
University of Trier
sachs@uni-trier.de, marina.schneider@uni-trier.de

MS92

Mathematical Modeling at Two Opposite Ends of the Scale Spectrum with the Same Objective in Mind: Improve Human Health

Osteoporosis is a common age related chronic disorder of the skeleton. It constitutes a considerable global public health problem currently affecting more than 200 million people worldwide. We will present an overview of osteoporosis, describe the bone remodeling process, and highlight the fundamental elements involve in determining bone strength in vivo. The talk will show why mathematical modeling is playing an increasingly relevant role in the research, diagnosis, and monitoring of osteoporosis. We will illustrate how modeling has been used at Merck in the development of novel osteoporosis treatments.

Antonio Cabal
Merck Research Laboratories
Quantitative Pharmacology and Pharmacometrics
antonio.cabal@merck.com

MS92

A Simultaneous Approach to Parameter Estimation with Ode Models: a Case Study with Viral Dynamics Models

The successful application of mathematical models and verifying their underlying hypotheses rely critically on our ability to estimate their parameters. Parameter estimation is usually performed using a sequential approach: an optimization algorithm calls the model independently to evaluate the agreement with data. Here, using an ODE-based viral dynamics model as an example, we explore and evaluate a simultaneous approach where the model is discretized and explicitly introduced as equality constraints to the optimization algorithm.

Khamir Mehta
Applied Mathematics and Modeling
Merck & Co, Inc.
khamir_mehta@merck.com

Junghoon Lee
Merck & Co, Inc
jung_hoon_lee@merck.com

MS92

Numerical Solutions of a Partial Differential Equations in a Pharmacometric Context

Pharmacometricians develop quantitative models of the efficacy, potency and safety of drug compounds under development. These models are often constrained to systems of ordinary differential equations (ODEs) because simulation and analysis tools are designed to work with ODEs. Emerging models are starting to clash with these constraints. For example, so-called age structured partial differential equations (PDEs) have potential utility in oncology and safety. We discuss these models and how they may be incorporated into current tools.

Jeffrey Saltzman
AstraZeneca
Research and Development Information Services
Jeffrey.Saltzman@astrazeneca.com

MS92

Applications of Modeling and Simulation in Drug Discovery and Development

Model based drug discovery has the potential to increase efficiency at various stages of drug discovery and development by reducing the cost and time required to make go/no-go decisions. Often times, a mathematical model is developed based on what is known about a particular

compound, such as its pharmacokinetic and pharmacodynamic properties. Such models can be used to simulate the behavior of a compound under multiple scenarios and understand various properties of the compound such as the parameters that drive efficacy, the efficacy/toxicity tradeoff or the most effective dosing regimen for the compound. The insight obtained using such models can be used to guide decisions such as the selection of the most promising pre-clinical compound or the selection of a dosing regimen in the pre-clinical or clinical space. The modeling techniques used in developing these models are usually tailored to the questions that arise at the different stages of drug discovery and development. For example, a system of ordinary differential equations can be used to represent the mechanistic model for a particular compound. Such a mechanistic model could then be analyzed to identify the parameters that affect the efficacy of a compound. In other cases, a simpler empirical model could be used to capture and assess the efficacy/toxicity tradeoff. In this talk I will present a few examples of how various modeling techniques can be used to answer important questions in the pre-clinical and clinical space of drug discovery and development.

Chandni Valiathan
Merck
chandni.valiathan@merck.com;

MS93

Extreme Scale Solution of Engineering Applications Using Uintah

Solving extreme scale problems using the Uintah computational framework was achieved with the development of a general runtime system capable of solving a broad class of problems using a graph-based approach. Important aspects of the solution process, adaptive approach involving dynamic out-of-order execution of tasks, the strengths and weaknesses of this approach are addressed, in particular with future architectures in mind. Examples of scaling to about 700K cores are shown on Blue Waters and Mira.

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

MS93

GPU for Adaptive Optics on Ground Based Astronomical Telescopes: Simulations and Real-Time Control

The European Extremely Large Telescope project (E-ELT) is one of Europe's highest priorities in ground-based astronomy. ELTs are built on top of several highly sensitive and critical astronomical instruments. Particularly, Adaptive Optics (AO), used to stabilize image quality, is essential for telescope routine operations. Designing and driving AO systems requires fast and high fidelity numerical simulations. We describe the simulation framework and highlight the extreme need for petascale computation to maintain real-time processing.

Damien Gratadour, Eric Gendron
Observatoire de Paris
damien.gratadour@obspm.fr, eric.gendron@obspm.fr

Hatem Ltaief
Extreme Computing Research Center
KAUST
hatem.Ltaief@kaust.edu.sa

MS93

Exascale: the Why and the How

Applications demand scale for resolution, dimension, multi-physics fidelity, isolation from artificial boundaries, parameter inversion, optimal control, uncertainty quantification, and the statistics of ensembles. Are extreme-scale systems able to effectively host state-of-art algorithms for these pursuits, however, or are advances in hardware and algorithms becoming than multiplicative? We examine this question of with respect to algorithmic premiums on uniform-thread concurrency, arithmetic intensity, asynchronicity, and fault-tolerance and from the perspective of recent Gordon Bell prize finalist computations.

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

MS93

From Optimal Algorithms to Fast Petascale Solvers

Scalability beyond petascale requires algorithms with optimal complexity, but optimal algorithms are not automatically fast. The power of current and future supercomputer can only be unleashed, when grid structure, discretization, solver, parallel implementations are carefully co-designed by using predictive performance models exploiting concurrency on all levels. We will describe the design of FE solvers for creeping flow solving a trillion degrees of freedom in around one minute and using up to a million parallel threads.

Björn Gmeiner
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
bjoern.gmeiner@informatik.uni-erlangen.de

Holger Stengel
Erlangen regional Computing Center, Germany
holger.stengel@fau.de

Christian Waluga
M2 Center of Mathematics
Technical University of Munich
waluga@ma.tum.de

Barbara Wohlmuth
M2, Centre for Mathematical Sciences,
Technische Universität München, Germany
wohlmuth@ma.tum.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

MS94

Affirmative Actions in Education Creating Division and Inefficiency Among Beneficiary Groups in India

The paper discusses the social and economic disparity within the group benefitting from affirmative actions in

educational sector of India and also tests this hypothesis-revealing that a major chuck of population is becoming inefficient because of these actions. With the help of mathematical tools it identifies negative cycle present in the society and calculates the economic cost due to this disparity. It concludes with a solution to this complex socio-political issue regarding affirmative actions.

Aprant Ajay
Delhi Techinical University, INDIA
aprant@gmail.com

MS94

System Architecture for a Cooperative Fleet of Autonomous Underwater Vehicles (AUVs)

The Eco-Dolphin system of AUVs has been engineered to deploy for the acquisition of coastal ecological data. Each member of the fleet has its own payload and objective in arriving at a predefined destination. The talk will detail the components of each vehicle, their characteristics, and what requirements are needed for the successful launch, maneuver, and recovery within the systems control center.

Stacey Joseph-Ellison, Qi Zhou
Embry-Riddle Aeronautical University
ellisos1@my.erau.edu, zhouq@my.erau.edu

Zhaoyang Fu
EmbryRiddle Aeronautical University
fuz@my.erau.edu

Zakaria Daud
Embry-Riddle Aeronautical University
daudz@my.erau.edu

Hong Liu
Department of Mathematics
Embry-Riddle Aeronautical University
liuho@erau.edu

MS94

Interfacial Motion by Mean Curvature in Liquid Crystals

Liquid crystals are mesogenic phases of matter between conventional solid and liquid phases. Nematic liquid crystals are anisotropic liquids with directional order. We use the gradient-flow model associated with the Landau-de Gennes free energy to study interfaces in nematic samples at the nematic-isotropic transition temperature, demonstrating that they propagate according to mean curvature flow in some asymptotic limits. We numerically compute dynamically metastable nematic configurations in cylinders and study their delicate dependence on initial conditions.

Amy Spicer
University of Bath, UK
A.Spicer@bath.ac.uk

Apala Majumdar
University of Bath
a.majumdar@bath.ac.uk

MS94

On Strategic Defense in Stochastic Networks

In this paper, we discuss a particular class of stochastic processes used for modeling the accumulation of damage to networks or systems experiencing a sequence of attacks or failures incapacitating random numbers of nodes (e.g. components), each with a random weight (e.g. a cost). Each component has threshold(s) whose crossing indicate the system entering a critical state. An operational calculus strategy is used to derive probabilistic information about the process within random vicinities of passage times.

Ryan White, J. H. Dshalalow
Florida Institute of Technology
ryan@chasnet.com, eugene@fit.edu

MS94

Theory and Computation for Bilinear Quadratures

We present a general framework for constructing numerical integration rules over an arbitrary domain in $\mathbb{R}^d$ that are exact when the integrand is the product of two functions belonging to prescribed subspaces. Such integrals are useful in Galerkin methods, for example. We prove that this framework reduces to Gaussian quadrature for univariate polynomials and trapezoid rule for trigonometric polynomials. We then describe a numerical procedure for constructing these integration rules and present some numerical results.

Christopher Wong
University of California, Berkeley
cawong@math.berkeley.edu

MS94

Survival Probability of Beneficial Mutations in Bacteria

Most novel mutations, even if they confer a benefit to the organism, are ultimately lost during the stochastic population growth process. Their survival probability is sensitive to the organism's life history details and the traits affected by the mutation. We develop a continuous time multitype branching process to predict the survival of initially rare beneficial mutations in bacteria. Predicting bacterial adaptation rates is critical to public health issues such as the evolution of drug resistance.

Anna Zhu, Lindi M. Wahl
Western University Canada
azhu4@uwo.ca, lwahl@uwo.ca

MS95

Asymptotic-Preserving Scheme for the Fokker-Planck-Maxwell System in the Quasi-Neutral Regime

Abstract not available at time of publication.

Stephane Brull
Institut de Mathématiques de Bordeaux UMR 5251
Université Bordeaux
stephane.brull@math.u-bordeaux1.fr

MS95

Solving Kinetic Equations to Model the Core-Collapse Supernova Explosion Mechanism

Core-collapse supernovae (CCSNe) are the explosive deaths of massive stars. CCSN explosions are driven by energy transfer from neutrinos to the stellar fluid. This neutrino

heating occurs under non-equilibrium conditions, and a kinetic description based on the Boltzmann equation is warranted. We describe our recent efforts to develop numerical methods to model neutrino transport in CCSNe with discontinuous Galerkin methods, including ongoing efforts to model neutrino-matter interactions.

Eirik Endeve, Cory Hauck
Oak Ridge National Laboratory
endevee@ornl.gov, hauckc@ornl.gov

Yulong Xing
Department of Mathematics
Univeristy of Tennessee / Oak Ridge National Lab
xingy@math.utk.edu

Tony Mezzacappa
ORNL
mezzacappa@phy.ornl.gov

MS95

Versions of Discontinuous Galerkin Algorithms for Diffusion and for Energy-Conserving Hamiltonian Dynamics

We present mixed continuous/discontinuous Galerkin schemes for solution of a class of kinetic Vlasov-Boltzmann problems in Hamiltonian Poisson-bracket form (plus collisions). These schemes conserve energy, and optionally the L_2 norm, exactly. Application to electromagnetic gyrokinetic problems requires novel extension to avoid the Ampere cancellation problem and significant time step limitations. There are subtle properties of commonly used DG schemes for second order derivatives, such as from diffusion, which we compare with a recovery-based approach.

Ammar Hakim
Princeton Plasma Physics Laboratory
ahakim@pppl.gov

Greg Hammett
Princeton Plasma Physics Laboratory
Princeton University
hammett@princeton.edu

Eric Shi, Ian Abel, Tim Stoltzfus-Dueck
Princeton University
eshi@pppl.gov, iabel@princeton.edu, tstoltzf@princeton.edu

MS95

Numerical Simulation of the Crookes Radiometer

The Crookes radiometer is a small mill enclosed in a glass bulb containing a partial vacuum. Its vanes rotate when exposed to light. This is due to the thermal transpiration, as explained by the kinetic theory of gases. In this talk, a numerical method to make full 3D simulations of this radiometer will be presented. It is based on a discretization of the Boltzmann equation by a cut-cell approach that allows to easily handle moving boundaries.

Guillaume Dechristé
Institut de Mathématiques de Bordeaux
guillaume.dechriste@math.u-bordeaux1.fr

Luc Mieussens
Institut de Mathematiques de Bordeaux
luc.mieussens@math.u-bordeaux1.fr

MS96

Swimming Dynamics of Microorganisms in Viscoelastic Fluids Near a Wall

Microorganisms swimming in viscoelastic fluids are ubiquitous in nature; this includes biofilms grown on surfaces, Helicobacter pylori colonizing in the mucus covering the stomach and spermatozoa swimming through cervical mucus. Previous studies have focused on the locomotion of microorganisms in unbounded viscoelastic fluids. However in many situations, microorganisms interact with solid boundaries. In this work, we numerically study the effect of solid boundaries on the swimming of an archetypal low-Reynolds number swimmer in viscoelastic fluids.

Gaojin Li
University of Notre Dame
gli5@nd.edu

Arezoo Ardekani
Notre Dame University
Department of Aerospace and Mechanical Engineering
ardekani@purdue.edu

MS96

Stabilizing the Collective Motion of Micro-swimmers using Confinement

Concentrated suspensions of swimming microbes and other forms of active matter are known to display intricate, self-organized spatiotemporal patterns on scales larger than those of the individual motile units. The collective dynamics of swimming microorganisms exhibits a complex interplay with the surrounding fluid: the motile cells stir the fluid, which in turn can reorient and advect them. This feedback loop can result in long-range interactions between the cells. We present a computational model that takes into account these cell-fluid interactions and cell-cell forces and that predicts counterintuitive cellular order driven by long-range flows. The predictions are confirmed by experiments with Bacillus Subtilis bacteria which measure the flagella bundle orientation and tracks the cells in the self-organized state. Simulations and experiments show that if the micro-swimmers are confined inside thin cylindrical chambers the suspension self-organizes into a stable swirling vortex. If the micro-swimmers are confined in thin racetracks, a persistent unidirectional stream can be observed. Both these phenomena emerge as a result of the complex interplay between the swimmers, the confining boundaries and the fluid flow. The study highlights the importance of models and simulations of micro-swimmers needing to correctly include both steric and hydrodynamic interactions in oder to capture the dynamics observed in experiments. Collaborators: Hugo Wioland and Raymond Goldstein, DAMTP, University of Cambridge, UK.

Enkeleida Lushi
Brown University
enkeleida_lushi@brown.edu

MS96

Dynamics of Micro-Swimmers Inside a Peristaltic Pump

Peristaltic pumping is a form of fluid transport along the

length of a tube containing liquid when the tube undergoes a contractile wave. While much is known about the peristalsis of Newtonian liquids, complex ones have received limited attention. There are many examples in nature where motile micro-particles or micro-swimmers (such as bacteria or spermatozoa) are suspended in the fluid inside a peristaltic micro-pump. We present a simulation method that accounts for the coupling of the dynamics of many micro-swimmers with each other, the pump, and the fluid flow. The pump and the fluid flow it drives can affect the swimmer dynamics in interesting ways. Moreover, the presence of the swimmers and their collective motion can affect the net transport and mixing in the pump. The efficiency of mixing the suspension for a variety of parameters will be discussed.

Adam Stinchcombe
University of Michigan
stinch@umich.edu

Enkeleida Lushi
Brown University
enkeleida_lushi@brown.edu

MS96

Flagellar Activity Influences Self-Organization in Confined Microswimmer Suspensions

Motile microorganisms are often subject to different types of boundary confinement in their natural environment, but the effects of confinement on their dynamics are poorly understood. We consider a model of microswimmers restricted to move in a 2D Hele-Shaw cell. In an unbounded periodic domain, we show that decreasing flagellar activity induces a hydrodynamically triggered transition in confined microswimmers from turbulentlike swimming to aggregation and clustering. We then impose two different types of boundary confinement: circular and sidewalls confinement and study how additional boundaries influence the emergence of global modes. In the case of circular confinement, the microswimmers can spontaneously organize into a single vortex, reminiscent to what have been observed in recent experiments of bacterial suspensions. In the case of sidewalls confinement, the microswimmers form density shock, via interaction with the sidewalls and background flow.

Alan Cheng Hou Tsang, Eva Kanso
University of Southern California
n/a, kanso@usc.edu

MS97

Threading Mesh Optimization Codes Using Transactional Memory

As threading becomes a necessity on multi- and many-core modern computers, Transactional Memory (TM) has received much attention as a thread synchronization mechanism that promises both coding elegance and efficiency. We study the effects of TM on mesh optimization algorithms using the IBM BG/Q and Intel Haswell platforms. Our preliminary results indicate that these iterative methods are good candidates for TM when the new metric of "time-to-convergence" is used instead of simple wall clock time.

Barna Bihari
Lawrence Livermore National Lab
bihari1@llnl.gov

Lori A. Diachin
Lawrence Livermore National Laboratory
diachin2@llnl.gov

Patrick M. Knupp
Sandia National Laboratories
pknupp@sandia.gov

MS97

An Array-Based Mesh Topological Representation That Effectively Supports General Mesh Modification

A parallel mesh data structure is presented which stores each part using a few arrays. It is flexible enough to allow constant-time insertion and removal of elements,enabling fully general mesh adaptation. Multiple element types defined at compile time can co-exist in one mesh. The memory usage is four times less than object-based structures with the same features. It can also be configured to represent reduced meshes, and various other adjacency-storage schemes.

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

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

MS97

Efficient Unstructured Mesh Traversal Methods Based on Array-Based Half Facets

Mesh-based discretization methods for solving PDE's depend on local mesh traversals for various purposes such as linear system formation, boundary conditions, etc. We present efficient traversal methods for unstructured meshes based on an array-based half-facet data structure. The half-facet representation is generalized from the half-edge/face representations in 2D/3D with support for non-manifold mixed meshes. We present the main constructs of the data structure and analyze several efficient entity traversal schemes for the discrete operator assembly.

Navamita Ray
Argonne National Labratory
nray@mcs.anl.gov

Xinglin Zhao
Stony Brook University
xinglin.zhao@stonybrook.edu

Vijay Mahadevan
Argonne National Laboratory
mahadevan@anl.gov

Xiangmin Jiao
Stony Brook University
xiangmin.jiao@stonybrook.edu

MS97

M3D-C1 Adaptive Loop Going from 2D Axisym-

metric to Full 3D

Tokamak fusion reactors are the experimental approach to study magnetically confined sustainable fusion reactions. The tokamak is a symmetric torus along the toroidal direction and the fusion material exists as the state of plasma in the torus. M3D-C1 simulates non-linear instabilities of plasmas in the tokamak. The on-going effort and results of M3D-C1 adaptive loop in coordination of SCOREC tools to change the simulation mode from axisymmetric to full 3D torus are presented.

E. Seegyoung Seol, Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
seols@rpi.edu, shephard@rpi.edu

Fan Zhang
RPI
fzhang@scorec.rpi.edu

MS98

Comparison of Continuous, Discontinuous and Hybrid Finite Element Methods for Accuracy and Efficiency

We compare and contrast continuous, discontinuous and hybrid (HDG/EDG) finite element methods for the scalar advection-diffusion problem. For each method we examine effectiveness of preconditioners for linear solves, an important consideration in the efficiency of each method as it will affect both convergence rates and overall run time. We also examine accuracy of both discretization and error estimates, the latter an important factor in the effectiveness of solution adaptation methods.

Steven R. Allmaras, Marshall Galbraith
Massachusetts Institute of Technology
allmaras@mit.edu, galbramc@mit.edu

David l. Darmofal
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
darmofal@mit.edu

MS98

Goal-Oriented Curved Mesh Optimization for High-Order Finite-Element Methods

High-order finite element methods generally require curved geometry representations, and hence curved elements, to maintain accuracy. While mesh validity is paramount, we take a further step and present a method for optimizing the shape of the curved elements to improve their approximation capabilities. The method optimizes the reference-to-global high-order mapping and yields curved elements ideally suited for representing a target solution, or for predicting outputs in a goal-oriented setting.

Krzysztof Fidkowski, Devina Sanjaya
University of Michigan
kfid@umich.edu, dsanjaya@umich.edu

MS98

Equivalence between the Energy Stable Flux Reconstruction and Filtered Discontinuous Galerkin Schemes

We demonstrate the equivalence between Discontinuous Galerkin (DG) schemes, in strong and weak form, and the Energy Stable Flux Reconstruction (ESFR) schemes through the application of a generalized filter to the penalized flux terms in the DG formulation. While the essence of the idea, the filtering of the highest mode of the correction functions, was given previously, an elegant extension to higher dimensional non-constant Jacobian formulations in arbitrary bases is presented here.

Philip Zwanenburg, Siva Nadarajah
McGill University
philip.zwanenburg@mail.mcgill.ca,
siva.nadarajah@mcgill.ca

MS98

Theoretical Aspects of High-Order Flux Reconstruction Schemes

In this talk I will discuss recent theoretical advances in the area of high-order Flux Reconstruction (FR) schemes. In particular I will discuss the influence of solution point placement on the stability and accuracy of FR schemes, and I will present a new extended family of energy-stable FR schemes for one-dimensional linear advection problems.

Peter E. Vincent
Department of Aeronautics
Imperial College London
p.vincent@imperial.ac.uk

MS99

The Method of Regularized Stokeslets: Motivation and Applications

Since its introduction in 2001, the method of regularized Stokeslets has been very useful in the simulation of many small scale fluid flows. Biological applications include swimming motions of microorganisms, cell growth, biofilm/fluid interactions, microfiltration for removing particulate matter, flagellar bundling, sperm motility, peristaltic pumping, ciliary motion, and other microscopic phenomena. We will describe the method, its motivation and the flexibility in choosing its components. We will also show applications to flagellar motion.

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

MS99

Bacteria Association with Ciliated Surfaces

Abstract not available at time of publication.

Eva Kanso
University of Southern California
kanso@usc.edu

MS99

Modeling Sperm Motility Using a Kirchhoff Rod Model

Sperm flagella have been observed to propagate different

waves of bending, depending on the fluid environment. In this talk, we will discuss modeling aspects of the relevant fluid environment and chemical concentrations, relating emergent waveforms and interactions to current experiments. The sperm flagellum is represented as a Kirchhoff rod and a regularized Stokes formulation will be used to solve for the local fluid flow. Results will be shown to describe emergent waveforms and swimming speeds.

Sarah D. Olson
Worcester Polytechnic Institute
sdolson@wpi.edu

MS99

The Effects of Rotation and Translation on Flagellar Synchronization

Many flagellated microswimmers exhibit some form of flagellar synchronization. For instance, the bundling and unbundling of flagella in *E. coli* is responsible for their run-and-tumble behavior. In this talk, we look in detail at the mechanisms responsible for phase synchrony in a pair of rigid side-by-side helices. Using an "end-pinned' model, we are able to isolate the effects of translation from those of rotation, finding synchrony in some cases and anti-synchrony in others.

Jonathan H. Tu
UC Berkeley
jhtu@berkeley.edu

Murat Arcak
University of California-Berkeley
arcak@eecs.berkeley.edu

Michel M. Maharbiz
UC Berkeley
maharbiz@eecs.berkeley.edu

MS100

Analysis and Control of Cascading Failures of Power Transmission Systems: a Macro View

Cascading failures of power systems are complex phenomena that are governed by a combination of existing low-level control mechanisms, physics, human input and exogenous factors (such as weather-related factors). In the event of a cascading failure, a goal is to arrest the failure with a minimum of demand lost. We describe ongoing research with a provably good control methodology that relies on a combinatorial algorithm, augmented with tools derived from the SAA (sample average approximation) approach for stochastic optimization. We describe our results in the context of simulations involving several large-scale grids, in particular the Eastern Interconnect.

Daniel Bienstock
Columbia University IEOR and APAM Departments
IEOR Department
dano@columbia.edu

Guy Grebla
Columbia University
guy@ee.columbia.edu

MS100

Visual Analytics for Detection and Tracking of Emergent Subgraphs in Social Networks

Social media networks, and more specifically Twitter retweet networks, are both very large and incredibly noisy. At the same time, analysis of Twitter networks has the potential to have both a predictive capability and detection of emergent trends. In this talk, we combine and co-optimize visual analytics with statistical graph analysis algorithms for detection and tracking of emergent subgraphs in dynamic graphs. A case study focused on movie name hashtag networks will be discussed.

Nadya Bliss, Ross Maciejewski
Arizona State University
nadya.bliss@asu.edu, rmacieje@asu.edu

MS100

Epidemic in Time and Space: Modeling Spatial Outbreak Dynamics

Most epidemic models represent the progression of the contagion process by estimating the number of infectious individuals per time unit. While this is a valid representation, it neglects to include the spatial progression of the epidemic. This presentation will focus on representing the spatio-temporal progression of an epidemic and highlight some of the factors that will effect the spatial path a disease outbreak will follow in the derived contact network.

Armin Mikler
Department of Computer Science & Engineering
University of North Texas
mikler@cs.unt.edu

MS100

Spectral Subgraph Detection in Noisy, Uncertain Networks

Network analysis in practice may involve considering relationships that must be observed through some noisy mechanism. This additional stochasticity compounds the uncertainty due to random connectivity fluctuations of the underlying graph. This presentation discusses various simple mechanisms for noise and uncertainty in network analysis and demonstrates their impact on the ability to detect small anomalous subgraphs using spectral techniques. We also demonstrate fusion of multiple noisy datasets to recover performance achievable on the uncorrupted data.

Benjamin A. Miller
MIT Lincoln Laboratory
bamiller@ll.mit.edu

MS101

Fast Computation of Orthonormal Bases for RBF Native Spaces

We proposed in a recent work the so called *WSVD basis*, which is strictly connected to the eigendecomposition of the integral operator (of Mercer's theorem) and allows one to overcome some problems related to the stability of the computation of the approximant for a wide class of radial kernels. Although effective, this basis is computationally expensive to compute. We discuss here a method to improve and compute in a fast way the basis using methods related to Krylov subspaces. After reviewing the connections between the two bases, we concentrate on the properties of the new one, describing its behavior by various

numerical tests.

Stefano De Marchi
University of Padova (Italy)
demarchi@math.unipd.it

Gabriele Santin
University of Padova
santin.gabriele@gmail.com

MS101

Meshless Vector Field Approximation with Radial Basis Functions

In this talk we will discuss customized radial basis function kernels for vector-field problems. Using shifts of a single scalar-valued kernel and its derivatives, one can easily construct approximation bases that are analytically divergence-free (or curl-free). Such bases lead to meshless methods that are suitable for problems involving electromagnetic fields, fluid flow, or other applications where a divergence-free or curl-free property must be preserved. Our focus will be on approximations designed to naturally "split" into divergence-free and curl-free components. We will present approximation rates of the resulting field-splittings, briefly consider a few applications, and share some numerical observations.

Edward Fuselier
High Point University
Department of Mathematics and Computer Science
efuselie@highpoint.edu

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

MS101

Beyond Quasi-unifomity: Kernel Approximation with a Local Mesh Ratio

Many theoretical results in RBF interpolation assume that data is sampled in a quasi-uniform way. This is somewhat supported by empirical evidence – as centers are badly arranged, the underlying interpolation matrices become ill-conditioned. In this talk we present a relaxed notion of quasi-uniformity by locally controlling the mesh-ratio. Under this type of assumption, and with the aid of recently developed localized bases for RBF approximation, centers may cluster or form gaps without giving rise to instability. At the same time, approximation error is controlled by a local density parameter, meaning that pointwise error responds to the local distribution of data.

Thomas C. Hangelbroek
Department of Mathematics
University of Hawaii at Manoa
hangelbr@math.hawaii.edu

MS101

Oversampling Near the Boundary and Improved Exponential Convergence Rates

We consider the reconstruction of smooth functions from scattered data. If the data are exact, then one theoretically expects exponential convergence rates for many kernel-based reconstruction processes. However, the associated schemes are notoriously ill-conditioned and often show strong boundary effects. We discuss a convergence analysis for a wide class of regularized reconstruction processes by means of sampling inequalities. In particular, we present improved convergence rates if the scattered data points are distributed more densely near the boundary. This is based on joint work with Christian Rieger (University of Bonn).

Barbara Zwicknagl
University of Bonn
zwicknagl@iam.uni-bonn.de

MS102

Parallel, Adaptive, Multi Block Methods for Cartesian Grids using ForestClaw

We will discuss recent multiblock and parallel capabilities in ForestClaw, an adaptive mesh refinement code based on the patched-based multi-rate Berger-Oliger algorithmic strategy coupled with quad/octree mesh refinement. The support for multiblock capabilities now allows for domains composed of quadrilaterals, (with up to four quads meeting at a vertex), each of which is refined as a quadtree. One important example of such a domain is the cubed sphere. We will also show results from parallel computations on the Blue Gene/Q supercomputer JUQUEEN, based in Julich, Germany. ForestClaw uses the dynamic grid management library p4est (C. Burstedde, Univ. of Bonn).

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

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

MS102

Using Explicit Filtering and Reconstruction to Improve Large-Eddy Simulation of the Atmosphere on Adaptive Grids

Large-eddy simulation (LES) and adaptive mesh refinement reduce the computational cost of turbulence modeling by restricting resolved length scales. Combining these techniques generates errors at grid refinement interfaces. This talk explores using the LES formulation to mitigate these errors. Explicitly filtering the advection term and the mixed model are compared to implicit filtering and the eddy viscosity model. Explicitly filtering the advection term reduces interpolation errors, and the mixed model decreases wave reflection.

Lauren Goodfriend, Fotini Katopodes Chow
UC Berkeley
lgoodfriend@berkeley.edu, tinakc@berkeley.edu

Marcos Vanella, Elias Balaras
George Washington University
mvanella@gwu.edu, balaras@gwu.edu

MS102

Local Time Stepping for Parallel Adaptive Mesh Refinement Simulation

In this talk we present a linear, multistep approach to local time stepping for parallel, adaptively refined discontinuous

Galerkin and multiblock finite difference methods. The scheme is efficiently initialized and, following mesh adaptation, restarted using a fixed number of Runge-Kutta steps; the number of Runge-Kutta steps is independent of the number of time levels. The coupled multiphysics problem of earthquake rupture dynamics is used to demonstrate the efficiency and accuracy of the method.

Jeremy E. Kozdon
Stanford University
Geophysics
jekozdon@nps.edu

Lucas Wilcox
Department of Applied Mathematics
Naval Postgraduate School
lwilcox@nps.edu

MS102

Progress in Parallel Adaptive Methods for Storm Surge Forecasting

Today the threat of coastal hazards such as storm surge and tsunamis has become increasingly critical to the sustainability of coastal infrastructure and communities. Leveraging adaptive mesh refinement approaches to these problems has proven to be an effective way to tackle these multi-scale problems, lowering the computational barrier substantially in some cases. In this talk progress in parallelizing these methods will be discussed including recent work on using many-core technologies.

Kyle T. Mandli
University of Texas at Austin
ICES
ktm2132@columbia.edu

MS103

Applying Object-oriented Programming to PDE Solutions

Object-oriented program design is very appealing for design of PDE solvers since many of the ideas fit well with top-down design, abstraction from the underlying discretisation, and high-level code re-use. However, too much data hiding and encapsulation, commonly espoused for object-oriented methods, leads to memory-use redundancy and inefficient codes. We will examine some lessons learned during application of object-oriented design to a spectral elementFourier NavierStokes solver.

Hugh Blackburn
Department of Mechanical and Aerospace Engineering
Monash University
hugh.blackburn@monash.edu

MS103

Architecting Spectral/HP Element Codes for Modern Hardware

Complex finite element software is necessary to meet the demands of some of today's most challenging computational problems. The design of software is therefore more important than ever to ensure correctness of the result and manageability and sustainability of the implementation. We illustrate the approach taken within the Nektar++ spectral/hp element framework to compartmentalise code to align with the mathematical description and produce a software solution which retains performance and platform portability.

Chris Cantwell, David Moxey
Imperial College London
c.cantwell@imperial.ac.uk, d.moxey@imperial.ac.uk

Mike Kirby
University of Utah
School of Computing
kirby@cs.utah.edu

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

MS103

High order DG Methods on Pyramidal Elements

High order time-explicit nodal discontinuous Galerkin (dG) methods have grown in popularity over the past decade for reasons both mathematical and computational in nature. Optimized Lagrange interpolation nodes and sharp trace inequalities with explicit constants allow for explicit expressions for optimal CFL and penalty constants. Finally, the computational structure of dG methods on simplices and hexahedra allows for efficient implementation on accelerators and graphics processing units. In this talk, we present extensions of these aspects of dG methods to high order pyramidal elements.

Jesse Chan
Rice University
jesse.chan@caam.rice.edu

MS103

What Makes Computational Open Source Libraries Successful?

While software is the backbone of scientific computing, we rarely publish information on best practices for writing large-scale, open source scientific software. In this talk, we discuss our observations for success for computational libraries. In particular, we talk about the roles of code, documentation, community, project management, testing, and licenses. The talk is based on our experience maintaining the finite element library deal.II (see www.dealii.org).

Timo Heister
Clemson University
Mathematical Sciences
heister@clemson.edu

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

MS104

Distributed Optimization in Directed Graphs: Push-Sum Based Algorithms

We consider distributed optimization by a collection of nodes, each having access to its own convex function, whose collective goal is to minimize the sum of the functions. The communications between nodes are described by a time-varying sequence of directed graphs, which is uniformly strongly connected. For such communications, as-

suming that every node knows its outdegree, we develop a broadcast-based algorithm, termed the subgradient-push, which steers every node to an optimal value under a standard assumption of subgradient boundedness. The subgradient-push requires no knowledge of either the number of agents or the graph sequence to implement. Our analysis shows that the subgradient-push algorithm converges at a rate of O(ln t/t). The proportionality constant in the convergence rate depends on the initial values at the nodes, the subgradient norms and, more interestingly, on both the speed of the network information diffusion and the imbalances of influence among the nodes.

Angelia Nedich
Industrial and Enterprise Systems Engineering
University of Illinois at Urbana-Champaign
angelia@illinois.edu

Alexander Olshevsky
University of Illinois at Urbana-Champaign
aolshev2@illinois.edu

MS104

Blessing of Scalability: A Tractable Dual Decomposition l-0 Approach for Large Graph Estimation

Estimating the topology of graphical models has been a critical problem in high-dimensional statistics. In large-scale graphs, the prior knowledge can be formulated as a total sparsity budget constraints. This induces a huge non-convex/combinatorial optimization problem involving l-0 constraint. An interesting observation is: as the graph size increases, the associated optimization problem becomes increasingly convex. This motives the use of the distributed dual decomposition method for estimating the graph topology. By relating the duality gap with certain Kullback-Leibler divergence associated with the graph estimation problem, we show that the estimator obtained by dual decomposition achieves asymptotically optimal statistical rate.

Mengdi Wang
Department of Operations Research and Financial Engineering
Princeton University
mengdiw@princeton.edu

MS104

On the O(1/k) Convergence of Asynchronous Distributed Alternating Direction Method of Multipliers

We consider a network of agents that are cooperatively solving a global optimization problem, where the objective function is the sum of privately known local objective functions of the agents and the decision variables are coupled via linear constraints. Recent literature focused on special cases of this formulation and studied their distributed solution through either subgradient based methods with O(1/ v k) rate of convergence (where k is the iteration number) or Alternating Direction Method of Multipliers (ADMM) based methods, which require a synchronous implementation and a globally known order on the agents. In this paper, we present a novel asynchronous ADMM based distributed method for the general formulation and show that it converges at the rate O(1/k).

Ermin Wei
MIT
erminwei@mit.edu

Asuman Ozdaglar
Massachusetts Institute of Technology
asuman@mit.edu

MS104

Distributed Optimization in Undirected Graphs: Gradient and EXTRA Algorithms

Decentralized optimization meets the future needs from mobile computing, self-driving cars, cognitive radios, and collaborative data mining, just to name a few. It allows optimization problems in a self-organizing network to be solved without a central computer. Compared to standard distributed computing with a central controller, a decentralized approach tolerates certain failed nodes or links, has better load balance, and allows each node to keep its data private during the computation. The setting of this talk is the same as the previous talk but assumes an undirected graph, which is more computationally friendly. We analyze existing first-order algorithms and their convergence rates and solution accuracies.

Wotao Yin
Department of Mathematics
University of California at Los Angeles
wotaoyin@math.ucla.edu

Qing Ling, Kun Yuan
University of Science and Technology of China
Department of Automation
qingling@mail.ustc.edu.cn, kunyuan@mail.ustc.edu.cn

MS105

Advanced Discretizations And Multigrid Methods For The Energy Minimization of Liquid Crystal Equilibrium Configurations

Abstract not available at time of publication.

David Emerson
Tufts University
tba@siam.org

MS105

A Parallel Volume Integral Equation Stokes Solver for Flows in Complex Geometries

Abstract not available at time of publication.

Dhairya Malhotra
University of Texas at Austin
tba@siam.org

MS105

Uncertainty Quantification in Control Problems for Flocking Models

Abstract not available at time of publication.

Mattia Zanella
University of Ferrara
tba@siam.org

MS106

Polynomial Approximation of Random PDEs by

Discrete Least Squares

We consider global polynomial approximation of the parameter-to-solution map for PDEs with random input parameters. The polynomial approximation is obtained by a discrete least squares approach based on noise-free pointwise evaluations of the map. We present results concerning the stability and optimality of the method for random as well as low discrepancy sequences of points, and various underlying probability measures and show high dimensional numerical tests supporting the theory.

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

Giovanni Migliorati
CSQI-MATHICSE, EPFL
giovanni.migliorati@epfl.ch

MS106

Combining Sparsity and Smoothness for Function Interpolation

Functions of interest are often smooth and sparse in some sense. Classical linear interpolation methods are effective under strong regularity assumptions, but cannot incorporate nonlinear sparsity structure. At the same time, nonlinear methods such as L1 minimization can reconstruct sparse functions from very few samples, but do not necessarily encourage smoothness. Here we show that weighted L1 minimization effectively merges the two approaches, promoting both sparsity and smoothness in reconstruction. More precisely, we provide specific choices of weights in the L1 objective to achieve rates for functions with coefficient sequences in weighted Lp spaces, $p <= 1$. We consider the implications of these results in particular for the multivariate setting.

Rachel Ward
Department of Mathematics
University of Texas at Austin
rward@math.utexas.edu

Holger Rauhut
Aachen University
rauhut@mathc.rwth-aachen.de

MS106

Quasi-optimal Polynomial Approximation of PDEs with Linear and Nonlinear Stochastic Coefficients

Abstract not available at time of publication.

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

Hoang A. Tran
Oak Ridge National Laboratory
Computer Science and Mathematics Division
tranha@ornl.gov

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

Ronald DeVore
Department of Mathematics
Texas A&M University
rdevore@math.tamu.edu

MS107

Runtime Systems for Fault Tolerant Computing

Since the number of hardware components is expected to grow by one or two orders of magnitude and the circuits' features will also decrease, reliability is expected to be a critical challenge for exascale computing. Specifically, exascale machines are expected to suffer hard faults, i. e., crashes, soft faults, i. e. silent data corruptions, plus performance slowdowns due to unexpected behavior of some hardware components. In this context, a runtime system able to tolerate the latency due to recovery mechanisms by overlapping them with algorithmic computations is required. Such system must dynamically adapt the workload depending on the hardware status and the prevalence of faults. In this talk, we will describe the general features that this kind of runtime software such have.

Marc Casas
Barcelona Supercomputing Center
marc.casas@bsc.es

MS107

Fenix: A Framework for Online Failure Recovery for Scientific Simulations Towards Exascale

In this talk we present Fenix, a framework for enabling recovery from process/node/blade/cabinet failures for MPI-based parallel applications in an online and transparent manner. Fenix provides mechanisms for transparently capturing failures, fixing failed communicators, restoring application state, and returning execution control back to the application. It relies on application-driven, diskless, implicitly coordinated checkpointing. Using the S3D combustion simulation we experimentally demonstrate Fenixs ability to tolerate high failure rates with low overhead.

Marc Gamell
NSF Cloud and Autonomic Computing Center, Rutgers Discovery
Informatics Institute, Rutgers University
mgamell@cac.rutgers.edu

Daniel Katz
Computation Institute, University of Chicago
Argonne National Laboratory, Chicago
d.katz@ieee.org

Keita Teranishi, Hemanth Kolla, Jacqueline Chen
Sandia National Laboratories
knteran@sandia.gov, hnkolla@sandia.gov,
jhchen@sandia.gov

Scott Klasky
Oak Ridge National Laboratory
Oak Ridge, TN
klasky@ornl.gov

Manish Parashar
Rutgers Discovery Informatics Institute

Department of Computer Science
parasparashar@rutgers.edu

MS107

DHARMA: Distributed asyncHronous Adaptive Resilient Management of Applications

DHARMA (Distributed asyncHronous Adaptive Resilient Management of Applications) is a new many-task programming model designed with resilience as a primary concern. It is a data-flow paradigm that emphasizes over-decomposition of work and exploits both task- and data-parallelism. These qualities enable dynamic and adaptable execution; yet introduce significant bookkeeping challenges to adequately address faults. In this talk we present an overview of DHARMAs runtime, discuss key design decisions, and present recent empirical results.

Hemanth Kolla, Janine C. Bennett, Jeremiah Wilke
Sandia National Laboratories
hnkolla@sandia.gov, jcbenne@sandia.gov,
jjwilke@sandia.gov

Nicole Slattengren
Sandia National Labs
nlslatt@sandia.gov

Keita Teranishi, John Floren
Sandia National Laboratories
knteran@sandia.gov, jfflore@sandia.gov

MS107

Understanding the Impact of Transient Faults at the Application Level in HPC

As new generations of microprocessors are created, soft error rates increase as a consequence of technology scaling and reduced energy consumption. Addressing the soft error problem has been identified as key to achieve exascale computing. In this talk, we will present our work on understanding the resilience problem with an application-level perspective: how they affect different components of HPC applications? and what mitigation strategies can be designed to overcome the problem?

Ignacio Laguna
Lawrence Livermore National Laboratory
ilaguna@llnl.gov

MS108

Numerical Solution of PDEs Posed on Graphs

There is currently considerable interest in a class of models (known as Quantum Graphs) which can be described in terms of PDEs posed on large and possibly complex graphs. Such models have found applications in quantum chemistry and solid state physics. The discretization of PDEs posed on graphs using finite element methods and implicit time stepping techniques leads to large, sparse systems of linear equations. This talk will address iterative methods and preconditioning techniques for solving these systems. This is joint work with Mario Arioli.

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

MS108

The Solution of Lyapunov Equations with Nonnormal Coefficients

Applications in control require solution of the Lyapunov equation $AX + XA^T = -BB^T$. When A is stable and B has low rank, the singular values of X often decay rapidly. Existing bounds on those singular values predict slow decay when A is far from normal. In contrast, we show that if the numerical range of A extends far into the right half plane, there must be a large difference between the extreme singular values of X.

Mark Embree
Department of Mathematics
Virginia Tech
embree@vt.edu

Jonathan Baker
Department of Computational and Applied Mathematics
Rice University
jpb7@rice.edu

John Sabino
The Boeing Company
john-paul.n.sabino@boeing.com

MS108

Indefinite Preconditioning of the Coupled Stokes-Darcy System

We propose the use of an indefinite (constraint) preconditioner for the iterative solution of the linear system arising from the finite element discretization of coupled Stokes-Darcy flow. We provide spectral and field-of-value bounds for the preconditioned system which are independent of the underlying mesh size. We present numerical results showing that the indefinite preconditioner outperforms both standard block diagonal and block triangular preconditioners both with respect to iteration count and CPU times.

Scott Ladenheim
Temple University
saladenh@temple.edu

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

Daniel B. Szyld
Temple University
Department of Mathematics
szyld@temple.edu

MS108

Exploiting Tropical Algebra in the Construction of Preconditioners

Recently it has been shown that tropical algebra can be usefully applied in numerical linear algebra, for example to approximate eigenvalues and singular values. The value of tropical algebra is often greater for unstructured matrices with entries that vary widely in magnitude, i.e., problems that are usually considered difficult. In this talk we show

how tropical algebra can be used to understand and design effective incomplete factorization preconditioners for Krylov subspace methods for solving linear systems.

James Hook
School of Mathematics
University of Manchester
james.hook@manchester.ac.uk

Jennifer Pestana
Mathematical Institute, Oxford University, UK
jennifer.pestana@manchester.ac.uk

Francoise Tisseur
University of Manchester
Department of Mathematics
ftisseur@maths.manchester.ac.uk

MS109

Error Decomposition and Adaptivity for Response Surface Approximations with Application to Bayesian Inference

We extend our work on error decomposition and adaptive refinement for response surfaces to focus on the development of a surrogate model for use in Bayesian inference. Estimation and adaptivity are driven by a quantity of interest. The desired tolerance in the error of the posterior distribution is used to establish a threshold for the accuracy of the surrogate model. Particular focus is paid to accurate estimation of evidence to facilitate model selection.

Corey M. Bryant, Serge Prudhomme
ICES
The University of Texas at Austin
corey.m.bryant@gmail.com, serge@ices.utexas.edu

Todd Oliver
PECOS/ICES, The University of Texas at Austin
oliver@ices.utexas.edu

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

MS109

A Scalable Measure-Theoretic Approach to the Stochastic Inverse Problem for Groundwater Contamination

We compute approximate solutions to inverse problems for determining parameters in groundwater contaminant transport models with stochastic data on output quantities. We utilize a measure-theoretic inverse framework to perform uncertainty quantification and estimation for these parameters. The inverse of the map from parameters to data defines a type of generalized contour map. Adjoint problems which are useful in determining a posteriori error estimates are developed and solved numerically.

Steven Mattis
University of Texas at Austin
steve.a.mattis@gmail.com

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

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

MS109

Quantifying Error in An Inadequate Model for Flow in a Porous Media

High fidelity models are prohibitively expensive to solve. For porous media flow, these are the Navier-Stokes equations at the pore-scale. To avoid the cost, we usually make simplifying assumptions leading to a cheaper inadequate model. We first give an overview of a general validation framework for these inadequate models. Then we give an explicit formulation of a model inadequacy for a porous media flow problem and explore how this inadequacy affects posterior uncertainty.

Teresa Portone
ICES
UT Austin
portone@ices.utexas.edu

Damon Mcdougall
Institute for Computational Engineering Sciences
The University of Texas at Austin
damon@ices.utexas.edu

Todd Oliver
PECOS/ICES, The University of Texas at Austin
oliver@ices.utexas.edu

Robert D. Moser
University of Texas at Austin
rmoser@ices.utexas.edu

MS109

A Scalable Computational Framework for Estimating Model Discrepancy

In this talk we will focus on devising a scalable computational framework for the estimation of model discrepancies in uncertainty quantification problems. We will answer questions like how do we include the choice of model selection in our parameter space, how does our data influence model selection, etc. We will answer this in the context of a measure theoretic framework for stochastic inverse problems.

Nishant Panda
Colorado State University
nishant.panda@gmail.com

MS110

Nested Iteration and Adaptive Finite Elements for Ice Sheet Models

This talk describes a First-order System Least-squares (FOSLS) formulation of the nonlinear Stokes flow used to model glaciers and ice sheets. A Nested Iteration (NI), Newton-FOSLS-AMG approach is used in which the majority of the work is done on coarse grids. A reformulation is described that avoids the difficulty of infinite viscosity in regions where ice is undergoing small deformations. Numerical tests are presented demonstrating good performance of the NI-Newton-FOSLS-AMG approach with

adaptive mesh refinement.

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

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

Harihar Rajaram
University of Colorado
Dept. Civil Engineering
hari@colorado.edu

MS110

Parametric Mixed Finite Elements for Two-Phase Flow Interface Problems

Optimal order convergence of a first-order system least squares method using lowest-order Raviart-Thomas elements is presented for domains with curved boundaries. Parametric Raviart-Thomas elements are introduced in order to retain the optimal order of convergence in the higher-order case. In particular, an estimate for the normal flux of the Raviart-Thomas elements on interpolated boundaries is derived. As an application, boundary values of forces are estimated in the Stokes problem.

Fleurianne Bertrand
University of Duisburg-Essen
fleurianne.bertrand@uni-due.de

Gerhard Starke
University of Duisburg-Essen
Fakultät für Mathematik
gerhard.starke@uni-due.de

Steffen Münzenmaier
Gottfried Wilhelm Leibniz University of Hannover
steffen.muenzenmaier@uni-due.de

MS110

Hybrid FOSLS/ll* for Nonlinear Systems of PDEs

This talk presents the extension of Hybrid FOSLS/LL* to nonlinear systems of PDEs. Hybrid FOSLS/LL* combines the best features of FOSLS and FOSLL*. It controls both the Operator Norm and the L^2 norm, approximately minimizing the Graph Norm of the error. It retains a locally sharp and asymptotically accurate *a posterior* error estimator that can be used with nested iteration and adaptive mesh refinement. The efficacy of this method is demonstrated on Navier/Stokes equation and Stokes equation with nonlinear viscosity, as used in ice sheet modeling.

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

Kuo Liu
Department of Applied Mathematics
University of Colorado at Boulder
liukuo99@gmail.com

Lei Tang
Department of Applied Mathematics
University of Colorado-Boulder
l6tang@gmail.com

John Ruge
University of Colorado at Boulder
ruge@colorado.edu

Chad Westphal
Wabash College
westphac@wabash.edu

MS110

A Least Squares Finite Element Method for Coupled Surface/Subsurface Flows

A coupled surface/subsurface flow is presented in this talk. The surface flow is modeled by shallow water equations and the subsurface flow by a generalized Darcy's law for the variably saturated case. Coupling is done along an interface and continuity of pressure and continuity of flux is enforced. A least squares finite element method for the spatial discretization is utilized and numerical experiments using adaptive refinement strategies will be examined.

Steffen Münzenmaier
Gottfried Wilhelm Leibniz University of Hannover
steffen.muenzenmaier@uni-due.de

MS111

Least Squares Shadowing for Adjoint Calculation of Chaotic and Turbulent Pdes

The adjoint method is an invaluable tool for scientific research and engineering design. However, it has been shown that the traditional adjoint method diverges and produces the wrong gradient for chaotic and turbulent PDEs. For these cases, a new approach, Least Squares Shadowing (LSS) has shown promising results. This talk discusses the method and its properties, including sources of error and convergence rates of the adjoint field. Applications including isotropic homogeneous turbulence are also presented.

Patrick Blonigan
MIT
blonigan@mit.edu

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

MS111

Actuator and Sensor Placement for Controlling High-Speed Jet Noise

The loudest source of high-speed jet noise, such as found on naval tactical fighters, appears to be unsteady wavepackets that are acoustically efficient but relatively weak compared to the main jet turbulence. These wavepackets can be usefully described by linear dynamics and connected to transient growth mechanisms. Through a component-wise structural sensitivity analysis of the turbulent jet baseflow, using both the equilibrium and time-average fields, estimates are given as to what location and kind of actuators and sensors are most effective, in a linear feedback context, to control the wavepackets to reduce their noise. Low and high frequency approaches are examined where the controlling mechanisms differ: the low-frequency control indirectly targets the slow variation of the mean on

which the wavepackets propagate while the high-frequency control targets the wavepackets themselves. The predicted control strategy is evaluated using direct numerical simulations on a series of Mach 1.5 turbulent jets.

Mahesh Natarajan, Daniel J. Bodony
University of Illinois at Urbana-Champaign
Department of Aerospace Engineering
nataraj2@illinois.edu, bodony@illinois.edu

MS111

Using Imperfect Outputs and Derivatives in Large-scale Optimization

The outputs of interest in many engineering systems are time averages of chaotic solutions. Relevant examples include the lift and drag on aerodynamic bodies, the energy produced by a fusion reactor, and the (phase-averaged) pressure in an internal-combustion engine. In practice, these outputs must be approximated over a finite integration period, which leads to noise-like errors in the parameter space. In addition, gradients of these outputs must be approximated using ensemble averages or other regularizations. Thus, both the outputs and their derivatives contain errors: they are "imperfect." In this talk we discuss a novel sampling strategy for gradient-based surrogate models that aims to address large-scale optimization problems in the context of imperfect data from computationally expensive simulations.

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

MS111

Parallel Bayesian Optimization of Massively Parallel Turbulent Flow Simulations

Abstract not available at time of publication.

Chaitanya Talnikar
MIT
talnikar@mit.edu

MS112

Software Productivity Application: Integrated Modeling for Fusion Energy

Integrated modeling of plasmas is a key scientific capability for the development of fusion as a power source. We will present our experience developing the Integrated Plasma Simulator (IPS), an environment which enables plasma physicists to be highly productive in developing and carrying out a broad range of coupled simulations, while improving supercomputer resource utilization. The design and implementation of the IPS is quite general, and it is also being used in other domains.

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

Wael R. Elwasif, Donald B. Batchelor
Oak Ridge National Laboratory
elwasifwr@ornl.gov, batchelordb@ornl.gov

MS112

Software Productivity Community Input: Concerns and Priorities

This session will start with an overview and history of the software productivity activities across the scientific computing community, followed by an open discussion of requirements, forums, and opportunities for further engagement.

Jeffrey C. Carver
Department of Computer Science
University of Alabama
carver@cs.ua.edu

Michael Heroux
Sandia National Laboratories
maherou@sandia.gov

MS112

Overview: Software Productivity Challenges for Extreme Scale Science

The Department of Energy has initiated research in software productivity for application development and software infrastructure for extreme-scale scientific computing. The eventual goal is to enable grand challenge science simulations that can survive and even leverage disruptive changes in extreme-scale computer architectures, and thus enable new frontiers in modeling, simulation, and analysis of complex multiscale and multiphysics phenomena. This session will give an overview of DOE activities, including workshops, participating communities, and pilot projects.

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

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

Michael Heroux
Sandia National Laboratories
maherou@sandia.gov

MS112

Software Productivity Challenges in Environmental Applications

Predictive simulations of environmental applications pose significant challenges for multiscale multiphysics frameworks. Problem complexity requires flexibility to compare different models or model features, to add new models, and to explore model coupling. We will present our experience in addressing these issues in two open-source codes, Amanzi and the Arctic Terrestrial Simulator. We will discuss the importance and challenge of leveraging existing frameworks and libraries as well as components of established codes.

David Moulton
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics

moulton@lanl.gov

Ethan T. Coon
Los Alamos National Laboratory
ecoon@lanl.gov

Carl Steefel
Lawrence Berkeley National Laboratory
cisteefel@lbl.gov

MS113

Modeling the Effects of Flow on Anticoagulant Therapy

Warfarin is a common anticoagulant used to treat blood clots. The success of anticoagulation treatment depends on *in vivo* dynamics that cannot be captured clinically. We combine computational models of warfarin dynamics and injury-initiated coagulation to determine the impact of blood flow and platelet dynamics on therapeutic outcomes. We highlight the differences between clinical assessment and *in vivo* measures of warfarin effectiveness. Results indicate that blood flow is a significant determinant of treatment success.

Erica J. Graham
North Carolina State University
Department of Mathematics
ejgraha2@ncsu.edu

Lisette dePillis
Harvey Mudd College
depillis@hmc.edu

Kaitlyn Hood
UCLA
ektuley1@gmail.com

Yanping Ma
Loyola Marymount University
ma@math.psu.edu

Julie Simons
Tulane University
Department of Mathematics
jsimons@tulane.edu

Ami Radunskaya
Pomona College
Mathematics Department
aer04747@pomona.edu

MS113

Navier Slip Condition for Viscous Fluids on a Rough Boundary

We study the effect of surface roughness on fluid flow over a solid surface. We are able to derive asymptotically an effective slip boundary condition (Navier slip condition) as a corrector to the no-slip condition on the surface.

Silvia Jimenez Bolanos
Department of Mathematics
Colgate University
sjimenez@colgate.edu

Bogdan M. Vernescu
Worcester Polytechnic Inst
Dept of Mathematical Sciences
vernescu@wpi.edu

MS113

Taming Targeted Drug Delivery: a Mathematical Model of Triggered Drug Delivery Across the Blood-Brain Barrier

The blood-brain barrier is necessary to protect the brain from microscopic pathogens and toxic molecules. However, this protective mechanism also poses a challenge to the delivery of pharmaceutical agents to the brain. Sonosensitive nanoparticles containing drug show promise as a way to overcome this challenge. Here we propose a schedule for the treatment of disorders of the brain based on a mathematical model of the circulation, triggered release and transport of encapsulated drugs.

Ami Radunskaya
Pomona College
Mathematics Department
aer04747@pomona.edu

MS113

Cooperative Swimming in Viscous Environments

Flagellated organisms exhibit different swimming behaviors, depending on their local environment. We consider the cooperative nature of swimming in populations in viscous environments, in an effort to understand sperm transport. Our approach is a numerical model that relies upon the method of regularized Stokeslets and is robust to three-dimensional effects. We investigate surface interactions as well as measures of cooperativity to elucidate why sperm have such a large variation in behavior across species.

Julie Simons, Lisa J. Fauci
Tulane University
Department of Mathematics
jsimons@tulane.edu, fauci@tulane.edu

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

MS114

A Study of the Entanglement in Polymer Melts

Polymer melts are dense systems of macromolecules. In such dense systems the conformational freedom and motion of a chain is significantly affected by entanglement with other chains which generates obstacles of topological origin to its movement. In this talk we will discuss methods by which one may quantify and extract entanglement information from a polymer melt configuration using tools from knot theory. A classical measure of entanglement is the Gauss linking integral which is an integer topological invariant in the case of pairs of disjoint oriented closed chains in 3-space. For pairs of open chains, we will see that the Gauss linking integral can be applied to calculate an average linking number. In order to measure the entanglement between two oriented closed or open chains in a system with three-dimensional periodic boundary conditions (PBC) we use the Gauss linking number to define the periodic linking number. Using this measure of linking to assess the extend of entanglement in a polymer melt we study the effect of CReTA (Contour Reduction Topolog-

ical Analysis) algorithm on the entanglement of polyethylene chains. Our results show that the new linking measure is consistent for the original and reduced systems.

Eleni Panagiotou
University of California Santa Barbara
lnpanagiotou@yahoo.com

MS114

A Fast Explicit Operator Splitting Method for a Multi-scale Underground Oil Recovery Model

In this talk, we propose a fast splitting method to solve a Multi-scale underground oil recovery model which includes a third order mixed derivatives term resulting from the dynamic effects in the pressure difference between the two phases. The method splits the original equation into two equations, one with flux term and one with diffusion term so that the classical numerical methods can be applied immediately. Two different spatial discretizations, second-order Godunov-type central-upwind scheme and WENO5 scheme, are used to demonstrate that higher order method provides more accurate approximation of solutions. The various numerical examples in both one and two dimensions show that the solutions may have many different saturation profiles depending on the initial conditions, diffusion parameter, and the third-order mixed derivatives parameter. The results are consistant with the study of traveling wave solutions and their bifurcation diagram. This is joint work with C.-Y. Kao, A. Kurganov, Z.-L. Qu.

Ying Wang
University of Oklahoma
wang@ou.edu

MS114

Computational Study of Dynamics and Transport in Vortex-Dipole Flows

A finite number of dipole interactions in free space are studied numerically in order to see how they give rise to a collective fluid flow pattern that is widely seen in ocean currents and clouds. The classical Lamb Dipole is used as our vortex unit. The computation is validated by comparing results with analytic solutions of a free-translating Lamb Dipole. The results have two ingredients. First, the intra- and inner-dipole kinematic pressure fields show distinct features and they relate the phenomenon of no mass exchange in dipoles interactions. Second, a general rule of the ultimate vortical flow pattern is observed based on dipoles interactions in three setups : head on collision, head-end marching, parallel marching.

Ling Xu
Georgia State University
xulingsue@gmail.com

MS114

A Stabilized Explicit Scheme for Coupling Fluid-structure Interactions

We develop a new stabilized explicit coupling partitioned scheme for the fluid-structure interaction problem, where the pressure and velocity are decoupled. Proper penalty terms are applied to control the variations at the interface. Using energy stability analysis, we show that the scheme is stable independent of the fluid-structure density ratio. Numerical examples are provided to show that although the penalty terms degrade the time accuracy, optimal accuracy is recovered by performing defect-correction subiterations.

Yue Yu
Brown University
yuy214@lehigh.edu

MS115

Reaction-diffusion and Electrical Signaling in Neurons (Rdesigneur): a System for Multiscale Modeling in MOOSE

MOOSE, the Multiscale Object-Oriented Simulation Environment, and Rdesigneur support coupled stochastic reaction-diffusion signaling embedded in electrical neuronal models. Rdesigneur coordinates loading of electrical and chemical models using standards like NeuroML and SBML. It then populates the complex neuronal geometry with the chemical models, and defines the interfaces between the electrical and chemical signaling. A set of modular, interchangeable numerical engines plug in to the system to carry out computations at the chosen level of detail.

Upinder Bhalla
National Center for Biological Sciences (NCBS), Bangalore
bahlla@ncbs.res.in

MS115

Interactive, Distributed Spatial Stochastic Simulation with PyURDME and MOLNs

Computational experiments have lead to new biological insights, but the complexity of managing the required distributed computation environments presents a barrier to the adoption. To address this need, we present MOLNs, a cloud computing appliance for distributed simulation of stochastic reaction-diffusion models. The appliance is based on IPython and a newly developed, spatial modeling and simulation package, PyURDME. MOLNs provides an interactive programming platform for development of sharable and reproducible distributed parallel computational experiments.

Brian Drawert
University of California Santa Barbara
bdrawert@cs.ucsb.edu

MS115

Gepetto/OpenWorm

Abstract not available at time of publication.

Stephen Larson
OpenWorm.org, MetaCell
stephen@openworm.org

MS115

Stochastic Simulation at Your Service

We present StochSS: Stochastic Simulation as-a-Service, an integrated development environment for simulation of biological systems with models ranging from ODE to spatial stochastic, where a user can build a simple model on a laptop and scale it up to increasing levels of complexity, deploying computing resources from the cloud with the push of a button when they are needed. StochSS is available for

download at www.stochss.org.

Linda R. Petzold
University of California, Santa Barbara
petzold@engineering.ucsb.edu

MS116

Convergence of Inverse Problems using Reduced Order Models

Inversion requires the repeated solution of expensive forward problems and the computation of Jacobians. In recent work, we have successfully used reduced order models to make nonlinear inversion cheaper. Our approach works well in practice but does not guarantee convergence. I will discuss the solution of a tomography problem using model reduction and an approach to guarantee convergence.

Eric De Sturler
Virginia Tech
sturler@vt.edu

MS116

From Data to Prediction Via Reduced Parameter-to-Observable Maps: Applications to Antarctic Ice Sheet Flow

Here we consider the following question: given a large-scale model containing uncertain parameters, (possibly) noisy observational data, and a prediction quantity of interest, how do we construct efficient and scalable algorithms to (1) infer the model parameters from the data, (2) quantify the uncertainty in the inferred parameters, and (3) propagate the resulting uncertain parameters through the model to issue predictions with quantified uncertainties? We present efficient and scalable algorithms for this end-to-end, data-to-prediction process under the Gaussian approximation and in the context of modeling the flow of the Antarctic ice sheet. We demonstrate that the work required is independent of the parameter and data dimensions. The key to achieving this is to exploit the fact that, despite their large size, the observational data typically provide only sparse information on model parameters.

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Toby Isaac
ICES
The University of Texas at Austin
tisaac@ices.utexas.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS116

Combined State and Parameter Reduction for the Inversion of Functional Neuroimaging Data

The evaluation of functional neuroimaging data such as EEG and fMRI requires the solution of large-scale inverse problems constrained by dynamical systems with high-dimensional state and parameter spaces. To accelerate the inversion, model order reduction can be employed to restrict the optimization to the dominant parameter subspace and the evaluation of the model to the dominant state subspace. With this combined state and parameter reduction of the underlying model, the inversion becomes less computationally costly. Two distinct methods for the combined reduction of states and parameters are presented. First, a gramian-based approach, which uses empirical gramians to reduce the state and parameter dimension based upon the associated controllability and observability. Second, an optimization-based approach that uses a data-driven greedy approach to assemble the reduced order model. Both methods are tested and compared on a dynamic causal model of a neuronal network.

Christian Himpe
Institute for Computational und Applied Mathematics
Universtiy of Muenster
christian.himpe@uni-muenster.de

Mario Ohlberger
Universität Münster
Institut für Numerische und Angewandte Mathematik
mario.ohlberger@uni-muenster.de

MS117

Imaging Biomarkers in Biopharmaceutical Industry

Increasingly biopharmaceutical industry is using imaging for understanding disease and for assessment of therapeutic effects in all stages of drug development. Different imaging modalities can visualize tissue or organ characteristics modulated by disease progression and/or therapy. Imaging enables development of predictive disease-based biomarkers that are translatable between preclinical and clinical stages of drug development by extracting and mining relevant quantitative information from qualitative imaging data in reliable, efficient and objective manner.

Belma Dogdas
Merck
belma.dogdas@merck.com

MS117

Simulation-Based Analysis of Complex Decision Options in Pharmaceutical Research and Development

I will present a multi-method modeling and simulation framework for decision analysis in dynamic resource-constrained situations under uncertainty. The framework was developed for, and applied to, nonlinear problems of strategic planning, portfolio management, asset valuation, process optimization, and what-if scenario analysis in the context of pharmaceutical research and development, and can be applied to a broader class of real world problems.

Otto Ritter
Independent Consultant
otto.ritter@gmail.com

MS117

Imaging Genomics for Pharmaceutical Applica-

tions

Imaging Genomics is an emerging field that integrates imaging and genomic data for identifying high-confidence genesets that relate to a phenotypic response measured from imaging. This talk will explore the use of Imaging genomics for improved diagnosis, patient stratification and assessment of therapeutic response in oncology. Analysis techniques described will encompass the areas of machine learning, image processing, and statistics to extract meaningful information from high-dimensional, multi-modality data for applications in drug discovery and development.

Sangeetha Somayajula, Chandni Valiathan
Merck
sangsom@gmail.com, chandni.valiathan@merck.com;

MS118

Petascale Simulations of Cloud Cavitation Collapse

We present our work on scaling workhorse computational fluid dynamics kernels and complex applications to the petascale domain. We show that by extreme algorithmic and code re-engineering we are able to increase sustained performance of CFD codes from the traditional single digit regime to a LINPACK like 74%. We employed our codes to study cloud cavitation collapse to unprecedented levels of fidelity.

Costas Bekas
IBM Research - Zurich
BEK@zurich.ibm.com

MS118

Petascale Simulation of Hurricane Sandy Using WRF Weather Model on Cray XE6 Blue Waters

The National Center for Atmospheric Research (NCAR) Weather Research and Forecasting (WRF) model has been employed on the largest yet storm prediction model using real data of over 4 billion points to simulate the landfall of Hurricane Sandy. Using an unprecedented 13,680 nodes (437,760 cores) of the Cray XE6 Blue Waters at NCSA at the University of Illinois, researchers achieved a sustained rate of 285 Tflops while simulating an 18-hour forecast.

Pete Johnsen
Cray Inc.
pjj@cray.com

Mark Starka
NCSA
m-straka@illinois.edu

Melvyn Shapiro, Alan Norton, Thomas Galarneau
NCAR
mshapiro@ucar.edu, alan@ucar.edu, tomjr@ucar.edu

MS118

Petascale Medical Simulations

We will present the usage of high performance computing systems for medical applications. The two examples presented are:

- Flow of blood in large arteries with a focus on abdominal aortic aneurysms
- Modelling and simulation of bone tissue using direct numerical simulation

We will present methods to integrate simulation into the medical treatment process covering the workflow from CT or MRI based imaging through simulation to visualization.

Michael Resch
HLRS, University of Stuttgart, Germany
resch@hlrs.de

Ralf Schneider
HLRS
schneider@hlrs.de

MS118

What Are the Priorities Beyond Petascale Computing?

Much of the focus in supercomputing has been on floating-point performance since maximizing FLOP/s and GF/W means minimizing time and energy to solution of simulations. However, does this apply to low-density problems and what are the right metrics for these problems? We discuss a detailed study of energy and time to solution performed on COSMO; a regional model used for climate weather forecasting, and derive effective parameters that should be considered in performance optimization.

Thomas C. Schulthess
Swiss National Supercomputing Center
schulthess@cscs.ch

MS119

Analysis of a Heterogeneous Multiscale Method for Poroelasticity

In this paper, we develop a highly parallelizable numerical method to solve the heterogeneous linear poroelasticity equations in multiple dimensions via operator splitting and a finite-volume based heterogeneous multiscale method for the linear elasticity and reaction diffusion equations. We demonstrate convergence both analytically and numerically, and analyze its computational complexity on high performance computers.

Paul M. Delgado
UTEP
pmdelgado2@utep.edu

Vinod Kumar, Son Young Yi
University of Texas at El Paso
vkumar@utep.edu, syi@utep.edu

MS119

Experimental Analysis of the Performance of GeoClaw, AnuGA and SurfWB-UC Numerical Models for the Simulation of Tsunami Inundation Phenomena

In this study we test and show the performance of three numerical codes for modeling tsunami inundation phenomena (GeoClaw, AnuGA and SurfWB-UC), by experimentally analyzing the numerical convergence to exact analytical solutions of the non-linear shallow water equations, and also, the speed-up and efficiency of their parallel implementation. This, in order to highlight the complexity of representing tsunami wave inundation in the presence of dry-wet

interfaces and shock-waves under variable bathymetry.

José Galaz
Pontificia Universidad Catolica de Chile, CHILE
jdgalaz@uc.cl

Rodrigo Cienfuegos
Hydraulic and Environmental Engineeering Department
Pontificia Universidad Catolica de Chile
racienfu@ing.puc.cl

MS119

Stabilization in Relation to Wavenumber in HDG Methods

We study the Hybrid Discontinuous Galerkin (HDG) method for complex wavenumber cases in acoustics and electromagnetics and show how the HDG stabilization parameter must be chosen in relation to the wavenumber. We show that the commonly chosen HDG stabilization parameter values are not appropriate for all complex wavenumbers, then discover a constraint on the stabilization parameter, dependent on the wavenumber, that guarantees unique solvability of both the global and the local HDG problems. We also perform a dispersion analysis for the Helmholtz case.

Nicole Olivares
Portland State University
nmo@pdx.edu

Jay Gopalakrishnan
Portland state university
gjay@pdx.edu

Liang Li
University of Electronic Science and Technology of China
plum_liliang@uestc.edu.cn

Ronan Perrussel
Université de Toulouse; CNRS;INPT, UPS;
Laboratoire Plasma et Conversion d'Energie (LAPLACE)
perrussel@laplace.univ-tlse.fr

MS119

Solving the Heat Equation with Wavelets

The numerical solution of parabolic time evolution problems such as the heat equation is required in numerous applictions. Solving this problem using the boundary element method (BEM) is an attractive alternative to traditional methods, such as Finite Elements combined with a time-stepping scheme. Using BEM generally leads to full systems, so we combine it with a wavelet method. This leads to sparse system matrices that can be solved in linear complexity.

Anne Reinarz
University of Reading, UK
a.reinarz@pgr.reading.ac.uk

Alexey Chernov
University of Bonn
Hausdorff Center for Mathematics
a.chernov@reading.ac.uk

MS119

Incompressible Flow and (Stabilised) Mixed Finite Element Methods on Highly Stretched Meshes

Anisotropic refinement is an interesting concept to resolve local features of solutions. Unfortunately, the stability of a mixed finite element method may depend on the aspect ratio and other mesh properties caused by the refinement. We show which part of the pressure space is responsible for the deterioration of stability. By imposing a minimal amount of constraints on the pressure, two mixed methods circumventing the behaviour arise. Numerical experiments confirm their stability.

Andreas Wachtel
University of Strathclyde, UK
andreas.wachtel@strath.ac.uk

Mark Ainsworth
Brown University
mark.ainsworth@brown.edu

Gabriel R. Barrenechea
University of Strathclyde
gabriel.barrenechea@strath.ac.uk

MS119

Sublinear Preconditioners for the 2D Helmholtz Equation

We present a new preconditioner for 2D Helmholtz equation in the high frequency regime based on domain decomposition, integral operators and fast algorithms. The preconditioner is designed to be seamlessly integrated in a high performance computing environment. The algorithm separates the computation in two parts, one expensive, but highly parallel, and a second one, sequential but with sublinear complexity; keeping a sub-linear overhead of communication per solve. We will discuss the new algorithm the supporting mathematics and prove that the new method has sub-linear complexity. Finally, we will demonstrate the sub-linear complexity numerically on examples from geophysics.

Leonardo Zepeda-Nñez
Massachusetts Institute of Technology
lzepeda@math.mit.edu

Laurent Demanet
Professor of Mathematics, MIT
laurent@math.mit.edu

MS120

A Fast Conservative Spectral Solver for the Nonlinear Boltzmann Collision Operator

We present a conservative spectral method for the fully nonlinear Boltzmann collision operator based on the weighted convolution structure in Fourier space developed by Gamba and Tharkabhushnanam. The novelty of this approach consists of factorizing the convolution weight on quadrature points by exploiting the symmetric nature of the particle interaction law. This procedure reduces the computational cost and storage of the method to $O(M^2N^4logN)$ from the $O(N^6)$ complexity of the original spectral method, where N is the number of velocity grid

points in each velocity dimension and M is the number of quadrature points in the factorization. We will present numerical results that exhibit the efficiency of this approach.

Jeffrey Haack
University of Texas at Austin
haack@math.utexas.edu

Jingwei Hu
The University of Texas at Austin
hu342@purdue.edu

Irene M. Gamba
Department of Mathematics and ICES
University of Texas
gamba@math.utexas.edu

MS120

Kinetic Equation in a Bounded Domain

Half-space problem is known as the key to understand the boundary layers for kinetic equations (Boltzmann equation, neutron transport equation etc.) in a bounded domain, where sharp transitions from kinetic to fluid regime present. In this talk, I will present a damping-recovery scheme to obtain the solution in a general setting: the equation is modified by a damping term, and a recovery process is proposed. The damped equation is proved to be well-posed, and a spectral method is designed to solve it.

Qin Li
Computing and Mathematical Sciences
CalTech
qinli@cms.caltech.edu

MS120

High-Order Semi-Lagrangian Discontinuous Galerkin Methods for Kinetic Plasma Models

We present a high-order operator split discontuous Galerkin (DG) method for solving the Vlasov-Poisson system. Our hybrid method relaxes strict CFL limitation by using semi-Lagrangian techniques, and we permit complicated geometries in configuration space with unstructured grids. We present 2D-2V results including the formation of a plasma sheath in the proximity of a cylindrical Langmuir probe, as well as the simulation of a single-species charged particle beam in a particle accelerator.

David C. Seal
Department of Mathematics
Michigan State University
seal@math.msu.edu

James A. Rossmanith
Iowa State University
Deparment of Mathematics
rossmani@iastate.edu

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

MS120

A Monte Carlo Method with Negative Particles for General Binary Collisions and Application to Coulomb Collisions

In this work we propose a novel negative particle method for the general bilinear collision operators in the spatial homogeneous case and apply it to the Coulomb collisions. This new method successfully reduces the growth of particle numbers from the numerical time scale to the physical time scale for the Coulomb collisions. We also propose a particle resampling method to reduce the particle number to further improve the efficiency. Various numerical simulations are performed to demonstrate the high accuracy and efficiency.

Bokai Yan
Department of Mathematics
UCLA
byan@math.ucla.edu

Russel Caflisch
Department of Mathematics
University of California, Los Angeles
caflisch@math.ucla.edu

MS121

OCC-Based Meshing for RGG Applications Using MeshKit

High fidelity simulations of physical phenomena described on complex geometries involve efficient generation of optimally resolved computational meshes. We present a completely open-source end-to-end workflow focused on nuclear engineering problems that provide components to describe the geometry representation, parallel mesh generation and in-situ visualization. This is made possible with a OCC-based geometry engine used in combination with a wide array of mesh generation algorithms implemented in MeshKit exposed through a GUI from Kitware called RGGNuclear.

Rajeev Jain
Argonne National Laboratory
jain@mcs.anl.gov

Jacob Becker
Kitware
jacob.becker@kitware.com

Vijay Mahadevan
Argonne National Laboratory
mahadevan@anl.gov

Patrick Shriwise
University of Wisconsin, Madison
shriwise@wisc.edu

Robert O'Bara
Kitware
bob.obara@kitware.com

Iulian Grindeanu
Argonne National Laboratory
iulian@mcs.anl.gov

MS121

High-Order Surface Reconstruction with Applications in Parallel Meshing and Finite Element Solvers

We present a method for reconstructing high-order surfaces

from a given unstructured surface mesh, based on weighted least squares polynomial fitting. The method can achieve third and even higher order accuracy. We present the theoretical framework and compare it with existing methods. We also present its applications in mesh refinement and finite element methods, and show that the meshes adapted or refined using the proposed method preserve the order of convergence of numerical discretizations.

Xiangmin Jiao
Stony Brook University
xiangmin.jiao@stonybrook.edu

Navamita Ray
Argonne National Labratory
nray@mcs.anl.gov

Cao Lu, Xinglin Zhao
Stony Brook University
clu@ams.sunysb.edu, xinglin.zhao@stonybrook.edu

MS121

Parallel Mesh Curving and Adaptation with High-Order Surface Continuity for High-Order Finite Element Simulations

This talk presents work on developing parallel curved meshing techniques for unstructured meshes where high-order surface continuity is maintained for the triangular element faces representing the curved domain surfaces. Optimal nodal placement methods are studied and applied to minimize interpolation error. Mesh modification operations are extended to deal with the complexities involved with adapting high-order curved meshes. Benefits of using curved meshes with high-order continuity are demonstrated in a set of CFD applications.

Qiukai Lu
Rensselaer Polytechnic Institute
luq3@rpi.edu

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

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

MS121

Parallel Meshing Technologies for Large Scale Adaptive Simulations

Scalable parallel simulation workflows require mesh generation and adaptation components to operate in parallel and interact with the analysis code using in-memory interfaces. Recent developments include improvements in scalability, and the ability to adapt anisotropic boundary layer elements in both tangential and normal directions. A procedure to distribute the geometric model to avoid maintaining the entire geometry on each process has been developed. Specific examples of in-memory interfaces with analysis codes will be presented.

Saurabh Tendulkar
Simmetrix Inc.
saurabh@simmetrix.com

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

Rocco Nastasia
Simmetrix Inc.
nastar@simmetrix.com

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

MS122

Modeling Active Flows and Stress Generation in Microtubule-Motor Networks

We describe a multi-scale theory for biologically-inspired soft active materials composed of microtubules crosslinked by molecular motors. Brownian dynamics simulations of microtubules with motile crosslinks reveal that activity-generated extensile stresses arise from both polarity sorting and crosslink relaxation. These simulations yield polarity-dependent active stress coefficients for a Doi-Onsager kinetic theory that captures activity-induced hydrodynamic flows. The model exhibits turbulent-like dynamics, and the continuous generation and annihilation of disclination defects associated with coherent flow structures.

Robert Blackwell, Meredith Betterton, Matthew Glaser
University of Colorado
robert.blackwell@colorado.edu,
meredith.betterton@colorado.edu,
matthew.glaser@colorado.edu

Michael J. Shelley
New York University
Courant Inst of Math Sciences
shelley@cims.nyu.edu

Tony Gao
Courant Institute of Mathematical Sciences
New York University
tgao@cims.nyu.edu

MS122

Computational Models of Cilia and Flagella in a Brinkman Fluid

The interaction between dynamic elastic structures and their surrounding fluid is important for sperm navigation and cilia beating within airways. We study a generalized Euler elastica immersed in a Brinkman fluid, a viscous fluid filled with a network of proteins. Regularized Greens functions for Brinkman flow are used to investigate emergent dynamics with preferred kinematics. Results are presented for swimming speeds, synchronization, and efficiency of flagella with planar waveforms in a Brinkman fluid.

Karin Leiderman
Applied Mathematics
UC Merced
kleiderman@ucmerced.edu

Sarah D. Olson
Worcester Polytechnic Institute

sdolson@wpi.edu

MS122

Fluid Coupling in Continuum Modeling of Microtubule Gliding Assays

Active networks are suspensions of actuated filaments obtained by mixing cytoskeletal filaments and motor protein complexes. We focus on gliding assays, where the molecular motors are anchored to a bottom plate. We present a continuum macroscopic model including the evolution of rigid filaments density, bound and free motors densities and fluid velocity. We focus on cumulative hydrodynamic effects and our numerical simulations show the emergence of ordered subregions.

Tamar Shinar
Computer Sciences
UC Riverside
shinar@cs.ucr.edu

Steven Cook
UC Riverside
scook005@cs.ucr.edu

Christel Hohenegger
University of Utah
Department of Mathematics
choheneg@math.utah.edu

MS122

Accurate Simulations of Complex Fluid Flow in Domains with Smooth Boundaries Using Fft-Based Spectral Methods

We present a method for simulating complex fluid flow in domains with smooth boundaries using simple FFT-based spectral methods to advance the stress evolution equation. Dirichlet conditions for the velocity are imposed by solving Stokes' equations using a novel high-order saddle-point method. Our numerical scheme automatically generates a smooth extension for the velocity field over the non-physical regions of the computational domain, allowing for straightforward coupling with the stress equation. We investigate the convergence of channel flow for an Oldroyd-B fluid to the analytical solution in the limit of vanishing artificial stress-diffusion.

David Stein
University of California Davis
dbstein@math.ucdavis.edu

Becca Thomases
University of California at Davis
thomases@math.ucdavis.edu

MS123

Further Developments in the Flux Reconstruction Method

While the Flux Reconstruction method can recover the nodal DG method, it also allows a wide number of alternatives by different choices of the correction function. The presentation will discuss some alternatives optimized to improve accuracy or to reduce complexity.

Antony Jameson
Professor, Department of Aeronautics & Astronautics
Stanford University
jameson@baboon.stanford.edu

MS123

Spectral Difference Method for Large Eddy Simulation Using Non-Conforming and Sliding Meshes

Recently, we have developed a simple, efficient, and high-order accurate sliding-mesh interface approach to the spectral difference method for 2D CFD simulations. The extension of the sliding interface spectral difference (SSD) method to solving unsteady turbulent flows requires careful verification and validation studies. This abstract reflects two aspects of our research. On the aspect of numerical methods, we report the development of the spectral difference method for a parallel solver of turbulent compressible flows on non-conforming and sliding meshes with all hexahedral elements. On the aspect of verification and validation, we report Large Eddy Simulation results of a turbulent Taylor Couette flow and compare to published Direct Numerical Simulation data.

Bin Zhang, Chunlei Liang
George Washington University
bzh@gwmail.gwu.edu, chliang@gwu.edu

MS123

High-Order Methods for Turbulent Flow Simulations on Deforming Domains

We present new high-order accurate methods for moving domains with large deformations. Unstructured moving meshes are generated by a sequence of entirely local operations. This produces high-quality meshes throughout the simulation, and provides a simple description of the mesh changes between each timestep. Using this information we can construct efficient numerical schemes, and we consider both space-time formulations and ALE/projection based methods. We demonstrate our methods on a range of problems involving complex domain deformations.

Per-Olof Persson
University of California Berkeley
Dept. of Mathematics
persson@berkeley.edu

MS123

On the Utility of High-Order Methods for Unstructured Grids: A Comparison Between PyFR and Industry Standard Tools

Conventional unstructured computational fluid dynamics solvers used by industry typically employ second-order accurate spatial discretizations on CPUs. In this work we assess the potential accuracy and efficiency benefits of high-order unstructured schemes implemented on GPUs, when compared to such industry standard tools. We compare accuracy and efficiency between the two methods for turbulent flow computations performed via direct numerical simulation and large eddy simulation.

Brian C. Vermeire
Imperial College London
b.vermeire@imperial.ac.uk

Freddie Witherden, Peter E. Vincent
Department of Aeronautics
Imperial College London

freddie.witherden08@imperial.ac.uk,
p.vincent@imperial.ac.uk

MS124

Massively Parallel Phase-Field Simulations using HPC Framework waLBerla

We present a massively parallel phasefield code, based on the HPC framework waLBerla, for simulating solidification processes of ternary eutectic systems. Various code optimizations are shown, including buffering strategies, vectorization and load balancing techniques. To reduce the effective domain size, a windowing mechanism is developed, such that only the region around the solidification front has to be simulated. Our simulations are run on SuperMUC, a supercomputer ranked number 12 in the Top500 list, using up to 32768 cores.

Martin Bauer
University of Erlangen-Nuremberg
Department of Computer Science
martin.bauer@fau.de

Harald Koestler
University Erlangen-Nürnberg
harald.koestler@fau.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

MS124

Large Scale and Massive Parallel Phase-field Simulations of Pattern Formations in Ternary Eutectic Systems

Different patterns are forming during directional solidification of ternary eutectics, depending on the physical parameters, with significant influence on the mechanical properties of the material. These patterns are studied with a thermodynamic consistent phase-field model based on the minimization of the grand potential difference, using the massive parallel framework waLBerla. We show structure formation on large scale domains, which give rise to spiral growth and compare them to experiments as well as analytic solutions.

Johannes Hötzer, Marcus Jainta, Philipp Steinmetz, Britta Nestler
Karlsruhe Institute of Technology (KIT)
Institute of Applied Materials (IAM-ZBS)
johannes.hoetzer@kit.edu, marcus.jainta@kit.edu, philipp.steinmetz@kit.edu, britta.nestler@kit.edu

MS124

Multi-Gpu Phase-Field Lattice Boltzmann Simulations for Growth and Moving of Binary Alloy Dendrite

Solidification microstructures are of great significance in materials science and engineering. The melt convection always occurs in casting and greatly affects the solidification microstructures. Meanwhile, the dendrite growth simulation in melt convection needs much computational cost. In this study, we accelerate the simulation by employing the phase-field lattice Boltzmann method and multiple graphics processing unit computational scheme. The dendrite growths accompanying movement and rotation are simulated by the GPU accelerated phase-field lattice Boltzmann scheme.

Tomohiro Takaki
Kyoto Institute of Technology
Mechanical and System Engineering
takaki@kit.ac.jp

Roberto Rojas
Graduate School of Science and Technology
Kyoto Institute of Technology
rcrmroberto_rojas@hotmail.com

Takashi Shimokawabe
Global Scientific Information and Computing Center
Tokyo Institute of Technology
shimokawabe@sim.gsic.titech.ac.jp

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

MS124

Large-scale Multi-Phase-Field Simulation of Abnormal Polycrystalline Grain Growth using TSUBAME2.5 GPU-Supercomputer

The multi-phase-field method has attracted attention as a very promising tool for simulating microstructure evolutions in polycrystalline materials. We developed a multiple-GPU computing technique that facilitate efficient 3D simulations. Using the technique developed, we performed large scale 3D multi-phase-field simulations of abnormal polycrystalline grain growth on a GPU-cluster and on the TSUBAME2.5 supercomputer at the Tokyo Institute of Technology.

Akinori Yamanaka
Tokyo Institute of Technology
Graduate School of Science and Engineering
a-yamana@cc.tuat.ac.jp

Masashi Okamoto
Tokyo University of Agriculture and Technology
50014643014@st.tuat.ac.jp

Takashi Shimokawabe
Global Scientific Information and Computing Center
Tokyo Institute of Technology
shimokawabe@sim.gsic.titech.ac.jp

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

MS125

Discussion on Future Directions of Noisy Networks

The organizers will lead a discussion, with participation from the other speakers and attendees, on the current state of research on analysis of networks with noise and uncertainty, and possible future directions. The discussion will focus on the implications of network noise from a theoretical perspective, and its impact in different applications, including the ones presented. A key point will be the interplay between theory and practice in this emerging subfield

of network science.

Sanjukta Bhowmick
Department of Computer Science
University of Nebraska, Omaha
sbhowmick@unomaha.edu

Benjamin A. Miller
MIT Lincoln Laboratory
bamiller@ll.mit.edu

MS125

Using Consensus to Inform Stochastic Graph Aggregation

Learning an appropriate graph representation from noisy, multi-source data is an area of increasing interest. We present a consensus-based framework to improve the stability and quality of graph representations learned by stochastic graph aggregation techniques. Furthermore, we use metrics of stability to quantify our confidence on which edges represent noise versus structure. We demonstrate the effectiveness of our framework using the Locally Boosted Graph Aggregation algorithm on a variety of synthetic and real datasets.

Layla Oesper
Brown University
layla_oesper@brown.edu

Rajmonda Caceres
Lincoln Laboratory
Massachusetts Institute of Technology
rajmonda.caceres@ll.mit.edu

MS125

Statistical Inference on Errorfully Observed Graphs

For statistical inference on graphs, the existence of edges is frequently based on imperfect assessment. Instead of observing a graph, for each potential edge we observe a "edge feature' which is used to assess the presence of an edge. Moreover, we face a quantity/quality trade-off: the proportion of assessed potential edges decreases with the edge features quality. For the stochastic blockmodel, we derive the optimal quantity/quality operating point for a specific inference task.

Carey Priebe
Johns Hopkins University
cep@jhu.edu

Daniel L. Sussman
Harvard University
dpmcsuss@gmail.com

Minh Tang
Johns Hopkins University
mtang10@jhu.edu

Joshua Vogelstein
Duke University
jo.vo@duke.edu

Vince Lyzinski, Donniell Fishkind, Nam Lee, Youngser Park
Johns Hopkins University
vincelyzinski@gmail.com, fishkind@ams.jhu.edu, nhlee@jhu.edu, youngser@jhu.edu

Avanti Athreya
Johns Hopkins University
Department of Applied Mathematics and Statistics
dathrey1@jhu.edu

MS125

Pockets of Instability in Network Centrality Metrics

We discuss how the sensitivity of centrality metrics alter as increasingly higher percentage of edges are altered in the network. Ideally the metric should change commensurately with the change in network. However, we see that instead of a monotonic change, there are certain pockets of vertices where the centrality values are stable and other pockets where they change significantly. We also discuss how the structure of the network leads to the formation of these pockets.

Vladimir Ufimtsev
University of Nebraska, Omaha
vufimtsev@unomaha.edu

Sanjukta Bhowmick
Department of Computer Science
University of Nebraska, Omaha
sbhowmick@unomaha.edu

MS126

Kernel-Based Image Reconstruction

In image reconstruction a central problem is the approximate inversion of certain integral transforms like the Radon transform, the spherical mean transforms, or in a more general setting the Funk transform. Popular techniques for the approximate reconstruction are kernel-based interpolation methods. This methods are easy to implement and fast algorithms are in many cases available. Nevertheless, the mathematical analysis of such methods is quite involved. We will discuss several problems regarding these procedures and compare some kernel-based methods for image reconstruction from spherical mean value data.

Frank Filbir
Institute of Biomathematics and Biometry
Helmholtz Center Munich
filbir@helmholtz-muenchen.de

MS126

A Numerical Study of the Accuracy of Divergence-Free Kernel Approximations

We present a numerical study of the accuracy of divergence-free radial basis function interpolants in 2D and 3D and preliminary numerical results indicating a polynomial flat limit ($\varepsilon \to 0$). When compared to standard interpolants, our results indicate that using a divergence-free basis improves accuracy of the derivatives in certain directions. In this talk, we explore strategies for improving approximations in these directions and compare accuracy of methods based on radial kernels and multivariate polynomials.

Arthur Mitrano
Arizona State University

arthur.mitrano@asu.edu

MS126

Meshfree Computations with SPH and Vortex Methods

We will provide a broad overview of Smoothed Particle Hydrodynamics (SPH) and Vortex Methods. This is followed by a brief discussion on some of the software tools that are being developed in our group. We will then look at recent comparisons of the vortex method and the SPH for incompressible fluid flow. This is followed by results from recent work on using SPH for gas dynamics, flood simulation, explosions, non-Newtonian fluids and Kelvin-Helmholtz instabilities.

Prabhu Ramachandran
Department of Aerospace Engineering
Indian Institute of Technology, Bombay
prabhu@aero.iitb.ac.in

MS126

Reproducing Kernels in Parametric Partial Differential Equations

In this talk, we address the problem of approximating the solution of a parametric partial differential equation. The number of parameters in the differential equation determines the dimensionality of the reconstruction problem. Without any further information such problems suffer from the curse of dimensionality. The physical model expressed in the partial differential equation, however, allows in many practical situations to identity a smaller set of important parameters. The identification of these important parameters build also the basis for many algorithms from machine learning. We will outline this connection in a specific situation using problem adapted reproducing kernels. Further, we will use sampling inequalities to show deterministic a priori (often exponential) convergence rates of a rather large class of regularized reconstruction schemes. This is partly based on joint work with M. Griebel and B. Zwicknagl (both Bonn University).

Christian Rieger
Universität Bonn
Collaborative Research Centre 1060
rieger@ins.uni-bonn.de

MS127

Recent Developments in Forest-of-octrees AMR

Forest-of-octrees AMR offers both geometric flexibility and parallel scalability and has been used in various finite element codes for the numerical solution of partial differential equations. Low and high order discretizations alike are enabled by parallel node numbering algorithms that encapsulate the semantics of sharing node values between processors. More general applications, such as semi-Lagrangian and patch-based methods, require additional AMR functionalities. In this talk, we present algorithmic concepts essential for recently developed adaptive simulations.

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

Toby Isaac
ICES
The University of Texas at Austin
tisaac@ices.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

MS127

An Implicit, High-Order Accurate, Incompressible Navier-Stokes Solver on Overlapping Grids

This talk describes an implicit, high-order accurate method for incompressible flow combining compact spatial discretizations with approximate factorization schemes and geometric multigrid on overlapping grids. Efficient implicit time discretization is achieved via a second order accurate approximately factored Crank-Nicolson method that incorporates the compact spatial approximations into a sequence of fast banded solves. When used with Overture's high-order accurate matrix-free multigrid for the pressure equation, our method provides an efficient high-resolution solver for LES applications.

Kyle Chand
Lawrence Livermore National Laboratory
chand1@llnl.gov

MS127

Exploring Astrophysical Flows with High-fidelity Large-scale Simulations

We present a systematic study of compressible flow problems found in many astrophysical environments using high-fidelity, large-scale and multiphysics numerical simulations. It covers topics like developing numerical models, designing simulation strategies according to the code scalability, and evaluating simulation results in comparisons with observation results. We will share our user experiences in using codes on DOE's and NSF's leading supercomputers and our interests in improving the efficiency of parallelization and memory usage.

Min Long
Department of Physics
Boise State University
minlong@boisestate.edu

MS127

Runtimes and Autotuning and Hybrid, Oh My! Chombo Navigates the Waters of Exascale

Adaptive Mesh Refinement (AMR) applications require a long-term sustained investment in software infrastructure to create scalable solvers that are capable of utilizing the full capabilities of the largest available HPC platforms. The scalable AMR framework Chombo provides an environment for rapidly assembling portable, high-performance AMR applications for a broad variety of scientific disciplines. In this talk, we present several levels of parallelism that can be exploited by Chombo to achieve scaling beyond petascale. These levels include threading the load handled sequentually by each MPI rank, fine-grain parallelism within the dimensional loops, and instruction-level parallelism models to make use of vector processing within

a larger threading model.

Brian Van Straalen
Lawrence Berkeley National Laboratory
Compuational Research Division
bvstraalen@lbl.gov

MS128

Nek5000: An Environment for Scalable Algorithm Development and Production Simulations

Nek5000 is an open source spectral element code for fluid flow and heat transfer. As Nekton 2.0, it was the first commercially available software for distributed memory computers and, with excellent strong scaling (to beyond a million processes) it is being used by hundreds of researchers in a variety of applications such as combustion, vascular flow modeling, stability analysis, and MHD. We discuss extensibility, scalability, and performance of Nek5000 in this context.

Paul F. Fischer
Argonne National Laboratory
fischer@mcs.anl.gov

MS128

Anisotropic Mesh Adaptation for the Many-core Era

Computing hardware is currently undergoing a rapid transformation from the historical trend of increasingly complex single processing units operating at ever increasing clock frequency, towards high numbers of low power processing units where computing throughput is achieved by increasing concurrency. The consequence for models is that their constituent algorithms must be amenable to parallelization, otherwise the (quasi-)serial section will quickly become the dominant computational cost. In short, algorithms without a high degree of parallelism are potentially facing extinction as the computational environment shifts. Many algorithms that are important for the finite element method, such as matrix-vector multiply, are readily parallelizable on many levels and lots of effort is going into optimal strategies on a range of different types of architectures. However, other important algorithms, such as mesh adaptation, have complex data interdependencies and irregular data access patterns. In such cases the overhead associated with parallelization can easily outweigh the gains. In this talk we describe the challenges, and solutions, for parallelizing anisotropic mesh adaptation. Previous work in this field has mostly been limited to domain decomposition methods implemented using MPI. As more finite element codes are capable of running in a mixed MPI-thread parallel mode, mesh it is necessary for mesh adaptation to also exploit thread level parallelism. We find that domain decomposition methods are not particularly successful for thread parallelism for this class of problem. Instead, a boarder range of parallelization strategies specifically designed for irregular computation must be adopted. Performance analysis shows that fast scalable performance is achieved but there are still room for improvement. We will consider what novel trends in parallel programming models and hardware might create new opportunities

Gerard J Gorman
Department of Earth Science and Engineering
Imperial College London
g.gorman@imperial.ac.uk

MS128

H-to-P Efficiently: a Nektar++ Update on Comparisons of Cg and Hdg

Since the inception of discontinuous Galerkin (DG) methods for elliptic problems, there has existed a question of whether DG methods can be made more computationally efficient than continuous Galerkin (CG) methods. Fewer degrees of freedom, approximation properties for elliptic problems together with the number of optimization techniques, such as static condensation, available within the CG framework made it challenging for DG methods to be competitive until recently. However, with the introduction of a static-condensation-amenable DG method, the hybridizable discontinuous Galerkin (HDG) method, it has become possible to perform a realistic comparison of CG and HDG methods when applied to elliptic problems. In this talk, we focus on embedded manifolds, which are considered a valid approximation for many scientific problems ranging from the shallow water equations to geophysics. We describe a comparison between a CG and an HDG numerical discretiztion in 2D, 3D and of an embedded two-dimensional manifold using high-order spectral/hp elements.

Mike Kirby
University of Utah
School of Computing
kirby@cs.utah.edu

MS128

Heterogeneous Computing with a Homogeneous Codebase

It is likely that finite element codes of the future will need to support a variety of hardware platformsincluding both conventional CPUs and accelerators such as GPUs. However, existing implementation strategies often result in feature disparity and quickly succumb to bit-rot. In this talk I will explain the approach we have taken in PyFR that promotes feature and performance parity across platforms with an emphasis on sustainability.

Freddie Witherden, Peter E. Vincent
Department of Aeronautics
Imperial College London
freddie.witherden08@imperial.ac.uk,
p.vincent@imperial.ac.uk

MS129

A Comparative Analysis of of Asynchronous Many Task Programming Models for Next Generation Platforms

Next generation platform architectures will require a fundamental shift in programming models due to a combination of factors including extreme parallelism, data locality issues (for managing both performance and energy usage), and resilience. The asynchronous, many-task programming model is emerging as a leading new paradigm to address these issues, with many variants of this new model being proposed. This talk surveys some of the leading proposed runtimes, highlighting their key design decisions, strengths, and weaknesses.

Janine C. Bennett
Sandia National Laboratories

jcbenne@sandia.gov

H. Kolla
SNL
hnkolla@sandia.gov

Jeremiah Wilke, Keita Teranishi
Sandia National Laboratories
jjwilke@sandia.gov, knteran@sandia.gov

Nicole Slattengren
Sandia National Labs
nlslatt@sandia.gov

Greg Sjaardema, Samuel Knight
Sandia National Laboratories
gdsjaar@sandia.gov, sknigh@sandia.gov

MS129

Structured Dagger: Supporting Asynchrony with Clarity

Abstract not available at time of publication.

Jonathan Lifflander
University of Illinois
jliff2@illinois.edu

MS129

Using Multiple Dags to Ensure Portability and Scalability in Large Scale Computations Using Uintah

Computational modeling of the hazards posed by thousands of explosive devices during a Deflagration to Detonation Transition (DDT),requires petascale computing resources to resolve the spatial and temporal scales present. The resulting scalable software is now capable of determining how the boosters, which should have just deflagrated, interacted to detonate. Preliminary results of the full scale simulation are shown and the broader scalability lessons for codes at this scale are discussed.

John A. Schmidt
SCI Institute
University of Utah
jas@sci.utah.edu

MS129

A DAG Approach to Tame Complexity in Multiphysics Software on Heterogeneous Architectures

Abstract Directed acyclic graphs (DAGs) provide an effective abstraction to effectively handle complexity arising from both hardware as well as physics. This talk addresses our usage of DAGs to dynamically assemble algorithms for highly complex, multiphysics problems. We also discuss how we are using DAGs to deploy these complex simulations on hybrid architectures such as CPU-GPU systems where overlapping computation with host-device transfers is of critical importance. Finally, we discuss other key aspects of our DAG approach including automated memory reuse, thread-pooling, etc.

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

Abhishek Bagusetty
Chemical Engineering
University of Utah
abhishek.bagusetty@utah.edu

MS130

Weno Finite Volume Methods for Embedded Boundary Grids

We discuss the discretization of hyperbolic conservation laws on Cartesian grids with embedded boundaries. For the regular part of the mesh (i.e., away from the cut cells), we use a high order accurate WENO finite volume method with Runge-Kutta time stepping. The cut cells are updated using an appropriate version of the h-box method.

Christiane Helzel
Department of Mathematics
Ruhr-University-Bochum
christiane.helzel@rub.de

MS130

Diffusion MRI on a Cartesian Grid with Immersed Interfaces

Diffusion MRI measures the diffusion of water in biological tissue. To simulate the diffusion MRI signal, it is important to accurately describe the geometry of the biological cells and cell membranes. We discretize this problem on a Cartesian grid and model cell membranes by interfaces that are not necessarily aligned with the computational grid. This results in a method that correctly accounts for the interface surface area, which has a strong influence on simulation results.

Khieu Van Nguyen
Neurospin, CEA, Saclay, France
nvkhieu89@gmail.com

Jing-Rebecca Li
INRIA Saclay
jingrebecca.li@inria.fr

Luisa Ciobanu
Neurospin, CEA, Saclay, France
luisa.ciobanu@gmail.com

MS130

High-Order Quadrature on Implicitly Defined Domains with Application to a High-Order Embedded Boundary Discontinuous Galerkin Method for Evolving Interface Problems

We present a high-order accurate (order 2p) numerical quadrature algorithm for evaluating integrals on implicitly-defined domains - suitable for, e.g., cut-cell finite element methods in which the mesh is implicitly generated by embedding the domain in a Cartesian grid. The algorithm naturally lends itself to a class of high-order discontinuous Galerkin methods ("tiny" cells are merged with neighbours) - we show some examples ranging from Poisson problems to multiphase incompressible fluid flow.

Robert Saye
Dept. of Mathematics
University of California, Berkeley

rsaye@lbl.gov

MS130

Terrain Following Versus Cut-Cells

A finite volume model is described with the same numerics for cut-cells and terrain-following layers. The model has curl-free pressure gradients which eliminate the horizontal pressure gradient error. Some clean comparisons between cut-cells and terrain-following layers are made and, on tests in which the flow interacts with the orography, the terrain-following layers give better accuracy whereas the cut-cells can excite the computational mode of the Lorenz grid.

Hilary Weller
University of Reading, UK
h.weller@reading.ac.uk

James Shaw
University of Reading
j.shaw@student.reading.ac.uk

MS131

Community feedback and discussion

This session will continue to solicit feedback and discussion from the CSE community on "Future Directions in CSE Education and Research".

Hans De Sterck
University of Waterloo
Applied Mathematics
hdesterck@uwaterloo.ca

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@MIT.EDU

MS131

The Future of CSE Research

This presentation will give an overview of the main points addressed in the draft white paper on "Future Directions in CSE Education and Research". It will comment on the rapid expansion of CSE since the beginning of the 21st century and the challenges the CSE field is encountering in the context of recent disruptive developments that include extreme-scale computing, data-driven discovery, and a comprehensive broadening of the application fields of CSE. The presentation will focus on the future of CSE research and will solicit feedback from the broad CSE community.

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

MS131

The Future of CSE Education

This presentation will continue the discussion of the main ideas presented in the draft white paper on "Future Directions in CSE Education and Research". It will focus on the future of CSE education, including expectations for graduate degree outcomes, workforce development, and changes in educational infrastructure. Feedback will be solicited from the broad CSE community.

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@MIT.EDU

Hans De Sterck
University of Waterloo
Applied Mathematics
hdesterck@uwaterloo.ca

MS132

Topological Sensitivity Analysis in Systems Biology

Mathematical models of natural systems are abstractions of much more complex processes. There are often many potential models consistent with our existing knowledge and experimental data, so it is critical to understand the impact of assumptions inherent to a selected model. Our method evaluates the dependence of inferences on the assumed model structure. Failing to consider this structural uncertainty, as is often done in practice, can give rise to misleading conclusions.

Ann C. Babtie
Theoretical Systems Biology
Imperial College London
ann.babtie12@imperial.ac.uk

Paul Kirk, Michael Stumpf
Imperial College London
Theoretical Systems Biology
paul.kirk06@imperial.ac.uk, m.stumpf@imperial.ac.uk

MS132

Bayesian Updating for Dynamic Systems Using Subset Simulation (Beck) and Active Model Selection (Busetto)

First part by Prof. Beck: A new approximate Bayesian computation algorithm, ABC-SubSim, is presented and applied for Bayesian updating of model parameters. It uses the Subset Simulation algorithm of Au and Beck (2001), a very efficient multi-level MCMC rare-event sampler. Second part by Prof. Busetto: A method for active approximate inference of nonlinear dynamical systems is introduced, and its concrete results are discussed. The method is computationally efficient and provides formal guarantees of near-optimal informativeness.

AlbertoGiovanni Busetto
UCSB
agb@ece.ucsb.edu

James Beck
Division of Engineering and Applied Science
California Institute of Technology
jimbeck@caltech.edu

MS132

Bayesian Inference of Chemical Kinetic Models from Proposed Reactions

We present a new framework for tractable Bayesian inference of chemical kinetic models in case of a large number of model hypotheses generated from a set of proposed reactions. The approach involves imposing point-mass mixture priors over rate constants and exploring the resulting posterior distribution using an adaptive Markov chain Monte Carlo method. We show that further gains in sampling efficiency can be realized by analyzing the chemical network structure in order to reduce the space of MCMC exploration.

Nikhil Galagali, Youssef M. Marzouk
Massachusetts Institute of Technology
nikhilg1@mit.edu, ymarz@mit.edu

MS132

Advanced Bayesian Computation for Challenging Problems in the Sciences and Engineering

Bayesian approaches to Uncertainty Quantification rely on efficient Markov chain Monte Carlo methods especially for models based on large systems of partial differential equations. This talk will present recent work on exploiting (1) Feynman-Kac identities defining the duality between numerical deterministic and stochastic (probabilistic) methods in obtaining solutions of certain classes of partial differential equations, and (2) surrogate geometric structures in defining Markov transition kernels on symplectic manifolds. An illustration with shallow water models for global climate models.

Mark Girolami
University College London
m.girolami@warwick.ac.uk

MS133

Multilevel Simulation of Mean Exit Times

Existing methods achieve ε RMS accuracy for the computation of mean exit times for very general Brownian diffusions, at a cost which is $O(\varepsilon^{-3})$ for a large class of SDEs using the Euler-Maruyama discretisation. We present a new multilevel Monte Carlo method which achieves the same result with a cost which is $O(\varepsilon^{-2}(\log|\varepsilon|)^3)$. This work relies heavily on theoretical results derived by E. Gobet and others, and is supported by numerical experiments.

Michael B. Giles
Mathematical Institute
Oxford University
Mike.Giles@maths.ox.ac.uk

Francisco Bernal
Instituto Superior Tecnico, Portugal
fco_bernal@hotmail.com

MS133

A Multilevel Stochastic Collocation Method for Pdes with Random Inputs

By employing a hierarchy of both spatial approximations and interpolations in stochastic parameter space, we develop a multilevel version of stochastic collocation methods for random partial differential equations, leading to a significant reduction in computational cost. We provide a convergence and cost analysis of the new algorithm, and demonstrate the gains possible on a typical random diffusion model problem.

Aretha Teckentrup
University of Bath
a.l.teckentrup@bath.ac.uk

Peter Jantsch
University of Tennessee
jantsch@math.utk.edu

Max Gunzburger
Florida State University
School for Computational Sciences
mgunzburger@fsu.edu

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

MS133

Hierarchical Acceleration of Stochastic Collocation Methods for PDEs with Random Input Data

We will present an approach to adaptively accelerate a sequence of hierarchical interpolants required by a multilevel sparse grid stochastic collocation (aMLSC) method. Taking advantage of the hierarchical structure, we build new iterates and improved pre-conditioners, at each level, by using the interpolant from the previous level. We also provide rigorous complexity analysis of the fully discrete problem and demonstrate the increased computational efficiency, as well as bounds on the total number of iterations used by the underlying deterministic solver.

Guannan Zhang, Clayton G. Webster
Oak Ridge National Laboratory
zhangg@ornl.gov, webstercg@ornl.gov

MS134

The Cost of Reliability: Iterative Linear Solvers and Reactive Fault Tolerance

We analyze several soft fault models and reactive approaches that may be used given a restarted solver or nested solver. Assuming some portion of the solver requires reliability, we analyze the costs of reliable and unreliable computations. We evaluate these costs given various parameters to the solvers' reactive fault tolerance.

James Elliott
North Carolina State University
jjellio3@ncsu.edu

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

Frank Mueller
North Carolina State University

mueller@cs.ncsu.edu

MS134

Inherent Error Resilience of a Complex Moment-Based Eigensolver

In this talk, we consider error resilience property of a complex moment-based parallel eigensolver for solving generalized eigenvalue problems. We show that the eigensolver with sufficient subspace size can achieve high accuracy for target eigenpairs, even if soft-errors like bit-flip occur in the most time-consuming part of the eigensolver. This property provides an inherent error resilience of the eigensolver that does not require checkpointing and replication techniques.

Akira Imakura, Yasunori Futamura, Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
imakura@cs.tsukuba.ac.jp, futamura@mma.cs.tsukuba.ac.jp, sakurai@cs.tsukuba.ac.jp

MS134

Analysis of Krylov Solver Resilience in the Presence of Soft-Faults

Convergence rates of Krylov iterative solvers are considered in the context of soft silent faults encountered in modern supercomputers. We propose an analytic hardware agnostic fault model based on selective reliability, that can provide rigorous answers to important resilience questions. We prove convergence for a class of Krylov methods, where the original iteration is coupled with periodic restarts (e.g., restarted GMRES). In addition, we derive optimal restart strategy that minimizes the resilience overhead.

Miroslav Stoyanov
Oak Ridge National Lab
mkstoyanov@gmail.com

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

MS134

On the Reliability of Soft Error Detection in CG-POP

Soft errors that are not detected by hardware mechanisms may be extremely complex to detect at the software layer. One option is to perform a full duplication of the computation (and data) and check on a regular basis that intermediate results are consistent. However, this mechanism may be prohibitive. In the context of CG solver, the most prohibitive operation to duplicate is SpMV. To avoid the duplication of this operation, checksum mechanisms may be employed. In this presentation, we investigate the reliability of such an approach in finite precision arithmetic. We illustrate our discussion with the CGPOP code, a miniapp for performing the CG within the Parallel Ocean Program (POP), which is a candidate for exascale climate simulations.

Agullo Emmanuel
Inria
emmanuel.agullo@inria.fr

Luc Giraud
Inria Bordeaux Sud-Ouest
Joint Inria-CERFACS lab on HPC
luc.giraud@inria.fr

Emrullah Fatih Yetkin
Inria
emrullah-fatih.yetkin@inria.fr

MS135

Model Calibration and Error Propagation for Large-Eddy Simulation of Turbulent Flows

We hypothesize that turbulence simulations of engineering applications can be made affordable on the design timescale by improving model calibration and error estimation. This study considers a Bayesian calibration approach followed by forward uncertainty quantification. Model parameters are calibrated with DNS of isotropic turbulence and are then tested in an LES turbulent channel flow configuration. Polynomial Chaos expansions are used for efficient propagation of uncertainties from input model parameters to output quantities of interest.

Myra Blaylock, Cosmin Safta
Sandia National Laboratories
mlblayl@sandia.gov, csafta@sandia.gov

Stefan P. Domino
Sandia. National Laboratories
spdomin@sandia.gov

John C. Hewson, Jeremy Templeton
Sandia National Laboratories
jchewso@sandia.gov, jatempl@sandia.gov

MS135

Predictive Rans Simulations Via Bayesian Model-Scenario Averaging

The turbulence closure model is the dominant source of error in Reynolds Averaged Navier-Stokes simulations, yet no reliable estimators exist for this error component. Here we develop a stochastic, *a posteriori* error estimate, based on variability in model closure coefficients across multiple flow scenarios, for multiple closure models. The variability is estimated using Bayesian Model-Scenario Averaging (BMSA), and used to obtain an stochastic solution estimate in an unmeasured (prediction) scenario.

Richard Dwight
Delft University of Technology
r.p.dwight@tudelft.nl

Wouter Edeling
TU Delft
w.n.edeling@tudelft.nl

Paola Cinnella
ENSAM, ParisTech
paola.cinnella@ensam.eu

MS135

Accounting for Model Error in the Calibration of Physical Models

It is important to account for model error in the fitting of physical models to data. In this talk, we discuss avail-

able Bayesian methods for accounting for model errors, highlighting the calibration of models of *physical* systems. We introduce a Bayesian calibration framework, relying on probabilistic embedding of the error within the model, allowing clear disambiguation of measurement errors and model structural errors. The method is demonstrated in the calibration of chemical kinetic rate parameters.

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

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

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

MS135

Bayesian Model Calibration Techniques That Incorporate Mixed Effects and Model Discrepancy

Measurement errors, model discrepancies, and variability due to differing experimental conditions can produce uncertainty in model parameters estimated through Bayesian model calibration techniques. In many cases, model discrepancies and variability among data sets are neglected during model calibration. However, this can yield non-physical parameter values and produce prediction intervals that are inaccurate in the sense that they do not include the correct percentage of future observations. In this presentation, we discuss techniques to quantify model discrepancies and mixed effects due to multiple data sets in a manner that yields physical parameters and correct prediction intervals. We illustrate aspects of the framework in the context of distributed models with highly nonlinear parameter dependencies.

Brian Williams
Los Alamos National Laboratory
brianw@lanl.gov

Kathleen Schmidt
North Carolina State University
klschmid@ncsu.edu

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

MS136

Resolvent Expansions for Higher-order Simulations of PDEs

By casting the solution to a PDE in terms of pseudo-differential operators, we leverage the use of resolvent expansions, in combination with successive convolution, to arrive at high order time accurate solutions for both linear and nonlinear PDEs. This class of solvers has the advantage that it naturally leads to a line-by-line approach that is well-suited to multi-core GPU computing. We consider several common examples to illustrate our method, such as the 2D Cahn-Hilliard, and the Fitzhugh-Nagumo equations.

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

Matthew Causley
Kettering University
matt.causley@gmail.com

Hana Cho
Michigan State University
chohana@math.msu.edu

MS136

Generalized Structure Additive Runge-Kutta Methods

This presentation discusses a general structure of the additively partitioned Runge-Kutta methods by allowing for different stage values as arguments of different components of the right hand side. An order conditions theory is developed for the new family of generalized additive methods, and stability and monotonicity investigations are carried out. The new family, named GARK, introduces additional flexibility when compared to traditional partitioned Runge-Kutta methods, and therefore offers additional opportunities for the development of flexible solvers for systems with multiple scales, or driven by multiple physical processes.

Adrian Sandu
Virginia Polytechnic Institute and
State University
sandu@cs.vt.edu

Michael Guenther
Bergische Universitaet Wuppertal
guenther@math.uni-wuppertal.de

MS136

Efficient Exponential Integrators: Construction, Analysis and Implementation

We will provide an overview of the latest advances in exponential integrators. In particular, we will discuss the exponential propagation iterative (EPI) methods framework and describe different classes of these integrators such as the unsplit, split, hybrid and implicit-exponential methods. Construction of the exponential methods using both classical and stiff order conditions will be discussed. We will also present the new software package that provides implementation of exponential schemes for serial and parallel computing platforms.

Mayya Tokman
University of California, Merced
School of Natural Sciences
mtokman@ucmerced.edu

MS136

K-Methods, An Extension of Exponential and Rosenbrock Time Integrators

We present a new class of time integration schemes, so called K-methods. We discuss the derivation of order conditions for Rosenbrock-K and exponential-K methods.

These schemes consider the time integration and the approximation of linear system solutions, in the case of Rosenbrock-K, or the approximation of matrix exponential vector products, in the case of exponential-K, as a single computational process. We also give some numerical results showing favorable scalability properties for parallel implementations.

Paul Tranquilli, Adrian Sandu
Virginia Polytechnic Institute and
State University
ptranq@vt.edu, sandu@cs.vt.edu

Ross Glandon
Virginia Tech
rossg42@vt.edu

MS137

Quantity-of-Interest Based Least-Squares Finite Element Methods

We present an approach to augment least-squares finite element formulation with a user-specified quantity of interest (QoI). The method inherits the global approximation properties of the standard least squares formulation with increased resolution of the QoI. We establish theoretical properties such as optimality and enhanced convergence under a set of general assumptions. We also present an adaptive approach that results in approximations which possess high accuracy in global norms as well as in the QoI. Several numerical experiments are presented to support the theory and highlight the effectiveness of our methodology.

Jehanzeb H. Chaudhry
Florida State University
jehanzeb.hameed@gmail.com

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Kuo Liu
Department of Applied Mathematics
University of Colorado at Boulder
liukuo99@gmail.com

Lei Tang
Department of Applied Mathematics
University of Colorado-Boulder
l6tang@gmail.com

MS137

An Energy-Minimization Finite-Element Approach for the Frank-Oseen Model of Nematic Liquid Crystals

We present an energy-minimization finite-element approach to the computational modeling of equilibrium configurations for nematic liquid crystals with applied electric fields. The method targets minimization of system free energy based on the electrically augmented Frank-Oseen free-energy model. We demonstrate the well-posedness of the associated intermediate, discrete, linearization systems. Numerical simulations involving heterogenous constant coefficients for both classical and complicated boundary conditions, relevant in ongoing research, are discussed and support the established theory.

David B. Emerson
Tufts University
Department of Mathematics
david.emerson@tufts.edu

James H. Adler
Tufts University
james.adler@tufts.edu

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

Timothy Atherton
Tufts University
Department of Physics
timothy.atherton@tufts.edu

MS137

Energy Laws and First-Order System Least Squares for MHD systems

Energy principles play a crucial role in understanding the interactions and coupling between different scales or phases in a physical system. This motivates one to investigate how well various discretization methods preserve the energy laws associated with PDEs. We discuss this question in the context of First-Order System Least Squares (FOSLS) finite element method applied to time-dependent heat equation and the equations of magnetohydrodynamics (MHD). Our study involves numerical experiments and some theoretical considerations.

Ilya Lashuk
Georgia Institute of Technology
ilya.lashuk@gmail.com

James H. Adler
Tufts University
james.adler@tufts.edu

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

Ludmil Zikatanov
Pennsylvania State University
ludmil@psu.edu

MS137

Nested Iteration and First-Order System Least Squares for Preconditioning a Two-Fluid Electromagnetic Plasma Model

A two-fluid plasma (TFP) model is presented both as a stand-alone solver and as the preconditioner to a fully im-

plicit particle-in-cell (PIC) simulation. The model couples fluid conservation equations for ions and electrons to Maxwell's equations. A Darwin approximation of Maxwell is used to eliminate spurious light waves. After scaling and modification, the TFP-Darwin model yields a nonlinear, first-order system of equations whose Fréchet derivative is shown to be uniformly H^1-elliptic. This system is addressed numerically by nested iteration (NI) and a First-Order System Least Squares (FOSLS) discretization. Numerical tests demonstrate the efficacy of this approach, yielding an approximate solution within discretization error in a relatively small number of computational work units (WU).

Chris Leibs
University of Colorado Boulder
leibs@colorado.edu

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

MS138

Hierarchically Accelerated Stochastic Collocation for Random PDEs

Stochastic collocation methods for partial differential equations with high-dimensional random inputs generate large collections of linear equations, and solving these linear systems is the dominant cost in the construction of an approximate solution. Interpolation on sequences of nested collocation nodes provides a natural multilevel hierarchy of sample points; thus, in this talk we present an accelerated method which utilizes this hierarchical structure to provide linear solvers with improved initial guesses and strong, cheap preconditioners. For a standard elliptic model problem we derive a priori estimates on the savings and computational cost of constructing an approximate solution using hierarchical acceleration.

Peter Jantsch
University of Tennessee
jantsch@math.utk.edu

Diego Galindo, Clayton G. Webster, Guannan Zhang
Oak Ridge National Laboratory
galindod@ornl.gov, webstercg@ornl.gov, zhangg@ornl.gov

MS138

Parametric Uncertainty Propagation in Resilient Domain Decomposition Methods

One challenging aspect of extreme scale computing concerns combining uncertainty quantification methods with resilient PDE solvers. We present and compare different approaches of propagating parametric uncertainty in resilient domain decomposition methods. In such methods, the PDE is solved on many subdomains whose boundary conditions become the unknowns of a global problem. We illustrate the implementation of the algorithms in light of results obtained for a diffusion equation with an uncertain diffusivity field.

Paul Mycek
Duke University
paul.mycek@duke.edu

Olivier P. Le Maitre
LIMSI-CNRS
olm@limsi.fr

Francesco Rizzi, Khachik Sargsyan, Karla Morris
Sandia National Laboratories
fnrizzi@sandia.gov, ksargsy@sandia.gov,
knmorri@sandia.gov

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

Omar M. Knio
Duke University
omar.knio@duke.edu

MS138

A Multilevel Solution Strategy for the Stochastic Galerkin Method for PDEs with Random Coefficients

The stochastic Galerkin method for solving partial differential equations (PDEs) with random coefficients yields highly accurate numerical solutions yet can be computationally demanding. In this talk, we present a multilevel approach to alleviate some of the prohibitive computational cost in the stochastic Galerkin method. Similar multilevel methods have been successfully applied to Monte Carlo approaches and stochastic collocation methods. We present numerical results for the proposed multilevel method compared to the standard single-level method.

Sarah Osborn
Texas-Tech University
sarah.osborn@ttu.edu

Victoria Howle
Texas Tech
victoria.howle@ttu.edu

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

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

MS138

Exploring Embedded Uncertainty Quantification Methods on Next-Generation Computer Architectures

We explore approaches for improving the performance of uncertainty quantification methods on emerging computational architectures. Our work is motivated by the trend of increasing disparity between floating-point throughput and memory access speed. We describe rearrangements of classical uncertainty propagation methods leading to improved memory access patterns and increased fine-grained parallelism. We then measure the resulting performance improvements on emerging multicore architectures in the context of computing solutions to PDEs with uncertain in-

put data.

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

H. Carter Edwards
Sandia National Laboratories
hcedwar@sandia.gov

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

MS139

Robust Optimization for Decision Making under Uncertainty

We present the derivative free trust-region algorithm NOWPAC for solving nonlinear constrained optimization problems. In this context we address optimization problems that are subject to two sources of uncertainties. First are uncertainties inherent to the constraints and objective function, yielding problems of robust optimization. Second are uncertainties stemming from computational inaccuracies in function evaluations; to detect these situations, we introduce a noise indication tool. We show results for a groundwater flow application.

Florian Augustin, Youssef M. Marzouk
Massachusetts Institute of Technology
fmaugust@mit.edu, ymarz@mit.edu

MS139

Comparison of Laminar Flame Models in the Presence of Uncertainty

In this work, we study one-dimensional and two-dimensional laminar flame models using Bayesian inference. In particular, existing experimental data is used to calibrate, in the Bayesian sense, parameters in chemical kinetics mechanisms used in the laminar flame models. Both kinetics and diffusion models are varied. The evidence is computed and the plausibility of the various models assessed.

Paul Bauman
Mechanical and Aerospace Engineering
University at Buffalo, State University of New York
pbauman@buffalo.edu

MS139

Liposome Vesicles in the Presence of Uncertainty

Liposome vesicles are artificially created vesicles and form a model system for more complicated biological cells, such as red blood cells. To date, models of vesicles have been deterministic in nature. Recent work has demonstrated that thermal fluctuations and the variability of material parameters play an important role in vesicle behavior. This talk will discuss the application of statistical approaches to vesicle simulations.

David Salac
University at Buffalo - SUNY
davidsal@buffalo.edu

MS139

Towards Experimental Design Strategies for Inadequate Models

Obtaining informative measurements is a fundamental problem when inadequate models are used to guide the design of experiments. The focus of this study is to develop a basic understanding of the impact that modeling errors have on experimental design strategies. Through a rigorous modeling of structural errors, new adaptive experimental design strategies can be obtained by exploiting structural uncertainty. The feasibility of the proposed methodology is demonstrated in the context of contaminant dispersion models.

Gabriel Terejanu
University of South Carolina
terejanu@cse.sc.edu

Xiao Lin
University of South Carolina
University of South Carolina
lin65@email.sc.edu

MS140

A Computational Model of Sperm Motility Through Viscoelastic Networks

Elastic polymers and filamentous networks within a fluid environment are ubiquitous. Mammalian sperm, for example, must navigate through the highly heterogeneous environment of layers of viscoelastic networks. We present a discrete model of such a network coupled to a Stokes flow using the method of regularized Stokeslets. The network consists of links made of springs and dashpots (typical viscoelastic elements). The results show the network effects on the swimming patterns of the microorganism.

Jacek Wrobel
Department of Mathematics, Tulane University
Center for Computational Science
jwrobel@tulane.edu

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

Lisa J. Fauci
Tulane University
Department of Mathematics
fauci@tulane.edu

MS140

Mathematical Modeling of Blood Clot Formation Under Flow

Vascular injury triggers two intertwined processes, platelet deposition and coagulation, and can lead to the formation of a blood clot that may grow to occlude a vessel. Formation of the clot involves complex biochemical, biophysical, and biomechanical interactions that are also dynamic and spatially-distributed, and occur on multiple spatial and temporal scales. We previously developed a spatial-temporal mathematical model of these interactions and looked at the interplay between physical factors

(flow, transport to the clot, platelet distribution within the blood) and biochemical ones in determining the growth of the clot. Recently, we extend this model to include reduction of the advection and diffusion of the coagulation proteins in regions of the clot with high platelet number density. The effect of this reduction, in conjunction with limitations on fluid and platelet transport through dense regions of the clot, can be profound. Our results suggest a possible physical mechanism for limiting clot growth.

Karin Leiderman
Applied Mathematics
UC Merced
kleiderman@ucmerced.edu

MS140

Experiment-Driven Surfactant Spreading Models

Surfactants are chemicals that lower surface tensions. They are used in many industrial applications as cleaners or stabilizers, but are also present in biological arenas such as the tear film of the eye and in the lungs. Surfactant spreading models often rely on an equation of state relating surfactant concentration to surface tension. To make mathematical analysis more tractable, models have often employed simple functional relationships. However, to model an experiment with a given fluid and surfactant, a physically meaningful equation of state can be derived from experimentally obtained isotherms. We compare model and experiment for NBD-PC lipid (surfactant) spreading on glycerol for an empirically-determined equation of state, and compare those results to simulations with traditionally employed functional forms. In particular we compare the timescales by tracking the leading edge of surfactant, the central fluid height and dynamics of the Marangoni ridge. We consider both outward spreading of a disk-shaped region of surfactant and the hole-closure problem in which a disk-shaped surfactant-free region self-heals.

Rachel Levy
Harvey Mudd College
levy@hmc.edu

MS140

A Multi-Moment Approach to Modeling the Onset of Vortex Merger

We use a low order model to understand how two corotating vortices transition from a quasi-steady distance from each other to convective merger. Experiments and computations have shown that this rapid phenomena occurs after diffusion causes the vortex core size to exceed some critical fraction of the separation distance. This model was derived from the recently developed Multi-Moment Vortex Method and provides several physical insights as well as pins down what causes the very initial onset of convective merger.

David T. Uminsky
University of San Francisco
Department of Mathematics
duminsky@usfca.edu

MS141

Efficient Deflated-Based Preconditioning for the Communication-Avoiding Conjugate Gradient Method

In this work, we demonstrate that deflation preconditioning can be applied in communication-avoiding Lanczos-based Krylov methods. We derive deflated communication-avoiding CG, which is mathematically equivalent to deflated CG of Saad et al. [SIAM J. Sci. Comput., 21 (2000), pp.1909–1926], but performs asymptotically less communication. Numerical examples and performance modeling reveal complex, problem- and machine-dependent tradeoffs between convergence rate and time per iteration. We discuss applications for which speedups can be obtained using this approach.

Erin C. Carson, Nicholas Knight
UC Berkeley
ecc2z@eecs.berkeley.edu, knight@cs.berkeley.edu

James W. Demmel
University of California
Division of Computer Science
demmel@berkeley.edu

MS141

Enlarged Krylov Subspace Methods for Reducing Communication

In this talk we will describe enlarged Krylov subspace methods. From a partitioning of the input matrix into t subdomains, these methods are based on adding t new vectors to the Krylov subspace at each iteration, instead of one vector in classic methods. The new enlarged search space contains the classical Krylov search space based on the initial residual, and hence the novel methods converge at least as fast as Classical CG in exact precision arithmetic. We will discuss parallel versions that reduce communication, and show that the methods converge at least as fast as Classical CG in exact precision arithmetic. The convergence results show that they also converge faster than CG in finite precision arithmetic.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Sophie Moufawad
LRI - INRIA Saclay, France
sophie.moufawad@inria.fr

Frederic Nataf
Laboratoire J.L. Lions
nataf@ann.jussieu.fr

MS141

Preconditioning Communication-Avoiding Krylov Methods

Krylov subspace projections methods are important for solving large-scale linear system of equations. Recent improvements in a communication avoiding s-step Krylov methods are important to the scalability of this approach in future architectures. A key missing piece of these improvements are robust preconditioners for the s-step Krylov methods. We present a preconditioner framework, based on domain decomposition, to precondition s-step Krylov methods without additional communication. We will present results on a hybrid CPU/GPU cluster.

Siva Rajamanickam
Sandia National Laboratories

srajama@sandia.gov

MS141

Hierarchical and Nested Krylov Methods for Extreme-Scale Computing

We present hierarchical Krylov methods and nested Krylov methods to overcome the scaling difficulties for eigenvalue problems on extreme-scale computers. The work is inspired by our previous work for linear systems and the arising applications. We demonstrate the impact at high core counts on the target applications. Because these algorithms can be activated at runtime via the PETSc library, application codes that employ PETSc can easily experiment with such techniques.

Hong Zhang
MCS, Argonne National Laboratory
hzhang@mcs.anl.gov

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

MS142

Anderson Acceleration: Convergence Theory and Numerical Experience

In this talk I will begin with a description of Anderson acceleration and some motivating applications. This part of the talk should get anyone up-to-speed on the topic of the mini symposium I will state some new convergence and close with a report on numerical experiments.

C.T. Kelley, Alex Toth
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu, artoth@ncsu.edu

MS142

On the Performance of Anderson Acceleration for Multiphysics Problems

Anderson Acceleration (AA) has recently garnered attention as an alternative to Picard and Newton-based nonlinear solution algorithms. This talk will discuss the efficiency and robustness of the method compared to the traditional solution techniques. Examples will be drawn from production simulation codes used for nuclear reactor core simulation, magnetohydrodynamics and ice sheet modeling. We will additionally discuss augmentations to the general algorithm to improve performance.

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Steven Hamilton
ORNL
hamiltonsp@ornl.gov

Mark Berrill
Oak Ridge National Laboratory
berrillma@ornl.gov

Alexander R. Toth
North Carolina State University
artoth@ncsu.edu

C.T. Kelley
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

MS142

Accelerating the EM Algorithm for Mixture-density Estimation

The EM algorithm is widely used for numerically approximating maximum-likelihood estimates in the context of missing information. This talk will focus on the EM algorithm applied to estimating unknown parameters in a (finite) mixture density, i.e., a probability density function (PDF) associated with a statistical population that is a mixture of subpopulations, using "unlabeled" observations on the mixture. In the particular case when the subpopulation PDFs are from common parametric families, the EM algorithm becomes a fixed-point iteration that has a number of appealing properties. However, the convergence of the iterates is only linear and may be unacceptably slow if the subpopulations in the mixture are not "well-separated" in a certain sense. In this talk, we will review the EM algorithm for mixture densities, discuss applying Anderson acceleration to improve the convergence of the iterates, and report on numerical experiments.

Joshua H. Plasse, Homer F. Walker
Worcester Polytechnic Institute
jhplasse@wpi.edu, walker@wpi.edu

MS142

Anderson Acceleration for Parallel Applications

Anderson acceleration has demonstrated significant benefits in accelerating fixed point solutions in a number of applications. The method, however, adds new synchronization points that can slow down its use in parallel. In this presentation, we will examine the parallel communication requirements of Anderson acceleration and discuss its performance for parallel application. In addition, we will discuss use of communication-avoiding ideas within the Anderson acceleration algorithm and show results on model problems as well as a large-scale application. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC.

John Loffeld
University of California
Merced
loffeld1@llnl.gov

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

MS143

Solution of the Full Waveform Inversion Problems

via Projection based Reduced Order Models

Often discrete measurements of transfer functions in the time or frequency domains can be equivalently transformed to projection based ROMs for the underlying PDEs. We use such ROMs for the numerical solution of the inverse hyperbolic problems. Justification of our approach is based on an intriguing connection of the ROMs with the discrete Krein-Marchenko-Gelfand-Levitan method. We show applications to the time-domain full waveform inversion on an example of Marmousi model of seismic exploration in 2D.

Vladimir L. Druskin
Schlumberger-Doll Research
druskin1@slb.com

Alexander V. Mamonov
University of Texas at Austin
mamonov@ices.utexas.edu

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

MS143

An Efficient Output Error Bound for Model Order Reduction of Parametrized Nonlinear Evolution Equations

We present an efficient a posteriori output error bound for model order reduction of parametrized (nonlinear) evolution equations. The error bound successfully avoids the accumulation of the residual in time, which is a common drawback in the existing error estimation for time-stepping schemes. The proposed error bound is applied to two kinds of parametrized instationary problems arising from chromatographic separation processes. Numerical experiments demonstrate the performance and efficiency of the proposed error bound.

Peter Benner, Lihong Feng
Max Planck Institute, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de,
feng@mpi-magdeburg.mpg.de

Yongjin Zhang
Max Planck Institute for Dynamics of Complex Systems,Germany
zhangy@mpi-magdeburg.mpg.de

MS143

Overlapping Clustering and Ldeim in Model Reduction for Nonlinear Inversion

Projection-based parametric model reduction is successfully employed in parameter inversion and optimization. However, efficient reduced model evaluation requires an affine parametrization of the system matrices. When the system matrices do not have this property, Discrete Empirical Interpolation Method (DEIM) can produce an approximate one. In this talk, we combine overlapping clustering algorithms and Local DEIM to generate a high-fidelity affine approximation with minimal on-line cost. A numerical example arising in diffuse optical tomography is presented.

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

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

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

Eric De Sturler
Virginia Tech
sturler@vt.edu

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

MS143

Efficiencies in Global Basis Approximation for Model Order Reduction in Diffuse Optical Tomography

We consider the nonlinear inverse problem of reconstructing parametric images of optical properties from diffuse optical tomographic data. Recent work shows MOR techniques have promise in mitigating the computational bottleneck associated with solving for the parameters. In this talk, we give an algorithm for efficiently computing the approximate global basis needed in MOR by utilizing a new interpretation of the transfer function and by capitalizing on Krylov recycling in a novel way.

Meghan O'Connell
Department of Mathematics
Tufts University
meghan.oconnell@tufts.edu

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

Eric De Sturler
Virginia Tech
sturler@vt.edu

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

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

MS144

Methods for Accurate and Efficient Computation of the Proper-Orthogonal-Decomposition with Large Data Sets

Methods for calculating the proper orthogonal decomposition are investigated with regards to efficiency and accuracy. The classical direct method, the snapshot method, and two new methods, one called "the deflation method" and the other called the "recursive snapshot method", are

compared. The sensitivity of the eigenvalue spectrum and POD modes to round-off errors and errors caused by using a reduced number of snapshots is investigated. Error bounds are given for these error sources.

Brian Helenbrook, Fariddudin Behzad
Clarkson University
helenbrk@clarkson.edu, farid.behzad@gmail.com

MS144

Hierarchical Bayesian Sampling for Image Reconstruction of X-Ray and Proton Radiographs

Reconstructing object densities by pulsing particles through a radially symmetric object and collecting them on a CCD results in an ill-posed Abel inversion problem. We present a Markov Chain Monte Carlo approach for solving Abel inversion, where we not only quantify uncertainty on the image reconstruction, but also quantify uncertainties on the prior image covariance matrix and the precision parameter from the noise model. The data presented were obtained from high-energy X-ray and proton radiography facilities.

Marylesa Howard
National Security Technologies, LLC
howardmm@nv.doe.gov

Michael Fowler
Mathworks
michael.fowler22@gmail.com

Aaron B. Luttman
National Security Technologies, LLC
luttmaab@nv.doe.gov

Margaret Hock
Columbia University
mhock1@slu.edu

MS144

An MCMC Approach to Quantifying Uncertainties in Neutron Tomography

Two of the most important properties characterizing pulsed fusion neutron sources are the time profile and the energy spectrum of the fusion neutrons. In this work we present a model for neutron creation as a function of time and energy, along with a Markov Chain Monte Carlo method for estimating the model parameters from neutron detector measurements. The formulation is demonstrated on real data from a U.S. Department of Energy Dense Plasma Focus fusion reactor.

Aaron B. Luttman
National Security Technologies, LLC
luttmaab@nv.doe.gov

Eric Machorro
National Security Technologies
machorea@nv.doe.gov

Daniel Lowe
National Security Technologies, LLC
lowedr@nv.doe.gov

MS144

Proper Orthogonal Decomposition Based Reduced Order Modeling for Real Time Monte Carlo Simulation

Monte Carlo simulations (MCS) are a powerful tool for modeling radiation transport (RT), but require significant computing resources to obtain accurate results. In this work, we develop a proper orthogonal decomposition (POD) based reduced order modeling (ROM) approach to reduce the number of MC particles that must simulated to obtain statistically significant results. POD typically is done in space, but here we use it to generate orthogonal basis functions to describe the radiation energy spectrum. We apply the POD to generate ROMs for terrestrial radiation detection scenarios and present numerical results to show the improvement in accuracy that can obtained using the ROM.

Indika G. Udagedara, Brian Helenbrook
Clarkson University
udagedig@clarkson, helenbrk@clarkson.edu

Aaron B. Luttman, Stephen Mitchell
National Security Technologies, LLC
luttmaab@nv.doe.gov, mitchese@nv.doe.gov

MS145

An Adaptable, Application-Aware Task-Centric Runtime System

Abstract not available at time of publication.

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

MS145

Task-Based Parallelization of the Fast Multipole Method on NVIDIA GPUs and Multicore Processors

Fast Multipole Methods are a fundamental operation for the simulation of many physical problems. In this talk, we present a new approach for implementing these methods that achieves high performance across many different computer architectures. Our method consists of expressing the FMM algorithm as a task flow and employing a state-of-the-art runtime system, StarPU, to process the tasks on the different computing units.

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

Emmanuel Agullo, Berenger Bramas, Olivier Coulaud
INRIA
emmanuel.agullo@inria.fr, berenger.bramas@inria.fr,
olivier.coulaud@inria.fr

Matthias Messner
Stanford University
burgerbua@gmail.com

Toru Takahashi

Nagoya University
ttaka@nuem.nagoya-u.ac.jp

MS145

Sparse Direct Solvers on Top of a Runtime System

To face the advent of multicore processors and the ever increasing complexity of hardware architectures, programming models based on DAG parallelism regained popularity in the high performance, scientific computing community. Modern runtime systems offer programming models and interfaces that comply with this paradigm and powerful engines for scheduling the tasks into which the application is decomposed. These tools have already proved their effectiveness on a number of dense linear algebra applications. This talk evaluates the usability of runtime systems for sparse matrix multifrontal factorizations which constitute irregular workloads, with tasks of different granularities and characteristics and with a variable memory consumption. Experimental results on real-life matrices show that it is possible to achieve the same efficiency as with an ad hoc scheduler which relies on the knowledge of the algorithm. This talk also shows that thanks to the effectiveness and expressiveness of these programming models, it is possible to implement more complex algorithms that achieve better performance.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Alfredo Buttari
CNRS-IRIT, France
alfredo.buttari@enseeiht.fr

Florent Lopez
Enseeiht
florent.lopez@enseeiht.fr

Abdou Guermouche
LaBRI-INRIA futurs
abdou.guermouche@labri.fr

MS145

A Task-Based Sparse Direct Solver Suited for Large Scale Hierarchical/heterogeneous Architectures

We study the benefits and limits of replacing the highly specialized internal scheduler of our parallel sparse direct solver PaStiX with two generic runtime systems PARSEC and STARPU. The analysis highlights that these generic task-based runtimes achieve comparable results on homogeneous platforms. Furthermore, they are able to significantly speed up the solver on heterogeneous environments by taking advantage of the accelerators while hiding the complexity of their efficient manipulation from the programmer.

Pierre Ramet
LaBRI, Univ Bordeaux, France
701868

MS146

Application of Lvy-Flight Firefly Algorithm in Solving Several Engineering Problems

There are a lot of optimization problems in engineering field. The problems are generally in nonlinear equations, multimodal and having a lot of constrains. To solve the problems we can use metaheuristic algorithm to get solution that is near the optimal solution in reasonable time. In this research, the author solves several engineering problems using Lvy-flights combination with the search strategy via the FIrefly Algorithm that is first developed by XinShe Yang in 2010. Advisor: Kuntjoro Adji Sidarto

Fauziah Andini Putri
Bandung Institute of Technology
fauziahandini16@gmail.com

MS146

Higher Dimensional Smooth Data Interpolation: Algorithmic Techniques from Computational Geometry

In this paper, we derive a construction for computing smooth interpolants for a given dataset of 3+ dimensions, and develop an algorithm for implementing our construction. The algorithm builds an n-dimensional cell complex using Delaunay triangulation, where each cell has an associated interpolation function that satisfies Lipschitz continuity for each internal point. This algorithm is implemented as a MATLAB package, and is the first of its kind that can be used on actual datasets. Advisor: Matthew Hirn, cole Normale Suprieure, Dpartement d'Informatique

Ariel Herbert-Voss
University of Utah
ariel@sci.utah.edu

MS146

A Bioinformatic Approach to Colorectal Cancer Research

My project involved using the R programming language to access metadata concerning colorectal cancer. We used epigenetic data to look at colon cancer in a new way, and came up with two novel pathways for its development, without ever doing a wet lab ourselves. Our results shifted current paradigms about colon cancer development. Advisor: Timothy Yeatman and Mingli Yang, Gibbs Cancer Center

Nicolas Limogiannis, Nick Napier
Wofford College/Gibbs Cancer Center
limogiannisna@email.wofford.edu,
napiernj@email.wofford.edu

MS146

Interpreting Twitter Data from World Cup Tweets

Cluster analysis is a field of data analysis that extracts underlying patterns in data. We clustered 30,000 tweets extracted from Twitter just before the World Cup started and compared the results of k-means, a commonly used clustering algorithm, and Non-Negative Matrix Factorization (NMF). The two algorithms gave similar results, but NMF proved to be faster and provided more easily interpreted results. Advisor: Carl Meyer, NC State

Carol Sadek, Caley Johns
Wofford College
sadekcw@email.wofford.edu, johnsc@email.wofford.edu

MS146

Valuation of American Options and E. Coli Muta-

tions

Numerical methods for the valuation of American options have long been an area of active research. Despite the amount of research in this area, there is not enough attention given to potential applications of the optimized numerical methods in use. Considering the valuation of American options intrinsic coupling with Brownian Motion, a Levy Process, this should be surprising. An area of promise seems to be in biological modeling, particularly E. Coli Mutations. Advisor: Liming Feng, UIUC

James A. Stronz
UIUC
jstronz2@illinois.edu

MS146

The Effects of Chronic Wasting Disease on Pennsylvania Deer and Coyote Populations

Chronic Wasting Disease (CWD) is prevalent in cervids, which can be transmitted directly and indirectly. Thus far, CWD results in death, and no treatments exist. This disease has captured the attention of officials and hunters. In this presentation, we focus on the effects of CWD directly on a Pennsylvania deer population and indirectly on the predatory coyote population. We present a dynamical system describing the relationship between deer and coyotes and some preliminary results. Advisor: Luis Melara, Shippensburg

Brandon D. Thrush
Shippensburg University
bt7856@ship.edu

MS147

Realizability in High-Order Numerical Solutions of Entropy-Based Moment Closures

Entropy-based moment closures for kinetic equations (colloquially known as M_N models) have attractive theoretical properties (hyperbolicity, entropy dissipation, and positivity) but are only defined in the set of realizable moment vectors, that is those which are consistent with a positive distribution. High-order numerical solutions do not always stay in this set, so we investigate the use of a limiter to handle nonrealizable moments in the implementation of a high-order discontinuous Galerkin method.

Graham Alldredge
RWTH AACHEN University
alldredge@mathcces.rwth-aachen.de

MS147

Kinetic Theory Molecular Dynamics

Electrons are weakly coupled in hot, dense matter that is created in high-energy-density (HED) experiments. They are also mildly quantum mechanical and the ions associated with them are classical and may be strongly coupled. In addition, the dynamical evolution of plasmas under these hot, dense matter conditions involve a variety of transport and energy exchange processes. Quantum kinetic theory is an ideal tool for treating the electrons but it is not adequate for treating the ions. Molecular dynamics is perfectly suited to describe the classical, strongly coupled ions but not the electrons. We develop a method that combines a Wigner kinetic treatment of the electrons with classical molecular dynamics for the ions. We refer to this hybrid method as "kinetic theory molecular dynamics,' or KTMD. The purpose of this paper is to derive KTMD from first principles and place it on a firm theoretical foundation. The framework that KTMD provides for simulating plasmas in the hot, dense regime is particularly useful since current computational methods are generally limited by their inability to treat the dynamical quantum evolution of the electronic component. Using the N-body quantum von Neumann equation for the electron-proton plasma, we show how this can be mapped to a classical Liouville equation for the ions coupled to a set of quantum kinetic equations for the 1-particle and 2-particle distribution functions.

Frank Graziani
Lawrence Livermore National Laboratory
graziani1@llnl

MS147

Control Strategies for Multi-Agent Games

We present an optimal control problem for a large system of interacting agents is considered using a kinetic perspective. As a prototype model we analyze a microscopic model of opinion formation under constraints. For this problem a Boltzmann–type equation based on a model predictive control formulation is introduced and discussed. The relation to meanfield is also explored and discussed. For numerical purposes a receding horizon strategy is introduced to embed the minimization of suitable cost functional into binary particle interactions. The corresponding Fokker-Planck asymptotic limit is also derived and explicit expressions of stationary solutions are given. Several numerical results showing the robustness of the present approach are finally reported.

Michael Herty
RWTH Aachen Universtiy
Department of Mathematics
herty@mathc.rwth-aachen.de

MS147

A High-Order / Low-Order Approach to Ocean Modeling

We examine a high-order/low-order approach for the free-surface ocean equations based on an implicit/explicit method. The two dimensional scalar continuity equation is treated implicitly with a preconditioned Jacobian-free Newton-Krylov method (JFNK) and the remaining three dimensional equations are subcycled explicitly within the JFNK residual evaluation with a method. The method is second-order accurate and scales algorithmically, with timesteps much larger than fully explicit methods. Moreover, the hierarchical nature of the algorithm lends itself readily to emerging architectures.

Chris Newman, Geoff Womeldorff, Dana Knoll, Luis Chacon
Los Alamos National Laboratory
cnewman@lanl.gov, womeld@lanl.gov, nol@lanl.gov, chacon@lanl.gov

MS148

Parameterized Reduced-Order Models for Shape Optimization of Flow Domains

The proper orthogonal decomposition is combined with derivatives of flow solutions with respect to geometric parameters to build parametric reduced-order models. The

effectiveness of these models is shown by solving optimal control and shape optimization problems and comparing the results to those obtained with full-order models.

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

MS148

Aeroelastic Design Optimization with Flutter Constraints and Local Rom Interpolation

Developing an efficient and fast algorithm in PDE-constrained optimization is an ongoing active research topic. One attempt is to replace PDE with a reduced order model (ROM). Unfortunately, ROMs are typically prone to parameter change, which is bad in optimization setting where a parameter space has to be explored. One remedy is to construct a database of ROMs and interpolate them for a parameter point that is not in the database. This approach is applied to a real application of aeroelastic wing optimization problem with flutter constraints. Detailed explanation of how to construct a database and sensitivities and numerical results will be illustrated in this talk.

Youngsoo Choi
Farhat Research Group, Stanford University
yc344@stanford.edu

David Amsallem, Charbel Farhat
Stanford University
amsallem@stanford.edu, CFarhat@stanford.edu

MS148

Projection-based ROMs for Parametrized PDE-constrained Optimization and Control Problems

Projection-based ROMs provide efficient strategies to tackle parametric optimization and parametrized control problems, where parameters are related to control/design variables, or to relevant features of the state system, respectively. In this talk we show how to construct projection-based ROMs in order to face the large computational costs arising in these cases, by discussing two general paradigms (optimize-then-reduce vs. reduce-then-optimize), and showing their performances by means of numerical examples.

Andrea Manzoni
EPFL, MATHICSE-CMCS
Switzerland
andrea.manzoni@epfl.ch

MS148

POD-G Reduced Order Models for Prediction and Control of Turbulent Flows

Model reduction has become an active area of scientific and engineering research in the past decade or so due to its ability to reduce the complexity of fluid dynamical systems to enable simulation, optimization design and control. Proper orthogonal decomposition Galerkin (POD-G) is one of the most commonly used methods to generate reduced-order models for turbulent flow systems. In reduced-order modeling, balancing the accuracy and efficiency is crucial for their success. In this talk, we propose an improvement to POD-G models to address this issue. A rigorous error analysis based on POD approximation theory and applications in realistic simulation and control problems will also be discussed.

S.S. Ravindran
University of Alabama in Huntsville
ravinds@uah.edu

MS149

Strategies for Reducing Setup Costs in Algebraic Multigrid

Algebraic multigrid (AMG) preconditioners are often employed in large- scale computer simulations to achieve scalability. AMG construction can be costly, sometimes as much as the solve itself. We discuss strategies for reducing expense through reuse of information from prior solves. The information type depends on the method. For smoothed aggregation AMG this includes aggregation data, whereas for energy minimization it includes sparsity patterns and prior interpolants. We demonstrate the effectiveness of such strategies in parallel applications.

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

Andrey Prokopenko
Sandia National Laboratories
aprokop@sandia.gov

MS149

Next Generation Sparse Symmetric Factorization

Sparse direct solvers are important in the solution of scientific problems but their parallel implementations typically exhibit poor scalability. Reducing synchronization and communication along the critical path is crucial to achieve good performance on future architectures. In this context, we are investigating a parallel implementation of sparse Cholesky factorization using the Fan-Both approach, and we will present our findings.

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

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

MS149

A Distributed CPU-GPU Sparse Direct Solver

We present the hybrid MPI+OpenMP+CUDA implementation of a distributed memory right-looking sparse LU factorization algorithm. The difficulty is that small problem sizes dominate the workload, making efficient GPU utilization challenging. We find ways to aggregate collections of small BLAS operations into larger ones; to schedule operations to achieve load balance and hide long-latency operations, such as PCIe transfer; and to exploit simultaneously all of a nodes available CPU cores and GPUs.

Xiaoye Sherry Li
Computational Research Division
Lawrence Berkeley National Laboratory

xsli@lbl.gov

Piyush Sao, Richard Vuduc
Georgia Institute of Technology
piyush3@gatech.edu, richie@cc.gatech. edu

MS149

New Developments in hypres Interfaces and Solvers

The hypre software library provides high performance preconditioners and solvers for the solution of large sparse linear systems on massively parallel computers via conceptual interfaces, which include a structured, a semi-structured, and a traditional linear-algebra based interface. We will discuss new algorithmic developments in hypre's solvers, as well as our efforts and plans to prepare hypre's interfaces and solvers for heterogeneous architectures.

Ulrike Meier Yang
Lawrence Livermore National Laboratory
yang11@llnl.gov

MS150

An Implicit Les Strategy for High Order Discontinuous Galerkin Discretizations

Due to their high scale-resolving capabilities per degree of freedom, high order Discontinuous Galerkin methods have been shown to be very accurate and effective for implicit LES simulations at moderate Reynolds number, if combined with a polynomial de-aliasing strategy for stability. In this work, we present an implicit LES strategy suitable for high order spectral element methods based on a locally and temporally adaptive de-aliasing, that ensures stability through physics-based indicators and models dissipation.

Andrea D. Beck, David Flad
University of Stuttgart
beck@iag.uni-stuttgart.de, flad@iag.uni-stuttgart.de

Claus-Dieter Munz
Institut für Aerodynamik und Gasdynamik (IAG)
munz@iag.uni-stuttgart.de

MS150

High-Order Finite-Volume Solution of Turbulent Aerodynamic Flows

Research in high-order unstructured solvers is encouraged by their accuracy advantages and flexibility for complex geometries. Traditionally, the finite-volume approach has been employed in computational aerodynamics because of intrinsic conservative properties. A major step towards the adoption of higher-order finite-volume methods in CFD solvers is turbulence modeling. We describe the integration of the Spalart-Allmaras model into a higher-order finite-volume solver and the treatments required to solve such a mathematically stiff problem on anisotropic meshes.

Alireza Jalali
PhD Candidate, University of British Columbia
arjalali@interchange.ubc.ca

Carl Ollivier-Gooch
University of British Columbia
cfog@mech.ubc.ca

MS150

A Theoretical and Computational Framework for Measure-Valued Solutions to Conservation Laws

Recent results cast doubts on the appropriateness of the entropy weak solution for nonlinear systems of conservation laws and it has been conjectured that the more general entropy measure-valued (emv) solutions might be the appropriate notion of solution. We proved that bounded solutions of an arbitrary high order space-time DG scheme combined with a nonlinear shock-capturing converge to an emv solution. The novelty in our work is that no streamline-diffusion terms are used for stabilization.

Zakerzadeh Mohammad
RWTH Aachen University
zakerzadeh@aices.rwth-aachen.de

Georg May
AICES
RWTH Aachen
may@aices.rwth-aachen.de

MS150

Understanding the Role of Spectral Vanishing Viscosity in High Reynolds Number Flows

To stabilize high Reynolds number simulations using Continuous Galerkin spectral/hp element discretisations stabilization is required, even after the application of dealiaising techniques, although apparently not in Discontinous Galerkin (DG) discretisations. For our complex geometry high Reynolds number simulations we are applying Spectral Vanishing Viscosity (SVV) for stabilization. In this presentation we discuss the dispersion analysis of the SVV approach and demonstrate that appropriate choices of the SVV parameters recover characteristics observed in DG methods.

Rodrigo Moura, Jean-Eloi Lombard, David Moxey, Yan Bao, Spencer Sherwin
Imperial College London
r.moura13@imperial.ac.uk, jean-eloi.lombard12@imperial.ac.uk, d.moxey@imperial.ac.uk, y.bao@imperial.ac.uk, s.sherwin@imperial.ac.uk

MS151

Resolving Uncomfortable Tradeoffs in Building Fast Boundary-Element Method Solvers: It's Not the Hows, It's the Whys

Fast-multipole methods (FMMs) keep pace with parallel computing, so where are the scalable, open-source fast boundary-element method (BEM) libraries? We suggest that a general BEM requires suitable strategies to account for the additional abstraction layer (a fast summation algorithm) between elementwise computations and the Krylov iteration. This perspective suggests an extensible architecture for general geometries, basis functions, collocation/Galerkin discretizations, scalar and vector problems, and interfaces to "downstream" applications, e.g. PDE-constrained optimization and coupling to FEM.

Jaydeep Bardhan
Northeastern University
j.bardhan@neu.edu

Matthew G. Knepley
University of Chicago
knepley@ci.uchicago.edu

MS151

BEM++ - Building Blocks for Galerkin Boundary Element Methods

The BEM++ boundary element library (www.bempp.org) is a versatile framework for the solution of boundary integral equations via Galerkin boundary element methods. The focus on this talk is on the design decisions and challenges of developing BEM++. This includes basic mesh representation, assembly of operators, FMM and H-matrix methods, preconditioning, linear algebra, and user interfaces in C++ and Python.

Timo Betcke, Simon Arridge, Elwin van't Wout
University College London
t.betcke@ucl.ac.uk, S.Arridge@cs.ucl.ac.uk, e.wout@ucl.ac.uk

MS151

Applications of Accelerated BEM in Aeronautics

Airbus Group Innovations is, inside Airbus Group, an entity devoted to research and development for the usage of Airbus Group divisions (Airbus Civil Aircraft, Airbus Defence & Space, Airbus Helicopters). The numerical analysis team has been working for now more than 20 years on integral equations and boundary element methods for wave propagation simulations, first in electromagnetism, later in acoustics and electrostatics. Since 2000, these BEM tools have received a multipole algorithm (called Fast Multiple Method) extension that allows to solve very large problems, with ten of millions of unknowns, in reasonable time on parallel machines. Recently, H-matrix technics have given access to fast direct solvers, able to solve with a very good accuracy problems with millions of unknowns without the problem induced by iterative resolution (no control on the number of iterations, difficulty to find a good preconditioned, etc.). The resulting software is used on daily basis in acoustics for installation effects computation, aeroacoustic simulation (in a coupled scheme with other tools), in electromagnetism for antenna siting, electromagnetic compatibility or stealth, and in electrostatics for fuel tank modelisation and lightning effects. The aim of this talk is to present a wide view of the realizations, to underline the recent developments (such as H-matrix) and to present the main perspectives and futures directions of research in a context of different physics and applicationsP

Nolwenn Balin, Benoit Lizé
Airbus
nolwenn.balin@airbus.com, benoit.lize@airbus.com

Guillaume Sylvand
EADS CCR, Centre de Toulouse
guillaume.sylvand@airbus.com

Isabelle Terrasse
Airbus
isabelle.terrasse@airbus.com

MS151

A Numerical Routine for Fast Spherical Grid Rotations

We present a numerical routine that interpolates function values on rotated spherical grids via hybrid nonuniform FFTs. This routine can be used for evaluating singular integral operators on smooth surfaces that are globally parametrized by spherical coordinates. Problems of this type arise, for example, in simulating Stokes flows with particulate suspensions and in multi-particle scattering calculations. The algorithm has a small complexity constant, and the cost of applying the quadrature rule is nearly-optimal $O(p^4 \log p)$ for a spherical harmonic expansion of degree p. This is joint work with Zydrunas Gimbutas (NIST).

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

MS152

Recovery Based a Posteriori Error Estimation for Finite Element Methods

The recovery-based Zienkiewicz-Zhu (ZZ) *a posteriori* error estimator has been widely used in the engineering practice due to the ease of implementation, generality, and accuracy. However, it is well known that the ZZ estimator is inefficient for non-smooth problems on relatively coarse meshes. In this talk, we will discuss three types of recovery-based *a posteriori* error estimators based on the L^2, the equilibrium, and the $H(\mathrm{div})$ recoveries, respectively. These estimators are efficient, reliable, or/and robust. Moreover, they preserve the mathematical structure of the underlying problem.

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

MS152

Robust a-Posteriori Error Estimation for Finite Element Approximation to H(curl) Problem

In this talk, we will discuss the recovery-type a posteriori error estimation for the conforming finite element formulation of the positive definite $\mathbf{H}(\mathbf{curl})$ problem. Both the $\mathbf{L}^2$-recovery and $\mathbf{H}(\mathbf{curl})$-recovery are discussed. The global recovery problems are localized through weighted averaging, and a partition of unity of the magnetizing field respectively.

Shuhao Cao
Pennsylvania State University
Department of Mathematics
scao@psu.edu

MS152

Localized H(div) Recovery-Based a Posteriori Error Estimators

In this talk, we present recovery-based error estimators for conforming finite element approximation of second-order elliptic problem. The flux is recovered in H(div) finite element subspaces by approximating equilibrium and constitutive equations simultaneously in a weighted H(div) norm. *A posteriori* error estimators are constructed locally on appropriate patches of triangular elements. Reliability and

efficiency of these local error estimators will be discussed. Numerical results are provided to demonstrate features of these error estimators.

Xu Zhang
Purdue University
xuzhang@purdue.edu

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

MS153

Continuum Spine Modeling with Application to Outer Retina Neurocircuitry

The continuum theory for dendritic spines is an extension of classical cable theory for which the distribution of spines is treated as a continuum. With the continuum theory, different spine morphologies, multiple populations of spines, and distributed physiological properties are represented compactly by relatively few differential equations. In this talk I will present a brief overview of continuum spine theory and discuss how to apply the theory to modeling neurocircuitry in the outer retina.

Steven M. Baer
Arizona State University
School of Mathematical & Statistical Sciences
steven.baer@asu.edu

MS153

Simulation of the Ephaptic Effect in the Cone-Horizontal Cell Synapse of the Retina

The drift-diffusion (Poisson-Nernst-Planck) model—including a numerical model for cell membranes that resolves surface charge boundary layers—is applied to the cone-horizontal cell synapse in the outer plexiform layer of the retina in a two-dimensional cross-section geometry. Numerical simulations reproduce the experimental calcium current-voltage curves for the goldfish retina in response to a bright spot, with and without an illuminated background. The ephaptic (electrical) effect is demonstrated by computing the shift in the IV curve for background off vs. background on.

Carl L. Gardner
Arizona State University
School of Mathematical and Statistical Sciences
carl.gardner@asu.edu

MS153

Modeling of Calcium-Induced Calcium Release

Calcium release through coordinated IP3 receptor and ryanodine receptor (RyR) openings are important signaling events in neurons in which Ca2+ release from one channel opens other nearby Ca2+ sensitive channels. This is modeled using extensive experimental data to compute how the open/closed probability of RyR changes with Ca2+ concentration and with RyR open/closed time. Simulations show how neighboring channels are recruited and also give insights into how this positive feedback mechanism terminates.

Dirk Gillespie
Rush University Medical Center
Department of Molecular Biophysics & Physiology
dirk_gillespie@rush.edu

MS153

Dendritic Coincidence Detection Enabling Wordspotting Computation

Dendritic computation is often ignored when building neuron models. However, dendrites have been shown to perform operations like non-linear filtering, spatial and temporal summation of synaptic inputs, coincidence detection, synaptic scaling and sequence detection. We show that a network of dendrites and a Winner-Take-all network is similar to a Hidden Markov Model(HMM) classifier often used for speech and pattern-recognition. We have developed a mathematical framework for Silicon dendrites, that models deep dendritic trees.

Jennifer Hasler
Electrical and Computer Engineering
Georgia Tech
jennifer.hasler@ece.gatech.edu

MS154

Fast Spectral PDE Solvers for Complex Structures: the Fourier-Continuation Method

We present fast spectral solvers for time-domain Partial Differential Equations. Based on a novel Fourier-Continuation (FC) method for the resolution of the Gibbs phenomenon these methodologies give rise to time-domain solvers for PDEs for general engineering problems and structures which enjoy a number of desireable properties, including spectral time evolution essentially free of pollution or dispersion errors for general PDEs in the time domain, with conditional/unconditional stability for explicit/alternating-direction methods. A variety of applications to linear and nonlinear PDE problems, including the Maxwell equations, the Navier-Stokes equations, the elastic wave equation demonstrate the significant improvements the new algorithms can provide over the accuracy and speed resulting from other approaches.

Oscar P. Bruno
California Institute of Technology
obruno@caltech.edu

MS154

Frame Theoretic Convolutional Gridding

This talk is about reconstructing images from non-uniform Fourier data. This problem is relevant in applications such as magnetic resonance imaging (MRI) and synthetic aperture radar (SAR). The non-uniform FFT algorithm provides a practical way to reconstruct images However, choosing the parameters can be difficult and in some cases the method may fail to converge. This talk provides a mathematical foundation, through the use of Fourier frames, for reconstructing functions from their non-uniform Fourier data. As a result, the parameters for the NFFT can be chosen to ensure numerical convergence under various non-uniform sampling schemes common in MRI and SAR.

Anne Gelb
Arizona State University
annegelb@asu.edu

Guohui Song

Clarkson University
gsong@clarkson.edu

MS154

Optimized Fourier Continuation Methods

Fourier continuation/extension (FC) methods, have been shown to achieve arbitrary orders of accuracy in generating Fourier approximations of smooth but non-periodic functions from evenly spaced data. The stability of PDE solvers based on FC methods is heavily dependent on the parameters of the continuation. Optimized FC methods will be presented that balance accuracy with sufficient conditions for stability within a FC based PDE solver. The inherent trade-off between accuracy and stability will be discussed in detail.

Mark Lyon
University of New Hampshire
mark.lyon@unh.edu

MS154

A New Radial Basis Functions (RBF)-based Frame Method to Bypass the Runge Phenomenon

Approximating nonperiodic analytic functions with spectral accuracy is not an easy task. RBFs, just like polynomials and other pseudospectral methods, are susceptible to suffer from the Runge phenomenon, spurious oscillations near the edges of the domain, and from the ill-conditioning of the method's associated system. Recently, Adcock et al have introduced a method which combines frames theory and Fourier extensions in [Adcock, Huybrechs, Martin-Vacquero, On the numerical stability of Fourier extensions, Found Comp Math 14, 635-687]. The method is not only spectrally convergent, it is also free from the Runge phenomenon even when the data is equispaced. We will use the fact that pseudospectral methods can be seen as particular cases of RBFs to construct an RBF-based frame method that generalized Adcocks Fourier extensions method.

Cecile M. Piret
Université catholique de Louvain
cmpiret@mtu.edu

MS155

Parallel and Adaptive Mantle Convection Simulation in Aspect

In this talk, we are presenting the open source mantle convection code ASPECT to solve large scale convection problems. The main numerical ingredients are higher order finite elements with adaptive mesh refinement, stabilization schemes, and linear/nonlinear solvers. The talk will also cover our software design approach to engineer a flexible, extensible, code with a healthy community that can support the project in the long run. Finally, we will talk about examples and show benchmark results.

Timo Heister
Clemson University
Mathematical Sciences
heister@clemson.edu

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

MS155

Large-Scale Forward and Inverse Numerical Simulations of Crustal and Lithospheric-Scale Deformation

Geological processes occur over long timescales and involve nonlinear constitutive relationships. We discuss a parallel 3D code to simulate such processes, based on a marker-and-cell technique with a staggered finite difference discretization and visco-elasto-plastic rheologies. We show a few examples of geological modelling applications. In addition, we demonstrate how the forward models can be combined with geophysical data and a MC Monte Carlo-based inversion approach to constrain the rheology and structure of mountains belts.

Boris Kaus
Johannes Gutenberg University Mainz
Institute of Geosciences
kaus@uni-mainz.de

Anton Popov, Tobias Baumann
Johannes Gutenberg University Mainz Institute of Geosciences
popov@uni-mainz.de, baumann@uni-mainz.de

MS155

HPC Finite Elements for Nonlinear Stokes Flow

Rock deformation over geological time scales can be expressed as a nonlinear, incompressible Stokes problem in which the viscosity is highly heterogeneous. Algorithmic and computationally scalable preconditioners are necessary to achieve high resolution 3D simulations. Here I present a mixed FE method employing a multilevel preconditioner which exploits matrix-free kernels. Trading flops for memory bandwidth, significant speedups compared to preconditioners requiring assembled operators are obtained. Performance is demonstrated using a simplified model of slab detachment.

Dave A. May
ETH Zurich
dave.may@erdw.ethz.ch

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory and CU Boulder
jed@jedbrown.org

MS155

Three-Field Block-Preconditioners for Models of Coupled Magma/mantle Dynamics

We discuss the iterative solution of a finite element discretisation of the magma dynamics equations. These equations share features of the Stokes equations, however, Elman-Silvester-Wathen (ESW) preconditioners for the magma dynamics equations are not optimal. By introducing a new field, the compaction pressure, into the magma dynamics equations, we have developed a new three-field preconditioner which is optimal in terms of problem size and less sensitive to physical parameters compared to the ESW preconditioners.

Sander Rhebergen

University of Oxford
Mathematics Institute and Department of Earth Sciences
sander.rhebergen@maths.ox.ac.uk

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

Andrew J. Wathen
Oxford University
Numerical Analysis Group
wathen@maths.ox.ac.uk

Richard F. Katz
University of Oxford
Richard.Katz@earth.ox.ac.uk

Laura Alisic, John Rudge
Bullard Laboratories
University of Cambridge
la339@cam.ac.uk, jfr23@cam.ac.uk

MS156

A Sub Cell Dynamics Based Closure Model for Multimaterial Arbitrary Lagrangian Eulerian Codes

Arbitrary Lagrangian Eulerian (ALE) codes introduce multimaterial cells to represent material interfaces that may undergo high deformation. A closure model is then required to close the governing equations, which are otherwise underdetermined, this defines how volume fractions and state variables of individual material components within these multimaterial cells evolve. The requirements for these models will be presented and a interface aware sub-scale-dynamics closure model will be described that has been developed to meet these requirements.

Andrew J. Barlow
Atomic Weapons Esatblishment, United Kingdom
Andy.Barlow@awe.co.uk

MS156

A High-Order/Low-Order Exponentially-Convergent IMC Method

It is well known that standard Monte Carlo converges as the inverse of the square root of the number of particle histories executed. We demonstrate that exponential convergence is possible using a defect approach in which each succeeding Monte Carlo batch estimates only the additive error in the solution from previous batches. This solution is obtained by projection of the Monte Carlo solution onto a FEM space. A refinement strategy is required to maintain exponential convergence.

Simon Bolding, Jim E. Morel
Texas A&M
bolding@tamu.edu, morel@tamu.edu

Robert B. Lowrie
Los Alamos National Laboratory
Los Alamos, NM 87545 USA
lowrie@lanl.gov

MS156

Multiphysics Lagrangian/Eulerian Modeling and deRham Complex Based Algorithms

Multiphysics modeling in an arbitrary Lagrangian/Eulerian (ALE) framework leads naturally to numerical algorithms matching the geometric meaning of the operators in the deRham complex. We discuss our experience with this approach relative to solid kinematics, magnetohydrodynamics and full Maxwell hydrodynamic modeling. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Allen C. Robinson
Sandia National Laboratories
acrobin@sandia.gov

MS156

Mimetic Finite-Difference Methods

The mimetic finite difference method (J. of Comput. Phys. 257(2014)1163) mimics fundamental properties of mathematical and physical systems including conservation laws, symmetry and positivity of solutions, duality and self-adjointness of differential operators, and exact mathematical identities of the vector and tensor calculus. We describe the major mimetic ideas and their relevance to academic and real-life problems. The supporting applications include diffusion, electromagnetics, fluid flow, and Lagrangian hydrodynamics problems on different type of meshes.

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

MS157

Representing Topography in Earth System Models with Porous Barriers

The ocean boundary (coasts and sea-floor) is fractal and exhibits features (e.g. ridges and passages) that are critical to shaping ocean circulation. These features are typically not resolvable in contemporary global earth system models. We explore an extension of cut-cells that represents the topography via porous barriers between cells. We provide an algorithm to calculate the porous barrier data and illustrate improved fidelity in coarse-resolution models using the porous barrier representation of topography.

Alistair Adcroft
Princeton
aadcroft@princeton.edu

MS157

A Higher-Order Cut Cell Finite Volume Method For Advection-Diffusion

We present a higher-order cut cell algorithm for the advection-diffusion equation in irregular domains. The finite volume operators are constructed using a unique weighted least-squares approach, which reduces the sensitivity to small cells despite the flux-difference conservation

form. A higher-order RK time discretization treats diffusion terms implicitly, so the overall algorithm is limited only by the advection term explicit CFL constraint, regardless of small cells. Accuracy and stability are demonstrated with several simple tests.

Dharshi Devendran
University of Chicago
pdevendran@gmail.com

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

MS157

A Mixed Explicit Implicit Time Stepping Scheme for Cartesian Embedded Boundary Meshes

We present a mixed explicit implicit time stepping scheme for solving the advection equation on Cartesian embedded boundary meshes. The implicit scheme is used to overcome the small cell problem and ensure stability at the cut cells. It is coupled to a standard explicit scheme which is used over most of the mesh. We present a theoretical result about the coupling, and show numerical results in one and more dimensions.

Sandra May
ETH Zurich
sandra.may@sam.math.ethz.ch

Marsha Berger
Courant Institute of Mathematical Sciences
New York University
berger@cims.nyu.edu

MS157

Inverse Lax-Wendroff Procedure for Numerical Boundary Conditions of Hyperbolic Equations

We develop a high order finite difference numerical boundary condition for solving hyperbolic Hamilton-Jacobi equations and conservation laws on a Cartesian mesh. The challenge results from the wide stencil of the interior high order scheme and the fact that the boundary may not be aligned with the mesh. Our method is based on an inverse Lax-Wendroff procedure for the inflow boundary conditions. Extensive numerical examples are provided to illustrate its good performance of our method. This is a joint work with Ling Huang and Mengping Zhang, and with Sirui Tan and Francois Vilar.

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

MS158

Sequentially Constrained Monte Carlo

Constraints in the parameter or model space typically make sampling from distributions more complex (although more interpretable) than their unconstrained counterparts. We define a Sequentially Constrained Monte Carlo algorithm connecting a simple distribution, to the target distribution by a path defined by the strictness of constraint enforcement. We show general applicability by expanding the usual definition of constraints to include adherence to a theoretical model governed by differential equations.

Dave A. Campbell
Simon Fraser University
dac5@sfu.ca

Shirin Golchi
Columbia University
golchi.shirin@gmail.com

MS158

On the Use of Particle Based Methods for the Real Time Identification and Control of Nonlinear Dynamical Systems

The field of structural dynamics is inherently related to the simulation, identification and control of structural systems. This task is not a straightforward one; firstly, due to potential nonlinear behavior; and secondly, due to uncertainties relating to erroneous modeling assumptions, imprecise sensory information, ageing effects, varying loads, and lack of a priori knowledge of the system itself. This talk discusses the implementation of methodologies capable of successfully simulating such systems by encompassing the aforementioned complexities.

Eleni Chatzi
Institute of Structural ENgineering
ETH Zurich
chatzi@ibk.baug.ethz.ch

MS158

Bayesian Uncertainty Quantification and Propagation for Molecular Dynamic Simulations in Nanoscale Fluid Mechanics

For five decades, molecular dynamics (MD) simulations, in synergy with experiments, have elucidated critical mechanisms in a broad range of physiological systems and technological innovations. However, in nanofluidics, the results of experiments and MD simulations may differ by several orders of magnitude. We show that experimental and large scale MD investigations can be consolidated through a Bayesian framework. Our findings indicate that it is essential to revisit MD simulations in the context of uncertainty quantification.

Petros Koumoutsakos
Chair of Computational Science, ETH Zürich
petros@ethz.ch

Panagiotis Angelikopoulos
Computational Science, ETH Zurich, Switzerland
panagiotis.angelikopoulos@mavt.ethz.ch

Costas Papadimitriou
University of Thessaly
Dept. of Mechanical Engineering
costasp@uth.gr

MS158

Computationally Efficient Tools for Bayesian Uncertainty Quantification and Propagation in Structural Dynamics

Bayesian uncertainty quantification, including Bayesian hierarchical modelling, are becoming standard tools in struc-

tural dynamics for model selection, model parameter calibration, uncertainty propagation and SHM using vibration measurements. For complex linear/nonlinear models, the computations involved in Bayesian asymptotic approximations and MCMC sampling tools may be excessive. Drastic reduction in computational effort in model intrusive and/or non-intrusive schemes is achieved by integrating model reduction techniques, adjoint methods, surrogate models, parallel computing and highly-parallelizable sampling schemes. Acknowledgement: This research has been implemented under the ARISTEIA Action of the Operational Programme Education and Lifelong Learning and was co-funded by the European Social Fund (ESF) and Greek National Resources.

Costas Papadimitriou
University of Thessaly
Dept. of Mechanical Engineering
costasp@uth.gr

Panagiotis Angelikopoulos
Computational Science, ETH Zurich, Switzerland
panagiotis.angelikopoulos@mavt.ethz.ch

Panagiotis Hadjidoukas
Computational Science
ETH Zurich, Switzerland
panagiotis.chatzidoukas@mavt.ethz.ch

Petros Koumoutsakos
Chair of Computational Science, ETH Zürich
petros@ethz.ch

MS159

H-Matrix Accelerated Second Moment Analysis for Potentials with Rough Correlation

We consider the efficient solution of strongly elliptic potential problems with stochastic Dirichlet data by the boundary integral equation method. The computation of the solutions two-point correlation is well understood if the two-point correlation of the Dirichlet data is known and sufficiently smooth. Unfortunately, the problem becomes much more involved in case of rough data. We will show that the concept of the H-matrix arithmetic provides a powerful tool to cope with this problem. By employing a parametric surface representation, we end up with an H-matrix arithmetic based on balanced cluster trees. This considerably simplifies the implementation and improves the performance of the H-matrix arithmetic. Numerical experiments are provided to validate and quantify the presented methods and algorithms.

Helmut Harbrecht, Juergen Doelz
Universitaet Basel
Departement of Mathematics and Computer Science
helmut.harbrecht@unibas.ch, juergen.doelz@unibas.ch

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

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

MS159

Adaptive Monte Carlo and Quasi-Monte Carlo Integration

Monte Carlo methods are used for computing means of random variables with complex distributions. Both Monte Carlo methods and quasi-Monte Carlo methods are also used for computing high dimensional integrals. A key question in these computations is when to stop, while ensuring that the desired accuracy is obtained. Adaptive (quasi-)Monte Carlo algorithms rely on data-based error bounds to answer this question. This talk describes recent work to construct such error bounds and ensure that they are trustworthy. We also derive upper bounds on the computational costs of our adaptive algorithms. A key idea is to consider cones of random variables or integrands. Our new algorithms have been implemented in the Guaranteed Automatic Integration Library (GAIL) https://code.google.com/p/gail/.

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

Lan Jiang
Dpet. of Applied Mathematics
Illinois Institute of Technology
ljiang14@hawk.iit.edu

Antoni Luís Jiménez Rugama
Department of Applied Mathematics
Illinois Institute of Technology
lluisantoni@gmail.com

MS159

Application of Quasi-Monte Carlo Methods to PDEs with Random Coefficients

PDEs with random coefficients are an important source of high dimensional problems. One example is the flow through porous medium: because of the near impossibility of modeling the microscopic channels through which water can flow in a porous layer, it is common engineering practice to model the porous medium as a random permeability field. The quantity of interest is therefore an expected value with respect to the random field, leading to a high dimensional integral where the number of variables is as high as the number of parameters needed to model this random field (it can be infinite). In this talk I will explain how quasi-Monte Carlo (QMC) methods can be tailored to a prototype of such integrals. I will discuss the fast construction of higher order QMC methods to improve the convergence rate and the use of multi-level techniques to improve the computational cost. The talk will touch on a number of joint works with Ivan Graham and Rob Scheichl (Bath), Dirk Nuyens (KU Leuven), Christoph Schwab (ETH Zurich), and Ian Sloan, James Nichols, Josef Dick and Quoc Le Gia (UNSW).

Frances Y. Kuo
School of Mathematics and Statistics
University of New South Wales

f.kuo@unsw.edu.au

MS159

Option Pricing and the Anova Decomposition of a Function of An Infinite Number of Variables

In this work we consider the ANOVA decomposition of the integrand in a continuous version of a path-dependent option pricing problem. We show that in the Brownian bridge (or Levy-Ciesielski) formulation, **every term in the (infinite) ANOVA decomposition is smooth**. With this result we are preparing for an error analysis of the cubature problem for the option pricing problem, in which the discrete-time problem is approximated by the continuous problem, and the error analysis then applied to the truncated ANOVA expansion, in which every term is smooth.

Ian H. Sloan
University of New South Wales
School of Mathematics
i.sloan@unsw.edu.au

Frances Y. Kuo
School of Mathematics and Statistics
University of New South Wales
f.kuo@unsw.edu.au

Michael Griebel
Universitat Bonn
Inst fur Angewandte Mathematik
griebel@ins.uni-bonn.de

MS160

Adaptive Determination of Optimal Multilevel Monte Carlo parameters in the Presence of Failures

In the multilevel Monte Carlo (MLMC) method numerous parameters have to be determined. Their choice is crucial for the work and error of the method. We propose to determine the number of samples per level, the finest and coarsest level by solving an integer optimization problem. Faults influence the MLMC levels to a different extent. We present a fault tolerant MLMC method that adapts to experienced failures without a priori knowledge of the failure distribution.

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

Stefan Pauli
ETH Zurich Computer Science Department
stefan.pauli@inf.ethz.ch

MS160

Hierarchical Resilience for Structured AMR

Structured AMR (SAMR) has hierarchy in spatial and temporal resolution that can be exploited for devising flexible and customizable resiliency strategies. We prototyped a hierarchical approach to resilience in Chombo, an SAMR library, using the Global View Resiliency (GVR) library that takes differentiated state snapshots for localized recovery. SAMR is used in various domains, so our strategy may benefit those applications, and perhaps an even broader class of applications with hierarchy.

Anshu Dubey
Lawrence Berkeley National Laboratory
adubey@lbl.gov

Hajime Fujita
Department of Computer Science
University of Chicago
hfujita@cs.uchicago.edu

Zachary Rubenstein
Department of Computer Science, The University of Chicago
zar1@uchicago.edu

Brian Van Straalen
Lawrence Berkeley National Laboratory
Compuational Research Division
bvstraalen@lbl.gov

Andrew A. Chien
The University of Chicago
Argonne National Laborator
achien@cs.uchicago.edu

MS160

Resilience Properties of Gossip-Style Algorithms

Gossip algorithms have interesting resilience properties which can potentially complement and extend state-of-the-art approaches towards algorithmic fault tolerance. We will review the latest developments in gossip-based all-reduce operations and some distributed linear algebra kernels built on top of them. Moreover, we will discuss how gossip-based approaches compare to state-of-the-art fault tolerant algorithms, such as ABFT methods. By combining gossip-based approaches with existing fault tolerant high performance algorithms, resilience at the algorithmic level can potentially be strengthened further.

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

Gerhard Niederbrucker, Michael Moldaschl, Karl Prikopa
University of Vienna
gerhard.niederbrucker@univie.ac.at,
michael.moldaschl@univie.ac.at,
karl.prikopa@univie.ac.at

MS160

Spatial Decomposition for Resilient Extreme-Scale Scientific Simulations

We present a probabilistic approach to PDEs, where computational results are viewed as data that is used in a resilient iterative fashion to update information about the solution until convergence. We rely on a domain decomposition method such that the original problem is reduced to solving the PDE on subdomains with uncertain boundary conditions. Resilience to both hard and soft errors is integrated within the algorithm. We show promising results for elliptic PDEs.

Francesco Rizzi, Khachik Sargsyan, Karla Morris, Cosmin

Safta
Sandia National Laboratories
fnrizzi@sandia.gov, ksargsy@sandia.gov,
knmorri@sandia.gov, csafta@sandia.gov

Paul Mycek, Omar M. Knio
Duke University
paul.mycek@duke.edu, omar.knio@duke.edu

Olivier LeMaitre
LIMSI-CNRS
olm@limsi.fr

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

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

MS161

Uncertainty in Reynolds Stress Closures for Turbulent Flow Calculations

Reynolds-averaged Navier-Stokes (RANS) simulations are a practical approach for solving complex multi-physics turbulent flows, but the underlying assumptions of the turbulence models introduce errors and uncertainties in the simulation outcome. We present a framework to characterize these uncertainty based on perturbations of the eigenvalues of the anisotropic Reynolds stress tensors which bias the model towards specific turbulent states. Results are presented for a broad class of turbulent flows and illustrate the potential for this approach to capture the salient uncertainties introduced by simplified modeling assumptions.

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

Michael A. Emory
Stanford University
memory@stanford.edu

Catherine Gorle
EMAT, UA
catherine.gorle@ua.ac.be

MS161

Eddy Viscosity Model Selection for Transonic Turbulent Flows Using Shrinkage Regression

Linear eddy viscosity models (LEVM), when used in k-epsilon models, often fail to capture complex turbulent phenomena. Our previous work has revealed the structural shortcomings of LEVMs in simulations of jet-in-crossflow interactions. We address this issue by calibrating a cubic eddy viscosity model using limited measurements from a jet-in-crossflow experiment. We use shrinkage regression to identify the terms that need to be retained. Improvement over LEVM predictions is quantified and presented.

Sophia Lefantzi
Sandia National Laboratories
slefant@sandia.gov

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Srinivasan Arunajatesan, Lawrence Dechant
Sandia National Laboratories
sarunaj@sandia.gov, ljdecha@sandia.gov

MS161

Formulation and Calibration of a Stochastic Model Form Error Representation for Rans

We develop a stochastic model inadequacy representation for the RANS equations for wall-bounded flows. The model takes the form of a stochastic PDE governing the error in the Reynolds stress computed according to an eddy-viscosity-based turbulence model. Further, the model is constructed to conform to available knowledge about the performance of typical turbulence models—e.g., that the models are constructed to capture the log layer. We test the model inadequacy representation using DNS data.

Robert D. Moser
University of Texas at Austin
rmoser@ices.utexas.edu

Todd Oliver
PECOS/ICES, The University of Texas at Austin
oliver@ices.utexas.edu

Bryan Reuter
University of Texas, Austin
bryan@ices.utexas.edu

MS161

Estimating a Model Discrepancy Term for the Community Land Model Using Latent Heat and Runoff Observations

We estimate three hydrological parameters of the CLM (Community Land Model) using runoff and latent heat flux measurements gathered from sites with a similar hydrological environment. The estimates are developed as probability density functions, using a CLM surrogate. Two structural error models were to used represent the discrepancy between model and data and selected by their ability to reproduce data. The sensitivity of the structural error model form to the two types of observations is also investigated.

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Laura Swiler
Sandia National Laboratories
Albuquerque, New Mexico 87185
lpswile@sandia.gov

Maoyi Huang, Zhangshuan Hou
Pacific Northwest National Lab
maoyi.huang@pnnl.gov, zhangshuan.hou@pnnl.gov

MS162

Optimal Convergence Rates of Co-Simulation Us-

ing Fine Structure Analysis

For coupled systems, where the monolithic simulation is difficult, co-simulation is an important methodology in time domain. In contrast to coupled ODEs, for coupled DAEs convergence can only be guranteed if certain contraction properties are given. Here, the fine structure analysis is used to verify that the corresponding recursion estimates are attained upper bounds on the convergence rate. Furthermore we verify that numerical convergence might not be detacted. The relation of iterations and order of the time integrations is discussed.

Andreas Bartel, Kai Gausling
University of Wuppertal
bartel@math.uni-wuppertal.de,
gausling@math.uni-wuppertal.de

Sebastian Schöps
Theorie Elektromagnetischer Felder (TEMF) and GSCE
Technische Universit
schoeps@gsc.tu-darmstadt.de

MS162

Trigonometric Integration Methods in Circuit Simulation

The numerical calculation of the limit cycle of oscillators with resonators exhibiting a high quality factor Q such as quartz crystals is a difficult task in the time domain. Time domain integration formulas may introduce numerical damping which leads to erroneous limit cycles. A novel class of adaptive multistep integration formulas is derived which circumvent the aforementioned problems. The results are compared with the well-known Harmonic Balance (HB) technique. Moreover the range of absolute stability and error estimates are derived for these methods.

Hans-Georg Brachtendorf, Kai Bittner
University of Applied Sciences Upper Austria
hans-georg.brachtendorf@fh-hagenberg.at, kai.bittner@fh-hagenberg.at

MS162

Modelling Transmission Power Systems with the Implicit DAE Solver, IDA

The main flow of power on a transmission power system can be modeled as a coupled set of differential-algebraic equations. Transmission models of interest, however, include limits on output and internal states as well as limits on state rates of change, deadbands, and positivity constraints. In this talk, we will outline how we are using a general DAE solver package, IDA, for solution of transmission grid models incorporating many of these conditions within our simulation. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC.

Philip Top
Lawrence Livermore National Laboratory
top1@llnl.gov

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

Alan Hindmarsh
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
alanh@llnl.gov

MS162

Fast Time-Domain Simulation for Reliable Fault Detection

Imperfections in manufacturing processes may cause unwanted connections (faults) that are added to the nominal, "golden", design of an electronic circuit. By fault simulation we simulate all situations: a huge number of new connections and each with many different values, up to the regime of large deviations, for the newly added element. We also consider "opens" (broken connections). A strategy is developed to efficiently simulate the faulty solutions until their moment of detection. We fully exploit the hierarchical structure of the circuit. Fast fault simulation is achieved in which the golden solution and all faulty solutions are calculated over a same time step.

E. Jan W. ter Maten
NXP Semiconductors, Corp IT, DMS-PDM-Mathematics
Jan.ter.Maten@math.uni-wuppertal.de

Bratislav Tasic
NXP Semiconductors
bratislav.tasic@nxp.com

Jos J. Dohmen
NXP Semiconductors, the Netherlands
jos.j.dohmen@nxp.com

Theo G.J. Beelen
NXP Semiconductors
theo.g.j.beelen@nxp.com

Rick Janssen
NXP Semiconductors, the Netherlands
rick.janssen@nxp.com

Wil Schilders
Eindhoven University of Technology
w.h.a.schilders@TUE.nl

Michael Guenther
Bergische Universitaet Wuppertal
guenther@math.uni-wuppertal.de

MS163

Sparse QR Factorization on Heterogenous Platforms with Multiple GPUs

We present a sparse multifrontal QR Factorization method on a heterogeneous CPU-GPU platform. Our method is extremely efficient for different sparse matrices sizes and structures even for matrices that don't fit in GPU memory. Our method benefits from both the highly parallel general-purpose computing cores available on a single GPU and from multiple GPUs on a single platform. This presentation focuses on parallelism over multiple GPU cores.

Mohamed Gadou
University of Florida
mgadou@cise.ufl.edu

Timothy A. Davis
Texas A&M
Computer Science and Engineering
davis@tamu.edu

Sanjay Ranka
CISE, University of Florida, USA
ranka@cise.ufl.edu

MS163

Sparse Communication Avoiding Pivoting and GPUs

Modern GPUs provide a window on future hardware platforms that we need to design for. In this talk we describe our experience in implementing direct methods for sparse symmetric indefinite matrices in CUDA, a topic that presents both mathematical and technical challenges. Whilst a number of communication-avoiding methods have been developed for dense linear algebra, their applicability to sparse linear algebra is limited as they ignore sparsity in selecting pivots, generating significant fill. We will describe heuristics that attempt to manage this complexity through a series of fall-back strategies that in the best case have similar communication properties to the unpivoted Cholesky factorization, but for numerically challenging problems deliver stability on a par with the best existing sparse symmetric indefinite direct solvers.

Jonathan Hogg
STFC Rutherford Appleton Lab, UK
jonathan.hogg@stfc.ac.uk

Jennifer Scott
Rutherford Appleton Laboratory
Computational Science and Engineering Department
jennifer.scott@stfc.ac.uk

MS163

GLU: LU Re-Factorization on the GPU

In realistic applications we often solve a sequence of similar linear systems that arise as part of a non-linear iteration process. In this talk we focus on the design and implementation of the parallel LU re-factorization on the GPU. It can be performed on the linear systems that have the same sparsity pattern, pivoting and reordering to minimize fill-in as the original system. We also discuss batched version of the algorithm and present relevant numerical experiments.

Maxim Naumov, Sharan Chetlur, Lung Sheng Chien
NVIDIA
maxim.a.naumov@gmail.com, schetlur@nvidia.com, lchien@nvidia.com

MS163

Accelerating the Supernodal Sparse Cholesky Factorization on GPUs

Achieving high performance sparse Cholesky factorization on a GPU is difficult due to irregular computation, limited GPU memory, and PCIe communication. Previous work showed how some PCIe communication could be hidden behind CPU and GPU computation. This work shows how a larger portion of the factorization can be accelerated on the GPU, and overall performance substantially improved, by moving entire branches of the elimination tree to the GPU and greatly reducing the required PCIe communication.

Steven C. Rennich
NVIDIA
srennich@nvidia.com

Timothy A. Davis
Texas A&M
Computer Science and Engineering
davis@tamu.edu

Darko Stosic
NVIDIA
ddstosic@bu.edu

MS164

A Fast N-body Algorithm for Kernel Sums in High Dimensions

In computational statistics kernelized methods require the rapid evaluation of kernel sums. We present a fast algorithm for such sums and introduce novel methods for pruning and for approximating the far field. The scheme is kernel-independent (does not use analytic expansions) and the pruning is combinatorial (not geometric). Its complexity depends linearly on the dimension. We present the structure of the algorithm and experimental results that demonstrate is performance. As a highlight, we report results for Gaussian kernel sums for one million points in 1000 dimensions and for problems in which the kernel has variable bandwidth.

George Biros, Bill March
University of Texas at Austin
biros@ices.utexas.edu, march@ices.utexas.edu

Bo Xiao
School of Computational Science and Engineering
Georgia Institute of Technology
boxiao@gatech.edu

MS164

Information-Theoretic Tools for Uncertainty Quantification of High Dimensional Stochastic Models.

We present mathematical tools for deriving optimal, computable bounds on sensitivity indices of observables for complex stochastic models arising in biology, reaction kinetics and materials science. The presented technique allows for deriving bounds also for path-dependent functionals and risk sensitive functionals. The use of variational representation of relative entropy also allows for error estimation and uncertainty quantification of coarse-grained models.

Paul Dupuis
Division of Applied Mathematics
Brown University
dupuis@dam.brown.edu

Markos Katsoulakis
University of Massachusetts at Amherst
markos.katsoulakis@gmail.com

Yannis Pantazis
University of Massachusetts, Amherst
Department of Mathematics and Statistics

pantazis@math.umass.edu

Petr Plechac
University of Delaware
Department of Mathematical Sciences
plechac@math.udel.edu

MS164

Use of Parallel MCMC Methods with the Community Land Model

We present the development of a parallel Delayed Rejection Adaptive Metropolis (DRAM) algorithm and its use with the Community Land Model for Bayesian calibration of model parameters. Bayesian calibration of expensive simulations is often done with emulators because of the computational cost of Monte Carlo Markov Chain (MCMC) sampling. In this case, we chose not to use emulators and instead use a parallel chain scheme directly on CLM to obtain reasonable run times.

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Laura Swiler
Sandia National Laboratories
Albuquerque, New Mexico 87185
lpswile@sandia.gov

Maoyi Huang
Pacific Northwest National Lab
maoyi.huang@pnnl.gov

Jason Hou
Pacific Northwest National Laboratory
zhangshuan.hou@pnnl.gov

MS164

Highly Scalable Hierarchical Sampling Algorithms for Gaussian Random Fields

The ability to generate samples of random fields with prescribed statistical properties is a key ingredient for many uncertainty quantification methods. In this talk, we will present a highly scalable multilevel hierarchical sampling technique that has linear complexity. Samples are generated on a coarse grid using the truncated Karhunen-Loéve expansion, and then extended to the finer levels by solving a stochastic partial differential equation. An application to Multilevel Monte Carlo simulations of subsurface flows will be presented to demonstrate the good scalability of our method.

Panayot Vassilevski
Lawrence Livermore National Laboratory
vassilevski1@llnl.gov

Umberto E. Villa
Center for Advanced Scientific Computing
Lawrence Livermore National Laboratory
villa13@llnl.gov

MS165

Construction of Gaussian Surrogate Process Using Numerical and Modeling Error Uncertainty

In this talk we will explore the use of adjoint based estimates of numerical error (upper bounds where available) in quantities of interest to aid in the construction of more effective surrogates. This will complement the vast body of work on constructing such surrogates for model parameter uncertainty characterization.

Hossein Aghakhani
University at Buffalo, SUNY
haghakha@buffalo.edu

Abani K. Patra
SUNY at Buffalo
Dept of Mechanical Engineering
abani.patra@gmail.com

Elaine Spiller
Marquette University
elaine.spiller@marquette.edu

MS165

Calibration of the Spalart-Allmaras Turbulence Model for Blunt Body Re-Entry Vehicle Flows Using DNS Data

The Spalart-Allmaras turbulence model is calibrated for use in predictions of transonic, chemically reacting boundary layers with favorable pressure gradients similar to those observed on blunt body re-entry vehicles. The calibration takes the form of a Bayesian update for the model parameters. We detail DNS cases specially designed for this calibration, uncertainty estimates for the resulting DNS data, a simple model inadequacy representation used for the calibration, and the posterior results for the model parameters.

Robert D. Moser
University of Texas at Austin
rmoser@ices.utexas.edu

Todd Oliver
PECOS/ICES, The University of Texas at Austin
oliver@ices.utexas.edu

Victor Topalian
PECOS/ICES
The University of Texas at Austin
topalian@ices.utexas.edu

Rhys Ulerich
The Univeristy of Texas at Austin
rhys.ulerich@gmail.com

MS165

Quantifying the Impact of Numerical Errors Along with Other Uncertainties on Probabilistic Hazard Mapping

A probabilistic hazard map requires: 1) a stochastic scenario model – incorporating data and expert opinion – that describes the system's aleatory variability, and 2) a physical model which lets one explore catastrophic hazards (inundation by flooding, landslides, tsunamis, etc). Geophysical modelers often assume that uncertainty >> numerical error and ignore impacts of the later. We have devised a surrogate-based strategy using adjoint error estimates to

quantify the effects numerical errors and other uncertainties on hazard probabilities.

Elaine Spiller
Marquette University
elaine.spiller@marquette.edu

Hossein Aghakhani
University at Buffalo, SUNY
haghakha@buffalo.edu

MS165

Predictive Uncertainty Quantification of An Ablating Entry Vehicle Heatshield

Simulations of hypersonic reacting flow require the integration of a number of complex physical models with uncertain parameters, evaluated at conditions for which experimental data may be sparse. We discuss calibration and forward propagation of uncertainty in the context of the atmospheric re-entry of a symmetric capsule body with an ablating heatshield, focusing on cycles of calibration and prediction for aerothermochemistry, surface ablation, turbulence, and other submodels. Multiple years' results will be presented.

Roy Stogner
University of Texas at Austin
roystgnr@ices.utexas.edu

MS166

Mathematical Modeling of Gliomas: Implications for Interpreting Therapeutic Efficacy Through Imaging

Glioblastoma is the most aggressive form of primary brain tumor. Traditional clinical images, such as magnetic resonance images, are limited in their ability to capture the full extent of the diffuse disease both prior to therapy and, even more, post therapy. In this talk, we will discuss how mathematical modeling can lend insights for interpreting these images on a patient-specific basis by estimating the extent of the disease not visible on imaging.

Andrea Hawkins-Daarud
Department of Neurological Surgery
Northwestern University
andrea.hawkins-daarud@northwestern.edu

Russell Rockne
Northwestern University
russell.rockne@northwestern.edu

David Corwin
LaunchPad Lab
dave@launchpadlab.com

Alexander R.A. Anderson
Moffitt Cancer Center
alexander.anderson@moffitt.org

Paul Kinahan
University of Washington
kinahan@uw.edu

Kristin R. Swanson
Northwestern University
kristin.swanson@northwestern.edu

MS166

Optimization of Computational Simulation Set for Quantification of Hurricane Surge Extreme-Value Statistics

Extreme-value hurricane flooding statistics are essential for effective coastal risk assessment. The joint probability method is a robust and reliable approach for quantifying these statistics and their uncertainty. Yet, the array of hurricane possibilities and complexity of physical processes that must be simulated result in an unmanageable computational burden. Here, optimal sampling to reduce this burden is achieved with physics-based, algebraic surge response functions to define continuous probability density functions for hurricane flood elevation.

Jennifer L. Irish
Virginia Polytechnic Institute and State University
jirish@vt.edu

MS166

Hurricane Storm Surge Risk Analysis for the US North Atlantic Coast

In this work, we use a physically based assessment to estimate the risk of hurricane storm surge in Narragansett Bay, RI, Jamaica Bay, NY, Atlantic City, NJ, and Norfolk, VA. Using a novel approach to risk analysis, we estimate storm surge recurrence intervals for the current and future climate by forcing a hydrodynamic model with thousands of synthetic hurricanes generated using data from observations and global climate models. Our results have been used to inform a multi-institutional, interdisciplinary research effort to develop "Structures of Coastal Resilience."

Talea Mayo, Ning Lin
Princeton University
tmayo@princeton.edu, nlin@princeton.edu

MS166

Representing model inadequacy: A stochastic operator approach

We investigate model form uncertainty for a reaction mechanism model in hydrocarbon combustion. In a typical reaction, the complete mechanism is either not well-understood, or too complex to effectively use as part of a larger combustion problem, necessitating a reduced model. To account for the discrepancy between the full model and its reduced version, we propose an additive, linear, probabilistic formulation. This representation is encoded in a random matrix, whose entries are calibrated using a hierarchical Bayesian scheme. In particular, this formulation is designed to respect certain physical constraints, but also be flexible enough to apply to multiple reactions.

Rebecca Morrison
UT Austin
rebeccam@ices.utexas.edu

Robert D. Moser
University of Texas at Austin

rmoser@ices.utexas.edu

MS167

Adaptive Reconnection-Based Arbitrary Lagrangian Eulerian Method

We describe a new h-adaptive ReALE (A-ReALE) method. The adaptive rezone of A-ReALE is proposed based on the equidistribution principle of centroidal Voronoi tessellations (CVTs). In the rezone phase, we introduce a quadtree based initialization to determine the initial guess of the cell generators. A new CVT is generated by minimizing the mesh energy of the Voronoi tessellation. For 2D examples of time-dependent simulations, A-ReALE shows better mesh convergence than ReALE.

Wurigen Bo
CCS-2, LANL
wbo@lanl.gov

Misha Shashkov
Los Alamos National Laboratory
misha@t7.lanl.gov

MS167

Multimaterial Simulation in Reale Framework

In the ReALE method presented here, connectivity of the mesh is allowed to change during the reconnection phase. The main idea is to define a new grid using specific movement of generators and formalism of Voronoi diagrams. This method leads to general polygonal mesh and allows to follow Lagrangian features of the mesh much better than for standard ALE methods. Furthermore, in the context of multimaterial computations using ReALE method (as for standard ALE method), grid and fluid move separately, and mixed cells containing two or more materials could appear. These mixed cells contain material interfaces, which need special treatment to be taken into account. This is done using the Moment of Fluid (MOF) method.

Jerome Breil
CELIA, UMR 5107 CNRS, University of Bordeaux and CEA, 351 Co
breil@celia.u-bordeaux1.fr

MS167

Triangular Metric-Based Mesh Adaptation for Compressible Multi-Material Flows in Semi-Lagrangian Coordinates

Lagrangian methods for multi-material flows are very attractive (precision, contact discontinuity preservation), but fail to calculate vertexes or shears. Classical remedy is ALE method which consists defining a "smoother" mesh without changing its connectivity and the remap the solution on it. The next step is mesh adaptation to improve robustness and precision. In this context, we present a metric-based triangular mesh adaptation method in 2D, using an interface reconstruction approach to address the numerical dissipation of material interfaces.

Stephane Del Pino, Isabelle Marmajou
CEA
stephane.delpino@cea.fr, isabelle.marmajou@cea.fr

MS167

Reconnection ALE in a Massively-Parallel, Staggered-Grid, Multi-Physics Code

The development of Reconnection-based Arbitrary Lagrangian-Eulerian (ReALE) methods has focused on cell-centered discretizations for Lagrangian hydrodynamics to greatly simplify the conservative remapping of fluid variables. Our approach is different, using both spatially-staggered and subzonal discretizations. This talk will focus on the benefits of this discretization in the context of multi-physics, the unique challenges it presents in the context of ReALE, and the novel techniques we have developed to address these challenges.

David Starinshak
Lawrence Livermore National Laboratory
starinshak1@llnl.gov

J. Michael Owen, Douglas S. Miller
Lawrence Livermore Nat. Lab
mikeowen@llnl.gov, dougmiller@llnl.gov

MS168

Interpolatory Model Reduction for Nonlinear Inversion

Model reduction by rational interpolation has been effectively used for producing accurate, in some cases, optimal reduced models in a numerically feasible way. In this talk, we will show that interpolatory methods can be also easily employed in nonlinear inversion problems to reduce the cost of the overall inversion process with little loss of accuracy. We will present the framework in the application setting of diffuse optical tomography.

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

MS168

Nonlinear Model Order Reduction Using Pod/DEIM 4-D Var with Trust Region Applied to a Spherical Shallow Water Equations Model

We address a very large-scale optimization problem with PDE constraints governed by a spherical shallow water equations as proxy for an atmospheric model.(4-D VAR) MOR with POD/DEIM and trust region model using both forward and adjoint snapshots opens opportunity of reducing computational cost of 4-D VAR data assimilation. Impact of number of POD basis functions, DEIM points along with choice of snapshots and optimization algorithm and finally CPU-speed-up will be investigated.

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Juan Du
IAP, Academia Sinica, Beijing, China
dujuan10@mail.iap.ac.cn

MS168

Reduced Order Modelling for Fluid-Structure Interaction Problems

We present some advances in Reduced Order Modelling applied to fluid structure interaction problems, modelled as inverse problems within a parametric setting. Special attention is devoted to the coupling between fluid and structure through the interface in order to apply reduced computational models to fluid and structure and manage the interface deformation with efficient algorithms. A special emphasis is devoted to the coupling between optimization and fluid-structure interaction problems.

Francesco Ballarin
Politecnico di Milano, Italy
francesco.ballarin@polimi.it

Gianluigi Rozza
SISSA, International School for Advanced Studies
Trieste, Italy
gianluigi.rozza@sissa.it

MS168

Aposteriori Error Estimates and Adaptive Reduced Order Modeling Data Assimilation

We will develop goal-oriented aposteriori estimators of the error in the strong constraint reduced order 4DVar solution. Specifically we will formulate goals as functions that depend on the analysis obtained by minimizing a standard data assimilation cost function. The novel framework will be used to guide an adaptive efficient choice of snapshots, location of discrete empirical interpolation points (DEIM) and matrix DEIM indexes leading to suboptimal reduced solutions that approximate well the full data assimilation analyses.

Razvan Stefanescu
Virginia Tech
rstefane@vt.edu

Adrian Sandu
Virginia Polytechnic Institute and
State University
sandu@cs.vt.edu

MS169

Numerical Stability Issues in H_2 approximation Methods

We discuss the fine details of numerical implementation of the IRKA and the VF algorithms for H_2 rational approximations of linear systems, i.e. turning the methods into reliable mathematical software. Our thesis is that such final step is by no means simple or routine task, even with the state of the art numerical software packages. We tackle the problems of the underlying Cauchy and Vandermonde structures, moment matching, as well as convergence issues in the appropriate function spaces.

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

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

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

MS169

Certified Reduced Basis Model Reduction for Maxwell's Equations

The Reduced Basis Method generates low-order models for the efficient evaluation of parametrized PDEs in many-query and real-time contexts. The approximation quality is certified by using rigorous error estimators. We apply the Reduced Basis Method to systems of Maxwell's equations arising from electrical circuits. Using microstrip models, the input-output behaviour of interconnect structures is approximated for a certain frequency range, parametrized geometry like distance between microstrips and material coefficients.

Martin W. Hess
Max-Planck-Institute for Dynamics of Complex Technical Systems Magdeburg
hessm@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS169

Parameter Estimation for Inverse Problems

As mathematical models continue to grow in size and complexity, the efficiency of the numerical methods used to solve their corresponding inverse problems becomes increasingly important. With differential equation models, for example, avoiding the computation of the forward solution is desirable. Constructing a "nearby" inverse problem avoids this computation to create a numerically robust approach while also providing parameter estimates suitable for the solution of the original inverse problem.

Justin Krueger
Department of Mathematics
Virginia Tech
kruegej2@vt.edu

MS169

Stochastic Approach to Nonlinear Inversion Combining Simultaneous Random and Deterministic Sources

We discuss parametric inversion for diffuse optical tomography using simultaneous random sources and detectors to drastically reduce the costs of the expensive solution of many large linear systems to be solved. Each linear system corresponds to a solving a 3D PDE. We compare this with inversion approaches using reduced order models. In addition, we consider methods to improve solution quality at low cost.

Selin Sariaydin

Department of Mathematics
Virginia Tech
selin@vt.edu

Eric De Sturler
Virginia Tech
sturler@vt.edu

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

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

MS170

Using Numerical Optimization Methods for Sampling in Inverse Problems

Many solution methods for inverse problems compute the maximum a posteriori (MAP) estimator, or equivalently, the regularized solution, by solving an optimization problem. Uncertainty quantification (UQ), on the other hand, typically requires sampling from the Bayesian posterior density function. In this talk, we bring these two ideas together and present posterior sampling methods that make use of existing algorithms for computing regularized solutions/MAP estimators. Theoretically correct samplers for both linear and nonlinear inverse problems will be presented.

Johnathan M. Bardsley
University of Montana
bardsleyj@mso.umt.edu

MS170

Point Spread Reconstruction from the Image of a Sharp Edge: Computation and Uncertainty Quantification

The blurring of an image is often modeled as convolution with a point spread function (PSF) specific to the imaging instrument. We present a method for estimating the PSF that is suitable for applications where it is possible to image a sharp edge. In our model, the PSF is a solution to a Fredholm integral equation of the first kind, known to be ill-posed. We employ optimization and sampling methods to compute several stable estimates of the solution.

Kevin Joyce
University of Montana
Department of Mathematical Sciences
kevin1.joyce@umontana.edu

Johnathan M. Bardsley
University of Montana
bardsleyj@mso.umt.edu

Aaron B. Luttman
National Security Technologies, LLC
luttmaab@nv.doe.gov

Peter Golubstov
Moscow State University
peter.golubstov@umontana.edu

MS170

Statistical Tests for Total Variation Regularization Parameter Selection

We explore three new algorithms for choice of scaling or regularization parameter in Total Variation denoising. TV regularization is viewed as an M-estimator and it is assumed to converge to a well defined limit even if the probability model is not correctly specified. The Discrepancy Principle, a modified version of the χ^2 method for Tikhonov regularization and an empirically Bayesian approach are implemented and compared on benchmark problems in digital image processing.

Jodi Mead
Boise State University
Department of Mathematics
jmead@boisestate.edu

MS170

Constrained Iterative Solver for Sparse Unmixing and Deblurring of Hyperspectral Images

Spectral unmixing involves the computation of fractional contributions of elementary spectra, called endmembers. The forward model involves a linear mixture of endmembers, with nonnegative sparse coefficients, and a wavelength dependent blurring operation. The inverse (reconstruction) problem requires solving a large scale, structured, constrained least squares problem. We show that by exploiting structure of the coefficient matrix, and using known properties of the endmembers, iterative methods can be used to efficiently reconstruct the fractional abundance coefficients.

Sebastian Berisha
University of Pennsylvania
berishas@mail.med.upenn.edu

James G. Nagy
Emory University
Department of Math and Computer Science
nagy@mathcs.emory.edu

Robert Plemmons
Wake Forest University
plemmons@wfu.edu

MS171

Task Based Programming with Pycompss: Leveraging Python in Parallel Platforms

StarSs is a family of task-based programming models which is based on the idea of writing sequential code which is executed in parallel at runtime taking into account the data dependences between tasks. COMPSs is an instance of StarSs, which intends to simplify the execution of Java applications in distributed infrastructures, including clusters and Clouds. The talk will focus in PyCOMPSs, a binding for the Python language which will enable a larger number of scientific applications in fields such as lifesciences and in the integration of COMPSs with new Big Data resource management methodologies developed at BSC.

Rosa M. Badia
Barcelona Supercomputing Center / CSIC

rosa.m.badia@bsc.es

MS171

Coarse Grained Task-Based Parareal Parallel-In-Time Applications in Fusion Energy

Parallelization in the temporal domain has recently emerged as a promising approach to utilize the massive concurreny at the heart of existing and emerging computing platforms. We present an event-driven task-based formulation of the parareal algorithm, implemented using a flexible, lightweight Python framework for coupled simulations. We describe recent applications of this framework to various problems in fusion modelling and simulation, including turbulance, advanced tokamak scenario analysis, and plasma edge dynamics.

Wael R. Elwasif
Oak Ridge National Laboratory
elwasifwr@ornl.gov

Debasmita Samaddar
CCFE
UK Atomic Energy Authority
dsamaddar@alaska.edu

MS171

A Task-Based Computational Astronomy Application

The European Extremely Large Telescope project is one of Europe's highest priorities in ground-based astronomy. The core implementation of the simulation lies in the intensive computation of a tomographic reconstructor, which is used to drive the deformable mirror in real-time from the measurements. A new task-based numerical algorithm is proposed to capture the actual experimental noise and to substantially speed up previous implementations by exposing more concurrency, while reducing the number of floating-point operations.

Hatem Ltaief
Extreme Computing Research Center
KAUST
hatem.Ltaief@kaust.edu.sa

MS171

Applications at Airbus Group of a Task-Based H-Matrix Solver

The Boundary Element Method (BEM) requires the solving of large, ill-conditionned, dense linear systems with a large number of righ-hand sides. A H-Matrix is a hierachical, approximate, data-sparse storage format for matrices that can be manipulated to produce a fast direct linear solver. We consider the H-Matrices for the BEM, with applications to Airbus Group industrial test cases ; and we present the parallelization of this algorithm in shared and distributed memory, using the StarPU runtime system. An almost optimal parallel efficiency is achieved in shared memory, along with promising results in distributed memory.

Guillaume Sylvand
EADS CCR, Centre de Toulouse
guillaume.sylvand@airbus.com

Benoit Lize
Airbus Group Innovations
benoit.lize@airbus.com

MS172

Estimation of Unmodeled Gravitational Wave Transients: an Application of Spline Based Regression and Particle Swarm Optimization

Detecting and estimating unmodeled transient gravitational wave (GW) signals in noisy data is a major challenge in GW data analysis. This paper explores a solution that combines spline regression with particle swarm optimization for knot placement and directional parameter estimation. Analyses of data from a network of detectors show fairly good directional estimates, with reasonable fidelity in the reconstruction of both GW polarization waveforms, at a signal to noise ratio capped at 15. Advisor: Soumya Mohanty, U Texas, Brownsville

Calvin Leung
Harvey Mudd College
cleung@g.hmc.edu

MS172

Persistent Random Walk of Microorganisms in a Porous Medium

We develop a persistent random walk model for the motion of swimming cells in an idealized lattice-like porous medium. The walk is described by a Markov chain in phase space, tracking both position and velocity. Physical parameters, including the overall geometry, bulk flow, and scattering laws, are incorporated into the memory-dependent transition amplitudes. We numerically compute first passage times in MATLAB to predict the effects of lattice structure on microbial transport. Advisor: Joern Dunkel (Massachusetts Institute of Technology)

Grace Lim
California State Polytechnic University, Pomona
gracelim@csupomona.edu

Aden Forrow
MIT
aforrow@mit.edu

MS172

Mean Squared Displacement and Mean First Passage Time in Fluids with Memory

Modeling the motion of passive particles in a viscous fluid is well studied and understood. Extensions to passive motion in a complex fluid which exhibits both viscous and elastic properties have been developed in recent years. However, questions remain on the characterization of mean-square displacement and mean first passage for different theoretical models. We present a statistically exact covariance based algorithm implemented in parallel C++ to generate particle paths to answer these questions. Advisor: Christel Hohenegger, Department of Mathematics, University of Utah

Michael Senter
University of Utah

u0822950@utah.edu

MS172

Pymethyl: A Bioinformatic Approach to Methylation Patterns and their Epigenetic Effects on Risk of Breast Cancer

The study looked at the genomes of two different groups of parous and nulliparous women and their risk for breast cancer. These genomes were sequenced for their genomic methylation patterns. The information was then processed by a high throughput, brute force algorithm to look for places of hypermethylation within the promoter regions or transcribable regions of specific genes. This information helps us gain information on the epigenetic effects to the risk of breast cancer. Advisor: Jose Russo, Fox Chase Cancer Center Mentor

Cody Watson
Wofford College
watsonca@email.wofford.edu

MS172

Modeling Bull Sperm Motility Using Image Processing

We use image processing to determine whether the chemical heparin may change bending in a bull sperm flagellum. Experimental movies for several sperm are analyzed to determine sperm location. From this data, curvature of bending and swimming speeds are analyzed. We also compute bending forces and resulting fluid flow using the method of regularized Stokeslets. We conclude that heparin does change bending in the neck region as well as changing hydrodynamic efficiency and energy. Advisor: Sarah D. Olson, Worcester Polytechnic Institute

Linan Zhang
Worcester Polytechnic Institute
lzhang2@WPI.EDU

MS172

Bounds on Electrical Fields in Two-Component Inhomogeneous Bodies

We investigate bounds on the maximum value of electrical field in two-component inhomogeneous bodies. We assume the conductivities of the inclusions and the surrounding body are known, with no assumptions on the precise geometry of the inclusions. Results are in terms of values of the voltage and current that can be attained from measurements on the boundary of the body. This research has potential applications in preventing electrical breakdown in composite materials. Advisor: Graeme Milton, University of Utah

Zoe Koch, Michael Primrose, Michael Zhao
University of Utah
kochzoe@gmail.com, tba, mz84092@gmail.com

MS173

On a New Class of Semi-Lagrangian Schemes for Kinetic Equations

In this talk we introduce a new class of semi-Lagrangian schemes for simulating kinetic type equations. In this approach we genealize the fast semi-Lagrangian scheme developed in [J. Comput. Phys., Vol. 255, 2013, pp 680-698] to the case of high order reconstructions. The first order accurate semi-Lagrangian scheme is supplemented with polynomial reconstructions of the distribution function and of the collisional operator leading to an effective high order accurate numerical scheme for all regimes, from extremly rarefied to highly collisional. The limitation of the reconstructions to avoid spurious oscillations is made a posteriori using polynomial order decrementation. This later is based on so-called detection criteria which filter problematic cells which need limitation from acceptable cells which are updated with an optimal accurate scheme.

Giacomo Dimarco
Università degli Studi di Ferrara
dmrgcm@unife.it

MS173

Practical Numerical Methods for Solving the Boltzmann Transport Equation in Nuclear Reactor Analysis

The numerical solution of the linearized form of the Boltzmann Transport Equation (BTE), as applied to the transport of neutrons within a "multiplicative" medium, results in the characterization of various macroscopic quantities of interest such as criticality, power distribution, and compositional changes in a nuclear reactor. For the purpose of practical reactor analysis and design, various physical approximations are introduced in order to reduce the size of the neutral-particle phase-space. The goal of this talk is twofold: first, to review current standard numerical methods used for the solution of BTE, and secondly, to present recent progress in the spatial discretization of the BTE in the context of lattice physics calculations. The Method of Characteristics (MoC), used to solve BTE in lattice physics, has become ubiquitous in design calculations. However, the order of accuracy of this scheme, as a function of mesh size, remains first-order. In order to relax this limitation, a spatially-linear scheme is introduced in the MoC scheme in order to obtain second-order accuracy. Although general high-order discretizations have been derived, the spatially-linear scheme remains an optimal compromise between run time and storage, for an equal level of accuracy. Numerical results will be presented to support the conclusions sketched above and realistic numerical simulations will presented and discussed.

Rodolfo Ferrer
Studsvik
rodolfo.ferrer@studsvik.com

MS173

A Deterministic-Particle Transport Solver for Scale-Bridging Simulation of Thermal Radiative Transfer

A high-order, low-order (HOLO) algorithm is a moment-based, scale-bridging algorithm for transport equations. HOLO algorithms have recently made significant efficiency and accuracy impacts in thermal radiative transfer (TRT) problems. The goal of this work is to develop an efficient, robust, Lagrangian, characteristic-based transport solver for TRT problems within the context of a HOLO algorithm. A preliminary example has shown that a new HO solver can achieve a factor of 100 reduction in computational effort compared to both Monte Carlo and Sn for and equivalent accuracy.

Hyeongkae Park

Los Alamos National Laboratory
hkpark@lanl.gov

MS173

A Hierarchy of Hybrid Numerical Methods for Multi-Scale Kinetic Equations

In this work, we construct a hierarchy of hybrid numerical methods for multi-scale kinetic equations based on moment realizability matrices, a concept introduced by Levermore, Morokoff and Nadiga. Following such a criterion, one can consider hybrid scheme where the hydrodynamic part is given either by the compressible Euler or Navier-Stokes equations, or even with more general models, such as the Burnett or super-Burnett systems. The efficiency of the method obtained can sometimes match the one of a fluid method, with the accuracy of a kinetic one.

Thomas Rey
University of Maryland
trey@cscamm.umd.edu

Francis Filbet
University of Lyon
filbet@math.univ-lyon1.fr

MS174

HJB-POD Feedback Control for Advection-Diffusion Equations

We present an algorithm for the approximation of an infinite horizon optimal control problem for advection-diffusion equations. The method is based on the coupling between a POD representation of the solution and a Dynamic Programming approximation scheme for the corresponding stationary Hamilton-Jacobi equation. We discuss several features regarding the selection of the snapshots for advection dominated terms. Some test problems are presented to illustrate the method.

Alessandro Alla
Department of Mathematics
University of Hamburg, Germany
alessandro.alla@uni-hamburg.de

Michael Hinze
Universität Hamburg
Department Mathematik
michael.hinze@uni-hamburg.de

MS174

Application of Discrete Empirical Interpolation Method to Reduced Order Modeling of Nonlinear Parametric Systems

We study two approaches of applying discrete empirical interpolation method (DEIM) in the finite element context, one applied to the assembled and the other to the unassembled form of the nonlinearity. We carefully examine how DEIM is applied in each case, and the substantial efficiency gains obtained by the DEIM. In addition, we demonstrate how to apply DEIM to obtain ROMs for a class of parameterized system that arises, e.g. in shape optimization.

Harbir Antil
George Mason University
Fairfax, VA
hantil@gmu.edu

MS174

Reduced Order Models for Nonlinear PDE-Constrained Optimization Problems in Fluid Dynamics

We propose a model order reduction framework for parametrized optimization problems constrained by nonlinear PDEs. First, we build – by means of either RB-greedy or POD methods – the reduced order model following a suitable all-at-once optimize-then-reduce paradigm. Then, combining rigorous error bounds with some heuristic computational strategies, we provide cheap yet reliable a posteriori error estimates. The methodology is applied to the boundary optimal control of parametrized Navier-Stokes equations.

Federico Negri
Mathematics Institute of Computational Science and Eng, EPFL
federico.negri@epfl.ch

Andrea Manzoni
EPFL, MATHICSE-CMCS
Switzerland
andrea.manzoni@epfl.ch

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

MS174

Reduced Basis Method for Hamilton-Jacobi-Bellman Equations

We consider the Hamilton-Jacobi-Bellman (HJB) equation of the form $-\partial_t v + \sup_\alpha \left\{ -a(\alpha;\mu)\Delta v - b(\alpha;\mu) \cdot \nabla v - f(\alpha;\mu) \right\} = 0$ for numerous different parameters $\mu \in \mathcal{D} \subset \mathbb{R}^p$ by applying the Reduced Basis Method. In our approach, the HJB equation is separated into an equation which determines v and another which yields α. To obtain an error estimator, we use a space-time formulation of these equations and apply the Brezzi-Rappaz-Raviart theory for quadratic non-linearities. The implementation shows that this error estimator is efficient for the chosen examples.

Sebastian Steck
Institute for Numerical Analysis
University of Ulm, Germany
sebastian.steck@uni-ulm.de

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

MS175

Application of Algebraic Multigrid (PETSc) for Block Structured Adaptive Mesh Refinement Applications (Chombo)

We report on progress in using algebraic multigrid (AMG) methods in the numerical library PETSc for challenging problems in block structured adaptive mesh refinement (BSAMR) applications that use the Chombo library.

Chombo's built-in geometric multigrid (GMG) solvers are fast for simple operators but for problems with complex geometry GMG is not effective and AMG is an effective solution. We discuss new capabilities in Chombo for constructing matrices, required for AMG, from BSAMR problems.

Mark Adams
Lawrence Berkeley Laboratory
mfadams@lbl.gov

MS175

Solvers and Error Control for Atmospheric Column Physics

Atmospheric column physics involves coupling of nonsmooth processes such as phase change and nonlocal/stiff processes like radiation and precipitation, along with strict positivity constraints. The splitting schemes currently used by CAM are not convergent in time and have been calibrated to compensate for systematic numerical errors. In this talk, we investigate the use of more rigorous ODE/DAE approaches using PETSc and Sundials on the basis of accuracy and efficiency.

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory and CU Boulder
jed@jedbrown.org

MS175

FASTMath Unstructured Mesh (MOAB) Solver (PETSc) Interactions

High fidelity applications typically rely on unstructured meshes representing complex geometry, and interoperable solver interfaces are critical for optimal mesh based traversal and operator assembly procedures. We present developments on a discretization manager *DM* implementation that utilizes array-based mesh database *MOAB* that exposes scalable mesh handling capabilities to simulators, resulting in reduced memory and communication overhead. Several scalar and multi-component solver demonstrations employing combinations of geometric multigrid and algebraic preconditioners is discussed.

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

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

MS175

Scalable Adaptive ImEx Integration with ARKode and HYPRE

Computational science increasingly focuses on simulations of complex systems, typically involving interacting physical processes and large data/computational requirements. In this talk, we consider the solution of a system of interacting advection-diffusion-reaction equations, decomposing these into stiff (reaction-diffusion) and non-stiff (advection) portions. For this problem, we apply two open-source solver libraries within the FASTMath project. The ARKode solver (part of the SUNDIALS library suite) enables high order and stable ImEx simulations, while HYPRE excels at massively-parallel linear systems and is used to precondition the implicit solves within ARKode.

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

MS176

Aerodynamic Simulations Using a High-Order Discontinuous Galerkin Solver

A parallel high-order Discontinuous Galerkin solver is developed to simulate aerodynamic problems. The solver capabilities include: mixed hybrid elements, hp-adaption, and an exact linearization used in an implicit preconditioned Newton-Krylov solver. Preconditioners include Gauss-Seidel relaxation, line implicit Jacobi, and ILU(0). Flows with strong shocks are simulated on simple geometries using a PDE based artificial viscosity. Also, turbulent flows are simulated on 2D airfoils and a full 3D aircraft using the new negative Spalart-Almaras RANS model.

Michael Brazell
University of Wyoming
mbrazell@uwyo.edu

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

MS176

Superconvergent HDG Methods with Symmetric Stress Approximations for Stokes Flow (and Linear Elasticity)

We present an abstract framework to obtain superconvergence of HDG methods with symmetric stress approximations (HDG-S methods) for the Stokes flow. Several 2D methods are constructed based on local enrichment of the stress spaces of existing (suboptimal) HDG-S methods. Numerical results are provided to show their superior performance. Extension of the results to linear elasticity with symmetric stress approximations is discussed. And challenges for constructing 3D superconvergent HDG-S methods are addressed.

Bernardo Cockburn, Guosheng Fu
School of Mathematics
University of Minnesota
cockburn@math.umn.edu, fuxxx165@umn.edu

MS176

Active Fluxes; A New High-Order Paradigm

We consider finite-volume schemes in which the interface fluxes are independent degrees of freedom, as proposed by van Leer in 1977. In the semi-discrete limit these are similar to, and in some cases identical to, Discontinuous Galerkin methods, but are naturally fully discrete and in that case have several advantages, being more accurate, faster and requiring less storage.

Philip L. Roe
Ann Arbor

philroe@umich.edu

MS176

Riemann-Solver-Free Space Time Discontinuous Galerkin Method for General Conservation Laws

We will talk about a Riemann-solver-free space-time discontinuous Galerkin method for general conservation laws. The method uses staggered space-time mesh to enforce space-time flux conservation in the DG framework. The resulting method is of high order and Riemann-solver free. The method is able to solve general conservation laws such as compressible Euler equations, shallow water equations and magnetohydrodynamics equations without the need of any type of Riemann solvers.

Shuang Z. Tu
Jackson State University
shuangzhang.tu@jsums.edu

MS177

ScalFMM: A Generic Parallel Fast Multipole Library

ScalFMM (Parallel Fast Multipole Library for Large Scale Simulations) offers all the functionalities needed to perform large parallel simulations while enabling an easy customization of the simulation components: kernels, particles and cells. We will present how we use our library on two kinds of application involving boundary integral representations of physical fields. The first one implements isotropic dislocation kernels for Dislocation Dynamics and the second a time dependent kernel for acoustic problems.

Pierre Blanchard
Inria
pierre.blanchard@inria.fr

Berenger Bramas, Olivier Coulaud
INRIA
berenger.bramas@inria.fr, olivier.coulaud@inria.fr

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

Laurent Dupuy
CEA-Saclay
laurent.dupuy@cea.fr

Arnaud Etcheverry
Inria
arnaud.etcheverry@inria.fr

Guillaume Sylvand
EADS CCR, Centre de Toulouse
guillaume.sylvand@airbus.com

MS177

PVFMM: A Parallel Fast Multipole Method for Volume Potentials

PVFMM is a scalable FMM library for computing solutions to constant coefficient elliptic PDEs (Poisson, Stokes, Helmholtz) on cubic domain with free-space and periodic boundary conditions. In this talk, we discuss performance optimizations including distributed and shared memory parallelism, integration with co-processors, vectorization and cache optimization. We also integrate our method with GMRES in PETSc to compute solutions to variable coefficient Stokes problems and demonstrate scalability to thousands of compute nodes on TACC's Stampede supercomputer.

Dhairya Malhotra
Institute of Computational Engineering and Sciences
The University of Texas at Austin
dmalhotra@ices.utexas.edu

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS177

Robust Implementation of Quadrature by Expansion (QBX)

QBX is a quadrature scheme for computing singular integrals that arise from the discretization of integral equations. In its standard (global) form, it fails to be robust for close to touching geometries. In this talk, we introduce a local variant of QBX. The local variant is computationally more expensive, and therefore should only be used judiciously when needed. We provide details of a robust hybrid implementation of global and local QBX which uses local QBX only in areas of complicated geometry.

Manas Rachh
Courant Institute
NYU
manasrachh@gmail.com

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

Michael O'Neil
New York University
oneil@cims.nyu.edu

Andreas Kloeckner
Department of Computer Science
University of Illinois at Urbana-Champaign
andreask@illinois.edu

MS177

ExaFMM – a Testbed for Comparing Various Implementations of the FMM

The software design space for the fast multipole method (FMM) is large, where various schemes exist for partitioning, communication, tree traversal, and translation operators. The optimal choice is problem and architecture dependent, and most FMM codes focus on a specific combination of the two. ExaFMM is a open source FMM code that provides the capability to explore this vast design space without sacrificing speed, scalability, or readability.

Rio Yokota
King Abdullah University of Science and Technology
rio.yokota@kaust.edu.sa

David E. Keyes

KAUST
david.keyes@kaust.edu.sa

Lorena A. Barba
Department of Mechanical and Aerospace Engineering
George Washington University
labarba@gwu.edu

MS178

A Posteriori Error Estimation in the Maximum Norm for Finite Element Methods

Adaptive finite element methods are designed to control the error in measuring a specific norm or functional of the solution. In some situations it is of interest to control the maximum error. In this talk we will discuss recent progress in developing a posteriori error estimates in the maximum norm for elliptic problems, with special focus on singularly perturbed reaction-diffusion equations. Time allowing, we will also discuss the sharpness of logarithmic factors which appear in our estimates.

Alan Demlow
Texas A&M University
demlow@math.tamu.edu

MS178

Robust Residual-Based a Posteriori Error Estimation for Interface Problems: Nonconforming Elements

A robust residual-based a posteriori error estimation for the non-conforming linear finite element approximation to the interface problem is studied. We introduce a new and direct approach, without using the Helmholtz decomposition, to analyze the reliability of the estimator. It is proved that our estimator is reliable with the constant independent of the jump of the interfaces, without the assumption that the diffusion coefficient is quasi-monotone. Numerical results are also presented.

Cuiyu He
Purdue University
he75@math.purdue.edu

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

Shun Zhang
City University of Hong Kong
shun.zhang@cityu.edu.hk

MS178

A PDE Approach to Fractional Diffusion: a Posteriori Error Analysis

We derive a computable a posteriori error estimator for the α-harmonic extension problem, which localizes the fractional powers of elliptic operators supplemented with Dirichlet boundary conditions. The derived estimator relies on the solution of small discrete problems on anisotropic cylindrical stars. It exhibits built-in flux equilibration and is equivalent to the error up to data oscillation. A simple adaptive algorithm is designed and numerical experiments reveal a competitive performance.

Abner J. Salgado
Department of Mathematics
University of Tennessee
asalgad1@utk.edu

MS178

Robust a Priori and a Posteriori Error Estimates for Diffusion Problems with Discontinuous Coefficients

For diffusion problems of discontinuous coefficients, the quasi-monotonicity assumption (QMA) is a very important condition to guarantee the robustness of problems independent of the coefficients. In this talk, new results of robust and optimal a priori and a posteriori error estimates for various finite element approximations of the diffusion problem with discontinuous coefficients with and without QMA are discussed. With new tools, we show that Conforming FEM is only optimal and robust in 2D with low regularity, while Mixed FEM, Nonconforming FEM, and DGFEM are both robust with respect to coefficients and optimal with respect to local regularities. For a posteriori error estimates, using nodal interpolations in stead of Clement interpolations, we show the robustness of CFEM, MFEM, NFEM, and DGFEM without GMA.

Shun Zhang
City University of Hong Kong
shun.zhang@cityu.edu.hk

MS179

Primal-Dual Newton Conjugate Gradients for Compressed Sensing Problems with Coherent and Redundant Dictionaries

Abstract not available at time of publication.

Kimon Fountoulakis
University of Edinburgh
k.fountoulakis@sms.ed.ac.uk

MS179

Augmented Lagrangian Methods for Large-Scale Nonlinear Optimization

Abstract not available at time of publication.

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

MS179

Column Generation Techniques for Large Mixed-Integer Programs

We study two-stage mixed-integer programs (MIP) that arise from co-generation in smart buildings. These MIP problems are beyond the capability of available MIP solvers due to a large number of binary variables in the second-stage problem. We develop a column generation approach that reduces problem size by building a coarser MIP model. The first-stage solutions are validated using a rolling horizon method. We test our approach using data sets from different types of buildings.

Fu Lin

Mathematics and Computer Science Division
Argonne National Laboratory
fulin@mcs.anl.gov

MS179

Convexification Methods for Sequential Quadratic Programming

Abstract not available at time of publication.

Elizabeth Wong
University of California, San Diego
Department of Mathematics
elwong@ucsd.edu

MS180

Gaussian-Localized Polynomial Interpolation (Hermite Function Interpolation) on a Finite Interval: Are Spectrally-Accurate Rbfs Obsolete?

Spectrally-accurate radial basis functions (RBFs) such as Gaussians often succeed where polynomial interpolation fails because of the Runge Phenomenon. Why? Comparative analysis of RBF and polynomial cardinal functions (nodal bases) shows that RBF cardinals are much more localized. This motivated us to abandon RBFs in favor of Gaussian-mollified polynomial cardinal functions: much more accurate, cheaper to construct and much better conditioned than RBFs. Using potential theory, we prove an exponential, geometric rate of convergence. On a uniform grid, the Runge Phenomenon is not completely annulled, but the Runge Zone is an order of magnitude smaller than for polynomial interpolation. Applications to solving differential equations, especially in complicated geometry, are in progress. A note was published as "Hermite Function Interpolation ..., *Appl. Math. Letters*, **26**, 10, 995-997 (2013).

John P. Boyd
University of Michigan
Ann Arbor, MI 48109-2143 USA
jpboyd@umich.edu

MS180

Automatic Multivariate Approximation

Fully automatic one-dimensional approximation can be implemented successfully using Chebyshev polynomial expansions, as demonstrated in the Chebfun software project. In higher dimensions, dimensionality and geometry are major new challenges. One proposal for automatic approximation in two and three dimensions using partition of unity, spectral methods, and radial basis functions will be presented.

Tobin Driscoll
University of Delaware
Mathematical Sciences
driscoll@udel.edu

MS180

A Windowed Fourier Method for Computations on the Sphere

A spectral method based on windowed Fourier approximations for computations on the sphere is presented. It relies on domain decomposition, such as in the cubed sphere, and is suitable for adaptive and parallel implementation. One of the advantages of this approach is that computations can be carried out using fast Fourier transforms on a nearly uniform grid. Approximations are obtained on overlapping domains and a global solution is obtained by weighted average.

Rodrigo B. Platte
Arizona State University
rbp@asu.edu

MS180

A Fast and Well-Conditioned Spectral Method

We describe a spectral method for the solution of linear ordinary differential equations that leads to well-conditioned matrices that are banded except for a few rows. We extend this to a spectral method for the solution of linear partial differential equations using low rank approximation and generalized Sylvester matrix equations. Techniques from automatic differentiation and preconditioning are employed to develop an automatic, adaptive, and spectrally accurate solver.

Alex Townsend
Oxford University
townsend@maths.ox.ac.uk

MS181

Recent Advances in Numerical Modelling of Thermo-Chemically Coupled Two-Phase Flow

Many processes in geodynamics involve a significant amount of active magmatism. Recent advances in computational magma dynamics have allowed formulating a novel multi-component two-phase flow model including a realistic visco-plastic rheology for host rock deformation and a basic thermodynamic disequilibrium model describing the simplified petrology of a multi-phase silicate material. The model has been implemented numerically using a parallel finite-difference method in 2-D for the use in studies of volatiles (H_2O, CO_2) in magmatic systems.

Tobias Keller
University of Oxford
Department of Earth Sciences
tobias.keller@earth.ox.ac.uk

MS181

Scalable Nonlinear Solvers for Magma Dynamics

We present new nonlinear solvers for magma dynamics problems based upon the nonlinear preconditioning framework in PETSc. We demonstrate their effectiveness on model problems, and examine a possible direction for convergence analysis.

Matthew G. Knepley
University of Chicago
knepley@ci.uchicago.edu

Richard F. Katz
University of Oxford
Richard.Katz@earth.ox.ac.uk

MS181

Multi-Scale Modelling of Granular Avalanches

It is important to be able to predict the distance to which a

hazardous natural flows travel. In dense flows the large particles segregate to the surface, where they are transported to the margins forming bouldery flow fronts. In natural flows these bouldery margins experience a much greater frictional force, leading to frontal fingering instabilities. A multiscale model for this fingering instability is compared in terms of cost and accuracy to full-scale discrete particle simulations.

Anthony R. Thornton
University of Twente
Multi-scale mechanics group
a.r.thornton@utwente.nl

MS182

Domain Specific Languages and Automated Code Generation: High Expressiveness and High Performance

Software should be carefully designed to closely follow the mathematical abstractions of the problem domain. We demonstrate this principle with FEniCS, which employs a domain-specific language, UFL, that mimics the mathematics of variational forms for finite element discretisations. By following this principle, we demonstrate how UFL and the FEniCS code generation techniques can be used to automatically derive extremely high-performance code for adjoints of PDEs, for use in sensitivity analysis, optimisation, and inverse problems.

Patrick E. Farrell
Department of Mathematics
University of Oxford
patrick.farrell@maths.ox.ac.uk

MS182

Moose: An Open Source Platform For Rapid Development of Multiphysics Simulation Tools

Many physical phenomena can be modeled with systems of partial differential equations. Some examples include nuclear reactors, geothermal flow, microstructural evolution, and fluid-structure interaction. Solving systems of nonlinear equations has typically been achieved with custom software or by combining existing simulation tools. The open-source Multiphysics Object-Oriented Simulation Environment (MOOSE; mooseframework.org) employs a different approach: it provides a generic, common platform for solving multiphysics problems. This technique allows scientists and engineers to focus on the physics while the framework manages the task of parallel, nonlinear solver development. A discussion of the platforms capabilities and software development model will be followed by several example applications in nuclear physics, geophysics, chemistry, and material science.

Derek R. Gaston
Idaho National Laboratory
friedmud@gmail.com

Cody Permann
Center for Advanced Modeling and Simulation
Idaho National Laboratory
cody.permann@inl.gov

David Andrs, John Peterson, Andrew Slaughter
Idaho National Laboratory
david.andrs@inl.gov, jw.peterson@inl.gov, andrew.slaughter@inl.gov

MS182

Project Jupyter: a Language-Independent Architecture for Cse, from Interactive Computing to Reproducible Publications

Project Jupyter is the evolution of the IPython interactive computing system into language-agnostic components to support all aspects of computational research. Jupyter abstracts the basic elements of interactive computing into openly documented file formats and protocols. Using these, "Jupyter kernels" can execute code in any language and communicate with clients that range from simple terminals to the rich web-based Jupyter Notebook that supports code, results, rich media, text and mathematics.

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

MS183

Bayesian Global Optimization of Expensive Functions with Low-Dimensional Noise

Motivated by the design of cardiovascular bypass grafts using computationally-expensive physics-based stochastic simulators, we consider optimizing an objective function that is an integral of some expensive-to-compute function over a low-dimensional space. While such problems can be solved using an optimization solver for expensive deterministic functions, or for expensive noisy functions, neither uses the low-dimensional structure of the noise. We provide a new Bayesian global optimization method that exploits this problem structure to improve performance.

Jing Xie
Cornell University
jx66@cornell.edu

Sethuraman Sankaran
Heartflow, Inc.
ssankaran@heartflow.com

Abhay Ramachandra
Department of Mechanical and Aerospace Engineering
University of California, San Diego
abbangal@ucsd.edu

Saleh Elmohamed
Cornell University
elmohamed@cornell.edu

Alison Marsden
Department of Mechanical and Aerospace Engineering
University of California, San Diego
amarsden@ucsd.edu

Peter I. Frazier
Cornell University
pf98@cornell.edu

MS183

Modeling An Augmented Lagrangian for Improved

Blackbox Constrained Optimization

We propose a combination of response surface modeling, expected improvement, and the augmented Lagrangian numerical optimization framework, allowing the statistical model to think globally and the augmented Lagrangian to act locally. We focus on problems where the constraints are the primary bottleneck, requiring expensive simulation to evaluate and substantial modeling effort to map out. In that context, our hybridization presents a simple yet effective solution that allows existing objective-oriented statistical approaches, like those based on Gaussian process surrogates and expected improvement heuristics, to be applied to the constrained setting with minor modification.

Robert Gramacy
University of Chicago
rbgramacy@chicagobooth.edu

Genetha Gray
Sandia National Laboratories
gagray@sandia.gov

Sebastien Le Digabel
GERAD - Polytechnique Montreal
Sebastien.Le.Digabel@gerad.ca

Herbert Lee
UC Santa Cruz
herbie@soe.ucsc.edu

Pritam Ranjan
Acadia University
pritam.ranjan@acadiau.ca

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

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS183

Topology Optimization of a Permanent-Magnet Synchronous Machine under Uncertainties

The modeling and simulation of a permanent-magnet synchronous machine requires the application of Maxwell's equations. Within a two-dimensional finite element model, the shapes of rotor poles, which are represented by zero-level sets, are optimized by a redistribution of an iron and a magnetic material using an iterative process. More specifically, we consider a stochastic problem, where random fields due to unknown variations of material properties are taken into account to model the propagation of uncertainties. Furthermore, in a robust optimization formulation, a multi-objective functional, which includes the mean and the standard deviation, is minimized subject to constraints. For this purpose, we apply methods dedicated to stochastic problems such as e.g. polynomial chaos expansions. Finally, the results of numerical simulations show that the proposed methods are robust and efficient.

Roland Pulch
University of Greifswald
pulchr@uni-greifswald.de

Piotr Putek
Bergische Universität Wuppertal
putek@math.uni-wuppertal.de

MS183

A Model and Variance Reduction Method for Computing Statistical Outputs of Stochastic Partial Differential Equations

We present a model and variance reduction method for computing statistical outputs of stochastic PDEs. We combine the hybridizable discontinuous Galerkin (HDG) and the reduced basis discretization of elliptic PDEs with a multilevel variance reduction method, exploiting the statistical correlation among the HDG and reduced basis approximations to accelerate the convergence of Monte Carlo simulations. We develop *a posteriori* error estimates for the statistical outputs, and propose an optimal selection of multilevel structure.

Ferran Vidal-Codina, Cuong Nguyen, Jaime Peraire
Massachusetts Institute of Technology
fvidal@mit.edu, cuongng@mit.edu, peraire@MIT.EDU

Michael B. Giles
Mathematical Institute
Oxford University
Mike.Giles@maths.ox.ac.uk

MS184

Accurate All-Electron Electronic Structure Theory for Large Systems

This contribution focuses on algorithmic developments towards efficient, accurate all-electron methods for computational simulations of materials and nanostructures from first principles. We demonstrate (i) accurate numeric atom-centered basis sets for ground-state density functional theory as well as perturbative methods such as GW or RPA, (ii) the open-source, massively parallel eigenvalue solver library ELPA and developments connected to it, essential for parallel scalability towards system sizes of ∼1000s of atoms, (iii) first steps towards a mixed, load-balanced CPU-GPU implementation real space integrals required in electronic structure theory. The methods described are implemented in the general-purpose all-electron electronic structure code FHI-aims [http://aims.fhi-berlin.mpg.de]

Volker Blum
Duke University
volker.blum@duke.edu

MS184

Enabling Large-Scale Hybrid Density Functional Theory Calculations

Hybrid density functional theory (DFT), in conjunction with a treatment of van der Waals/dispersion interactions, provides a substantial improvement over popular DFT functionals based on the generalized gradient approximation, in the modeling of condensed-phase molecular systems such as liquid water, albeit with a prohibitively large computational cost. In this work, we will demonstrate how novel theoretical and algorithmic developments, coupled with efficient utilization of supercomputer architectures, enable large-scale hybrid DFT calculations on condensed-

phase molecular systems of interest.

Robert A. DiStasio, Jr.
Princeton University
distasio@princeton.edu

MS184

Finite Elements for Large, Accurate Quantum Mechanical Materials Calculations: from Classical to Enriched to Discontinuous

We discuss recent developments in finite-element (FE) based methods for the solution of the Kohn-Sham equations that have made possible smaller basis sets and larger calculations than possible by current state-of-the-art planewave (PW) based methods, in some cases by an order of magnitude or more. We begin with classical FE based approaches, then discuss recent enriched partition-of-unity FE (PUFE) methods, which build known atomic physics into the basis while retaining strict locality and systematic improvability. By incorporating known physics, these bases can achieve the required accuracies with an order of magnitude fewer degrees of freedom (DOF) than traditional PW based methods. However, with such enrichment comes more expensive quadrature and some degree of ill-conditioning. By incorporating not only local-atomic but also environmental physics into the basis, recent Discontinuous Galerkin (DG) based approaches can achieve larger reductions in DOFs still, while retaining strict locality and systematic improvability. Most notably, the DG formulation allows for orthonormality as well, alleviating conditioning issues and allowing for the solution of standard rather than generalized discrete eigenproblems in the critical N^3 scaling step of the Kohn-Sham solution. Accurate quantum mechanical forces and molecular dynamics have been demonstrated. Incorporation of Pole Expansion and Selected Inversion (PEXSI) has been undertaken to eliminate the need for diagonalization entirely. We conclude with an outlook and applications interests going forward.

John Pask
Lawrence Livemore National Lab
pask1@llnl.gov

MS184

Enabling Large Scale LAPW DFT Calculations by a Scalable Iterative Eigensolver

In LAPW-based methods a sequence of dense generalized eigenvalue problems appears. Traditionally these problems were solved using direct eigensolvers from standard libraries like ScaLAPACK. We developed a subspace iteration method pre-conditioned with Chebyshev polynomials of optimal degree (ChFSI). This algorithm is consistently competitive with direct eigensolvers and greatly enhance performance and scalability. ChFSI is included in the FLEUR software and improves its scaling behaviour for calculations of large physical systems on modern supercomputers.

Daniel Wortmann
Institute for Asvanced Simulation
Forschungszentrum Juelich
d.wortmann@fz-juelich.de

Edoardo A. Di Napoli
Juelich Supercomputing Centre
e.di.napoli@fz-juelich.de

Mario Berljafa
School of Mathematics
The University of Manchester
m.berljafa@maths.man.ac.uk

MS185

Solving Optimal Feedback Control Problems for Partial Differential Equations Using Adaptive Sparse Grids

An approach to solve finite time horizon optimal feedback control problems for partial differential equations using adaptive sparse grids is proposed. A semi-discrete optimal control problem is introduced and the feedback control is derived from the corresponding value function. The value function can be characterized as the solution of an evolutionary Hamilton-Jacobi Bellman (HJB) equation which is defined over a state space whose dimension is equal to the dimension of the underlying semi-discrete system. Besides a low dimensional semi-discretization it is important to solve the HJB equation efficiently to address the curse of dimensionality. We propose to apply a semi-Lagrangian scheme using spatially adaptive sparse grids. Sparse grids allow the discretization of the high(er) dimensional value functions arising in the numerical scheme since the curse of dimensionality of full grid methods arises to a much smaller extent. For additional efficiency an adaptive grid refinement procedure is explored. We present several numerical examples studying the effect of the parameters characterizing the sparse grid on the accuracy of the value function and optimal trajectories. Furthermore we analyze the behaviour of the trajectories in case of noise.

Jochen Garcke
University of Bonn
garcke@ins.uni-bonn.de

Axel Kroner
RICAM, Austria
axel.kroener@ricam.oeaw.ac.at

MS185

Dimension-Independent, Likelihood-Informed Mcmc Sampling Algorithms for Bayesian Inverse Problems

Many Bayesian inference problems require exploring the posterior distribution of high-dimensional parameters, which in principle can be described as functions. By exploiting the intrinsic low dimensionality of the likelihood function, we introduce a suite of MCMC samplers that can adapt to the complex structure of the posterior distribution, yet are well-defined on function space. Posterior sampling in nonlinear inverse problems arising from various partial differential equations and also a stochastic differential equation are used to demonstrate the efficiency of these dimension-independent likelihood-informed samplers.

Kody Law
SRI UQ Center, CEMSE, KAUST
kodylaw@gmail.com

Tiangang Cui, Youssef M. Marzouk
Massachusetts Institute of Technology
tcui@mit.edu, ymarz@mit.edu

MS185

Numerical Solution of Elliptic Diffusion Problems

on Random Domains

In this talk, we provide regularity results for the solution to elliptic diffusion problems on random domains. Especially, based on the decay of the Karhunen-Loeve expansion of the domain perturbation field, we establish rates of decay which imply the tractability of the Quasi-Monte Carlo method. By taking into account only univariate derivatives, the regularity results can considerably be sharpened in order to show also the applicability of the stochastic collocation method and related rates of convergence. We moreover employ parametric finite elements to compute the solution of the diffusion problem on each particular realization of the domain generated by the perturbation field. This simplifies the implementation and yields a non-intrusive approach. The theoretical findings are complemented by numerical examples.

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

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

MS185

Data and Uncertainties: Representation of High-Dimensional Dependencies Using Adaptive Sparse Grids

Wherever simulations depend on parameters, high-dimensional problems arise, and the representation of high-dimensional dependencies based on data and/or uncertainties is required. Where the dependencies are non-smooth, adaptive refinement and basis functions with local support are able to avoid Gibbs phenomenon. Sparse grids, together with suitable refinement strategies, provide such a hierarchical and incremental approach. We will discuss some of their properties and give best practice examples. Examples stem from uncertainty quantification, data mining, and density estimation.

Dirk Pfluger
University of Stuttgart
Dirk.Pflueger@ipvs.uni-stuttgart.de

Fabian Franzelin
Universität Stuttgart
fabian.franzelin@ipvs.uni-stuttgart.de

MS186

Incorporating Error Detection and Recovery into Hierarchically Semi-Separable Matrix Operations

Hierarchically semi-separable (HSS) matrix factorizations are promising components for high-performance linear solvers. To construct a fault-tolerant HSS solver, we have identified several checksum relationships that are preserved during HSS matrix-vector multiplication. Whenever these checksum assertions fail, the Containment Domains library restores the faulty arrays and re-executes the necessary region of the code. The timing profile of our algorithm is used to parametrize a Markov model and determine the optimal granularity of the checksum tests.

Brian Austin
NERSC
Lawrence Berkeley National Laboratory
baustin@lbl.gov

Alex Druinsky, Xiaoye Sherry Li
Computational Research Division
Lawrence Berkeley National Laboratory
adruinsky@lbl.gov, xsli@lbl.gov

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

Eric Roman
Lawrence Berkeley National Laboratory
eroman@lbl.gov

MS186

Parallel Spectral Element-Based Agglomeration Algebraic Multigrid for Porous Media Flow

We present an element based algebraic multigrid strategy for scalable parallel simulation of porous media flow. Coarse basis functions in the multigrid algorithm are adapted to the high contrast coefficients by solving sparse or dense local eigenvalue problems. An advantage of the approach is that setup and application are based on sparse matrix-vector and matrix-matrix products, so that any resilient implementations of these components will be inherited by the entire algorithm.

Andrew T. Barker
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
barker29@llnl.gov

Delyan Kalchev
University of Colorado
delyank@gmail.com

Panayot Vassilevski
Lawrence Livermore National Laboratory
vassilevski1@llnl.gov

Umberto E. Villa
Center for Advanced Scientific Computing
Lawrence Livermore National Laboratory
villa13@llnl.gov

MS186

Attaining High Arithmetic Intensity in Finite-volume Methods through High-order Quadratures

Finite-volume discretizations in production codes are typically only second-order accurate, resulting in methods with low arithmetic intensity. The widening performance gap between processor and memory motivates developing schemes that do more work for less data motion. We present an arithmetic intensity analysis for finite-volume methods with high-order flux approximations. The model suggests the flops-to-byte ratio of the methods increases rapidly with order of accuracy. We also discuss performance measurements of the methods on current multicore hardware. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Lawrence Livermore National Security, LLC.

John Loffeld
University of California
Merced
loffeld1@llnl.gov

Jeffrey A. Hittinger
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
hittinger1@llnl.gov

MS186

Controlling Numerical Error in Particle-In-Cell Simulations of Collisionless Dark Matter

Particle-in-cell (PIC) methods are commonly used to simulate the gravitational evolution of collisionless dark matter in cosmological settings. However, such methods are known to be prone to numerical error on scales below the mean inter-particle separation. An understanding of these errors is crucial, both for correctly interpreting simulation results, and for designing higher-order PIC methods. We discuss the use of two techniques: regularization and adaptive phase-space remapping, to improve the accuracy of cosmological PIC simulations.

Andrew Myers
Computational Research Division
Lawrence Berkeley National Laboratory
atmyers2@gmail.com

Brian Van Straalen
Lawrence Berkeley National Laboratory
Compuational Research Division
bvstraalen@lbl.gov

Colella Phillip
LBNL
pcolella@lbl.gov

MS187

Accelerating the Solution of Inverse Problems Using Reduced-Order Models

Inverse problems are a special class of PDE-constrained optimization problems for which model parameters are estimated. This inverse problem often needs to be solved in real-time based on measurements obtained on the field. The computational complexity associated with the optimization problem of interest generally prevents its fast solution. To accelerate its solution, a strategy based on a database of pre-computed reduced-order model is here chosen. Special attention will be given to the training of such a database in the case of high-dimensional parameter spaces.

David Amsallem
Stanford University
amsallem@stanford.edu

MS187

Reduced Basis Method for Uncertainty Quantification Problems: A Recent Update

In this talk, we present a recent update of the development and application of reduced basis method for forward and inverse uncertainty quantification (UQ) problems. Specific topics include reduced basis method for time-dependent Bayesian inverse problem, time-dependent optimal control problems, convergence of reduced basis method for high-dimensional or possibly infinite-dimensional problems, etc.

Peng Chen
ETH Zurich
peng.chen@sam.math.ethz.ch

Gianluigi Rozza
SISSA, International School for Advanced Studies
Trieste, Italy
gianluigi.rozza@sissa.it

MS187

Recent Advances in Reduced Order Modelling in Computational Fluid Dynamics within EU-MORNET COST Activities

We provide some recent updates about reduced order modelling techniques in computational fluid dynamics with a special focus on flow instabilities and bifurcations for time dependent non-linear flows. A special attention is devoted to parameter sampling techniques and approximation stability. A posteriori error bounds are recalled. A cardiovascular application is presented as benchmark example. This work wants to provide a general overview on some recent tasks of the EU-MORNET COST activities (European Union MOdel Reduction NETwork, Cooperation in Science and Technology).

Gianluigi Rozza
SISSA, International School for Advanced Studies
Trieste, Italy
gianluigi.rozza@sissa.it

Giuseppe Pitton
SISSA, International School for Advanced Studies Trieste
Italy
gpitton@sissa.it

Annalisa Quaini
Department of Mathematics, University of Houston
quaini@math.uh.edu

MS187

A Minimum-residual Mixed Reduced Basis Method: Exact Residual Certification and Simultaneous Finite-element Reduced-basis Adaptive Refinement

We present a reduced basis method for parametrized PDEs certified by a dual-norm bound of the residual computed not in the typical finite-element "truth' space but rather in an infinite-dimensional function space. The bound, which admits an efficient offline-online decomposition, builds on a minimum-residual mixed finite-element method and an associated reduced-basis approximation. The method, combined with a spatio-parameter adaptation strategy, yields an online system that meets any user-specified residual tolerance for any parameter value.

Masayuki Yano
Massachusetts Institute of Technology
myano@mit.edu

MS188

A New Look at Global Error Estimation in Differ-

ential Equations

Global error represents the actual discretization error resulting after solving a system of differential equations. I will introduce new time-stepping methods with built-in global error estimates for ordinary and differential algebraic equations. Calculating and controlling a posteriori errors is an expensive process, and therefore in practice only the (local) error from one step to the next is used to estimate the global errors. However, local error estimation is not always suitable and may lead to error underestimation. I will review several strategies for a posteriori error estimation and discuss new approaches that generalizes the classical strategies.

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS188

On the Construction of Robust Additive Runge-Kutta Methods

Space discretization of some PDEs gives rise to systems of ODEs in additive form whose terms have different stiffness properties. Sometimes, the solution to these PDEs have qualitative properties which are relevant in the context of the problem. In these cases, it is convenient to preserve numerically these properties. In this talk we show how robust schemes can be constructed for this class of problems. We will also show their performance on nontrivial problems.

Inmaculada Higueras
Departamento de Ingenieria Matematica e Informatica
Universidad Publica de Navarra
higueras@unavarra.es

MS189

A Sparse Multiresolution Regression Framework for Uncertainty Quantification

A novel nonintrusive, i.e., sampling-based, framework is presented for approximating stochastic solutions of interest admitting sparse multiresolution expansions. The coefficients of such expansions are computed via greedy approximation techniques that require a number of solution realizations smaller than the cardinality of the multiresolution basis. The effect of various random sampling strategies is investigated. The proposed methodology is verified on a number of benchmark problems involving nonsmooth stochastic responses.

Daniele E. Schiavazzi
University of California, San Diego
dschiavazzi@ucsd.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

MS189

Adaptive Compressive Sensing Method for Uncertainty Quantification

The standard generalized polynomial chaos (gPC) method in uncertainty quantification (UQ) selects basis functions based on the randomness of the input. However, this may not be an efficient way. We propose a new method to construct the basis functions based on the output of the system. With the new gPC expansion, the system has a sparser representation, hence sparse recovery approaches, e.g., compressive sensing technique, can be more effective.

Xiu Yang
Pacific Northwest National Laboratory
xiu.yang@pnnl.gov

Xiaoliang Wan
Louisiana State University
Department of Mathematics
xlwan@math.lsu.edu

Huan Lei, Guang Lin
Pacific Northwest National Laboratory
huan.lei@pnnl.gov, guang.lin@pnnl.gov

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS189

Stochastic Collocation Methods Via L1 Minimization

In this talk, we discuss the stochastic collocation methods via L1 minimization. The motivation it to construct polynomial approximations for parametric functions. Two sampling strategies, that is, the random sampling and the deterministic sampling methods will be introduced. We shall also discuss how to handle derivative information in this framework.

Tao Zhou
Institute of Computational Mathematics, AMSS
Chinese Academy of Sciences
tzhou@lsec.cc.ac.cn

Ling Guo
Department of Mathematics, Shanghai Normal
University, China
lguo@shnu.edu.cn

Dongbin Xiu
Purdue University
dxiu@purdue.edu

MS190

Development of a Time-Dependent Ice Flow Model Adjoint and Its Applications

The adjoint approach is quite popular in fitting glacial ice flow models to observations. For the most part, however, the approaches are time-independent. As remote sensing data becomes more frequent and more complete, it is fair to

expect that time-dependent models of ice dynamics might be constrained with transient data, providing better capture of transient state, and a better representation of unknown properties, such as bed topography or sub-ice shelf melt.

Daniel Goldberg
University of Edinburgh
dan.goldberg@ed.ac.uk

Patrick Heimbach
Massachusetts Institute of Technology
heimbach@mit.edu

MS190

Parallel 4D Variational Data Assimilation

A parallel algorithm to solve the inverse problem associated with the 4D Var Data assimilation problem is described. The inverse problem is solved in the variational framework. The optimization required to evaluate the analysis is performed using augmented Lagrangian technique. The evaluation of the cost function and the derivative information required to perform the optimization is computed in parallel using forward and adjoint sensitivity analysis.

Vishwas Hebbur Venkata Subba Rao
Virginia Polytechnic Institute and State University
visrao@vt.edu

Adrian Sandu
Virginia Polytechnic Institute and
State University
sandu@cs.vt.edu

MS190

An Adjoint Based Analysis of the Physical Drivers of Uncertainty in Air-Sea Exchange and Ocean Draw Down of Co2

We describe the use of an innovative adjoint approach to quantify and attribute uncertainty in model estimate of air-sea CO2 exchange. The adjoint approach provides a computationally efficient way to capture spatio-temporal variability in the sensitivity of net air-sea exchange to physical processes represented in a model. In this study we focus on diapycnal mixing in the upper ocean, which can vary widely in space and time. The processes by which the ocean responds to disequilibrium with atmospheric CO2 are key attributes of the Earth system decadal-centenial-millenial response to changes in CO2 emissions. The study provides quantitative insight into where and how diapycnal mixing processes contribute to that response. The study makes extensive use of the open-source automatic differentiation tool, OpenAD.

Chris Hill
Earth, Atmospheric and Physical Sciences
M.I.T.
cnh@mit.edu

Oliver Jahn
MIT
jahn@mit.edu

Jean Utke
Argonne National Laboratory
utke@mcs.anl.gov

MS190

Improving the Efficiency of the Adjoint of Fixed-Point Iterations

Reverse mode automatic differentiation (AD) computes the adjoints of codes precisely and efficiently. The reverse mode employs checkpointing to store data between its forward and the reverse computational sweeps. When the code contains fixed point iterations, unecessary checkpointing can result in excessive memory or disk usage. Reformulating the adjoint of the fixed point iteration, drastically reduces the usage. An implementation of the reformulation and its use in an ice sheet model will be presented.

Sri Hari Krishn Narayanan
Argonne National Laboratory
snarayan@mcs.anl.gov

Daniel Goldberg
University of Edinburgh
dan.goldberg@ed.ac.uk

Paul D. Hovland
Argonne National Laboratory
MCS Division, 221/C236
hovland@mcs.anl.gov

MS191

Multivariate Weighted Least-squares using Monte Carlo Samples

We propose and analyze an algorithm for computing polynomial discrete least-squares approximations to high-dimensional parameterized systems. Our sample grid for discrete least-squares is generated using Monte Carlo samples, and we show that our method exhibits several asymptotically optimal properties. Stability and accuracy can be established when the number of samples scales log-linearly with the dimension of the approximation space. This analysis relies on results from the fields of orthogonal polynomial and pluripotential theory. Our method is straightforward and efficient to apply, addresses approximations for general random variables with bounded and unbounded range, and easily extends to cases of epistemic uncertainty when the distribution of the parametric random variable is not known. We validate the algorithm on several low- and high-dimensional examples that are common in the uncertainty quantification community.

Akil Narayan
University of Massachusetts Dartmouth
akil.narayan@umassd.edu

MS191

Title Not Available at Time of Publication

Abstract not available at time of publication.

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

MS191

Local Polynomial Chaos Methods for High Dimen-

sional SPDE

We present a localized polynomial chaos expansion for PDE with random inputs. We focus on problems with input random fields of short correlation length, resulting in high dimensional random inputs. The local polynomial chaos method employs a domain decomposition technique to approximate the stochastic solution locally and independently with very low dimensions and then obtains the correct global solution via sampling. The drastic dimensional reduction makes the method highly efficient for practical problems.

Yi Chen
Purdue University
chen411@purdue.edu

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Xueyu Zhu, Dongbin Xiu
University of Utah
xzhu@sci.utah.edu, dongbin.xiu@utah.edu

MS191

Uncertainty Propagation Using Infinite Mixture of Gaussian Processes and Variational Bayesian Inference

Uncertainty propagation (UP) is a very challenging mathematical and computational problem. Among other things, UPs difficulty is due to the limited number of model evaluations, the curse of dimensionality, discontinuities, and multivariate responses with non-trivial correlations. In order to deal with all these problems simultaneously, we develop an infinite mixture of multi-output Gaussian process model. We train the model using variational Bayesian inference and we obtain highly competing results in porous flow problems.

Peng Chen, Nicholas Zabaras
Cornell University
pc423@cornell.edu, nzabaras@gmail.com

Ilias Bilionis
Purdue University
ebilionis@gmail.com

MS192

The Collective Impact of Social Factors and Interventions on the Dynamics of Reported Narcotic-Related Criminal Cases in the Community Areas of Chicago

Abstract not available at time of publication.

Maryam Khan
SAL Mathematical Computational Modeling Science Center
Arizona State University
maryam.khan.1@asu.edu

MS192

An Effective Community-based Approach to Mitigate Sybil Attacks in Online Social Networks

The popularity of online social networks (OSNs) has resulted in increasing Sybil attacks where an adversary forges many fake identities (called Sybils) to disrupt or control the normal functioning of the system. Proposed attack mitigation schemes primarily work by computing the landing probability or statistical distribution of the visiting frequency of short random walks and are dependent on accurate trusted nodes identification. In this talk, we will present SybilExposer, an algorithmic framework which addresses these limitations. Our experiments on several large graphs validate the effectiveness of SybilExposer.

Satyajayant Misra
Department of Computer Science
New Mexico State University
misra@cs.nmsu.edu

MS192

The Dynamics of Co-Evolution of Health Behaviors in College Population

Abstract not available at time of publication.

Anuj Mubayi
Arizona State University
anuj.mubayi@asu.edu

MS192

Analysis of Information Diffusion on Social Networks

Abstract not available at time of publication.

Daniel Romero
School of Information
University of Michigan
drom@umich.edu

MS193

Fluctuating Hydrodynamics of Suspensions of Rigid Particles

We develop a computational method for fluid-structure coupling at small Reynolds numbers that consistently includes the effects of thermal fluctuations. This is important in problems involving Brownian rigid and semi-rigid structures immersed in a fluid. We couple an immersed-boundary Lagrangian representation of rigid bodies to a fluctuating finite-volume fluid solver. We handle complex rigid (e.g., synthetic nanorods) and semi-rigid (e.g., short DNA segments) bodies by composing each structure from a collection of spherical particles constrained to move (semi)rigidly. The underlying fluctuating hydrodynamics formulation automatically ensures the correct translational and rotational Brownian motion.

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
aleks.donev@nyu.edu

MS193

Boltzmann's State of Motion: Phenomenological Modeling of Chemical and Ecological Systems

We carry out a mathematical analysis, *à la* Helmholtz's and Boltzmann's 1884 studies of monocyclic Newtonian mechanics, for chemical reaction and ecological systems containing oscillatory dynamics. One of the important

features of the systems considered, absent in the classical mechanical model, is a natural stochastic dynamic formulation of which the deterministic differential equation is the infinite population limit. We separate the conserved dynamics from the dissipative models and show how the conservation law along a single trajectory can be extended to incorporate both variations in model parameters and in the initial conditions: Helmholtz's theorem establishes a broadly valid conservation law in a class of chemical and ecological dynamics. Further analysis identifies an entire orbit as a stationary behavior, and establishes the notion of an "equation of behavioral state'. Studies of the stochastic dynamics shows the conserved dynamics as the robust, fast cyclic underlying behavior. The mathematical narrative provides a novel way of capturing long-term dynamical behavior with an emergent conservative motion.

Yian Ma
University of Washington
yianma@uw.edu

Hong Qian
Department of Applied Mathematics
University of Washington
hqian@u.washington.edu

MS193

A Nano Pore-Scale Model for the Nanostructured Cathode of Lithium-Oxygen Batteries

We propose a nano pore-scale model that bridges the gap between continuum models and atomistic models to study the impact of cathode microstructure and rate-dependent morphology of Li_2O_2 on the discharge performance of Li-O_2 batteries. The model captures the micro-scale resolved nanostructure of cathode with varying porosities, pore size distributions and surface-to-volume ratios, the diffusion-limited transport of oxygen across the porous structure of the cathode, and the morphology of Li_2O_2 growth at different current densities.

Wenxiao Pan
Pacific Northwest National Laboratory
wenxiao.pan@pnnl.gov

MS193

SPH Model for Landau-Lifshitz Navier-Stokes and Advection-Diffusion Equations

We propose a novel Smoothed Particle Hydrodynamics (SPH) discretization of the fully-coupled Landau-Lifshitz-Navier-Stokes (LLNS) and advection-diffusion equations. The accuracy of the SPH solution is demonstrated by testing the scaling of velocity variance and self-diffusion coefficient with kinetic temperature and particle mass, the spatial covariance of pressure and velocity fluctuations, the so-called giant fluctuations of the front between light and heavy fluids with and without gravity, and the effect of thermal fluctuation on the Rayleigh-Taylor instability.

Alexander Tartakovsky
Pacific Northwest National Laboratory
alexandre.tartakovsky@pnnl.gov

Jannes Kordilla
Geoscientific Centre
University of Goettingen
jkordil@gwdg.de

Wenxiao Pan
Pacific Northwest National Laboratory
wenxiao.pan@pnnl.gov

MS194

Computing Exactly and Eficiently Arbitrarily-High-Order Response Sensitivities to Model Parameters

This work presents a new methodology for computing exactly and (most) efficiently response sensitivities, of arbitrarily high-order, to model parameters. This new method generalizes the adjoint method for sensitivity analysis of nonlinear systems originated by Cacuci (1981), requiring at most O(NK-1) large-scale adjoint computations per response to obtain exact Kth-order sensitivities for a system with N parameters. In contradistinction, conventional methods require O(NK) large-scale computations to obtain approximate values for the Kth-order sensitivities.

Dan G. Cacuci
University of South Carolina
cacuci@cec.sc.edu

MS194

Toward New Applications of the Adjoint Tools in 4D-Var Data Assimilation

The development of efficient methodologies to assess and improve the information content ('value') of high-resolution atmospheric measurements is an imperative task. This talk presents recent advances in the adjoint-based evaluation of the model forecast sensitivity to observations, error covariance parameter specification, and impact estimation in a four-dimensional variational data assimilation system. The practical significance and further research directions are discussed together with preliminary experiments and the current status of implementation at numerical weather prediction centers.

Dacian N. Daescu
Portland State University
Department of Mathematics and Statistics
daescu@pdx.edu

Ricardo Todling
NASA/GSFC Global Modeling and Assimilation Office
ricardo.todling@nasa.gov

Rolf Langland
Naval Research Laboratory
Monterey, CA
rolf.langland@nrlmry.navy.mil

Austin Hudson
Portland State University
Portland, OR, USA
hudsonaustins@gmail.com

MS194

Dealing with Nonsmoothness in Data Assimilation

The theory of adjoints and their use in inverse problems usually assumes that the state equations and error norms are smooth. In practice this assumption may be violated through the use of nonsmooth norms or due to nonsmooth modification of the state equations, e.g. the use of flux limiters in spatial discretizations. We analyze the theoret-

ical and practical effects and discuss suitable algorithmic modifications.

Andreas Griewank
HU Berlin, Germany
griewank@math.hu-berlin.de

MS194

Second Order Analysis in Variational Data Assimilation

The problem of Variational Data Assimilation, considered in the framework of Optimal Control Theory, leads to an Optimality System containing all the available information : model, data and statistics, its solution gives a necessary condition for optimality. Introducing a second order adjoint a second order analysis can be carried out with important applications such as: -Improvement of the optimization algorithms - Sensitivity Analysis - Estimation of a posteriori statistics on the solution. Theory and applications will be presented.

Francois-Xavier L. Le-Dimet
University Joseph Fourier
ledimet@imag.fr

M Yousuff Hussaini
Florida State University
myh@csit.fsu.edu

Ha Tran Thu
Vietnamese Academy of Sciences
fxld@yahoo.com

MS195

Data-Driven Model for Solar Irradiation Based on Satellite Observations

We construct a data-driven model for solar irradiation based on satellite observations. The model yields probabilistic estimates of the irradiation field every thirty minutes starting from two consecutive satellite measurements.The probabilistic nature of the model captures prediction uncertainties and can therefore be used by solar energy producers to quantify the operation risks. The dynamics are represented in a nonlinear, nonparameteric way by a recursive Gaussian process. The predictions of the model are compared with observed satellite data as well as with a similar model that uses only ground observations at the prediction site. We conclude that using satellite data in an area including the prediction site significantly improves the prediction compared with models using only ground observation site data.

Ilias Bilionis
Purdue University
ebilionis@gmail.com

Emil M. Constantinescu, Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov, anitescu@mcs.anl.gov

MS195

Two-Stage Adaptive Robust Unit Commitment Using Scenarios Induced Uncertainty Set

In this talk, we propose a new robust optimization framework for unit commitment under uncertainty. We present a new type of uncertainty set induced by correlated scenario samples to capture spatial-temporal correlations of the uncertainties. Furthermore, for large-scale problems, we propose an efficient approximation of the Scenarios Induced Uncertainty (SIU) set using Principal Component Analysis. Computational results demonstrate the economic benefits of using SIU sets and the efficacy of the proposed solution approach.

Richard L. Chen
Sandia National Laboratories
Livermore, CA
rlchen@sandia.gov

Cosmin Safta
Sandia National Laboratories
csafta@sandia.gov

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

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

Ali Pinar
Sandia National Labs
apinar@sandia.gov

MS195

Economic Impacts of Wind Covariance Estimation on Power Grid Operations

We study the impact of capturing spatiotemporal correlations between multiple wind supply points on economic dispatch procedures. Using a simple dispatch model, we first show analytically that over/underestimation of correlation leads to positive and negative biases of dispatch cost, respectively. A rigorous, large-scale computational study for the State of Illinois transmission grid with real topology and physical constraints reveals similar conclusions. For this study, we use the Rao-Blackwell-Ledoit-Wolf estimator to approximate the wind covariance matrix from a small number of wind samples generated with the numerical weather prediction model WRF and we use the covariance information to generate a large number of wind scenarios. The resulting stochastic dispatch problems are solved by using the interior-point solver PIPS-IPM on the BlueGene/Q (Mira) supercomputer at Argonne National Laboratory. We find that strong and persistent biases result from neglecting correlation information and indicate to the need to design a market that coordinates weather forecasts and uncertainty characterizations.

Cosmin G. Petra
Argonne National Laboratory
Mathematics and Computer Science Division
petra@mcs.anl.gov

Victor Zavala
Argonne National Laboratory
vzavala@mcs.anl.gov

Elias Nino-Ruiz
Virginia Polytechnic Institute and State University

enino@vt.edu

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

MS195

An Efficient Approach for Stochastic Optimization of Electricity Grid Operations

Stochastic unit commitment optimization problems typically handle uncertainties in forecast demand by considering a finite number of random realizations from a stochastic process model for uncertain loads. In this paper we propose a more efficient approach using Polynomial Chaos representations valid over the range of the forecast uncertainty. We demonstrate the approach for several power grid models with and without transmission constraints.

Cosmin Safta
Sandia National Laboratories
csafta@sandia.gov

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

Richard L. Chen
Sandia National Laboratories
Livermore, CA
rlchen@sandia.gov

Ali Pinar
Sandia National Labs
apinar@sandia.gov

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

MS196

On the Coupling of Far-Field Wind-Wave Simulation and Near-Field Free-Surface Flow Simulation

We develop a computational framework for simulating in a coupled manner the interaction of large-scale ocean waves and winds with the presence of complex floating structures. We employ an efficient large-scale model to develop offshore wind and wave environmental conditions, which are then incorporated to a high resolution two-phase flow solver with fluid-structure interaction (FSI). The far-field/near-field coupling algorithm is validated under several test cases involving simple wave trains as well as three-dimensional directional waves.

Antoni E. Calderer
Saint Anthony Falls Lab, Civil Engineering department
University of Minnesota
cald0075@umn.edu

Xin Guo
UC Berkeley
xinguo@ieor.berkeley.edu

Fotis Sotiropoulos
Saint Anthony Falls Lab, Civil Engineering department
University of Minnesota
fotis@umn.edu

Lian Shen
University of Minnesota
shen@umn.edu

MS196

Topological Change with a Cut Cell based Sharp Interface Method for Multi-phase Flows

In the conservative, consistent, all-speed, sharp-interface method presented by Chang, Deng & Theofanous 2013, a free interface is represented by cut faces and evolved with the help of a level set function. It shows good performance on simulating multi-phase flow problems with clear interfaces, but cannot deal with topological change. By switching between this cut cell/level set mixed method and a pure level set method, topological change can be realized smoothly.

Xiao-Long Deng
Beijing Computational Science Research Center
xiaolong.deng@csrc.ac.cn

MS196

High Resolution PDE Solvers on Octree Grids and Parallel Architectures

Abstract not available at time of publication.

Frederic G. Gibou
UC Santa Barbara
fgibou@engineering.ucsb.edu

MS196

A Second Order Virtual Node Algorithm for NavierStokes Flow Problems with Interfacial Forces and Discontinuous Material Properties

We present a numerical method for the solution of the Navier-Stokes equations in three dimensions that handles interfacial discontinuities due to singular forces and discontinuous fluid properties such as viscosity and density. We discretize the equations using an embedded approach on a uniform MAC grid to yield discretely divergence-free velocities that are second order accurate. The method leads to a discrete, symmetric KKT system for velocities, pressures, and Lagrange multipliers. We also present a novel simplification to the standard combination of the second order semi-Lagrangian and BDF schemes for discretizing the inertial terms. Numerical results indicate second order spatial accuracy for the velocities (L^∞ and L^2) and first order for the pressure (in L^∞, second order in L^2).

Joseph Teran
UCLA
jteran@math.ucla.edu

MS197

Asynchronous Preconditioning on Accelerators

This talk focuses on how preconditioners can be constructed using asynchronous algorithms. Such algorithms are fine-grained and can tolerate memory latency, making them suitable for modern architectures, particularly accelerators. Preconditioners including incomplete factoriza-

tions and sparse approximate inverses may be constructed this way. We discuss convergence of various types of asynchronous algorithms applied to constructing preconditioners, necessary implementation differences on GPUs and Intel Xeon Phi, and inner-outer asynchronous iterations designed to reduce communication volume from low levels of the memory hierarchy.

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

MS197

Overview and Contrast of Modern Computer Architectures Including the Intel Phi

State-of-the-art distributed-memory clusters today contain multi-core CPUs with 8 to 12 cores, massively parallel GPUs with thousands of computational cores, and many-core accelerators such as the 60-core Intel Phi, connected by high-performance networks such as InfiniBand interconnect. This talk will give an overview of their features, show sample code how to use GPUs and the Intel Phi processor in hybrid mode, and contrast performance results for basic test problems.

Jonathan Graf, Samuel Khuvis, Xuan Huang
Department of Mathematics and Statistics
University of Maryland, Baltimore County
jongraf1@umbc.edu, khsa1@umbc.edu, hu6@umbc.edu

Matthias K. Gobbert
University of Maryland, Baltimore County
Department of Mathematics and Statistics
gobbert@umbc.edu

MS197

Offloading Computational Kernels in Long-Time Simulations to the Intel Phi

Simulating calcium waves in a heart cell requires both fine meshes and large final times to match laboratory experiments. This is a perfect example of a problem that can profit from offloading numerical kernels to accelerators such as the 60-core Intel Phi processor on each hybrid compute node, combined with pooling memory from several nodes to allow for the desired spatial resolutions. We report results for special-purpose code for this problem.

Samuel Khuvis, Xuan Huang, Jonathan Graf
Department of Mathematics and Statistics
University of Maryland, Baltimore County
khsa1@umbc.edu, hu6@umbc.edu, jongraf1@umbc.edu

Matthias K. Gobbert
University of Maryland, Baltimore County
Department of Mathematics and Statistics
gobbert@umbc.edu

MS197

The HPCG Benchmark Using Intel Phi Accelerators

The High Performance Preconditioned CG (HPCG) Benchmark, developed by Sandia National Laboratories, uses a preconditioned conjugate gradient method to solve the Poisson equation, since this may be more relevant to many applications than the LINPACK benchmark. After introducing this new benchmark that reported rankings for the first time in June 2014 we show its implementation and performance on CPUs only and on hybrid nodes with Intel Phi accelerators, connected by a QDR InfiniBand interconnect.

Adam Cunningham, Gerald Payton, Jack Slettebak
University of Maryland, Baltimore County
cunning5@umbc.edu, gpayton1@umbc.edu,
slet1@umbc.edu

Jordi Wolfson-Pou
Department of Physics
University of California, Santa Cruz
jswolfso@ucsc.edu

MS198

Creating the Largest Decoys Database to Improve Scoring Functions Using Machine Learning

In the protein structure prediction field it is absolutely vital to develop reliable scoring functions that will allow researchers to efficiently identify the very best models out of hundreds of thousands and sometimes even millions of predictions. Such scoring functions have proven to be difficult to produce. The protein folding community believes that machine learning techniques will advance the capabilities of scoring functions. Datasets of computationally generated models known as decoys are used to train and test the scoring functions that are created. While there are decoys datasets available to perform these tests, these datasets are very small in comparison to the amount of protein models that are created during the Critical Assessment of protein Structure Prediction (CASP) competitions. Larger datasets are essential to create more accurate scoring functions. The WeFold collaboration mediated by the homonymous gateway gathers an enormous amount of models, which are shared by its users. We are creating the largest database of decoys to make available to the machine learning community. The database does not only store models, but also numerical features that describe important characteristics of protein models. MongoDB is being used to store all of this information and an interface is being created to integrate MongoDB with the WeFold gateway. Users will be able to submit queries to the database and download all models created by all groups, all models created by specific groups, specific features for specific models, and even more. Speakers: Ricardo Ferreira and Christopher Cook

Ricardo Ferreira, Christopher Cook
Mountain View Community College
ri.ferreira2020@gmail.com, cc112591@gmail.com

MS198

WeFold: a Collaborative and Educational Experiment

Determining protein structure is key to advances in science and medicine. To advance this field a social-media based consortium of labs worldwide was created to bring together researchers from all walks and disciplines. Recently, the project added an educational component by bringing a group of students and faculty under one roof to discuss protein folding. We will discuss the results of WeFold both as a new way to conduct research and engage students to learn computational science. Speakers: Silvia

Crivelli, John Hatherill, and Jesse Fox.

Silvia N. Crivelli
Lawrence Berkeley National Laboratory
sncrivelli@lbl.gov

John Hatherill
Del Mar Community College
jhatherill@delmar.edu

Jesse Fox
Mountain View Community College
jfox@dcccd.edu

MS198

Reducing the Data Complexity with Filtering and Clustering

Wefold maintains hundreds of thousands of models obtained from participants like Foldit. Model computation is time intensive due to NP complete nature, thus this team sought smaller accurate protein conformation datasets. Preliminary results from our filter concluded the filtration achieved a decrease in set size with an accuracy increase. However, this set retains unworkable complexity. Different parallel clustering algorithms and metrics were analyzed to find accurate clusters and representatives to finally reduce the set. Speakers: Davis, Ogden, and Raiyyani

Rachel A. Davis
Drake University
rachel.davis@drake.edu

Jennifer Ogden
Saint Mary's College of California
jennogden95@gmail.com

Rehan Raiyyani
University of California, San Diego
rehan.raiyyani@gmail.com

MS198

The Maintenance of the WeFold Gateway for CASP11

The protein structure prediction and refinement methods generated by the WeFold community in the context of the CASP11 competition were executed as pipelines of components contributed by different labs. Twenty labs worldwide including more than 100 researchers participated in this collaboration. Because of deadlines to submit predictions to the competition, it was important to facilitate the execution of the pipelines. I will discuss the maintenance of the WeFold gateway as well as the results achieved.

Anthony Lopez
Del Mar Community College
lopez52891@yahoo.com

MS199

An Implcit, Conservative Vlasov-Darwin Pic Solver in Multiple Dimensions

We propose a new, implicit 2D-3V particle-in-cell algorithm for non-radiative, electromagnetic kinetic plasma simulations. Local charge, global energy, and canonical momentum are exactly conserved. The Darwin-Vlasov equations are discretized with particles on a grid, using a time-centered, implicit scheme, which is converged nonlinearly. The nonlinear solver is accelerated by a moment-based preconditioner. We demonstrate the properties of the algorithm with several verification examples, including a Weibel instability and a Kinetic Alfven wave ion-ion streaming instability.

Guangye Chen, Luis Chacon
Los Alamos National Laboratory
gchen@lanl.gov, chacon@lanl.gov

MS199

A Moment Model for the Vlasov Fokker Planck Equation

We present a moment method based on spherical harmonics for the Vlasov equation coupled with some common collision operators. A transformation in phase space is used to enable the moments to give a perturbation around the equilibrium solution in phase space. Furthermore, we give numerical results which show the difference between using/not using a transformation in phase space.

Charles K. Garrett, Cory Hauck
Oak Ridge National Laboratory
garrettck@ornl.gov, hauckc@ornl.gov

MS199

Modeling Non-Ideal Plasmas: a Hyrbid Quantum Hydrodynamics and Molecular Dynamics Approach

Abstract not available at time of publication.

Michael Murillo
Los Alamos National Laboratory
murillo@lanl.gov

MS199

iFP: An Optimal, Fully Conservative, Fully Implicit, Vlasov-Fokker-Planck Solver

We introduce a new, exactly conservative (mass, momentum, and energy), and fully nonlinearly implicit solver for a multi-species 1D-2V Vlasov-Rosenbluth-Fokker-Planck system. The new solver optimizes mesh resolution requirements by 1) adapting the velocity-space mesh based on the species' local thermal-velocity, and 2) treating the cross-species collisions exhibiting disparate thermal velocities by an asymptotic formulation. We demonstrate the efficiency and accuracy properties of the approach with challenging numerical examples.

William T. Taitano, Luis Chacon, Andrei Simakov
Los Alamos National Laboratory
taitano@lanl.gov, chacon@lanl.gov, simakov@lanl.gov

MS200

Current Challenges in Mesh Partitioning for Physics Simulations

For modern and complex physics simulations on very large data, achieving a good load balancing of the computations at every time step is still a hard problem. For mesh based simulations, load balancing is achieved by performing mesh partitioning. In this talk, we shall define what can be the goals of mesh partitioning. We shall also present main

problems we are currently working on at CEA: - Multi criteria (re-)partitioning, in order to deal with multi physics simulations; - Memory constrained (re-)partitioning, to model layers of ghost cells which can fill up the available memory.

Cédric Chevalier
French Alternative Energies & Atomic Energy Commission (CEA)
cedric.chevalier@cea.fr

MS200

The Zoltan2 Toolkit: Partitioning, Task Placement, Coloring, and Ordering

Zoltan2 is a toolkit of partitioning, coloring, and ordering algorithms for use in parallel scientific applications. A successor to the Zoltan toolkit, Zoltan2 is designed to address current issues in scientific computing: extreme dataset sizes, parallel computers built of multicore processors, and sustainable software design. Zoltan2 is built on the templated software stack in the Trilinos solvers framework, and, thus, provides more seamless integration with Trilinos' data structures. This presentation provides an introduction to the design and capabilities of Zoltan2.

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

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

Siva Rajamanickam, Lee Ann Riesen
Sandia National Laboratories
srajama@sandia.gov, lriesen@sandia.gov

Mehmet Deveci
The Ohio State University
mdeveci@bmi.osu.edu

Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
umit@bmi.osu.edu

MS200

Unstructured Mesh Partitioning to over 500k Parts

Parallel unstructured mesh-based applications running on the latest petascale systems require partitions with over 500k parts. Methods combining the most powerful graph based and geometric methods with diffusive methods directly operating on the unstructured mesh will be discussed. Initial partitioning results on meshes with over 12 billion elements and over 512k parts indicate that these methods maintain quality while keeping run times to a small fraction of application run time.

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

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

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

MS200

Zoltan2 for Extreme-Scale Data Partitioning

As the number of cores in modern high performance computing architectures has increased, data partitioning algorithms have struggled to scale. As part of the DOE SciDAC FastMATH Institute, we have been attempting to address this problem in our load-balancing software Zoltan2. Specifically, we have be focusing on developing partitioners to create high quality partitions for 100k-1M cores. In this talk, we discuss many of the techniques that we have found useful in achieving this goal.

Michael Wolf
Sandia National Laboratories
mmwolf@sandia.gov

MS201

Block-Structured AMR: Applications Using BoxLib

BoxLib is a publicly available software framework for building parallel block-structured AMR applications using both MPI and OpenMP. It supports grid-based and particle-mesh operations on adaptive hierarchical meshes with optional subcycling in time. This talk focuses on the common components that form the basis of existing codes in astrophysics, cosmology, combustion and porous media, and allow construction of new application codes and extension to AMR of existing single-level codes, in a variety of application areas.

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

MS201

Preconditioners for Implicit Atmospheric Climate Simulations in the Community Atmosphere Model

High resolution, long-term climate simulations are essential to understanding regional climate variation on the decade scale. To maintain accuracy with time steps commensurate with the physical processes of interest, the spectral element dynamical core of the Community Atmosphere Model (CAM-SE) implements fully implicit schemes utilizing FASTMath Trilinos solvers. We discuss the development of preconditioners to reduce the computational cost of ancillary linear system solves within each time step to improve solver efficiency and scalability.

David J. Gardner
Lawrence Livermore National Laboratory
gardner48@llnl.gov

Katherine J. Evans
Oak Ridge National Laboratory
evanskj@ornl.gov

Aaron Lott

D-Wave Systems, Inc.
palott@gmail.com

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

MS201

Gyrokinetic Poisson Equation Solvers with Explicit Flux Surface Averaging in XGC1 with PETSc

An accurate representation of the flux surface averaging term in the gyrokinetic Poisson equation is explored for the first time, to our knowledge, in a global extreme-scale Tokamak code XGC1. We use the FieldSplit preconditioner capability in PETSc, which provides access to a wide variety solver algorithms, including direct, block near direct, Schur complement, Gauss-Seidel, and additive and multiplicative solver algorithms.

Mark Adams
Lawrence Berkeley Laboratory
mfadams@lbl.gov

Seung-Hoe Ku
Princeton Plasma Physics Laboratory
sku@pppl.gov

MS201

Algebraic Multigrid Solvers for Lattice QCD in the Hypre Software Library

Algebraic multigrid (AMG) has been instrumental in many simulation codes for solving the requisite linear systems in a scalable manner. In quantum chromodynamics (QCD), the development of effective AMG methods has been fairly recent so their use in QCD simulations has not yet become commonplace. In this talk, we discuss work to develop a variant of the bootstrap AMG method in the hypre library and efforts to interface it with the USQCD software stack.

Evan Berkowitz
Lawrence Livermore National Laboratory
berkowitz2@llnl.gov

James Brannick
Department of Mathematics
Pennsylvania State University
brannick@psu.edu

Robert Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

Chris Schroeder
Lawrence Livermore National Laboratory
schroeder1@llnl.gov

MS202

Optimization Approach for Tomographic Inversion from Multiple Data Modalities

Fluorescence tomographic reconstruction can be used to reveal the internal elemental composition of a sample while transmission tomography can be used to obtain the spatial distribution of the absorption coefficient inside the sample. In this work, we integrate both modalities and formulate an optimization approach to simultaneously reconstruct the composition and absorption effect in the sample. By using multigrid-based optimization framework (MG/OPT), significant speedup and improvement of accuracy has shown for several examples.

Zichao Di
Argonne National Lab
wendydi@mcs.anl.gov

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS202

Multigrid Preconditioning for Space-time Distributed Optimal Control Problems Constrained by Parabolic Equations

We present some recent results regarding multigrid preconditioning of the linear systems arising in the solution process of space-time distributed optimal control problems constrained by parabolic equations. While the construction of the preconditioners is based on ideas extracted from optimal control problems constrained by elliptic equations, in the parabolic-constrained case the multigrid preconditioners exhibit a suboptimal behavior, namely they approximate the operators to be inverted by half an order less than optimal.

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

Mona Hajghassem
Department of Mathematics and Statistics
University of Maryland, Baltimore County
mona4@umbc.edu

MS202

About Some Smoothers for Saddle-point Problems

In this work, we focus on the multigrid solution of saddle point problems. A key ingredient in designing an optimal multigrid solver is the choice of the smoother. The smoothers that we consider are mainly the obtained via a coupled relaxation approach and they are tuned specifically for the indefinite matrices corresponding to the saddle point problems of interest. This is done guided by the local Fourier analysis (LFA), which allows us to estimate the spectral radius of the k-grid operator in order to obtain quantitative measures for the error reduction and to accurately predict the asymptotic convergence factor of the method. The LFA of such smoothers requires non-standard techniques, and they are used to adjust the parameters of the underlying multigrid method and increase its efficiency.

Carmen Rodrigo, Francisco José Gaspar, Francisco Lisbona
University of Zaragoza
carmenr@unizar.es, fjgaspar@unizar.es, lisbona@unizar.es

Ludmil Zikatanov
Pennsylvania State University
ludmil@psu.edu

MS202

Multigrid in Chaos

Dynamics of molecules, fluid flows, the climate, and other chaotic systems appear in many science and engineering applications. The long-time behavior of these dynamical systems, and their sensitivity to perturbations, is interesting to both scientists and engineers. The multigrid method may be the key to finding such sensitivity. Sensitivity of chaotic dynamical systems can be computed using a new method known as Least Squares Shadowing, which requires solving a space-time system of elliptic nature. Although multigrid shows great promise, this application also brings new challenges to multigrid methods.

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

Patrick Blonigan
MIT
blonigan@mit.edu

MS203

Fast Solvers for Hierarchical Matrices

Recent years have seen the emergence of novel fast direct methods to solve linear systems with hierarchical matrices. These methods are based on a type of direct elimination of the unknowns, in a manner similar to the LU factorization; their accuracy is therefore less sensitive to the condition number and distribution of eigenvalues of the matrix than with iterative solvers. We will present a novel class of fast direct solvers for sparse matrices.

Amirhossein Aminfar
Stanford University
aminfar@stanford.edu

Sivaram Ambikasaran
Department of Mathematics
Courant Institute of Mathematical Sciences
sivaram@cims.nyu.edu

Mohammad Hadi Pour Ansari
Stanford
hadip@stanford.edu

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

MS203

Rank-Structured Preconditioners for Two and Three-Dimensional Integral and Differential Equations

The inverses and LR factors of stiffness matrices appearing in the context of partial differential equations and integral equations are frequently *rank-structured*, i.e., they contain submatrices that have low numerical rank. This property can be exploited to efficiently construct approximations of these matrices and thereby find preconditioners for a wide range of problems. This talk presents algorithms for carrying out algebraic operations (matrix multiplication, factorization, inversion) approximately exploiting and preserving the low-rank structure. The algorithms rely on sequences of low-rank updates applied to $\mathcal{H}^2$*-matrices*, a relatively general class of rank-structured matrices that have efficient data-sparse representations.

Steffen Börm, Knut Reimer
Universitaet Kiel
sb@informatik.uni-kiel.de, knr@informatik.uni-kiel.de

MS203

A New Integral Formulation and Fast Direct Solver for Periodic Stokes' Flow

Many solution techniques have recently been developed to accurately and efficiently numerically model vesicle flow. The introduction of a periodic confining geometry adds further complications to such simulations. This talk presents a new integral formulation which avoids the use of the periodic Green's function. Additionally, a fast direct solver for the discretized confining geometry is presented. This solver allows for efficient time stepping by decoupling the static geometry from the moving vesicles.

Adrianna Gillman
Dartmouth College
Department of Mathematics
adrianna.gillman@rice.edu

Alex H. Barnett
Department of Mathematics
Dartmouth College
ahb@math.dartmouth.edu

Shravan Veerapaneni, Gary Marple
Department of Mathematics
University of Michigan
shravan@umich.edu, gmarple@umich.edu

MS203

Practical and Efficient Direct Solvers for BIEs

The discretization of integral equations leads to dense linear systems. For large problems, these systems have traditionally been solved using iterative solvers, often in combination with accelerated techniques for the dense matrix-vector multiplication, such as, e.g., the Fast Multipole Method. However, in the last several years a number of *direct solvers* with very high practical efficiency, and favorable (often linear) asymptotic complexity have been developed. This talk will describe techniques aimed at improving performance and simplifying coding by using randomized methods to accelerate certain structured matrix computations. The direct solvers used will be based on the hierarchical merging of discrete analogs of *scattering matrices*.

Gunnar Martinsson
Univ. of Colorado at Boulder

martinss@colorado.edu

MS204

A One-Stage High-Resolution Constrained Transport Method for Magnetohydrodynamic Equations

Failure to maintain zero divergence of magnetic field has been known to cause instability of numerical methods for solving magnetohydrodynamic equations. In this talk we will present a high-order conservative finite difference WENO method that uses constrained transport to enforce a divergence free magnetic field. This method is based on the Picard Integral Formulation of the PDE, which allows us to construct a single-stage single-step scheme. This method is high-resolution, scalable to large clusters, and amenable to Adaptive Mesh Refinement (AMR).

Xiao Feng
Department of Mathematics
Michigan State University
fengxia2@msu.edu

MS204

A Hybrid Weno Reconstruction on Unstructured Mesh

The weighted essentially non-oscillatory (WENO) schemes are a popular class of high order numerical methods for hyperbolic partial differential equations (PDEs). In this talk, we will first talk about a hybrid approach for WENO reconstructions on the unstructured meshes and then discuss the maximum-principle-preserving/positivity-preserving properties of the schemes. Numerical examples coupled with third order Runge-Kutta time integrator are reported.

Yuan Liu
Department of Mathematics
Michigan State University
yliu7@math.msu.edu

MS204

A New RKDG Method with Conservation Constraint to Improve CFL Condition for Solving Conservation Laws

Abstract not available at time of publication.

Zhiliang Xu
University of Notre Dame
zxu2@nd.edu

MS204

High Order WENO Method for Steady State Problems

High order accurate shock capturing schemes such as WENO schemes often suffer from difficulties in their convergence towards steady state solutions. In this talk, I shall present our recent results on improving the convergence of fifth order WENO scheme for solving steady state hyperbolic conservation laws by an explicit Gauss-Seidel sweeping framework combined with different new techniques.

Liang Wu, Yongtao Zhang
University of Notre Dame
lwu2@nd.edu, yzhang10@nd.edu

Shuhai Zhang
China Aerodynamics Research and Development Center
China
zhang_shuhai@tom.com

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

MS205

Constraint Aggregation Methods for PDE-Constrained Optimization

Many engineering design optimization problems are formulated with a bound on a physical quantity over a domain of interest. Aggregation methods approximate this infinite-dimensional constraint in a differentiable manner. However, many classical methods are non-conservative and it is often difficult to assess the error incurred through aggregation. To address these issues, we present a new class of aggregation technique for PDE-constrained optimization and present a post-optimality method to assess the aggregation error. We present the results of these new techniques on large-scale structural optimization problems.

Graeme Kennedy
Georgia Institute of Technology
School of Aerospace Engineering
graeme.kennedy@aerospace.gatech.edu

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

MS205

Large-Scale PDE-Constrained Fluid-Structure Optimization

A multi-physics fluid-structure optimization problem with millions of state variables and thousands of design variables is presented. The fluids domain is described by the compressible Reynolds-Averaged Navier–Stokes equations and the structural domain is described by the equations of linear elasticity. Parallel solution techniques for the coupled, non-linear multidisciplinary system and the linearized adjoint system are discussed.

Gaetan Kenway
University of Michigan
Aerospace Engineering
gaetank@gmail.com

Joaquim R. R. A Martins
University of Michigan
jrram@umich.edu

MS205

A Krylov-Based Iterative Solver for Equality-Constrained Non-Convex Quadratic Subproblems

Conventional matrix-based constrained-optimization algorithms are not well suited to reduced-space PDE-constrained optimization, because, even for modest-sized primal and dual spaces, the cost of evaluating the Hessian and Jacobian can be prohibitive. This makes matrix-free optimization algorithms attractive for these applications. Some challenges for reduced-space matrix-free algorithms

include Hessians that are nonconvex in the null space of the Jacobian, as well as finding effective matrix-free preconditioners for the primal-dual system. Several matrix-free algorithms have been proposed to handle nonconvexity by adapting existing iterative solvers. In this work, we address nonconvexity within the primal-dual iterative method itself by modifying the subspace problem solved by flexible GMRES. The approach is surprisingly simple and effective. We also present preliminary investigations of a multigrid matrix-free preconditioner.

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

Pengfei Meng
Rensselaer Polytechnic University
mengp2@rpi.edu

MS205

Aerodynamic Shape Optimization with Goal-Oriented Error Estimation and Control

We investigate the control of discretization error in gradient-based aerodynamic shape optimization through use of adaptive mesh refinement. The approach makes dual use of the adjoint method first in the computation of the objective function gradient and second in the estimation of discretization error. We focus on progressive optimization, where the depth of mesh refinement is systematically increased as the design improves. The approach is demonstrated on several challenging problems, including three-dimensional inverse design for sonic-boom minimization and the optimization of a flexible wing for transport aircraft. This is an important step in our research toward dynamic error control to improve automation and minimize cost.

Marian Nemec
Science & Technology Corp.
NASA Ames Research Center
marian.nemec@nasa.gov

Michael Aftosmis
NASA Ames Research Center
michael.aftosmis@nasa.gov

MS206

A Parallel Local Timestepping Runge-Kutta Discontinuous Galerkin Method with Applications to Coastal Ocean Modeling

Abstract not available at time of publication.

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

MS206

Development and Validation of DG Wave: a Discontinuous Galerkin-Based Numerical Wave Prediction Model

Abstract not available at time of publication.

Ethan Kubatko
The Ohio State University
ekubatko@gmail.com

MS206

Physically Based Assessment of Hurricane Surge Threat under Climate Change

Abstract not available at time of publication.

Ning Lin
Princeton University
nlin@princeton.edu

MS206

Understanding Coastal Hydrodynamic Processes and Mitigating Risk Through High Fidelity Computer Simulations

Abstract not available at time of publication.

Joannes Westerink
Department of Civil Engineering and Geological Sciences
University of Notre Dame
jjw@nd.edu

MS207

FinAT: A Mathematical Structure-Preserving Library of Finite Elements

Currently, automated finite element software generates fairly traditional algorithms based on black box tabulation of basis functions. FInAT is a new project that will exploit the algebraic structure of basis functions to generate more sophisticated algorithms that exploit tensor product structure. It replaces arrays of tabulated values, with specialized recipes for performing evaluation and integration. This will create efficient implementations of product elements, vector elements, Bernstein polynomials, and elements such as Morley and Hermite with complicated pullbacks.

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

Rob Kirby
Baylor University
robert_kirby@baylor.edu

MS207

Finite Element Geometric Multigrid Solvers from High-Level Problem Descriptions

We present an implementation of geometric multigrid solvers in the Firedrake finite element framework. Our approach combines composable building blocks – appropriate smoothers and intergrid transfer operators – with a symbolic representation of the PDE to simplify development of solvers for users: for many problems they need only provide the fine grid discretisation. Our approach is agnostic to how the operators are applied, allowing simple switching between matrix-free and assembled operators.

Lawrence Mitchell
Department of Computing
Imperial College London
lawrence.mitchell@imperial.ac.uk

Eike H. Mueller
University of Bath
e.mueller@bath.ac.uk

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

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

MS207

Towards a Unified Framework for Automated a Posteriori Error Estimation and Adaptivity in Space-Time

Finite element systems that operate at a sufficiently high level of abstraction allow for efficient implementation of a number of automated features. Here, we will discuss how to efficiency extend previous work on automated goal-oriented error control and adaptivity to space-time finite element methods. We'll present a unified framework for automated derivation of problem-targeted error estimates and indicators, and high level interfaces, efficient data structures and algorithms for arbitrary dimension tensor product finite element methods.

Marie E. Rognes
Simula Research Laboratory
meg@simula.no

Anders Logg
Chalmers University of Technology
logg@chalmers.se

Benjamin Kehlet
Simula Research Laboratory
benjamik@simula.no

MS207

Multicore Parallelism for Common Finite Element Operations

As the number of cores on a chip and in a node increases, the importance of multithreading increases. We identify a design pattern where an embarrassingly parallel operation on every cell is followed by a reduction. This design pattern appears in finite element code through matrix assembly, estimating discretization errors, post-processing, etc. We also show a scalable implementation of this design pattern.

Bruno Turcksin
Texas A&M University
turcksin@math.tamu.edu

Martin Kronbichler
Technische Universitat Munchen
kronbichler@lnm.mw.tum.de

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

MS208

Thermodynamically Consistent and Meta-Stable Equation of State Models for Hydro and Solid Dynamics

Simulations of real engineering interest require the use of thermodynamic and constitutive models that go beyond simple analytic equations of state. In this talk we will discuss the complications that arise from using general equations of states for fluids and solids. Issues that need to be addressed include equilibrium and non-equilibrium treatments for mixed and separated species and robust treatments for incompatible mixtures such as metals in tension and gases.

John W. Grove
Computational Physics & Methods, CCS-2
Los Alamos National Laboratory
jgrove@lanl.gov

MS208

Modelling of Fabric Surface for Parachute Inflation through Front Tracking

We use the front tracking library on a triangulated spring system to model the dynamic evolution of parachute canopy surface. The model is numerically convergent under the constraints that the summation of point masses is constant and that both tensile and angular stiffness of the spring conform with the material's Young modulus and Poisson ratio. This assembly is coupled with the fluid solver through the impulse method to compute parachute inflation for comparison with experiments.

Xiaolin Li
Department of Applied Math and Stat
SUNY at Stony Brook
xiaolin.li@stonybrook.edu

MS208

Overlapping BEM on FEM computations

In this talk I will review some theoretical and computational work on how to deal with time-harmonic wave transmission problems by superposing solutions of problems in unbounded domains (discretized with BEM) and bounded domains with homogeneous properties (discretized with FEM).

Francisco J. J. Sayas
Department of Mathematical Sciences
University of Delaware
fjsayas@math.udel.edu

Victor Dominguez
Universidad Publica de Navarra
Spain
victor.dominguez@unavarra.es

Matthew Hassell
Department of Mathematical Sciences
University of Delaware

mhassell@math.udel.edu

MS208

Fractional Schrödinger Dynamics

We study the dynamics of the Schrödinger equation with a fractional Laplacian $(\Delta)^{\alpha}$. Analytically, we find equations describing the dynamics of the expected position and expected momentum in the fractional Schrödinger equation, equations that are the fractional counterpart of the Newtonian equations of motion for the non-fractional Schrödingier equation ($\alpha = 1$). We also propose a numerical method for to study the dynamics of fractional Schrödinger equation and find that the nonlocal interaction from the fractional Laplacian introduces decoherence into the wave function.

Yanzhi Zhang
Missouri University of Science and Technology
zhangyanz@mst.edu

MS209

Towards Ab-Initio Simulations of Nanoelectronic Devices

To accurately simulate always smaller devices and predict their "current vs. voltage" characteristics prior to fabrication, the usage of a quantum transport (QT) solver has become a must. Also, empirical models such as tight-binding cannot be considered as fully reliable when it comes to nanostructures so that in many applications they should be replaced by $ab-initio$ solutions. Here, we will therefore present the core algorithms of a DFT-based approach that combines the CP2K community code with an existing QT simulator. As key feature the resulting tool can treat electronic transport in 2-D and 3-D systems with more than 10,000 atoms.

Mathieu Luisier, Mauro Calderara, Sascha Brueck
Integrated Systems Laboratory
ETH Zurich
mluisier@iis.ee.ethz.ch, mcalderara@iis.ee.ethz.ch,
bruecks@iis.ee.ethz.ch

Hossein Bani-Hashemian, Joost VandeVondele
Nanoscale Simulations
ETH Zurich
hossein.banihashemian@mat.ethz.ch,
joost.vandevondele@mat.ethz.ch

MS209

Truly Scalable O(N) Approach for First-Principles Molecular Dynamics (FPMD) of Non-Metallic Systems

We present a scalable O(N) FPMD algorithm based on a non-orthogonal localized orbitals formulation of DFT. A scalable strategy is used to approximately compute selected elements of the inverse of the associated Gram matrix. The algorithm exploits sparsity and uses nearest neighor communication only for excellent scalability. Accuracy is controlled by the mesh spacing of the discretization, the size of the localization regions confining the orbitals, and a cutoff radius for the Gram matrix. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Jean-Luc Fattebert
Lawrence Livermore National Lab.
fattebert1@llnl.gov

Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
oseikuffuor1@llnl.gov

MS209

Ab Initio Quantum Monte Carlo in Computational Materials Science and Chemistry

Ab initio Quantum Monte Carlo is an electronic structure method that is highly accurate, well suited to large scale computation, and potentially systematically improvable in accuracy. The method has recently been applied to transition metal oxides and to the prediction of defect properties and thermodynamics of materials, where established electronic structure methods have difficulty reaching the accuracies desired to inform experiment and the design and selection of materials. Nevertheless there remain significant challenges to achieving a methodology that is robust and systematically improvable in practice, e.g., where the statistical nature of the method precludes the use of algorithms developed for density functional theory or quantum chemical techniques. In this talk I will outline the current state of the art and describe several opportunities for advances in the mathematical foundation and numerical implementation of key components of the method.

Paul Kent
Oak Ridge National Laboratory
kentpr@ornl.gov

MS209

Using Next-generation Architectures to Model Large and Complex Molecular Environments

Modelling large and complex molecular environments requires fast, computationally scalable and efficient computational tools utilizing novel algorithmic approaches on the latest hardware technologies. Results of optimization efforts of the high-accuracy coupled cluster triples and the solid state planewave modules in the NWChem computational chemistry software for Intel Knights Corner and Landing multicore architectures combined with algorithmic advances will be discussed.

Bert de Jong, Hongzhang Shan
Lawrence Berkeley National Lab
wadejong@lbl.gov, hshan@lbl.gov

Leonid Oliker
Lawrence Berkeley National Laboratory
loliker@lbl.gov

MS210

Segmental Refinement: A Multigrid Technique for Data Locality

We investigate a technique, segmental refinement (SR), proposed by Brandt in the 1970s as a low memory multigrid method. The technique is attractive for modern computer architectures because it provides high data locality, minimizes network communication, is amenable to loop fusion, and is naturally highly parallel and asynchronous. The

network communication minimization property was recognized by Brandt and Diskin in 1994; we continue this work by developing a segmental refinement method for a finite volume discretization of the 3D Laplacian on massively parallel computers. An understanding of the asymptotic complexities, required to maintain textbook multigrid efficiency, are explored experimentally with a simple SR method. A two-level memory model is developed to compare the asymptotic communication complexity of a proposed SR method with traditional parallel multigrid. Performance and scalability are evaluated with a Cray XC30 with up to 64K cores. We achieve modest improvement in scalability from traditional parallel multigrid with a simple SR implementation.

Mark Adams
Lawrence Berkeley Laboratory
mfadams@lbl.gov

MS210

Comparative Performance Analysis of an Algebraic Multigrid Solver on Leading Multicore Architectures

We present a comparative performance analysis of a novel element-based algebraic multigrid method combined with a robust coarse-grid solution technique based on HSS low-rank sparse factorization. Our test datasets come from the SPE Comparative Solution Project for oil reservoir simulations. We contrast the performance of the code on one 12-core CPU of a Cray XC30 with that on a 60-core Intel Xeon Phi. The steps we required to obtain top performance are described in detail.

Alex Druinsky
Computational Research Division
Lawrence Berkeley National Laboratory
adruinsky@lbl.gov

Brian Austin
NERSC
Lawrence Berkeley National Laboratory
baustin@lbl.gov

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

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

Eric Roman, Samuel Williams
Lawrence Berkeley National Laboratory
eroman@lbl.gov, swwilliams@lbl.gov

MS210

Iterative Performance of Monte Carlo Linear Solver Methods

Stochastic linear solvers based on Monte Carlo Synthetic Acceleration are being studied as a potentially resilient alternative to standard methods. Our work has shown that for certain classes of problems, these methods can also demonstrate performance competitive with modern techniques. In this talk we will review recent developments in parallelizing Monte Carlo solvers for linear systems. Our work targets large, distributed memory systems by adapting parallel algorithms developed in part by the radiation transport community.

Massimiliano Lupo Pasini
Department of Mathematics and Computer Science
Emory University
massimiliano.lupo.pasini@emory.edu

MS210

Parallel Algorithms for the Monte Carlo Synthetic Acceleration Linear Solver Method

Stochastic linear solvers based on Monte Carlo Synthetic Acceleration are being studied as a potentially resilient alternative to standard methods. Our work has shown that for certain classes of problems, these methods can also demonstrate performance competitive with modern techniques. In this talk we will review recent developments in parallelizing Monte Carlo solvers for linear systems. Our work targets large, distributed memory systems by adapting parallel algorithms developed in part by the radiation transport community.

Stuart Slattery
Computer Science and Mathematics Division
Oak Ridge National Laboratory
slatterysr@ornl.gov

Tom Evans
Oak Ridge National Laboratory
evanstm@ornl.gov

Steven Hamilton
ORNL
hamiltonsp@ornl.gov

MS211

Energy-based Inner Products for POD/Galerkin Model Reduction for Compressible Flows

The focus of this talk is the development of energy-based inner products for POD/Galerkin model reduction for compressible flows. An energy stability analysis reveals that the inner product employed in the Galerkin projection step of the model reduction dictates the ROMs stability. Following the review of a symmetry inner product that gives rise to a stable formulation for linearized compressible flow, a new energy-based inner product for nonlinear compressible flow is derived and evaluated.

Jeffrey Fike, Irina K. Tezaur, Matthew Barone, Srinivasan Arunajatesan
Sandia National Laboratories
jafike@sandia.gov, ikalash@sandia.gov, mbarone@sandia.gov, sarunaj@sandia.gov

MS211

Data-driven Optimal Rational Approximation via Numerical Quadrature

Iterative Rational Krylov Algorithm (IRKA) has proved very successful in producing optimal rational approximations. The main cost for IRKA is the need to evaluate the underlying transfer function and its derivatives at new points at every step. In this talk, we present a quadrature-based framework for IRKA that will remove this computa-

tional cost. Transfer function will be evaluated only in an off-line phase at selected quadrature nodes and iteration will not need new function evaluations.

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

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

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

MS211

Reduced Order Modeling of Geophysical Flows

The reduced order models (ROMs) are frequently used in the simulation of complex flows to overcome the high computational cost of direct numerical simulations. The proper orthogonal decomposition (POD), as one of the most commonly used tools to generate ROMs, has been utilized in many engineering and scientific applications. For many complex flows, however, its original promise of computationally efficient, yet accurate approximation of coherent structures still remains to be fulfilled. To balance the low computational cost required by ROMs and the complexity of the targeted flows, appropriate closure modeling strategies need to be employed. Several closure models for the POD-ROMs of structurally dominated flows are carefully derived and numerically investigated in the context of the quasi-geostrophic equations, which model the large scale wind-driven ocean flows.

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MS211

Efficient Reduced Basis Methods for Contact and Related Problems

We present online-efficient RB methods for contact and related problems. Existing methods are inefficient since the online cost to compute the error estimates depends on the dimension of the FE problem. Furthermore, the resulting error bounds are often pessimistic. We present two alternative, online-efficient approaches. The first approach introduces an additional problem which enables the computation of sharp(er) error bounds. The second approach recasts the nonlinearity to allow treatment by the Empirical Interpolation Method (EIM).

Karen Veroy-Grepl, Zhenying Zhang, Eduard Bader
Graduate School AICES
RWTH Aachen University
veroy@aices.rwth-aachen.de, zhang@aices.rwth-aachen.de, bader@aices.rwth-aachen.de

Mark Kaercher
RWTH Aachen University
kaercher@aices.rwth-aachen.de

MS212

Optimal Explicit Strong Stability Preserving RungeKutta Methods with High Linear Order and Optimal Nonlinear Order

The search for high order strong stability time-stepping methods with large allowable strong stability coefficient has been an active area ofresearch over the last two decades. This research has shown that explicit SSP RungeKutta methods exist only up to fourth order. However, if we restrict ourselves to solving only linear autonomous problems, the order conditions simplify and this order barrier is lifted: explicit SSP RungeKutta methods of any linear order exist. These methods reduce to second order when applied to nonlinear problems. In the current work we aim to find explicit SSP RungeKutta methods with large allowable time-step, that feature high linear order and simultaneously have the optimal fourth order nonlinear order. These methods have strong stability coefficients that approach those of the linear methods as the number of stages and the linear order is increased. This work shows that when a high linear order method is desired, it may be still be worthwhile to use methods with higher nonlinear order.

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

Zachary J. Grant
University of Massachusetts at Dartmouth
zgrant@umassd.edu

Daniel L. Higgs
University of Massachusetts, Dartmouth
DanielHiggs@gmail.com

MS212

Strong Stability Preserving General Linear Methods

We describe the construction of strong stability preserving (SSP) general linear methods (GLMs) for ordinary differential equations. This construction is based on the monotonicity criterion for SSP methods. This criterion can be formulated as a minimization problem, where the objective function depends on the Courant-Friedrichs-Levy (CFL) coefficient of the method, and the nonlinear constrains depend on the unknown remaining parameters of the methods. The solution to this constrained minimization problem leads to new SSP GLMs of high order and stage order. This is a joint work with Giuseppe Izzo from the University of Naples.

Zdzislaw Jackiewicz
Arizona State University
jackiewicz@asu.edu

MS212

Stability-Optimized Time Integrators for WENO Discretizations

We present numerical time integrators designed for use with WENO and compact-WENO spatial discretizations of hyperbolic PDEs. These integrators have been optimized for, in order of priority

1. Absolute stability with respect to the WENO semi-discretizations (using the convex optimization technique of Ketcheson & Ahmadia)
2. Strong stability preservation
3. Accuracy

We present theoretical and numerical evidence that these new methods allow the use of larger time steps without sacrificing accuracy.

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Debojyoti Ghosh
Mathematics and Computer Science Division
Argonne National Laboratory
ghosh@mcs.anl.gov

MS212

Implicit-Explicit General Linear Methods

Implicit-explicit (IMEX) time stepping methods can efficiently solve differential equations with both stiff and non-stiff components. In this work we study new implicit-explicit methods of general linear type. We develop an order conditions theory for high stage order partitioned GLMs that share the same abscissae, and show that no additional coupling order conditions are needed. Consequently, GLMs offer an excellent framework for the construction of multi-method integration algorithms. Next, we propose a family of IMEX schemes based on diagonally-implicit multi-stage integration methods and construct practical schemes of order up to five with numerically optimized stability regions. The new methods have similar stability properties as IMEX Runge-Kutta methods, but they do not suffer from order reduction, and are superior in terms of accuracy and efficiency. Numerical experiments with two and three dimensional test problems illustrate the potential of the new schemes to speed up complex applications.

Hong Zhang
Virginia Tech
zhang@vt.edu

MS213

Sampling Strategies for L1 Minimization

One of the biggest challenges of uncertainty quantification is the inability to 'densely' sample the model input space due to high-dimensionality and/or the immense computational expense of a high-fidelity model. The need to 'densely' sample can be often be alleviated by building a sparse Polynomial Chaos Expansion (PCE) using l1-minimization. In this talk I will present sampling, preconditioning and basis adaptive strategies that allow one to accurately approximate high-dimensional compressible functions from limited data using any orthonormal polynomial basis.

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Akil Narayan
University of Massachusetts Dartmouth
akil.narayan@umassd.edu

MS213

Reweighted Minimization Method for Uncertainty Quantification of Microscopic Modeling

In this talk, reweighted minimization method for uncertainty quantification of microscopic modeling will be discussed. Fast and accurate surrogate model is very useful in mescoscopic modeling. It helps to evaluate the quantity of interest, calibrate the model e?ciently (with Bayesian framework), to study the uncertainty in the system quickly, etc. In this work, reweighted minimization method is employed for microscopic modeling when the system includes i.i.d. uniform random variables. The proposed method enhances our ability of exploiting information from limited source of experiments or simulations. We employ the new method to practical problems in mesoscopic modeling in physical chemistry, electrochemistry, biophysics, biochemistry, etc., to help to quantify the uncertainty, calibrate model, design experiments, etc. In conclusion, this work will provide a ?exible approach of e?ciently utilizing limited experimental or simulation results to construct an accurate surrogate model.

Guang Lin, Xiu Yang, Huan Lei
Pacific Northwest National Laboratory
guang.lin@pnnl.gov, xiu.yang@pnnl.gov, huan.lei@pnnl.gov

MS213

Least Square Methods for Low-Rank Approximations with Sparsity Inducing Regularization

Approximation of high dimensional stochastic functions using functional approaches is often limited by the so called curse of dimensionality. In literature, approximation methods often rely on exploiting particular structures of high dimensional functions. One such structure which is increasingly found to be applicable in these functions is sparsity on suitable basis. Also, these functions exhibit optimal low rank representation and can be approximated in suitable low rank tensor subsets. In this work, we exploit sparsity within low rank representation to approximate high dimensional stochastic functions in a non intrusive setting using few sample evaluations. The proposed method can also be combined with clustering and classification approaches to approximate high dimensional irregular and discontinuous functions.

Prashant Rai
LUNAM Universite, Ecole Centrale Nantes, CNRS, GeM
prashant.rai@ec-nantes.fr

Mathilde Chevreuil
LUNAM Universite, Universite de Nantes, CNRS, GeM
mathilde.chevreuil@univ-nantes.fr

Loïc Giraldi
GeM, Ecole Centrale de Nantes
loic.giraldi@ec-nantes.fr

Anthony Nouy
LUNAM Universite, Ecole Centrale Nantes, CNRS, GeM

anthony.nouy@ec-nantes.fr

MS213

Interpolation Via Weighted L1 Minimization

Functions of interest are often smooth and sparse in some sense. Classical linear interpolation methods are effective under strong regularity assumptions, but cannot incorporate nonlinear sparsity structure. At the same time, nonlinear methods such as L1 minimization can reconstruct sparse functions from very few samples, but do not necessarily encourage smoothness. Here we show that weighted L1 minimization effectively merges the two approaches, promoting both sparsity and smoothness in reconstruction. We consider the implications of these results for spherical harmonic and polynomial interpolation, in the univariate and multivariate setting. Along the way, we extend concepts from compressive sensing such as the restricted isometry property and null space property to accommodate weighted sparse expansions; these developments should be of independent interest in the study of structured sparse approximations and continuous-time compressive sensing problems.

Rachel Ward
Department of Mathematics
University of Texas at Austin
rward@math.utexas.edu

Holger Rauhut
Aachen University
rauhut@mathc.rwth-aachen.de

MS214

Nonlinear Image Registration with a Sliding Motion Deformation Model

Common medical image registration approaches are enforcing a global continuity of the estimated deformation field. However, for deformations like e.g. the sliding of the lung along the ribcage a continuous deformation estimation is not feasible. Therefore we present in this talk a registration framework that models discontinuities in the deformation field along organ boundaries described by arbitrary orientable submanifolds. The incorporated methods involve constrained nonlinear registration in the Lagrange frame and a finite element discretization.

Alexander Derksen
Institute of Mathematics and Image Computing
University of Luebeck
alexander.derksen@mic.uni-luebeck.de

MS214

Efficient Algorithms for Physically Constrained Diffeomorphic Image Registration

We treat image registration as a problem of optimal control. The deformation is represented by its velocity. This leads to a constrained variational optimization problem, where the constraints are partial differential equations (PDEs) . We augment standard H^1 and H^2 smoothing regularization with differential constraints on the velocity that both encapsulate prior knowledge and explicitly control the determinant of the deformation gradient. The associated optimality conditions are a system of space-time non-linear multi-component PDEs that is challenging to solve in an efficient way. We solve for the first-order optimality conditions using a matrix-free preconditioned Gauss-Newton-Krylov method for the Schur complement of the velocity field. We use a spectral Galerkin method in time to reduce the number of unknowns. Also, we use a spectral discretization in space, which in turn allows for an efficient preconditioning of the Hessian.

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

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS214

Efficient Algorithms for High-Resolution Diffusion-Weighted MRI

Diffusion-weighted Magnetic Resonance Imaging (DW-MRI) allows the acquisition of functional information in-vivo and has important clinical applications, for instance, related to Ischemia and Alzheimer's disease. In this talk, I will present computational methods to overcome two limitations currently inherent in DW-MRI: Geometrical distortions due to measurement artifacts and low spatial resolution. The presented 3D correction and super-resolution method is based on a physical prior and employs a slice-wise parallelization to increase computational efficiency.

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

MS214

Constrained Optimal Control Approaches in Large Deformation Diffeomorphic Metric Mapping

The large deformation diffeomorphic metric mapping (LDDMM) algorithm, which addresses image or shape registration, can be interpreted as an optimal control problem in some suitable high-dimensional space of shapes. It has led to a large number of applications in the domain of computational anatomy, in which the focus is set on the analysis of anatomical variations in relation with disease. In some cases this algorithm can be invoked with additional constraints, which leads to interesting new applications and challenging implementation problems. We will review some of these examples in this talk, including, in particular, issues related to the combined registration of multiple shapes, and to the evaluation of deformations with atrophy constraints.

Laurent Younes
Center for Imaging Science , The Johns Hopkins University
laurent.younes@jhu.edu

MS215

Preconditioned MCMC and Adaptive Posterior Refinement Leveraging Sparse PCE

The process of performing Bayesian inference is frequently hindered by expense and by a lack of reliability in the MCMC sampling used for computing the posterior distribution. We present a recent prototype for adaptive

emulator-based inference that employs ℓ_1-regularized regression for sparse polynomial chaos emulation, exploits analytic derivatives from the emulator to inform the proposal density, and then refines the emulator by performing new model evaluations for points with high posterior density.

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Laura Swiler
Sandia National Laboratories
Albuquerque, New Mexico 87185
lpswile@sandia.gov

MS215

Accelerated Bayesian Inference with Transport Maps

The essential challenge in Bayesian computation for UQ is one of characterizing the posterior distribution, which often is high-dimensional and non-Gaussian. We will show that the use of transport maps can dramatically accelerate Bayesian inference in this setting. First, we use a combination of optimal transport and the Metropolis-Hastings rule to yield a new adaptive MCMC approach for exact inference. Second, we develop new conditioning techniques, with tunable accuracy, that can utilize massively parallel offline computation.

Matthew Parno, Youssef M. Marzouk
Massachusetts Institute of Technology
mparno@mit.edu, ymarz@mit.edu

MS215

Sparse, Adaptive Smolyak Quadrature Algorithms for Stochastic Inverse Problems

Deterministic, dimension-adaptive quadrature methods for Bayesian inverse problems of parametric operator equations with distributed uncertain inputs are proposed. For data from a sparsity class, the parametric, deterministic density of the Bayesian posterior belongs to the same sparsity class. Dimension-independent convergence rates for dimension-adaptive Smolyak and QMC integration algorithms are shown. In the vanishing observation noise variance limit the posterior concentrates near the MAP estimate. We give asymptotic expansions of the Bayesian estimate w.r. to observation noise variance. A "variable metric' rescaling of the posterior near concentration points "preconditions' quadrature for vanishing observation noise covariance. Numerical results in agreement with the theory are presented. Supported by SNF and ERC under AdG247277.

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

Claudia Schillings
ETH Zurich
SAM
claudia.schillings@sam.math.ethz.ch

MS215

Quasi Optimal Sparse-Grid Approximation of Random Elliptic PDEs

Solutions of PDEs depending on parameters/random coefficients can be conveniently approximated by polynomial expansions over the parameter space. However, these approximations suffer from a performance degradation as the number of random parameters increases ("curse of dimensionality' effect). In this talk we will propose an "a-priori/a-posteriori profit approach' to minimize such effect for sparse grids approximations. The efficiency of the proposed technique will be supported by theoretical convergence results and numerical tests.

Lorenzo Tamellini
EPF-Lausanne, Switzerland,
lorenzo.tamellini@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

MS216

Approaches to Evaluate Interactions in Collaborative Groundwater Management

Abstract not available at time of publication.

Joseph Amaya
Texas A&M University - Kingsville
joseph.amaya@tamuk.edu

MS216

Managing Surface Water Resources in Data Sparse Regions

Abstract not available at time of publication.

Felipe Estrada
Texas Tech University
felipe.estrada@ttu.edu

MS216

Climate Change and Water Scarcity

Abstract not available at time of publication.

Donna Mitchell
Texas Tech University
donna.m.mitchell@ttu.edu

MS216

Application of Simulation-Optimization for Water Management in Hydraulic Fracturing Operations

Abstract not available at time of publication.

Elma A. Uddameri
Texas Tech University

annette.hernandez@ttu.edu

MS217

Interface Resolved Numerical Method to Study Electrokinetic Particle Assembly in Microdevices

A hybrid immersed interface-immersed boundary method is developed to study electric field driven particle assembly where both electric and hydrodynamic forces are calculated with interface-resolved approach instead of commonly used point-particle method. In this study, the Maxwell stress tensor is used to calculate the electric force acting on particles by considering the physical effect of particles in the computational domain. Thus, this method eliminates the approximations used in point dipole methods for calculating forces. A comparative study between Maxwell stress tensor and point dipole methods for computing electric forces will be presented to elucidate the shortcoming of the latter method. Next, electric field driven particle motions and related fluid flow phenomena will be presented for charged, uncharged and bipolar particles. We will particularly demonstrate the particle chaining phenomena for similar and dissimilar type particles using applied electric fields.

Prashanta Dutta
Washington State University
prashanta@wsu.edu

MS217

Classical Density Functional Theory of Charged Fluids at Interfaces

Classical density functional theory (DFT) is a statistical mechanical theory for inhomogeneous fluids, based on the minimization of a free energy functional. Here I will describe the application of DFT to charged systems and in particular the ionic correlations included by the DFT that go beyond the Poisson-Boltzmann treatment. I will present results for macroion interactions in electrolyte solutions and the behavior of charged particles in nanochannels.

Amalie Frischknecht
Sandia National Laboratories
alfrisc@sandia.gov

MS217

Efficient Parallel Implementation of Implicit SPH/MLS using LAMMPS and Trilinos

We present an efficient parallel implementation of 3D implicit mesh-free particle methods for incompressible flows. Mesh-free methods are widely used for solving complex problems that mesh-based methods cannot handle easily. However, implicit time methods are rarely applied to large-scale particle simulations due to the prohibitive computational cost solving linear systems of equations. Our implementation is based on LAMMPS and adopts Trilinos packages for implicit time integration. Numerical experiments are presented to demonstrate the parallel scalability.

Kyungjoo Kim
Sandia National Laboratories
kyukim@sandia.gov

Mauro Perego
CSRI Sandia National Laboratories
mperego@fsu.edu

Nathaniel Trask
Brown University
Department of applied mathematics
nathaniel_trask@brown.edu

Michael L. Parks
Sandia National Laboratories
mlparks@sandia.gov

MS217

Meshless Methods for the Mesoscale - High Order Implicit ALE Schemes using Collocated MLS

Meshless methods are ideal for simulating flows occurring at the mesoscale, since they trivially handle both deformation of complex boundaries and interface tracking for flows involving multispecies phenomena. Classical meshless methods tend to either be expensive or inconsistent and have mainly gained traction in the past as a low order technique. We present a collocation formulation of moving least squares (MLS) that allows high order discretization in space, and a fully implicit ALE projection scheme that is demonstrated to provide up to third order accuracy in time. When the resulting system is solved using algebraic multigrid preconditioners, the scheme provides a scalable, flexible framework for simulating multiphysics at the mesoscale.

Nathaniel Trask
Brown University
Department of applied mathematics
nathaniel_trask@brown.edu

Kyungjoo Kim
Sandia National Laboratories
kyukim@sandia.gov

Mauro Perego
CSRI Sandia National Laboratories
mperego@fsu.edu

MS218

A Sampling Filter for Non-Gaussian Data Assimilation

Current operational ensemble-based filters like Ensemble Kalman Filter (EnKF), and Maximum Likelihood Ensemble Filter (MLEF), usually fail in case of non-linear observations or non-Gaussian distributions. We propose a general ensemble-based data assimilation method that works by sampling directly from the posterior distribution following a Hybrid Monte Carlo (HMC) ap- proach. The proposed filter was tested on Lorenz-96 model with several observation operators. Results show that this filter is capable of handling non-linear as well as linear observations.

Ahmed Attia
Virginia Tech
attia@vt.edu

Adrian Sandu
Virginia Polytechnic Institute and
State University

sandu@cs.vt.edu

MS218

Bayesian Nonlinear Smoothing and Adaptive Sampling

New schemes are presented for optimal Bayesian nonlinear state estimation and adaptive sampling of large nonlinear fluid and ocean dynamical systems, both forward and backward in time. The Bayesian nonlinear smoothing combines reduced-order Dynamically-Orthogonal (DO) equations with Gaussian Mixture Models (GMMs), extending linearized backward pass updates to a Bayesian nonlinear setting. Bayesian nonlinear adaptive sampling schemes are then derived to predict the observations to be collected that maximize information about variables of interest. When combined with rigorous time-optimal path planning, we obtain efficient coordinated swarms of autonomous ocean sampling systems.

Pierre F.J Lermusiaux
Massachusetts Institute of Technology
pierrel@mit.edu

Tapovan Lolla
MIT
ltapovan@mit.edu

MS218

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

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

Allison Lewis
North Carolina State University
allewis2@ncsu.edu

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

Brian Williams
Los Alamos National Laboratory
brianw@lanl.gov

MS218

Displacement Data Assimilation

We are developing a data assimilation methodology specifically geared to problems in which the preservation of features and topological structures is crucial. The application of this methodology is aimed at improving estimates of such things as hurricane tracks and the Lagrangian tracks of tracers in a flow. The basic strategy combines conventional nonlinear/non-Gaussian data assimilation on the state variables as well as on the underlying kinematics of the flow. In addition to describing the methodology we show comparisons of displacement data assimilation to more conventional assimilation in the context of pure advection of a tracer in a flow constrained kinematically by area preserving maps. This comparison will demonstrate that our methodology minimizes phase errors and thus delivers better estimates of the topological features of advected tracers.

Juan M. Restrepo
Departments of Mathematics, Atmospheric Physics, and Physics
University of Arizona
restrepo@math.oregonstate.edu

Steven Rosenthal
University of Arizona
srosenthal@gmail.com

Shankar C. Venkataramani
University of Arizona
Department of Mathematics
shankar@math.arizona.edu

Arthur Mariano
RSMAS/MPO
University of Miami
amariano@rsmas.miami.edu

MS219

Statistical Metrics for Assessing Quality of Scenarios for Unit Commitment and Dispatch

Wind power scenarios for use in stochastic unit commitment should accurately represent the stochastic process for available wind power. We employ statistical evaluation metrics to assess whether a scenario set possesses properties that are expected to minimize expected cost. A new mass transportation distance rank histogram assesses calibration of unequally likely scenarios according to their bias, variability and autocorrelation. Energy scores, rank histograms, and Brier scores are applied to alternative sets of wind power scenarios.

Sarah M. Ryan, Didem Sari
Iowa State University
smryan@iastate.edu, dsari@iastate.edu

MS219

Adaptive Robust Optimization with Dynamic Uncertainty Sets for Power System Operations

In this talk, we present an adaptive robust optimization model with a new type of dynamic uncertainty sets for power system operations under significant renewable generation uncertainty. We introduce a data-driven framework that fuses statistical estimation with uncertainty set construction to model temporal and spatial correlations of renewable generation. Computational results show the advantages of the proposed robust model in terms of operational cost and system reliability over existing deterministic and robust models.

Andy Sun
Georgia Tech
College of Engineering
andy.sun@isye.gatech.edu

Alvaro Lorca
Georgia Institute of Technology
Industrial and Systems Engineering
alvarolorca@gatech.edu

MS220

Reconstructed Discontinuous Galerkin (RDG) Method for Multi-Material Flows on Unstructured Meshes

In this work, we discuss extension of the Reconstructed Discontinuous Galerkin method, developed previously for high-order spatio-temporal discretization of single-fluid flows on hybrid unstructured meshes, to fluid flow applications with multi-material interfaces. The focus is placed on challenging issues of hierarchical WENO based limiting of discontinuous flow solutions, modal orthogonal basis functions, interface tracking using combination of volume tracking and level set methods, sharp treatment of interfacial jump conditions, and combination with Arbitrary Lagrangian-Eulerian (ALE) techniques.

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

MS220

An Eulerian Projection Method for Quasi-Static Elastoplasticity

A well-established numerical approach to solve the Navier–Stokes equations for incompressible fluids is Chorin's projection method, whereby the fluid velocity is explicitly updated, and then an elliptic problem for the pressure is solved, which is used to project the velocity field to maintain the incompressibility constraint. In this talk, a mathematical correspondence between Newtonian fluids in the incompressible limit and elastoplastic solids in the slow, quasi-static limit will be presented. Using this correspondence, a new fixed-grid, Eulerian numerical method for simulating quasi-static elastoplastic solids will be developed, whereby the stress is explicitly updated, and then an elliptic problem for the velocity is solved, which is used to project the stress to maintain the quasi-staticity constraint. Numerical tests of the method will be given, and a number of correspondences between incompressible fluid mechanics and quasi-static elastoplasticity will be shown, creating possibilities for translating other numerical methods between the two classes of physical problems.

Chris H. Rycroft
Harvard SEAS
chr@seas.harvard.edu

MS220

A Robust and Efficient Solver for Interfacial Multiphase Flows on Unstructured Grids

We present a robust, accurate and efficient numerical framework for multiphase interfacial fluid dynamics on unstructured grids of arbitrary shapes of elements. The multi-moment finite volume method is used as a reliable and practical solver for complex flows in presence of complex geometries, which well balances numerical accuracy and computational cost. The free interfaces are computed by the THINC scheme which is an accurate algebraic VOF method well-suited for unstructured grids.

Feng Xiao, Bin Xie, Sun Ziyao
Tokyo Institute of Technology
xiao@es.titech.ac.jp, xie.b.aa@m.titech.ac.jp,
sun.z.ab@m.titech.ac.jp

MS220

A New Incompressibility Discretization for a Hybrid Particle Mac Grid Representation with Surface Tension

We take a particle based approach to incompressible free surface flow motivated by the fact that an explicit representation of the interface geometry and internal deformations gives precise feedback to an implicit solver for surface tension. Methods that enforce incompressibility directly on the particles are typically numerically inefficient compared to those that utilize a background grid. However, background grid discretizations suffer from inaccuracy near the free surface where they do not properly capture the interface geometry. Therefore, our incompressibility discretization utilizes a particle based projection near the interface and a background MAC grid based projection for efficiency in the vast interior of the liquid domain as well as a novel method for coupling these two disparate projections together. We show that the overall coupled elliptic solver is second order accurate, and remains second order accurate when used in conjunction with an appropriate temporal discretization for parabolic problems. A similar second order accurate discretization is derived when the MAC grid unknowns are located on faces (as opposed to cell centers) so that Navier-Stokes viscosity can be solved for implicitly as well. Finally, we present a fully implicit approach to surface tension that is robust enough to achieve a steady state solution in a single time step. Beyond stable implicit surface tension for our novel hybrid discretization, we demonstrate preliminary results for both standard front tracking and the particle level set method.

Wen Zheng, Bo Zhu
Stanford University
zhw@stanford.edu, boolzhu@stanford.edu

Byungmoon Kim
Adobe Systems Inc.
bmkim@adobe.com

Ronald Fedkiw
Stanford University
fedkiw@cs.stanford.edu

MS221

Applications of Distributed Methods to Non-Traditional Linear Systems

We investigate solutions to multi-dimensional linear systems using single machine parallelization techniques. The tri-color-channel radioisty method is used as a test case where the Jacobi method is applied to solve the radiosity linear equation. The multi-dimensionality of the system is represented through a structure of arrays and an array of structures, where the latter was found to be optimal. Work distribution was implemented with CUDA and OpenMP where OpenMP threading was found to be optimal.

Julian Gilyard
Department of Computer Science
Wake Forest University
gilyja12@wfu.edu

Thomas Stitt

Pennsylvania State University
Department of Electrical Engineering
tms54@psu.edu

Oluwapelumi Adenikinju, Joshua Massey
University of Baltimore Maryland County
Department of Computer Science and Electrical Engineering
adenik2@umbc.edu, masseyj1@umbc.edu

MS221

Multigrid Solvers on Heterogeneous Architectures

HHG is a multigrid finite element solver that scales to a million threads and that can solve in excess of a trillion (10^{12}) unknowns in about a minute compute time. This is achieved by a careful co-design of grid structure, discretization, multigrid components, and the hybrid parallelization, all driven by a meticulous performance engineering process. HHG is now being extended to use GPU kernels and we will report on the acceleration achieved.

Björn Gmeiner
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
bjoern.gmeiner@informatik.uni-erlangen.de

Daniel Iuhasz
FAU Erlangen
danieliuhasz@gmail.com

Sebastian Kuckuk
Universität Erlangen-Nürnberg
kuckuk@i10.informatik.uni-erlangen.de

Markus Stuermer, Harald Koestler
University of Erlangen-Nuremberg
markus.stuermer@informatik.uni-erlangen.de,
harald.koestler@informatik.uni-erlangen.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

MS221

Speeding Up Sparse Triangular Solution on Multicores and GPUs

Multicores have complex memory hierarchy and increased parallelism; GP-GPUs have tremendous amount of thread-based parallelism. We examine structures to expose the underlying parallelism present in sparse-matrices to utilize multicore NUMA architecture and GPU. Specifically we have developed CSR-k, a multilevel form of the traditional compressed sparse row format that can be mapped to NUMA hierarchy. We formulate sparse triangular solve in terms of CSR-k with coloring and demonstrate how this perform with CPU and GPU.

Humayun Kabir
Department of Computer Science and Engineering
The Pennsylvania State University
hzk134@psu.edu

Joshua D. Booth
The Pennsylvania State University
jdb5132@cse.psu.edu

Padma Raghavan
The Pennsylvania State Univ.
Dept of Computer Science Engr.
raghavan@cse.psu.edu

MS221

General SpMV and SpMM for AMG on GPUs

Scientific computations on sparse graphs mainly require two computational patterns: processing of neighbors (general SpMV) and processing of all neighbors of a set of nodes, often all neighbors of neighbors (general SpMM). These two patterns have different properties and requirements for massively parallel implementations. The main challenge for the former is the different number of neighbors and indirection in access, while the latter has the dominating problems of index matching and varying output sizes. Once these problems are solved in parallel an algebraic multigrid solver (AMG) can be efficiently assembled on the GPU.

Robert Strzodka
NVIDIA
strzodka@mpi-inf.mpg.de

MS222

Industrial Mathematics Education at Worcester Polytechnic Institute

Central to the philosophy of learning at Worcester Polytechnic Institute (WPI) is the principle of project-based learning. In the Department of Mathematical Sciences through our Center for Industrial Mathematics and Statistics (CIMS) we have a strong focus on industry-sponsored projects for students. These projects are based on real-world problems that come directly from business or government. Students work on these projects under the guidance of a faculty advisor and an industrial liaison. These projects are used in many settings, including senior projects for the Bachelor's degree, Master's degree projects, and summer Research Experience for Undergraduates (REU) projects. The focus of this talk will be about the project process and its benefits to our WPI. Examples will be provided for illustration.

Marcel Blais
Worcester Polytechnic Institute
myblais@wpi.edu

MS222

PIC Math: Preparation for Industrial Careers in Mathematical Sciences

PIC Math is a new program to prepare students in the mathematical sciences to succeed in careers in business, industry, and government (BIG). Funded by a 2 million dollar NSF grant, this program (a) helps students be aware of their choices for non-academic careers and opportunities for internships, (b) helps faculty be more fully aware of non-academic career options for their students, make connections with people working for local BIG organizations, and develop internship opportunities for their students, (c) offers students the opportunity to have a research experience related to real-world problems from BIG during a spring semester course, and (d) provide training to students and faculty in how to successfully work on problems

from BIG and develop the needed communications skills. To accomplish these objectives, we are developing a set of educational and informative videos, conducting summer training workshops for faculty, and preparing materials for a semester-long course in which students learn skills and work on research problems from BIG.

Michael Dorff
Brigham Young University
mdorff@math.byu.edu

MS222

A New Curriculum in Applied and Computational Mathematics

We present BYU's new undergraduate curriculum in Applied and Computational Mathematics. We highlight the main features of this program and discuss how we have been able to overcome numerous pedagogical, institutional, and cultural challenges, particularly as we've launched an applied math degree in a predominantly pure math department. We also show how we have been able to teach a number of advanced topics to undergraduates, and we share our approach to recruiting and retaining students.

Jeffrey Humpherys
Brigham Young University
jeffh@math.byu.edu

MS222

A Student Perspective on Industrial Capstones at Harvey Mudd College

This talk will give a student's perspective on Harvey Mudd's Mathematics Clinic program, in which seniors majoring in mathematics work in teams to complete industrial projects. These projects give ample opportunity not only to hone mathematical skills, but also to practice effective collaboration, communication, and project management. As a student, the combination of faculty mentorship and student-driven work provides an engaging transition into industrial careers in mathematics.

Elizabeth Schofield
Harvey Mudd College
eschofield@g.hmc.edu

MS223

Multilevel Projection Method for Nonlinear Radiative Transfer Problems

This talk presents a deterministic computational method for solving coupled radiative transfer and energy balance equations. It uses a multilevel system of equations consisting of the high-order radiative transfer equation, multigroup low-order quasidiffusion (LOQD) and grey LOQD equations defined for moments of the specific intensity. The energy balance equation is coupled to the grey LOQD equations. Temperature is evaluated in a projected space of smallest dimensionality. We study discretization and linearization methods for coupled multiphysics equations.

Dmitriy Y. Anistratov
North Carolina State University
anistratov@ncsu.edu

MS223

A Multigrid Method for Two-Dimensional Discrete-Ordinates Radiation-Transport Calculations

We present a multigrid method for discrete-ordinates radiation-transport calculations, in particular for the Self-Adjoint Angular Flux (SAAF) form of the transport equation in two-dimensional Cartesian geometry. For smoothing, we employ multistage cellwise block Jacobi iteration. The use of the SAAF equation allows straightforward determination of optimal multistage parameters for small spatial cells; single-stage cellwise block Jacobi iteration is not effective in this regime. On coarse grids, we apply a Galerkin discretization based on bilinear interpolation.

Jeffery D. Densmore, Daniel Gill, Justin Pounders
Bettis Atomic Power Laboratory
jeffery.densmore@unnpp.gov, daniel.gill@unnpp.gov, justin.pounders@unnpp.gov

MS223

A Conservative High-Order / Low-Order Method Based Upon a Non-Conservative High-Order Least Squares S_n Formulation

We have developed a high-order/low-order (HO/LO) scheme for neutronics eigenvalue calculations. The HO equation is a second-order least-squares form of the S_n equations that contains scattering and fission sources from the LO equation. The LO equation is an angular moment equation in drift-diffusion form containing closure information from the HO equation. The HO equation generally yields a non-conservative solution, while the LO equation always yields a conservative solution adequate for reactor analysis.

Jacob Peterson
Department of Nuclear Engineering
Texas A&M
jrp314@neo.tamu.edu

Jim E. Morel
Texas A&M
morel@tamu.edu

MS223

Multilevel Monte Carlo Methods for Kinetic Equations

The adaptation of the Multilevel Monte Carlo (MLMC) method - introduced by Giles for SDEs in finance - to kinetic equations will be presented. MLMC introduces multiple time-steps and uses correlations between them for variance reduction, thereby reducing the complexity of achieving RMS error ε from $O(\varepsilon^{-3})$ to $O(\varepsilon^{-2})$. Application to the spatially homogeneous Landau-Fokker-Planck model of Coulomb collisions will be presented, as well as progress toward an MLMC-type scheme for the inhomogeneous Vlasov equation.

Lee F. Ricketson
UCLA
lfr224@cims.nyu.edu

MS224

Greedy Algorithms for Parametric Eigenvalue Problems

Some greedy algorithms will be presented to compute the solution of parametric eigenvalue problems. This work is

a continuation of the results presented in [1] where the authors presented algorithms for non-parametric eigenvalue problems. The principle of the method is to compute a reduced-order model for the solution of a parametric eigenvalue problem as a sum of pure tensor-product functions, each term being computed in an iterative way which will be precised in the talk. Some theoretical results on the convergence of these algorithms will be presented and their numerical behaviour will be illustrated on some simple test cases. [1] E. Cancs, V. Ehrlacher and T. Lelivre. "Greedy algorithms for high-dimensional eigenvalue problems", accepted for publication in Constructive Approximation.

Virginie Ehrlacher
CERMICS - Ecole des Ponts Paristech / INRIA
ehrlachv@cermics.enpc.fr

MS224

Semi-Supervised Robust Matrix Completion for Dynamic Subspace Estimation and Tracking

Recent SVD-free matrix factorization formulations have enabled rank minimization for systems with millions of rows and columns, paving the way for matrix completion for extremely high dimensional data. In this talk, we discuss a robust matrix completion and subspace tracking algorithm that uses factorized matrix decomposition with a pre-specified rank to detect and track a low rank subspace from incomplete measurements and in the presence of sparse noise. We demonstrate the performance of our algorithm for video background subtraction.

Hassan Mansour
Mitsubishi Electric Research Laboratories
mansour.hassan@gmail.com

MS224

Low-rank Approximation of Matrices and Tensors for Dynamical and Optimization Problems

In this talk we will consider several topics that are connected through the ideas of low-rank approximation. First, we propose a new method for the approximate solution of the Lyapunov equation. Second, we will discuss how the low-rank methods help to solve reaction-diffusion equations with time-dependent potential. Third, we will describe the ideas of efficient multidimensional sampling via cross approximation methods apply them to the problems of global optimization.

Ivan Oseledets
Institute of Numerical Mathematics
Russian Academy of Sciences, Moscow
ivan.oseledets@gmail.com

Denis Kolesnikov, Mikhail Litsarev
Skolkovo Institute of Science and Technology
denis.kolesnikov@skolkovotech.ru,
m.litsarev@skolkovotech.ru

MS224

Preconditioned Riemannian Optimization for Low-Rank Tensor Equations

The solution of very large linear systems is a challenging task often encountered as a core ingredient when solving partial differential equations on high-dimensional domains. In these cases, the degrees of freedom in the linear system grow exponentially with the number of dimensions, making classic approaches unfeasible. Approximation of the solution by low-rank tensor formats often allows us to avoid this *curse of dimensionality* by exploiting the underlying structure of the linear operator. We propose a new algorithm that performs a preconditioned gradient method on the manifold of tensors of fixed rank. In particular, we focus on tensors represented in the *Tensor Train* (TT) / *Matrix Product States* (MPS) format. We demonstrate the flexibility of our algorithm by comparing different approximations of the Riemannian Hessian as preconditioners for the gradient directions. Finally, we compare the efficiency of our algorithm with other tensor-based approaches such as the *Alternating Linear Scheme* (ALS). This is joint work with Michael Steinlechner and Daniel Kressner (EPF Lausanne).

Bart Vandereycken
Department of Mathematics
Princeton University
bartv@math.princeton.edu

MS225

Accurate Adaptive Loops for Finite Deformation Plasticity in Albany

The Parallel Albany Adaptive Loop with Scorec software (PAALS) provides an automated framework for adaptive finite element simulations on massively parallel machines. Within the context of finite deformation plasticity in PAALS, remapping the material state during adaptation, maintaining high quality element shapes, and updating the reference configuration are necessary to produce accurate solution results. Methods to handle these situations will be discussed and results will be shown that emphasize the interaction of these components.

Brian Granzow
Rensselaer Polytechnic Institute
granzb@rpi.edu

Glen Hansen
Sandia National Laboratories
gahanse@sandia.gov

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

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

MS225

Massively Parallel Flow Simulation using PETSc

Implicit computational fluid dynamics solvers have recently been shown to scale to the full machine at several of the largest computer facilities using linear equation solvers that were coded to match the data structures of the equation discretization. In this work, we describe efforts to apply PETSCs broad class of preconditioners to flow simulations at large core count to understand the tradeoff between improved preconditioners and the cost of the required additional data structure translation.

Michel Rasquin, Benjamin Matthews

University of Colorado Boulder
michel.rasquin@colorado.edu,
benjamin.a.matthews@colorado.edu

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

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

MS225

Variational Multiscale Analysis of Stochastic Partial Differential Equations in Albany

We present the variational multiscale (VMS) method for stochastic PDEs to compute an accurate solution in a coarse physical and stochastic space while accounting for the missing scales through a model term. The model term is based on an algebraic approximation of the fine-scale stochastic Green's function and is rational in the random variables. We consider computationally efficient approximations of this term and demonstrate their efficacy using the Albany code.

Onkar Sahni
Rensselaer Polytechnic Institute
sahni@rpi.edu

Jason Li, Jayanth Jagalur-Mohan
Department of Mechanical, Aerospace and Nuclear Engineering
RPI
lij14@rpi.edu, jagalj@rpi.edu

Assad Oberai
Department of Mechanical, Aerospace and Nuclear Engineering
Rensselaer Polytechnic Institute
oberaa@rpi.edu

MS225

Albany: A Trilinos-based code for Ice Sheet Simulations and other Applications

Albany is a finite element application code that incorporates capabilities through interoperable software components. By leveraging independently developed mathematical libraries, Albany-based applications have ready access to advanced technologies, spanning: discretizations, automatic differentiation, linear and nonlinear solvers, adjoint-based optimization and UQ, performance-portable kernels, mesh adaptivity, and dynamic load balancing. This strategy enables rapid development of sophisticated new codes, including the Albany/FELIX Ice Sheet application. Albany incorporates libraries from Trilinos, Dakota, and PUMI, some with FASTMath support.

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

Glen Hansen, Irina K. Tezaur
Sandia National Laboratories
gahanse@sandia.gov, ikalash@sandia.gov

Mauro Perego
CSRI Sandia National Laboratories
mperego@fsu.edu

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

MS226

Fast Algorithms for Shape Analysis of Planar Objects

Effective computational tools for shape analysis are needed in many areas of science and engineering. We address this problem and propose a fast iterative algorithm to compute the elastic geodesic distance between boundaries of planar objects. This algorithm is the most important component for data mining of shapes, and a fast algorithm is essential for large-scale shape analyses.The key to our algorithm is the decoupling of the optimization for the starting point and rotation and the optimization for the reparameterization function in the distance formulation. Moreover, we develop a fast dynamic programming algorithm and a nonlinear constrained optimization algorithm that work in tandem to compute optimal reparameterizations fast.

Gunay Dogan
Theiss Research, NIST
gunay.dogan@nist.gov

MS226

Deflation-based Domain Decomposition Methods

Domain decomposition methods are widely used in applied mathematics and regarded as highly scalable algorithms that can be seen as specific cases of multigrid methods. Projection operators are one of the essential tools for achieving scalability: they are used for building deflation preconditioners. A C++ framework will be presented accompanied by theoretical results to show how effective it can be to solve problems arising from Darcy's law, elasticity, or Helmholtz equation.

Pierre Jolivet
Laboratoire Jacques-Louis Lions
jolivet@ann.jussieu.fr

Frederic Nataf
Laboratoire J.L. Lions
nataf@ann.jussieu.fr

Christophe Prud'homme
University of Strasbourg
France
prudhomme@unistra.fr

MS226

Multiscale Methods for Networks

Networks are a widely used type of abstraction for complex data. Optimization of different quantitative objectives on networks often plays a crucial role in network science, not only when a practical solution is needed, but also for a general understanding of structural and statistical features of networks. We present multiscale approaches for two problems: optimal response to epidemics, and network generation. Both approaches are inspired by AMG scheme

reinforced by the algebraic distance connectivity strength.

Ilya Safro
School of Computing
Clemson University
isafro@g.clemson.edu

MS226

The Auxiliary Space Solvers and Its Applications

We talk about the mathematically optimal multigrid solvers for general unstructured grids which are robust and easy to use in practice based on the methodology of Fast Auxiliary Space Preconditioning. a new parallel unsmoothed aggregation algebraic multigrid (UA-AMG) method has also been developed based the idea of FASP. It provides (nearly) optimal load balance and predictable communication patterns factors that make our new algorithm suitable for parallel computing.

Lu Wang
Penn State University
wang_l@math.psu.edu

MS227

Robust Algorithms for Periodic Problems and Evaluation of Layer Potentials

I overview methods developed recently by our group for solving time-harmonic acoustics and Maxwell scattering from media with 10^3 layers, doubly-periodic obstacles in 3D (with efficiency for axisymmetric obstacles), and Stokes flow in periodic pipes and doubly-periodic porous media. In each case we use only the free-space Greens function, an auxiliary field, and a small linear system, to circumvent the traditional but problematic periodic Greens function. We also develop quadrature components such as 3D QBX and close evaluation for Stokes.

Alex H. Barnett
Department of Mathematics
Dartmouth College
ahb@math.dartmouth.edu

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

Adrianna Gillman
Dartmouth College
Department of Mathematics
adrianna.gillman@rice.edu

Min Hyung Cho
Department of Mathematics
Dartmouth College
min.h.cho@dartmouth.edu

Lin Zhao, Yuxiang Liu
Dartmouth College
lin.zhao.gr@dartmouth.edu,
yuxiang.liu.gr@dartmouth.edu

Bowei Wu
University of Michigan
bobbielf2@gmail.com

Gary Marple
Department of Mathematics
University of Michigan
gmarple@umich.edu

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

MS227

Adaptive Boundary Element Methods

One particular strength of the boundary element method is that it allows for a high-order pointwise approximation of the solution of the related partial differential equation via the representation formula. We propose an adaptive mesh-refining algorithm and discuss recent results on its quasi-optimal convergence behavior with respect to the point error in the representation formula. Numerical examples for the weakly-singular integral equations for the Laplacian underline our theoretical findings.

Michael Feischl, Thomas Führer, Gregor Ganter,
Alexander Haberl, Dirk Praetorius
Vienna University of Technology
michael.feischl@tuwien.ac.at,
thomas.fuehrer@tuwien.ac.at,
gregor.gantner@tuwien.ac.at,
e0925297@student.tuwien.ac.at,
dirk.praetorius@tuwien.ac.at

MS227

Fast Algorithms for the Evaluation of Layer Potentials using 'Quadrature by Expansion'

Quadrature by Expansion, or 'QBX', is a systematic, high-order approach to singular quadrature that applies to layer potential integrals with general kernels on curves and surfaces. Being based on a scheme for close evaluation due to Barnett, the scheme provides a unified evaluation capability for layer potentials. This talk discusses algorithmic options for using QBX within a variant of the Fast Multipole Method. A method is presented that preserves accuracy, generality and close evaluation capability while only requiring a relatively modest increase in computational cost in comparison to a point-to-point FMM. In addition, an optionally GPU-accelerated set of open-source software libraries is discussed that implements the proposed method.

Andreas Kloeckner
Department of Computer Science
University of Illinois at Urbana-Champaign
andreask@illinois.edu

MS227

Title Not Available at Time of Publication

Abstract not available at time of publication.

Denis Zorin
Computer Science Department
Courant Institute, New York University

dzorin@cs.nyu.edu

MS228

Compact-Reconstruction WENO on Non-uniform Meshes

Compact-Reconstruction WENO (CRWENO) schemes use solution-dependent combinations of compact schemes to produce a higher-order scheme that is non-oscillatory and has the superior spectral resolution of compact schemes. Previous work on these schemes has used uniform grids exclusively. In this talk we present a generalization of the fifth-order CRWENO scheme to non-uniform meshes in one space dimension, with some results concerning the effects of unequal spacing on the scheme's accuracy and spectral resolution.

Kilian Cooley
University of Maryland
kcooley2@math.umd.edu

James Baeder
Department of Aerospace Engineering
University of Maryland
baeder@umd.edu

MS228

Superconvergence Properties of Discontinuous Galerkin Methods Based on Upwind-Biased Fluxes for Linear Hyperbolic Equations

Superconvergence properties of discontinuous Galerkin (DG) methods using purely upwind fluxes under the assumptions of periodic boundary conditions and a uniform mesh for solving linear hyperbolic equations with smooth solutions have been studied through various approaches: pointwise a posteriori spacial discretization error estimates; spectral analyses based on the Fourier approach; and the smoothness increasing accuracy conserving (SIAC) filtered solution. We analyze via these approaches DG methods using upwind-biased fluxes, illustrating the discussion with numerical experiments.

Daniel Frean
University of East Anglia
d.frean@uea.ac.uk

MS228

A Compact-Reconstruction WENO Scheme with Semi-Implicit Time Integration

Weighted, nonlinear compact finite difference schemes are well suited for flows with a large range of length scales such as atmospheric flows. Semi-implicit time integration methods allow for efficient solutions by treating the stiff components of the hyperbolic flux implicitly. We propose a high-order finite-difference algorithm for atmospheric simulations based on the CRWENO scheme and implicit-explicit time stepping. The performance and scalability of the algorithm are evaluated for benchmark flow problems.

Debojyoti Ghosh
Mathematics and Computer Science Division
Argonne National Laboratory
ghosh@mcs.anl.gov

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS229

Parallel Preconditioning for Time-Dependent PDE-Constrained Optimization

All-at-once schemes aim to solve all time-steps of time-dependent PDE-constrained optimization problems in one coupled computation, leading to exceedingly large linear systems requiring efficient iterative methods. We present a new block diagonal preconditioner which is both optimal with respect to the mesh parameter and parallelizable over time, thus can provide significant speed-up. We will present numerical results to demonstrate the effectiveness of this preconditioner.

Eleanor McDonald
University of Oxford
eleanor.mcdonald@new.ox.ac.uk

Andy Wathen
Oxford University, UK
wathen@maths.ox.ac.uk

MS229

Preconditioning of Active-Set Newton Methods for PDE-Constrained Optimal Control Problems

We address the problem of preconditioning saddle point linear systems arising in the solution of PDE-constrained optimal control problems via active-set Newton methods, with control and (regularized) state constraints. We present two preconditioners based on a full block matrix factorization of the Schur complement of the Jacobians matrices where the active-set blocks are merged into the constraint blocks. The robustness of the new preconditioners is discussed and exhaustive numerical experiments are presented.

Margherita Porcelli
Dipartimento di Matematica
Università di Bologna
margherita.porcelli@unibo.it

Valeria Simoncini, Mattia Tani
Universita' di Bologna
valeria.simoncini@unibo.it, mattia.tani2@unibo.it

MS229

HPC Methods for Structured Inverse Modeling in Diffusive Processes

In this talk we present a method which combines high performance computing and large scale applications with an optimized, inverse shape identification procedure. We introduce a limited memory BFGS approach for optimizing the shape of the distribution of a jumping permeability in diffusive flow processes. These techniques are utilized to fit a model of the human skin to data measurements and not only estimate the permeability coefficients but also the shape of the cells.

Martin Siebenborn
University of Trier
siebenborn@uni-trier.de

Volker H. Schulz
University of Trier

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

MS229

Nonstandard Sobolev Spaces for Preconditioning Mixed Methods

In mixed methods for elliptic boundary value problems auxiliary variables are used to reformulate the involved differential equation as a system of differential equations of lower order. These methods can often be written as PDE-constrained optimization problems, whose analysis motivates the introduction of nonstandard Sobolev spaces for the auxiliary variables and leads to efficient solution techniques. The approach is exemplified for the Hellan-Herrmann-Johnson method applied to biharmonic boundary value problems.

Walter Zulehner
University of Linz, Austria
zulehner@numa.uni-linz.ac.at

Wolfgang Krendl
Johannes Kepler University Linz, Austria
wolfgang.krendl@dk-compmath.jku.at

MS230

Strengthening the Hurricane Wave and Surge Forecast Guidance provided to Coastal Communities in North Carolina

North Carolina is sensitive to hurricane waves, storm surge and flooding. A computational modeling system is utilized operationally to provide daily forecast guidance for coastal waves and inundation (http://nc-cera.renci.org/). This guidance has been expanded beyond Web-based delivery to include additional formats that are targeted to the needs of users within the state, are representative of the guidance at high levels of resolution, and are portable to widely-used geographic information systems.

Rosemary Cyriac
North Carolina State University
Dept. of Civil, Construction and Environmental Engineering
rcyriac@ncsu.edu

J. Casey Dietrich
University of Texas at Austin
jcdietrich@ncsu.edu

Jason Fleming
Seahorse Coastal Consulting
Morehead City, NC
jason.fleming@seahorsecoastal.com

Brian Blanton
University of North Carolina at Chapel Hill
Renaissance Computing Institute
bblanton@renci.org

Rick Luettich
University of North Carolina - Moorehead City
rick_luettich@unc.edu

MS230

Three-Dimensional Coupled Wind-Wave and Cohesive Sediment Transport Modeling in South San Francisco Bay

Abstract not available at time of publication.

Oliver Fringer
Environmental Fluid Mechanics Laboratory
Stanford University
fringer@stanford.edu

MS230

Discontinuous Galerkin Methods for Spectral Wave/Circulation Modeling

On large geographic scales, waves are represented in a spectral sense via the action balance equation. We present a computational spectral wave model developed using discontinuous Galerkin (DG) methods. DG methods allow for the use of unstructured meshes and adaptive, higher-order approximations, which we show leads to increased accuracy and efficiency. We loosely couple the new DG spectral wave model to the DG-Shallow Water Equation Model (DG-SWEM), an existing circulation model.

Jessica Meixner
University of Notre Dame
jdmeixner@gmail.com

MS230

Computational Modeling of Storm Surge in Galveston Bay

Abstract not available at time of publication.

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

MS231

Simulating Coupled Pressure-Temperature Equations for Trace Gas Sensors Using FEniCS and PETSc

Trace gas sensors are currently used in many applications from leak detection to national security and may some day help with disease diagnosis. These sensors are modelled by a coupled system of complex elliptic partial differential equations for pressure and temperature. Solutions are approximated using the finite element method which requires the development of custom block preconditioners. Finite element solutions are approximated using FEniCS and PETSc4py.

Brian W. Brennan
Baylor University
b_brennan@baylor.edu

MS231

Mesh-Independent Convergence for PDE-Constrained Optimisation Solvers in Dolfin-Adjoint

A key feature of dolfin-adjoint is its high-level framework for solving PDE-constrained optimisation problems using the finite element environment FEniCS. In this talk we discuss the mesh-independent convergence of the optimisation framework. This is achieved by formulating standard optimisation algorithms in a Hilbert space setting. We discuss

the implementation, show its effectiveness and compare it to preconditioning approaches. Thanks to the high-level problem formulation in FEniCS the resulting framework requires minimal user input, is fully parallel and scales to large problems.

Simon W. Funke, Magne Nordaas
Center for Biomedical Computing
Simula Research Laboratory
simon@simula.no, magneano@simula.no

MS231

Spectral/HP Element Modelling in Nektar++

In this talk we will highlight the high-level aspects of the Nektar++ high order finite element framework, which enables rapid development of high-performance parallel solvers. Using a straightforward linear PDE as an illustrative example, we will demonstrate how the library can be utilised at a variety of levels to fit the requirements of the end-user, and how simulation parameters such as the time-stepping scheme can be changed without detailed technical knowledge of the underlying methods.

David Moxey, Chris Cantwell, Spencer Sherwin
Imperial College London
d.moxey@imperial.ac.uk, c.cantwell@imperial.ac.uk,
s.sherwin@imperial.ac.uk

Mike Kirby
University of Utah
School of Computing
kirby@cs.utah.edu

MS231

Supporting Modern HPC Hardware in the DUNE Framework

Upcoming exa-scale computers will exhibit multiple levels of concurrency. Specialized implementations, exploiting all levels of parallelism, can gain a significant part of the peak performance. For high flexibility the C++ framework DUNE [Bastian et.al., 2008] defines generic interface, which allow various different implementations. The EXA-DUNE project follows this approach to enable all levels of concurrency for general DUNE applications. We present recent results and discuss implications on the interface design.

Christian Engwer, Fahlke Jorrit
INAM, University of Münster
christian.engwer@uni-muenster.de,
jorrit.fahlke@uni-muenster.de

Steffen Müthing
Heidelberg University
steffen.muething@iwr.uni-heidelberg.de

MS232

Transitioning from Game Design to Simulation Using Agent-Based Modeling

The objective of the Scalable Game Design project is to create a pipeline of people to work in computer-related fields by introducing them to game design. Hopefully, they will then transfer their interests and skills from game design to science simulation. Over the past five years several thousand students have participated in the project, with encouraging results. AgentSheets and AgentCubes are visual programming agent-based modeling applications that are ideally suited to this projects goals.

Fred Gluck
AgentSheets, Inc.
fred.gluck@comcast.net

MS232

Applying Run-Modify-Build Templates for Agent-Based Models

Students can grow in competence and confidence in modeling and simulation through well-scaffolded templates of working models. Instead of typing code or building structures, students can learn any given syntax in the context of real models, focusing on learning the science and improving the model. This is especially true for agent models that require more detailed descriptions of behavior. Examples of complex model templates across the computational sciences and for different tools will be shown.

Robert M. Panoff
President and Executive Director
Shodor
rpanoff@shodor.org

MS232

Teaching Freshman Science Using Agent-Based Computational Laboratories

With improved computational abilities and explosion of amounts of data, scientists routinely implement computation to test hypotheses and guide their research. We have developed a course that employs computational approaches to investigate scientific questions. Students explore science concepts, and using computational tools and algorithmic thinking, implement the scientific method to understand the natural world. Satisfying a science requirement, the course is designed so faculty from any science or computer science department can teach it.

George W. Shiflet
Wofford College
shifletgw@wofford.edu

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

MS232

NetLogo in the Secondary Life Science Classroom

In response to the adoption of new middle school science textbooks which incorporated NetLogo models as well as other visualization software, NetLogo models were modified or created for use in middle school life science and environmental lessons as well as in high school biology classrooms. Large group professional development, individual instructional coaching and classroom co-teaching were used to help teachers and students interact with the simulations successfully.

Charlotte M. Trout
WCPS-Retired

scitrout@icloud.com

MS233

Microstructure for Free Surface Flows

Numerical simulations of turbulent mixing in the Large Eddy Simulation regime are nonunique, with the selection of solution microstructure regulated by subgrid turbulence models. Validation, ie comparison to experiments, is then essential. We give examples of successfully validated Front Trecking simulations with extrapolation to flows outside of the experimental range. Appllcation to the selection of atomic vs. chunk mix for an inertial confinement fusion context will be given.

James G. Glimm
Stony Brook University
Brookhaven National Laboratory
glimm@ams.sunysb.edu or glimm@bnl.gov

MS233

Volume-Preserving Adaptive Moment-of-Fluid Method for Interface Tracking

We have developed a method to preserve the volume of a cell when the vertices of the cell move freely in a divergence-free velocity field. An optimization procedure is proposed to minimize the total variation of the node positions of the vertices after they are advanced according to the velocity field. A matrix-free Newton Krylov method is used to solve the nonlinear system generated by the optimization. The method has been applied to the multi-material moment-of-mluid interface tracking method with adaptive mesh refinement. Examples are provided to demonstrate the effectiveness of our numerical algorithm.

Shengtai Li
Los Alamos National Laboratory
sli@lanl.gov

Hyung T. Ahn
University of Ulsan
htahn@ulsan.ac.kr

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

MS233

A Fictitious Domain Method with a Hybrid Cell Model for Simulating Motion of Cells in Fluid Flow

In this work, we develop a hybrid model to represent membranes of biological cells and use the distributed-Lagrange multiplier/fictitious-domain (DLM/FD) formulation for simulating the fluid/cell interactions. The hybrid model representing the cellular structure consists of a continuum representation of the lipid bilayer, from which the bending force is calculated through energetic variational approach, a discrete cytoskeleton model utilizing the worm-like chain to represent network filament, and area/volume constraints. For our computational scheme, a formally second-order accurate fractional step scheme is employed to decouple the entire system into three sub-systems: a fluid problem, a solid problem and a Lagrange multiplier problem. Numerical results compare favorably with previously reported numerical and experimental results, and show that our method is suited to the simulation of the cell motion in flow. This work is in collaboration with W. Hao, C. Liu and G. Lin.

Zhiliang Xu
University of Notre Dame
zhiliangxu@nd.edu

MS233

Dissipation and Dispersion Errors of Discontinuous Galerkin Method and Its Application to Level Set Equations

The discontinuous Galerkin (DG) method is known to provide high resolution properties, especially when applying after long time run. In this talk, we consider analysing the error behaviour of the DG method for linear hyperbolic equations. Through Fourier analysis we observe, with P2 quadratic polynomial approximations, the dissipation error is on the order of 5 and the dispersion error is on the order of 6. The part of the error that grows linearly in time is on the order of 6. When solving interface problems in a complex incompressible flow, the DG method is shown to dramatically improve the mass conservation property of the level set method. Numerical examples demonstrate the high order accuracy of the scheme and the high resolution property especially when the interface undergoes large topological changes.

Jue Yan
Dept. of Math.
Iowa State University
jyan@iastate.edu

MS234

Massively Parallel GW Calculations for Current and Next-generation HPC

The traditional GW-Bethe-Salpeter (BSE) approach has, in practice, been prohibitively expensive on systems with large numbers of atoms. We show that through a combination of methodological and algorithmic improvements, the standard GW-BSE approach can be applied to systems with thousands of atoms. We will discuss the massively parallel GW-BSE implementation in the BerkeleyGW package (on-top of common DFT packages) including the importance of hybrid MPI-OpenMP parallelism, parallel IO and library performance. We will discuss optimization strategies for and performance on many-core architectures.

Jack Deslippe
National Energy Research Scientific Computing Center
jrdeslippe@lbl.gov

MS234

Recent Progress on Quantum Mechanics Embedding Theory

Abstract not available at time of publication.

Chen Huang
Theoretical Division, T-1
Los Alamos National Laboratory
chuang3@fsu.edu

MS234

Towards Predictive Modeling of Correlation Effects

in Many-electron Systems

In this presentation we will discuss the development of scalable and unique computational capabilities for modeling quasi-degenerate systems using implementations of multi-reference coupled-cluster (MRCC) methods in NWChem. We will discuss novel parallel algorithms for several MRCC methodologies, which are capable of taking advantage of existing peta-scale architectures. In this context, the emergence of the heterogeneous computer architectures offers a unique chance to advance accurate yet expensive MRCC formalisms. We will also outline the recent development of coupled-cluster Green function formalism geared towards efficient inclusion of higher-order correlation effects.

Karol Kowalski
Pacific Northwest National Laboratory
karol.kowalski@pnnl.gov

MS234

A Parallel Orbital-Updating Approach for Electronic Structure Calculations Based on Singularity Decompositions

In this presentation, we will introduce an orbital-based parallelization algorithm for electronic structure calculations and demonstrate the efficiency of our algorithm by numerical experiments. This algorithm is based on our understanding of the single-particle approximation equations of independent particles that move in an effective potential and a singularity decomposition of the effective potential. With this algorithm, the solution to the single-particle equation can be reduced to some solutions of several independent linear algebraic systems and a small scale algebraic eigenvalue problem. This presentation is based on some joint works with Xiaoying Dai, Xingao Gong, and Xin Zhang.

Aihui Zhou
Chinese Academy of Sciences
azhou@lsec.cc.ac.cn

MS235

Algorithmic Adaptations for Scalable Community Detection on the Tilera Many-Core Architecture

As power constraints and data movement costs become significant barriers for high-end computing, the Tilera many-core architecture offers a low-power platform exhibiting key characteristics of future systems: a large number of simple cores, a sophisticated network-on-chip, and fine-grained control over memory locality. We implemented a graph community detection application using platform-aware memory layouts and scheduling techniques resulting in speedups of up to 46x on the TileGX36 and comparable result quality and performance to x86 platforms.

Daniel Chavarria, Howard Lu, Mahantesh Halappanavar
Pacific Northwest National Laboratory
daniel.chavarria@pnl.gov, luhowardmark@hotmail.com, hala@pnnl.gov

Ananth Kalyanaraman
Associate Professor
School of EECS, Washington State University
ananth@eecs.wsu.edu

MS235

Using Performance Tools to Assist Porting to New Platforms

Performance tool play a critical role in porting applications to new platforms. They not only enable identifying hot spots; the first targets for acceleration, but can also show the context of these hot spots. Tracing and profiling tools can be used in combination to provide different resolution of detail and amounts of runtime perturbation. Tools like Score-P can capture all parallel activity (MPI/OpenMP/pthreads/CUDA/OpenCL/OpenACC) and display it to the user using analysis engines like Vampir.

Guido Juckeland
TU-Dresden
guido.juckeland@tu-dresden.de

MS235

Algorithmic Selection, Autotuning, and Scheduling for Accelerator-Based Codes for Numerical Linear Algebra

Often, a rewrite of the scientific software is unsustainable in the long run as the complexity increases to respond to the more parallel and heterogeneous hardware. In this talk, I will give practical perspective on how algorithmic selection, software autotuning at installation time, and various forms of dynamic runtime scheduling for multicore systems with accelerators can alleviate the aforementioned problems in the context of widely used numerical linear algebra libraries.

Piotr Luszczek
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
luszczek@eecs.utk.edu

MS235

Toward Heterogeneous Memory Systems for HPC

Compute nodes equipped with a variety of memory technologies such as scratchpad memory, on-chip 3D-stacked memory, or NVRAM-based memory, apart from traditional DRAM, are already a reality. Careful use of the different memory subsystems is mandatory in order to exploit the potential of such supercomputers. I will present our view on upcoming heterogeneous memory systems, which comprises exposing the different memory subsystems as first-class citizens to efficiently exploit their capabilities.

Antonio J. Peña
Argonne National Lab
apenya@mcs.anl.gov

MS236

Reduced Order Models for Patient-Specific Haemodynamics of Coronary Artery Bypass Grafts

In this talk we present a POD-Galerkin reduced order model to simulate the haemodynamics of coronary artery bypass grafts. Clinically relevant physical (inlet flow rates) and geometrical (stenoses severity, anastomoses geometry)

quantities are considered in a parametrized setting for unsteady incompressible Navier-Stokes equations. The proposed ROM is applied to some patient-specific cases. A presentation of a comparison, with respect to such parameters, and a clinical discussion of the resulting flow patterns follows.

Francesco Ballarin
Politecnico di Milano, Italy
francesco.ballarin@polimi.it

Elena Faggiano
Politecnico di Milano
Italy
elena.faggiano@polimi.it

Sonia Ippolito
Luigi Sacco Hospital
Milan, Italy
ippolito.sonia@hsacco.it

Andrea Manzoni
EPFL, MATHICSE-CMCS
Switzerland
andrea.manzoni@epfl.ch

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

Gianluigi Rozza
SISSA, International School for Advanced Studies
Trieste, Italy
gianluigi.rozza@sissa.it

Roberto Scrofani
Luigi Sacco Hospital
Milan, Italy
scrofani.roberto@hsacco.it

MS236

pyMOR - A New Model Order Reduction Software Framework

Over the past years, projection-based model order reduction (MOR) techniques, such as the reduced basis method (RB) ([1] and references therein), have become well-established tools for efficient numerical solution of partial differential equations. However despite the fact that RB-methods are very generic in nature, most software solutions still provide only ad-hoc implementations which are tailor-made for specific application problems and are tied to a given software ecosystem. In this talk we present pyMOR [2], a new open source library for building MOR-applications with the Python programming language. pyMOR's main focus lies on the easy application of RB-methods to parameterized partial differential equations solved by external high-dimensional discretization packages. We give a brief overview of pyMOR's main components and highlight the design philosophy and choices that allow us to arrive at a set of completely generic reduction algorithms which can be readily incorporated with any external solver connected to pyMOR via its lightweight operator and vector interfaces. As a result of this abstraction, well-proven MOR-algorithms can be easily applied to new real-world problems. At the same time, different reduction algorithms can be compared for a given problem without modification of the high-dimensional solver. We conclude our presentation by illustrating these benefits by the example of our current work on the reduction of microscale lithium-ion battery models. [1] B. Haasdonk, M. Ohlberger, Reduced basis method for finite volume approximations of parametrized linear evolution equations, M2AN Math. Model. Numer. Anal., 42 (2008), 277–302. [2] R. Milk, S. Rave, F. Schindler, pyMOR - Model Order Reduction with Python, http://pymor.org.

Rene Milk
University of Muenster
rene.milk@uni-muenster.de

Mario Ohlberger
Universität Münster
Institut für Numerische und Angewandte Mathematik
mario.ohlberger@uni-muenster.de

Stephan Rave
University of Muenster
stephan.rave@wwu.de

Felix Schindler
WWU
felix.schindler@wwu.de

MS236

Reduced Basis Methods for Option Pricing

We present a reduced basis method for pricing European and American options based on the Black-Scholes and Heston model. To tackle each model numerically, we formulate the problem in terms of a time dependent variational equality or inequality and propose a method that combines POD-Greedy and Angle-Greedy procedures for the construction of the reduced spaces. In addition, we obtain a posteriori error estimators. Numerical examples illustrate the efficiency of our approach.

Julien Salomon
CEREMADE, Universite Paris-Dauphine
julien.salomon@dauphine.fr

Olena Burkovska
Technical University of Munich
burkovsk@ma.tum.de

Bernard Haasdonk
University of Stuttgart
haasdonk@mathematik.uni-stuttgart.de

Barbara Wohlmuth
Technical University of Munich
barbara.wohlmuth@ma.tum.de

MS236

Adaptivity and Reduced Basis Methods

Usually, the computations in a Reduced Basis Method (RBM) are split into an offline and an online phase. In the offline phase, expensive computations are used to construct a Reduced Basis in terms of snapshots of a detailed model. This has a number of consequences that are also possible criticisms on the RBM itself: (1) The snapshots need to based upon the same discretization (at least for parts of the parameter range); (2) If the problem itself is challenging, offline costs may be significant; (3) The separation into offline and online phase hides the overall complexity of the problem. At an extreme, one could pre-compute

'everything' offline and just pick the right data online. On the other hand, several adaptive methods are known and proven to be asymptotically optimal for a wide class of operator equations. Asymptotically optimal means that the rate of convergence is comparable to a best-possible benchmark and the computational cost is linear in the required number of degrees of freedom. This is the motivation to consider the combination of adaptive numerical methods and the RBM. We show some consequences of using adaptive computations both for the snapshot generation as well as the evaluation of standard RB error estimates. The talk is based upon joint work with Kristina Steih (Ulm).

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

MS237

An Accelerated Domain Decomposition Method for Time Dependent Problems

We consider the time integration of the system of ODEs which results from the semi-discretization of PDEs. The solution components are assumed to evolve on different time scales and hence some sort of multi-rate time integration method which allows different time steps for different components may be suitable. Here we consider two algorithms for such problems: a multi-rate Schwarz Waveform relaxation hybrid algorithm and a multi-rate accelerated Schwarz Waveform relaxation method.

Ronald Haynes
Memorial University of Newfoundland
St. John's, NL, Canada
rhaynes74@gmail.com

MS237

Developing a Custom Time Integrator for the Nonlinear Schrödinger Equation for An Application in Paraxial Laser Propagation

I will report on the construction of a temporal integration scheme designed to be used with a higher-order finite-difference spatial approximation to the nonlinear Schrödinger equation with perfectly matched layer boundary conditions. The temporal method is based on a semi-implicit spectral deferred corrections approach coupled with an alternating direction implicit (ADI) approximation to the (linear) implicit terms. The accuracy and efficiency of methods up to tenth-order in both time and space will be presented as well as some preliminary results on space-time parallelization of the method.

Michael Minion
Stanford University
mlminion@lbl.gov

MS237

Spatially Partitioned Embedded Runge-Kutta Methods

We study spatially partitioned embedded Runge-Kutta schemes for partial differential equations (PDEs), in which each of the component schemes is applied over a different part of the spatial domain. Such methods may be convenient for problems in which the smoothness of the solution or the magnitudes of the PDE coefficients vary strongly in space. We focus on embedded partitioned methods as they offer greater efficiency and avoid the order reduction that may occur in non-embedded schemes. We demonstrate that the lack of conservation in partitioned schemes can lead to non physical effects and propose conservative additive schemes based on partitioning the fluxes rather than the ordinary differential equations. A variety of schemes are presented, including an embedded pair suitable for the time evolution of fifth-order weighted non-oscillatory (WENO) spatial discretizations. Numerical experiments are provided.

Steven Ruuth
Mathematics
Simon Fraser University
sruuth@sfu.ca

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Colin B. Macdonald
Oxford University
macdonald@maths.ox.ac.uk

MS237

A Massively Parallel Solver for the Incompressible Navier–Stokes Equations

In 2011, Guermond and Minev proposed a massively parallelizable directional splitting method for the time discretization of the incompressible Navier–Stokes equations. It was shown to scale well in a distributed memory environment of up to 1024 CPUs. We modified this method to take advantage of the high computational throughput afforded by (multiple) GPUs as well as the efficiency of variable time steps. We demonstrate that the modified method exhibits substantial performance gains over the original.

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

MS238

Coherence Motivated Sampling of Polynomial Chaos Expansions

We investigate solution recovery, particularly regarding methods for sampling basis functions associated with Hermite and Legendre Polynomial Chaos Expansions. Recent results have identified interesting quantities related to the basis functions and their sampling which are computable, and analytically fruitful. Through these parameters we analyze asymptotically motivated samplings, and construct sampling methods with statistical optimality. We demonstrate these methods on examples motivated from Uncertainty Quantification.

Jerrad Hampton
University of Colorado, Boulder
jerrad.hampton@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder

Alireza.Doostan@Colorado.EDU

MS238

Sparse Solutions to Large-Scale Nonlinear Subsurface Flow Inverse Problems

We present sparse inversion formulations for nonlinear subsurface flow model calibration problems. By promoting sparsity, the inversion is transformed into a feature selection problem that is flexible and robust against prior uncertainty. Specifically, in addition to effective low-rank representation, the formulation allows for discriminating against alternative plausible prior geologic scenarios based on nonlinear multiphase flow data from scattered well locations. Using several case studies, we demonstrate the advantages of sparse nonlinear inversion for geoscience application.

Benham Jafarpour
Viterbi School of Engineering
University of Southern California
behnam.jafarpour@usc.edu

Reza Khaninezhad
Electrical Engineering
University of Southern California
m.khaninezhad@usc.edu

MS238

An Efficient Method for the Computation of the Stochastic Galerkin Projections by Means of Tensor Format Representations

The stationary diffusion problem with an uncertain permeability coefficient is solvedin a low-rank tensor representation without changes of the available deterministic code. The presented uncertainty is modeled by a random field, which is discretized by KLE as well as by PCE. It was shown in previous works [Espig et al 2012, 2013] under which conditions the stochastic Galerkin operator can be approximated/represented in a low-rank tensor format. The obtained symmetrical and positive defined linear system is solved by the new iterative method. The iteration matrix as well as the whole computations are done also in that low-rank tensor format. The choice of the tensor format is not crucial here. The computational cost and storage (independent on the tested tensor formats) are linear in the stochastic dimension.

Alexander Litvinenko
SRI-UQ Center, King Abdullah University of Science and Techn
alexander.litvinenko@kaust.edu.sa

Mike Espig
Max Planck Institute for Mathematics in the Sciences
mike.espig@alopax.de

Matthies HERMANN G.
Instiute of Scientific Computing
TU Braunschweig
h.matthies@tu-bs.de

MS238

Tensor Approximation Methods for Stochastic Problems

Spectral stochastic methods have gained wide acceptance as a tool for efficient modelling of uncertain stochastic systems. The advantage of those methods is that they provide not only statistics, but give a direct representation of the measure of the solution as a so-called surrogate model, which can be used for very fast sampling. Especially attractive for elliptic stochastic partial differential equations is the stochastic Galerkin method, since it preserves essential properties of the differential operator. One drawback of the method is, however, that it requires huge amounts of memory, as the solution is represented in a tensor product space of spatial and stochastic basis functions. Different approaches have been investigated to reduce the memory requirements, for example, model reduction techniques using subspace iterations to reduce the approximation space or methods of approximating the solution from successive rank-1 updates. See more in my dissertation [1]. [1] Elmar K. Zander, Tensor Approximation Methods for Stochastic Problems, Dissertation, TU Braunschweig, 2013.

Elmar Zander
TU Braunschweig
Inst. of Scientific Computing
e.zander@tu-bs.de

MS239

Novel Priors and Algorithms for 4d Tracking and Classification of Cells

In biomedical imaging efficient reconstruction and tracking methods play a fundamental role. Particularly in cell biology and medical tomography innovative spatio-temporal imaging models are of strongly growing interest. The aim of this talk is to highlight novel priors and related convex optimization for reconstruction and flow quantification in 4D image sequences. Motivated by the success of L1 and higher-order (e.g. TGV) regularization for static imaging we will focus on spatio-temporal priors for the optimal transportation of cells.

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

MS239

Parameter Estimation for Malignant Brain Tumors

Determining patient specific tumor parameters is important in choosing optimal treatment plans. We present an inverse problem formulation to approximate the parameters of interest using PDE-constrained optimization algorithms. The choice of the optimization method is important as solving this problem has a high computational cost. We will present a reduced space formulation, novel preconditioners to speed up the solution process, and numerical experiments on a synthetic dataset for preliminary evaluation of the method.

Amir Gholaminejad
PhD Student, Institute for Computational Engineering and Sciences, University of Texas at Austin
i.amirgh@gmail.com

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS239

Incorporating Uncertainty in MR Images of

Glioblastoma when Leveraging Models to Interpret Therapeutic Efficacy

There is a lack of response metrics that can quickly and accurately predict the therapeutic efficacy for glioblastoma. Recently, a mathematical model of glioma growth has been used to define a response metric based on patient-specific parameters of proliferation and invasion calibrated from pre-treatment imaging which was shown to be prognostic for overall survival. This talk will focus on how data uncertainty affects the calibrated parameters, the response metric, and the implications for clinical translation.

Andrea Hawkins-Daarud
Department of Neurological Surgery
Northwestern University
andrea.hawkins-daarud@northwestern.edu

Kristin R. Swanson
Northwestern University
kristin.swanson@northwestern.edu

MS239

Platform-independent Description of Image Registration Algorithms

In the last years real-time imaging and more complex models for the various imaging applications like image registration became feasible mostly because of the progress made in parallel computer architectures like found in GPUs and modern multi-core CPUs. However, the implementation effort increases, if one wants to achieve good performance. A solution to this problem is to formulate the image registration algorithms in an abstract way in a domain-specific language (DSL) and then automatically generate efficient C++ or CUDA code. A multi-layered approach is sketched that allows users to describe applications in a natural way from the mathematical model down to the program specification.

Harald Koestler
University of Erlangen-Nuremberg
harald.koestler@informatik.uni-erlangen.de

MS240

Computational Complexity of Stochastic Galerkin and Collocation Methods for PDEs with Random Coefficients

We will present a rigorous cost metric, used to compare the computational complexity of a general class of stochastic Galerkin methods and stochastic collocation methods, when solving high-dimensional stochastic PDEs. Our approach allows us to calculate the cost of preconditioning both the Galerkin and collocation systems, as well as account for the sparsity of the Galerkin projection. Theoretical complexity estimates will also be presented and validated with use of several computational examples.

Nick Dexter
University of Tennessee
ndexter@utk.edu

Clayton G. Webster, Guannan Zhang
Oak Ridge National Laboratory
webstercg@ornl.gov, zhangg@ornl.gov

MS240

QoI Basis Adaptation

We demonstrate the efficiency and scope of the basis adaptation methodology developed for Gaussian polynomial chaos. The approach is based on adapting the basis in L_2 by applying an isometry to its Gaussian base. We discuss error estimates associated with this approach, and the the concept of an optimal isometry. We also present implementation details and applications to large scale problems from across science and engineering. We also explore the importance of specific quantities of interest to the adaptation process, and extend the approach to situations defined by multiple quantities of interest.

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

Kenny Chowdhary
Sandia National Laboratories
kennychowdhary@gmail.com

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

MS240

Numerical Methods for SPDEs with Levy Jump Processes SPDES: Stochastic and Deterministic Approaches

We present numerical results on SPDEs with multiple Levy jump processes of various depen- dence structures by gPC (stochastic approach) and by joint PDF from the generalized Fokker- Planck equations for spatial modes to decompose the SPDEs into SODEs (deterministic ap- proach). As examples, we solve SPDEs with TaS jump processes, with tempered fractional PDEs as joint density functions. We will demonstrate that we largely improve the accuracy and efficiency of the traditional MC method!

Mengdi Zheng
applied math, brown university
mengdi_zheng@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS240

Optimal Least-Squares Projection: Applications to Uq

In this talk, we discuss the problem of approximating a multivariate function by discrete least-squares projection onto a general polynomial space, using a random chosen sub-grid of the tensor grid of Gaussian points. The independent variables of the function are assumed to be random variables, and thus, the framework provides a non-intrusive way to construct the generalized polynomial chaos expansions, stemming from the motivating application of Un-

certainty Quantification (UQ). We prove the stability and an optimal convergence estimate, provided the number of points scales linearly (up to a logarithmic factor) with the dimension of the polynomial space. The framework includes both the bounded measures such as the uniform and the Chebyshev measure, and the unbounded measures which include the Gaussian measure. Several numerical examples are given to confirm the theoretical results.

Tao Zhou
Institute of Computational Mathematics, AMSS
Chinese Academy of Sciences
tzhou@lsec.cc.ac.cn

Akil Narayan
University of Massachusetts Dartmouth
akil.narayan@umassd.edu

Dongbin Xiu
University of Utah
dongbin.xiu@utah.edu

MS241

N-Body Algorithms for Matrices with Decay: Multiplication, Projection, Inverse Factorization & Fock-Exchange

We report on several fast algorithms for operations on dense matrices with decay. Based on the SpAMM algorithm [arXiv:1203.1692 and 1403.7458] for matrix-matrix multiplication, we develop methods for spectral projection and inverse factorization with reduced, O(N) complexity as well as error control in accordance with estimates. We further demonstrate the generalization of these methods to the higher dimensional problem of Fock-exchange, involving multi-level multipole approximation with nested SpAMM contraction.

Matt Challacombe
Los Alamos National Laboratory
Theoretical Division
MChalla@lanl.gov

Nicolas Bock
Los Alamos National Laboratory
nicolasbock@freeon.org

Terry Haut
University of Colorado, Boulder
terry.haut@colorado.edu

MS241

Linear-Cost Storage and Computation with Kernel Matrices

Kernel matrices embrace a rich structure that enables more efficient storage and computation than does a usual dense matrix. The FMM and tree code methods are prominent examples exploiting such a structure: far-field low-rank interaction. In this talk, we lay the fundamental O(n) data structure in the context of linear algebra and present O(n) algorithms for performing matrix operations, including matrix-vector multiplication, matrix inversion, and determinant calculation. We demonstrate numerical stability, linear scaling, and applications in statistical data analysis.

Jie Chen
Argonne National Laboratory
jiechen@mcs.anl.gov

MS241

Parallel Randomized Structured Multifrontal Method for Sparse Direct Solutions

We present a fast parallel direct solver for general sparse matrices based on a fully structured randomized multifrontal method. Distributed memory is used. A sequence of innovative parallel strategies are designed for randomized compression, hierarchical structures, and structured multifrontal solution. Multilevel parallelism is built upon two hierarchical tree layers. This gives a significant new direction for designing fast scalable sparse direct solvers.

Zixing Xin, Jianlin Xia
Department of Mathematics
Purdue University
zxin@purdue.edu, xiaj@math.purdue.edu

MS242

Fluctuating Hydrodynamics Methods for Electrokinetics and Capillary Electrophoresis of Charged Colloids

Motivated by problems in the design and operation of recent microfluidic and nanofluidic devices, we present fluctuating hydrodynamic methods for confined electrokinetics. We investigate the capillary electrophoresis of charged colloids. We show that the colloid counter-ion layers exhibit deviations from Poisson-Boltzmann theory and interesting phenomena arising from discrete ion effects, hydrodynamic interactions, and overlap with ion layers near the channel walls.

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

MS242

Mesoscale Models for Molecular Solvation: Funny Business at the Solute-Solvent Interface

Electrostatic interactions between biomolecules (proteins) and the surrounding aqueous environment are crucial for correct function. Continuum electrostatic models are significantly faster than atomistic molecular-dynamics simulations, but trade off realism: e.g., at protein surfaces, solvent is not a bulk material. Although multiscale models resembling gradient-elasticity theories improve accuracy somewhat, non-standard boundary conditions capture relevant interfacial-solvent effects much better. These findings motivate the development of rigorous methods for obtaining boundary conditions between multiscale materials.

Jaydeep Bardhan
Northeastern University
j.bardhan@neu.edu

Matthew G. Knepley
University of Chicago
knepley@ci.uchicago.edu

MS242

Mesoscopic Modeling of Temperature-dependent Properties in Non-isothermal Fluid Systems

We develop an energy conserved dissipative particle

dynamics model to capture the correct temperature-dependent properties of fluids. We demonstrate that the proposed model can predict correctly the temperature dependence of the diffusivity, viscosity and thermal conductivity of liquid water, which is consistent with available experimental data of water at various temperatures. Subsequently, this mesoscopic model is applied to investigations of the nonisothermal hydrodynamic flows and the phase transition of thermoresponsive polymers.

Zhen Li
Brown University
Brown University
Zhen_Li@brown.edu

Yu-Hang Tang, Bruce Caswell
Brown University
yuhang_tang@brown.edu, bruce_caswell@brown.edu

George E. Karniadakis
Division of Applied Mathematics
Brown University
george_karniadakis@brown.edu

MS243

Ensemble Kalman Filters Without Tuning for Large Applications

Ensemble Kalman filters are widely used for large geophysical data assimilation problems but generally require significant tuning of inflation and localization to produce good results in the presence of sampling error from small ensembles. Methods to reduce the tuning required by estimating and correcting for errors in ensemble sample covariances are presented. The feasibility of treating ensemble filters as a black box for general applications is discussed.

Jeffrey Anderson
National Center for Atmospheric Research
Institute for Math Applied to Geophysics
jla@ucar.edu

MS243

Orthogonal Transformations for the Ensemble Kalman Filter

Data assimilation of complex physical systems is commonplace in many areas of science in order to predict phenomena of interest. In this talk we use orthogonal transformations in the ensemble Kalman filter to capture the dominant correlations within a model and reduce noise. A series of assimilation experiments are presented for a ionosphere-thermosphere model using real observational data. The results show a reduction in forecast error using orthogonal transformation compared with assimilation without transformation.

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

Earl Lawrence
Los Alamos National Laboratory
earl@lanl.gov

Dave Higdon
Los Alamos National Laboratory
Statistical Sciences
dhigdon@lanl.gov

MS243

Multi-Scale Data Assimilation for Fine-Resolution Models

The commonly used data assimilation algorithms are based on optimal estimation theory, in which error covariance is of fundamental importance. It is shown that the standard optimal estimation algorithm is inherently ineffective when it is applied to fine resolution models, and the ineffectiveness arises from its filtering nature. We propose a multi-scale data assimilation algorithm, in which the cost function is decomposed for a set of distinct spatial scales. Data assimilation is implemented sequentially from large to fine scales. Results are presented to demonstrate the algorithm.

Zhijin Li
Jet Propulsion Laboratory
California Institute of Technology
zhijin@jpl.nasa.gov

MS243

Unified Ensemble-Variational Data Assimilation System

In this presentation we present a development of a unified ensemble-variational data assimilation system that combines best properties of ensemble and variational DA methods, as well as filter and smoother methodologies. The main idea is to create a practical, yet theoretically advanced system for general data assimilation applications. The new system has a complete feedback in terms of uncertainties and optimal states. Recent progress and results will be presented.

Milija Zupanski
Cooperative Institute for Research in the Atmosphere
Colorado State University
Milija.Zupanski@colostate.edu

MS244

An Adaptive Augmented Lagrangian Method for Large-Scale Constrained Optimization

We present augmented Lagrangian (AL) methods for solving nonlinear optimization problems. AL methods are advantageous as they can be implemented matrix-free, and so are scalable for solving large scale problems. However, they often suffer from poor initial choices of the penalty parameter and Lagrange multipliers. Our methods overcome this disadvantage by incorporating dynamic updates for the penalty parameter while minimizing the AL function. Numerical experiments with LANCELOT software illustrate the benefits of our proposed methods.

Frank E. Curtis
Industrial and Systems Engineering
Lehigh University
frank.e.curtis@gmail.com

Nicholas I.M. Gould
Rutherford Appleton Laboratory
UK
nick.gould@stfc.ac.uk

Hao Jiang

Department of Applied Mathematics and Statistics
Johns Hopkins University
hjiang13@jhu.edu

Daniel Robinson
Northwestern University
daniel.p.robinson@gmail.com

MS244

A Data-Driven Approach to PDE-Constrained Optimization Under Uncertainty

I present an approach for incorporating data in PDE-constrained optimization. First, I develop a data-driven discretization for probability measures of uncertain PDE parameters and prove rigorous error bounds. I then formulate a robust optimization problem that accounts for the uncertainty in the estimated probability measures. Finally, I propose an algorithm to solve the resulting semidiscretized minimax problem. This algorithm employs trust regions and permits inexact derivative computations. I conclude with preliminary numerical results.

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

MS244

PDE-Constrained Optimization Under Uncertainty for Convection-Diffusion-Reaction Systems

We present an overview of algorithms for large-scale optimization of partial differential equations (PDEs) with uncertain coefficients. Our algorithms minimize risk-based objective functions using sparse-grid discretizations. We demonstrate our approach on large scale source inversion for multi-spieces convection-diffusion-reaction dynamics.

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

Harriet Li
MIT
kameeko@mit.edu

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

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

MS244

Inexact Primal-Dual Interior Point Filter Method

Abstract not available at time of publication.

Victor Zavala
Argonne National Laboratory
vzavala@mcs.anl.gov

MS245

A Finite-Volume Based Formulation for Viscoelastic Two-Phase Flows

We present computational schemes for the study of viscoelastic two-phase flows. A new formulation that transforms the constitutive equation for viscoelastic stress into a conservative form is developed and combined with a finite-volume discretization of the Navier-Stokes equations. The numerical method is based on a volume of fluid algorithm for tracking the interface. Numerical examples are presented to demonstrate the higher order accuracy of developed schemes.

Shahriar Afkhami
Department of Mathematical Sciences
New Jersey Institute of Technology
shahriar.afkhami@njit.edu

MS245

A Moment-of-Fluid Method for Computing Solutions to Multi-Phase Flows

The Moment-of-Fluid (MOF) reconstruction algorithm is volume preserving, generalizable to arbitrary number of materials, and uses only information local to the computational cell. These properties enable one to develop a multiphase flow algorithm that accurately computes compressible or incompressible multiphase flows consisting of any number of materials without the aid of Riemann solvers and no acoustic time step constraint required. An outline of the MOF multiphase flow algorithm will be explained, and example multiphase flow simulations with applications in energy and medicine will be shown.

Mark Sussman
Department of Mathematics
Florida State University
sussman@math.fsu.edu

MS245

A Time Splitting Projection Scheme for Compressible Two-Phase Flows : Application to the Interaction of Bubbles and Droplets with Ultrasound Waves

We will present a time-splitting scheme for pressure and velocity in order to account with surface tension effects in the framework of projection method for compressible two-phase flows. Several benchmarks, based on the Rayleigh-Plesset theory (volume oscillation of a bubble), will be presented in order to demonstrate the efficiency of this new time-splitting scheme in comparisons with others existing methods (HLLC solver, Turkel preconditionning) in situations involving a low Mach number.

Sebastien Tanguy
Universite de Toulouse
sebastien.tanguy@imft.fr

MS245

Fourth-Order Interface Tracking and Curvature Estimation for An Arbitrary Number of Materials in Two Dimensions

We present the iPAM method as the first fourth-order interface tracking method in two dimensions, prove its convergence rates by mapping and adjusting regular semi-algebraic sets, and propose fourth-order algorithms for estimating the unit normal and the curvature. Exploiting algorithms and theories in computational geometry, general

topology, and algebraic geometry, iPAM directly applies to both structured and unstructured grids, both incompressible and compressible flows, and an arbitrary number of phases without any algorithmic modifications.

Qinghai Zhang, Aaron L. Fogelson
University of Utah
qinghai@math.utah.edu, fogelson@math.utah.edu

MS246

Large-Scale 3D Electromagnetic Imaging Using Julia

In this talk we discuss the implementation of large scale electromagnetic inverse problems with multiple sources and frequency using Julia. We discuss how asynchronous optimization and regularization techniques can be implemented in an efficient way and propose a paradigm for general data rich inverse problems that invoke with nontrivial simulations.

Eldad Haber
Department of Mathematics
The University of British Columbia
haber@math.ubc.ca

MS246

JuMP: Algebraic Modeling of Optimization Problems in Julia

We present JuMP, an open-source algebraic modeling language for mathematical optimization (cf. AMPL, GAMS, CVX, YALMIP). We describe how JuMP takes advantage of Julia's advanced technical features like just-in-time compilation and metaprogramming to achieve competitive performance with commercial tools, including a reimagined implementation of Automatic Differentiation (AD) techniques for computing exact Hessian matrices for nonlinear optimization. JuMP supports a number of commercial and open-source solvers for linear, mixed-integer, quadratic, conic, and nonlinear optimization.

Miles Lubin
Operations Research Center
Massachusetts Institute of Technology
mlubin@mit.edu

Iain Dunning
Operations Research Center
Massachusetts Insitute of Technology
idunning@mit.edu

Joey Huchette
Operations Research Center
Massachusetts Institute of Technology
huchette@mit.edu

MS246

Distributed and Parallel Computing for Pde Constrained Optimization in Julia

This talk presents a Julia framework for the solution of large-scale PDE constrained optimization problems. It is based on a discretize-then-optimize approach, uses a (projected) Gauss-Newton method, and provides interfaces state-of-the-art linear solvers (both explicit and iterative). The framework uses Julia's potential for parallel and distributed computation. Being written in a dynamic language, it is easily extendable and yet fast as will be outlined for a large electromagnetic test problem.

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

MS246

An Extensible Test Matrix Collection for Julia

Matrix Depot is an open source project aiming to provide a rich collection of test matrices. This software can be easily extended and therefore facilitate exchange of matrices between researchers in numerical linear algebra community. We will describe the design and implementation of this software. Matrix Depot supports a variety of languages, including Fortran and Python, but in this talk, we will focus on how to use Matrix Depot in Julia.

Weijian Zhang, Nicholas Higham
School of Mathematics
The University of Manchester
weijian.zhang@postgrad.manchester.ac.uk,
nick.higham@manchester.ac.uk

MS247

Meet Informally with the CSE15 Co-Chairs and Several Invited Speakers

Meet Informally with the CSE15 Co-Chairs and Several Invited Speakers.

Hans De Sterck
University of Waterloo
Applied Mathematics
hdesterck@uwaterloo.ca

Christopher Johnson
University of Utah
Department of Computer Science
crj@sci.utah.edu

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

MS248

Approximation of the Boltzmann Equation with the Method of Moments for Low Speed Gas Flow

Gas flows in micro-electro-mechanical systems (MEMS) are usually at a low speed and in the transition regime. The Navier-Stokes-Fourier equations are no longer adequate to capture the non-equilibrium phenomena in MEMS and the Boltzmann equation is necessary to describe the flow field correctly. However, the full solution of the Boltzmann equation is complicated and expensive. Approximation of the Boltzmann equation with the moment method is introduced and the accuracy and validity range of the method will be discussed.

Xiaojun Gu
Scientific Computing Department
STFC Daresbury Laboratory
xiaojun.gu@stfc.ac.uk

David Emerson, Jianping Meng
Scientific Computing Department STFC Daresbury Laboratory
david.emerson@stfc.ac.uk, jianping.meng@stfc.ac.uk

MS248

A Framework on Moment Model Reduction for Kinetic Equation

Model reduction of kinetic equation turns a high dimensional problem to a low dimensional quasi-linear system, which not only provides further understanding of the problem, but also essentially improves the efficiency of the numerical simulation. As a quasi-linear system with Cauchy data, the well-posedness of the model deduced is required to be hyperbolic. In the existed models, some of them are hyperbolic, and some of them may be regularized to be hyperbolic, while there are seldom progress on the else models. Studies indicate that the hyperbolicity is possibly related with H-theorem, preserving of certain physics, etc. In this talk, I will reveal the underlying reason why some existed models are hyperbolic and the others are not, and how the hyperbolic regularization works. It is pointed out that the hyperbolicity is not related to H-theorem, conservation, etc, at all. The even fascinating point is, with only routine calculation, symmetric hyperbolic models can always be deduced with any ansatz for generic kinetic equation by the framework we proposed. By this framework, existing models are re-presented and brand new models are discovered. Even if the study is restricted in the scope of the classical Grad's 13-moment system, a new model with global hyperbolicity can be deduced.

Ruo Li
School of Mathematical Science
Peking University
rli@math.pku.edu.cn

MS248

Numerical Solution of a Fourteen-Moment Closure for Non-equilibrium Gases

Recently, a hyperbolic fourteen-moment closure for the Boltzmann equation from gaskinetic theory has been proposed that is based on an interpolation technique between local equilibrium and realizability boundaries using entropy maximization as a guide. This closure has proven to be very accurate for moderately non-equilibrium flows, however a singularity in the closing flux presents challenges to numerical solution. This talk discusses the handling of these challenges and presents solutions to canonical flow problems.

James McDonald, Amir Baradaran
Department of Mechanical Engineering
University of Ottawa
james.mcdonald@uottawa.ca, abara061@uottawa.ca

Boone Tensuda
University of Toronto
Institute for Aerospace Studies
boone.tensuda@mail.utoronto.ca

Clinton P. Groth
University of Toronto Institute for Aerospace Studies
Canada
groth@utias.utoronto.ca

MS248

Theoretical and Computational Investigations of the Non-linear Coupled Constitutive Relations (NCCR)

In a classical framework, the Navier-Stokes-Fourier (NSF) equations can be obtained via linear uncoupled thermodynamic flux-force relations which guarantee the non-negativity of the entropy production. It is commonly accepted that the conventional thermodynamic description is only valid when the Knudsen number is sufficiently small. Here, we will show that the range of validity of the NSF equations may be extended further by considering the nonlinear coupling between thermodynamic fluxes and forces. The resulted Nonlinear Coupled Constitutive Relations (NCCR) can capture many interesting rarefaction effects, such as Knudsen paradox, transpiration flows, thermal stress, heat flux without temperature gradients, etc. We will also derive a set of phenomenological boundary conditions for NCCR which respect the second law of thermodynamics. We will explore some boundary value problems by comparing the NCCR and the DSMC simulations.

Anirudh Singh Rana
GNU Building 404 Room 305
Dept: ACML, GNU, South Korea
anirudh@uvic.ca

Rho Shin Myong
School of Mechanical and Aerospace Engineering
Gyeongsang National University, South Korea
myong@gnu.ac.kr

MS249

Exploiting Active Subspaces for Nonlinear Programming

The active subspace of function $f : \mathbb{R}^m \to \mathbb{R}$ is the span of $n < m$ vectors constructed so that input perturbations along the active subspace change f more *on average* than input perturbations orthogonal to the active subspace. Such directions may help optimize complex nonlinear functions of several variables.

Paul Constantine
Colorado School of Mines
Applied Mathematics and Statistics
pconstan@mines.edu

David F. Gleich
Purdue University
dgleich@purdue.edu

MS249

Using Stochastic Optimization Methods for the Polyadic Decomposition of Large-Scale Tensors

Large-scale higher-order datasets pose many challenges in terms of computational and storage costs. As these tensors may not fit into the computer's memory, most canonical polyadic decomposition algorithms cannot be used. In this talk, we propose to use stochastic optimization methods, as these kind of algorithms require only a small amount of entries in memory in every iteration. The power of these methods will be illustrated on both synthetic and real life

data.

Nico Vervliet
Department of Electrical Engineering (ESAT)
KU Leuven
nico.vervliet@esat.kuleuven.be

Lieven De Lathauwer
Katholieke Universiteit Leuven
lieven.delathauwer@kuleuven-kulak.be

MS249

Non-Convex Low-Rank Matrix and Tensor Recovery

In this talk, I will present algorithms for solving non-convex low-rank matrix and also tensor recovery models. Global convergence to stationary point or in terms of first-order optimality condition is given. Although we cannot guarantee global optimal solutions, numerical experiments demonstrate that our algorithms can give more faithful solutions than those for solving convex models on both synthetic and real-world data.

Yangyang Xu
Rice University
yangyang.xu@rice.edu

MS249

Towards an Optimal Scalability in Computing Extreme Eigenpairs of Large Matrices

SVD and/or eigen-decomposition are fundamental problems in scientific and engineering computing. For large-scale data and on modern computers, classic algorithms have encountered scalability bottlenecks. In this work, we study two block methods that are based on, respectively, the Gauss-Newton and the power methods. We present theoretical and numerical results to demonstrate their promises. In particular, we show that the proposed multi-power method can frequently achieve an optimal scalability.

Yin Zhang
Rice University
Computational and Applied Mathematics
yzhang@rice.edu

Zaiwen Wen
Peking University
Beijing International Center For Mathematical Research
wendouble@gmail.com

Xin Liu
Academy of Mathematics and Systems Science
Chinese Academy of Sciences
liuxin@lsec.cc.ac.cn

MS250

Towards a Recursive Graph Bipartitioning Algorithm for Well Balanced Domain Decomposition

In the context of hybrid sparse linear parallel solvers based on Schur complement approaches, getting a domain decomposition tool leading to a good balancing of both internal and interface nodes for all the domains is a critical point for parallel efficiency. In this presentation, we introduce several variations of the existing algorithms in the multi-level Scotch partitioner that take into account these multiple criteria. We illustrate the results on graphs used for numerical scientific applications.

Astrid Casadei
INRIA Bordeaux Sud-Ouest & CNRS (LaBRI UMR 5800)
Bordeaux University
astrid.casadei@gmail.com

Pierre Ramet
LaBRI, Univ Bordeaux, France
701868

Jean Roman
INRIA
Jean.Roman@inria.fr

MS250

Handling Multiple Communication Metrics for Hypergraph Partitioning

We investigate connectivity-based partitioning methods, specifically hypergraph partitioning-based methods, for task mapping problem in order to efficiently parallelize the communicating tasks. A good partitioning method should divide the load among the processors as evenly as possible and minimize the inter-processor communication overhead. The total communication volume is the most popular communication overhead metric which is reduced by the existing state-of-the-art hypergraph partitioners. However, other metrics such as the total number of messages, the maximum amount of data transferred by a processor, or a combination of them are equally, if not more, important. We propose a directed hypergraph model to capture multiple communication metrics and a one-phase approach where all the communication cost metrics can be effectively minimized in a multi-objective setting. We wrapped the proposed model and methods in a multi-objective, multi-level hypergraph partitioner called UMPa. The partitioner takes various prioritized communication metrics into account, and optimizes all of them together. Compared to the state-of-the-art methods which only minimize the total communication volume, we show on a large number of problem instances that UMPa produces better partitions in terms of several communication metrics.

Mehmet Deveci
The Ohio State University
mdeveci@bmi.osu.edu

Kamer Kaya
The Ohio State University
Department of Biomedical Informatics
kamer@bmi.osu.edu

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

Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
umit@bmi.osu.edu

MS250

Complex Objective Partitioning of Small-World

Networks Using Label Propagation

We present PULP, a parallel and memory-efficient graph partitioning method specifically designed to partition low-diameter networks with skewed degree distributions. Partitioning determines the in-memory layout of a graph, which affects locality, inter-task load balance, communication time, and overall memory utilization of graph analytics. A novel feature of our method PULP (Partitioning using Label Propagation) is that it optimizes for multiple objective metrics simultaneously, while satisfying multiple graph constraints. Using our method, we are able to partition a web crawl with billions of edges on a single compute server in under a minute. For a collection of test graphs, we show that PULP uses 8-39x less memory than state-of-the-art partitioners, and is up to 14.5x faster, on average, than alternate approaches (with 16-way parallelism). We also achieve better partitioning quality results for both the multi-constraint as well as multi-objective scenarios.

George Slota, Kamesh Madduri
Pennsylvania State University
gslota@psu.edu, madduri@cse.psu.edu

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

MS250

Load Balancing Multiscale Simulations

Adaptive procedures and associated load-balancing operations are requirements for a high-performance parallel simulations to take full advantage of computational resources. Multiscale simulations introduce additional sensitivities and costs into adaptive processes which are not present in the standard single-program multiple-data (SIMD) parallel model used by the overwhelming majority of single-scale codes. The Adaptive Multiscale Simulation Infrastructure (AMSI) supports the implementation and execution of general multiscale simulations, and provides methods for global and local scale-sensitive load balancing and adaptive operations. An implementation of a multiscale simulation of soft-tissue mechanics using AMSI is discussed along with its usage of the scale-sensitive adaptive and load-balancing procedures provided by AMSI.

William R. Tobin, Daniel Fovargue
Rensselaer Polytechnic Institute
tobinw2@rpi.edu, dfovargue@gmail.com

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

MS251

High-Order Accurate Numerical Methods for Elliptic and Parabolic Interface Models

Designing numerical methods with high-order accuracy for problems with interfaces (for example, models for composite materials or fluids, etc), as well as models in irregular domains is crucial to many physical and biological applications. In this talk we will discuss recently developed efficient numerical schemes based on the idea of the Difference Potentials for elliptic and parabolic composite domain/interface problems. Numerical experiments to illustrate high-order accuracy and the robustness of the developed methods will be presented as well.

Yekaterina Epshteyn
Department of mathematics
University of Utah
epshteyn@math.utah.edu

MS251

A High-Order Adaptive Finite Volume Solver for Steady Euler Equations

In this talk, we will present our work on high-order adaptive finite volume methods for steady Euler equations. There are two main components in our solver. The first one is a high-order finite volume solver for steady Euler equations. To reach the high-order accuracy, the k-exact reconstruction is used. The Newton iteration is applied to linearize the equations, then the linear system is solved by a geometrical multigrid method. In the second component, we use the framework of the goal oriented a posteriori error estimation, and develop some new and efficient techniques to generate the error indicator. Then a h-adaptive method will be introduced to optimize the distribution of the mesh grids, and the improvement of the efficiency on the algorithm implementation can be expected. The numerical results verify our theoretical results, and the high-order behavior of our method will be demonstrated by the benchmark examples.

Guanghui Hu
University of Macau
garyhu@umac.mo

MS251

A Seventh Order Hybrid Weighted Compact Scheme Based on WENO Stencil for Hyperbolic Conservation Laws

In this paper, we propose a hybrid seventh order weighted compact scheme for shock capturing. In smooth region, present scheme recovers a seventh linear compact scheme. Near discontinuities, fifth order weighted compact scheme is used. To guarantee convergence condition, the seventh order scheme and fifth order scheme are coupled by a new smooth indicator based on WENO weights. Present scheme only uses five points which is the same as fifth order WENO scheme.

Jun Peng, Yiqing Shen
State Key Laboratory of High Temperature Gas Dynamics
Institute of Mechanics, Chinese Academy of Science
pengjun@imech.ac.cn, yqshen@imech.ac.cn

MS251

Maximum Principle and Positivity Preserving Flux Limiters for High Order Schemes

Abstract not available at time of publication.

Zhengfu Xu
Michigan Technological University
Dept. of Mathematical Sciences
zhengfux@mtu.edu

MS252

PDE-constrained Optimization with Local Control

and Boundary Observations: Robust Preconditioners

We consider PDE-constrained optimization problems with control functions defined on a subregion of the domain of the state equation. The main purpose of this paper is to define and analyze robust preconditioners for KKT systems associated with such optimization tasks. That is, preconditioners that lead to iteration bounds, for the MINRES scheme, that are independent of the regularization parameter α and the mesh size h. Our analysis addresses elliptic control problems, subject to Tikhonov regularization, and covers cases with boundary observations only and locally defined control functions. A number of numerical experiments are presented

Ole Løseth Elvetun
Department of Mathematical Sciences and Technology
Norwegian University of Life Sciences
ole.elvetun@nmbu.no

Bjørn F. Nielsen
Simula Research Laboratory
bjorn.f.nielsen@nmbu.no

MS252

Robust Preconditioners for PDE-Constrained Optimization with Limited Observation Data

Regularization robust preconditioners have been successfully developed for some PDE-constrained optimization problems. These methods, however, typically assume that observation data is available throughout the entire domain of the state equation. For many inverse problems, this is an unrealistic assumption. We propose and analyze preconditioners for PDE-constrained optimization problems with limited observation data, e.g. observations are only available at the boundary of the computational domain. Our methods are robust with respect to both the regularization parameter and the mesh size.

Magne Nordaas
Center for Biomedical Computing
Simula Research Laboratory
magneano@simula.no

Kent-Andre Mardal
Simula Research Laboratory and Department of Informatics
University of Oslo
kent-and@simula.no

Bjørn F. Nielsen
Simula Research Laboratory
bjorn.f.nielsen@nmbu.no

MS252

All-at-once Approach to Optimal Control Problems Constrained by PDEs with Uncertain Inputs

We present an efficient approach to simulate optimization problems governed by partial differential equations involving random coefficients. This class of problems leads to prohibitively high dimensional saddle point systems with Kronecker product structure, especially when discretized with the stochastic Galerkin finite element method. Here, we derive robust Schur complement-based block diagonal preconditioners for solving the resulting stochastic Galerkin systems with all-at-once low-rank solvers. Finally, we illustrate the effectiveness of our solvers with numerical experiments.

Akwum Onwunta, Peter Benner
Max Planck Institute, Magdeburg, Germany
onwuntajunior@gmail.com,
benner@mpi-magdeburg.mpg.de

Martin Stoll
Max Planck Institute, Magdeburg
stollm@mpi-magdeburg.mpg.de

MS252

Accelerated Source-Encoding Full-Waveform Seismic Inversion with Additional Constraints

We present a semismooth Newton-PCG method for full-waveform inversion governed by the elastic wave equation that can handle additional constraints on the material. Source-encoding substantially reduces the computational costs compared to conventional approaches that consider every source individually. In particular, we accelerate the minimization of a sample average approximation model by using inexact Hessian information based on mini-batches of the samples. Furthermore, we compare the performance with preconditioned stochastic descent schemes.

Michael Ulbrich
Technische Universitaet Muenchen
Chair of Mathematical Optimization
mulbrich@ma.tum.de

Christian Boehm
Technische Universität München
boehm@ma.tum.de

MS253

Improved Understanding of Atmospheric Stability Effects on Wind Farm Performance Using Large-Eddy Simulation

Recent evidence from operational wind farms has shown that atmospheric stability, especially heat fluxes at the surface, can have a profound effect on turbine-atmosphere interactions. We use two LES codes to study atmospheric stability effects in wind farms: the Simulator for Offshore/Onshore Wind Farm Applications (SOWFA), a publicly-available, finite-volume code developed in C++ by the National Renewable Energy Laboratory, and the Wind Turbine and Turbulence Simulator (WiTTS), a new, finite-difference, Fortran code developed in-house. Comprehensive results will be presented for both single- and multiple-turbine cases under a variety of initial conditions.

Cristina L. Archer
University of Deleware
carcher@udel.edu

Shengbai Xie, Niranjan Ghaisas
University of Delaware
sxie@udel.edu, nghaisas@udel.edu

MS253

Characterizing Turbulence in Wind Turbine Wake:

Role of Stratification

Abstract not available at time of publication.

Kiran Bhaganagar
University of TExas, San Antonio
kiran.bhaganagar@utsa.edu

MS253

Les Study of a Large Wind Farm Within a Diurnal Atmospheric Boundary Layer

Abstract not available at time of publication.

Marc Calaf
University of Utah
marc.calaf@utah.edu

MS254

Bootstrap and Adaptive Methods

In this talk, I will highlight several recent advances in the development and analysis of AMG coarsening algorithms. I will discuss various strategies for selecting the coarse variables and defining interpolation, in both the adaptive AMG and Bootstrap AMG settings. Numerical experiments of the proposed techniques applied to various applications will also be provided.

James Brannick
Department of Mathematics
Pennsylvania State University
brannick@psu.edu

MS254

Algebraic Multigrid for H-hermitian Matrices

We develop an algebraic multigrid method for solving linear systems of equation $Ax = b$ with non-Hermitian matrices A that possess a simple symmetrizing operator, e.g., Saddle-point problems, Hamiltonian matrices. That is there exists a simple matrix H such that HA is hermitian. We observe that by carefully constructing the intergrid transfer operators it is possible not only to transfer this non-standard symmetry to coarse scales, but also to reduce a general Petrov-Galerkin to a Galerkin coarse grid construction, namely for HA. We show that by using this construction and a Kaczmarz smoother we obtain a method that yields the same iterates when applied to $Ax = b$ or $HAx = Hb$. To demonstrate the applicability of this approach we develop a method for the Wilson discretization of the 2-dimensional Dirac equation. The proposed approach uses a bootstrap setup algorithm based on a multigrid eigensolver. It computes test vectors which define the least squares interpolation operators by working mainly on coarse grids, leading to an efficient and integrated self learning process for defining algebraic multigrid interpolation. The algorithm is motivated by the γ_5-symmetry of the Dirac equation, which carries over to the Wilson discretization. This discrete γ_5-symmetry is used to reduce a general Petrov Galerkin bootstrap setup algorithm to a Galerkin method for the Hermitian and indefinite formulation of the Wilson matrix. Kaczmarz relaxation is used as the multigrid smoothing scheme in both the setup and solve phases of the resulting Galerkin algorithm. Extensive numerical results are presented to motivate the design and demonstrate the effectiveness of the proposed approach.

Karsten Kahl
Bergische Universität Wuppertal
Department of Mathematics
kkahl@math.uni-wuppertal.de

MS254

Root-Node Based Algebraic Multigrid

Recent approaches to improving multigrid convergence through modified coarsening and enhanced interpolation have shown to be effective in a general setting (e.g. complex, non-Hermitian, and indefinite). Yet, the the resulting multigrid hierarchies may exhibit higher complexities than necessary. In this talk we outline a root-node based approach to multigrid, which can be viewed as a hybrid of classical and aggregation based multigrid methods. This allows both point-wise decisions in the setup while retaining the framework of aggregation. We give an overview of root-node multigrid using interpolation based on energy minimization and show how the complexity of the multigrid cycle is controlled through selective filtering by utilizing a root-node. The method yields improved interpolation (and convergence), while limiting the total work of the cycle with minimal tuning of parameters. We present some numerical results in support and discuss directions for further theoretical and numerical development.

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

Jacob B. Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

MS254

Algebraic Multigrid Method for Implicit Smoothed Particle Hydrodynamics

Recently, a Lagrangian particle model based on smoothed particle hydrodynamics (SPH) is proposed to numerically solve the coupled system that describes the electrokinetic phenomena. We generalize the aggregation based Alegebraic multigrid (AMG) method to solve the large-scale linear system of equations discretized from implicit SPH. Auxiliary grid approach is used to improve the efficiency and reduce the computational complexity. Numerical experiments for modeling electroosmotic flow in microchannels and flow throgh charged membranes are presented to demonstrate the effectiveness of the proposed AMG method.

Xiaozhe Hu
The Pennsylvania State University
xiaozhe.hu@tufts.edu

Wenxiao Pan
Pacific Northwest National Laboratory
wenxiao.pan@pnnl.gov

Jinchao Xu
Pennsylvania State University
xu@math.psu.edu

Hongxuan Zhang
Penn State University

zhxsxuan@gmail.com

MS255

Designing Visualizations for Biological Research

Advances in measurement devices in the last decade have given rise to an explosion of scientific data. In biology, access to massive amounts of quantitative data has fundamentally changed how discoveries are made, and now an important component of the scientific process is making sense of this data using visualization methods. For most biologists, however, their toolbox is made up of only broadly-available tools that were designed for over-arching problems, often leaving them without answers to their specific questions. A growing trend in the visualization community is to develop tools that focus on specific, real-world problems. Called a design study, the process of developing these tools relies on a close collaboration with end-users as well as the use of methods from design. In this talk I'll present several design studies that target complex, biological data analysis, from discovering trends in molecular networks to understanding the results of comparative genomics algorithms.

Miriah Meyer
University of Utah
miriah@cs.utah.edu

MS255

Block-Based Analysis of Scientific Data

Block-based analysis is the division of an analysis problem into blocks (not processes) that communicate. By allowing blocks to be configurable size and number, flexibly assigning blocks to processes, and migrating blocks among different levels of memory/storage hierarchy, an analysis infrastructure can scale and adapt to current and future science applications and HPC platforms. We will discuss these fundamental concepts and their implementation in a reusable data movement library designed for block-based scientific data analysis.

Tom Peterka
Argonne National Laboratory
tpeterka@mcs.anl.gov

MS255

Exascale Scientific Data Analytics and Visualization

Among many existing feature descriptors, statistical information derived from data samples is a promising approach to taming the big data avalanche because data distributions computed from a population can compactly describe the presence and characteristics of salient data features with minimal data movement. The ability to computationally summarize and process data using distributions also provides an efficient and representative capture of information that can adjust to size and resource constraints, with the added benefit that uncertainty associated with results can be quantified and communicated. In this talk, I will discuss applying distribution-based approaches for visualization and data analytics, including multivariate analysis, vector/scalar field analysis, and histogram compression and query.

Han-Wei Shen
The Ohio State University
hwshen@cse.ohio-state.edu

MS255

Exploring Big Urban Data

For the first time in history, more than half of the world's population lives in urban areas. Given the growing volume of data that is being captured by cities, the exploration of urban data will be essential to inform both policy and administration, and enable cities to deliver services effectively, efficiently, and sustainably while keeping their citizens safe, healthy, prosperous, and well-informed. Urban data analysis is a growing research field that will not only push computer science research in new directions, but will also enable many others, including urban planners, social scientists, transportation experts. We have been working on methods and systems that support urban data analysis, with a focus on spatio-temporal aspects. We will describe these efforts, in particular our work on analyzing the NYC taxi dataset, which contains information about over 850 million yellow cab trips. Supported by NSF, Google, Moore-Sloan Data Science Environment, IBM, and CUSP.

Claudio T. Silva
NYU
csilva@nyu.edu

MS256

Probability Density Methods for the Analysis of Power Grids Under Uncertainty

We present the probability density function (PDF) method for the analysis of power grids under uncertainty due to random renewable energy inputs. We derive deterministic equations for the evolution of the PDF of the power systems using various closures. The resulting PDEs are solved numerically, and results are compared against reference Monte Carlo simulations to assess the accuracy of the proposed closures.

David A. Barajas-Solano
Pacific Northwest National Lab
david.barajas-solano@pnnl.gov

Alexander Tartakovsky
Pacific Northwest National Laboratory
alexandre.tartakovsky@pnnl.gov

Debojyoti Ghosh
Mathematics and Computer Science Division
Argonne National Laboratory
ghosh@mcs.anl.gov

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

Shrirang Abhyankar
Argonne National Laboratory
abhyshr@mcs.anl.gov

Zhenyu Huang
Pacific Northwest National Laboratory

zhenyu.huang@pnnl.gov

MS256

Distributed Optimization Algorithms for Wide-Area Oscillation Monitoring in Power Systems

In this talk we will present three distributed optimization algorithms based on Alternating Directions Multiplier Method (ADMM) for estimating the electro-mechanical oscillation modes of large power system networks using Synchrophasors. Both synchronous and asynchronous communications will be considered. Each architecture integrates a centralized Prony-based algorithm with several variants of ADMM. We will discuss convergence and resiliency properties of each architecture using analytical results as well as simulations of IEEE prototype power system models.

Aranya Chakrabortty
North Carolina State University
achakra2@ncsu.edu

Seyedbehzad Nabavi
NC State University
snabavi@ncsu.edu

MS256

Exploiting Network Laplacian Structure in Power Grid Dynamics

The electromechanical dynamics of the power grid possess a structure that is "nearly" Hamiltonian, with the gradient an underlying scalar function capturing much of the information needed to reconstruct the vector field. This work will go further to exploit added structure in this scalar potential function, demonstrating that its network related portion can be written as a Hermitian form in the complex variables of the voltage phasors (i.e., the fundamental component for a windowed Fourier transform of the nearly sinusoidal bus voltages of the power grid). Moreover, the defining matrix of this Hermitian from is a weighted Laplacian, defined by the topology and electrical characteristics of the transmission grid. We show how this structure can be exploited to build tractable dynamic models of cascading failure phenomena, in which overload and protective disconnection of one transmission line may induce subsequent overloads and disconnects of further links, with the risk of cascade into system-wide failure.

Christopher DeMarco, Honghao Zheng
University of Wisconsin-Madison
demarco@engr.wisc.edu, hustfreedom.zheng@gmail.com

MS256

Fast Algorithms for Synchrophasor Computations

There is an urgent need to develop fast real-time algorithms that can extract operator friendly information out of large-scale high speed synchronized measurement devices being implemented in power systems all over the world. This talk will highlight the stability monitoring algorithms developed recently in this context while emphasizing the computational challenges in realizing the theoretical methods. Distributed algorithms as well as centralized formulations will be compared.

Vaithianathan Venkatasubramanian, Tianying Wu, Seyed Arash Sarmadi, Ebrahim Rezaei
Washington State University
mani@eecs.wsu.edu, lilyrosejam@hotmail.com, arash.sarmadi@gmail.com, rezaee.ebrahim@gmail.com

MS257

A multi-physics Domain Decomposition Method for Navier-Stokes-Darcy Model

This presentation discusses a multi-physics domain decomposition method for solving the coupled steady-state Navier-Stokes-Darcy system with the Beavers-Joseph interface condition. The wellposedness is first showed by using a branch of singular solutions. Robin boundary conditions on the interface are constructed based on the physical interface conditions to decouple the Navier-Stokes and Darcy parts. Then a domain decomposition method is developed and analyzed. Numerical examples are presented to validate this method.

Xiaoming He
Department of Mathematics and Statistics
Missouri University of Science and Technology
hex@mst.edu

MS257

Improvements in the Level Set Method with a Focus on Curvature-Dependent Forcing

One particular subset of moving interface problems is when the interface exerts a curvature-dependent force on the surrounding material. Such a force can arise, for example, when the interface is resistant to bending or has surface tension. Stable numerical computation of the interface curvature can be difficult as it is often expressed as high-order derivatives of either marker particle positions or of a level set function. Focusing on the latter, the level set method is modified to track not only the interface position, but the curvature as well. The definition of the signed-distance function that is used in this approach is also used to develop an interpolation-free, closest-point method for solving surface PDEs.

Chris Vogl
University of Washington
cvogl@u.washington.edu

MS257

Optimal Energy Conserving Discontinuous Galerkin Methods for the Wave Propagation Problems in Heterogeneous Media

Solving wave propagation problems within heterogeneous media has been of great interest and has a wide range of applications in physics and engineering. The design of numerical methods for such general wave propagation problems is challenging because the energy conserving property has to be incorporated in the numerical algorithms in order to minimize the phase or shape errors after long time integration. In this talk, we will discuss multi-dimensional wave problems and consider linear second-order wave equation in heterogeneous media. We will present an LDG method, in which numerical fluxes are carefully designed to maintain the energy preserving property and accuracy. We propose compatible high order energy conserving time integrators and prove the optimal error estimates and the energy conserving property for the semi-discrete methods. Our numerical experiments demonstrate optimal rates of convergence, and show that the errors of the numerical solution do not grow significantly in time due to the energy

conserving property.

Yulong Xing
Department of Mathematics
Univeristy of Tennessee / Oak Ridge National Lab
xingy@math.utk.edu

MS257

Surface Phase Separation Mediated by Nonlocal Interactions

Motivated by the lipid phase separation on membrane mediated by electrostatic interactions, we propose here a general Cahn-Hilliard equation with nonlocal electrostatic interactions. A C-0 interior penalty discontinuous Galerkin method is adapted to solve the coupled system of PDEs on arbitrary surfaces. Phase separation of smaller scales are found compared to scenario without nonlocal interaction, indicating that the nonlocal interaction may generate lipid rafts.

Yongcheng Zhou
Department of Mathematics
Colorado State University
yzhou@math.colostate.edu

MS258

The Role of Autotuning Compiler Technology

Autotuning empirically evaluates a search space of possible implementations of a computation to identify the implementation that best meets its optimization criteria (e.g., performance, power, or both). Autotuning compilers generate this search space of different implementations either automatically or with programmer guidance. This talk will explore the role of compiler technology in achieving very high levels of performance, comparable to what is obtained manually by experts. It will focus on the optimizations required for specific domains: geometric multigrid, stencils, and tensor contraction computations.

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

MS258

Active Harmony: Making Autotuning Easy

Active Harmony is an auto-tuning framework for parallel programs. In this talk, I will describe how the system makes it easy (sometime even automatic) to create programs that can be auto-tuned. I will present example from a few applications and programming languages.

Jeffery Hollingsworth
Computer Science Department
University of Maryland
hollings@cs.umd.edu

MS258

Towards Auto-tuning in the Era of 200+ Thread Parallelisms — FIBER Framework and Minimizing Software Stack —

Currently, parallelism of 200+ threads is pervasive by many-core processors, such as the Xeon Phi. In the processors, we need careful optimizations with respect to execution conditions at run-time, such as problem sizes and hybrid MPI/OpenMP. In this talk, we introduce FIBER framework of Auto-tuning (AT) to optimize codes for the many-core processors. We also show effect of the AT with multi nodes of the Xeon Phi with an application of Finite Difference Method.

Takahiro Katagiri
Information Technology Center, The University of Tokyo
katagiri@cc.u-tokyo.ac.jp

Satoshi Ohshima
The University of Tokyo
Information Technology Center
ohshima@cc.u-tokyo.ac.jp

Masaharu Matsumoto
Information Technology Center
The University of Tokyo
matsumoto@cc.u-tokyo.ac.jp

MS258

A Framework for Separation of Concerns Between Application Requirements and System Requirements

An HPC application is usually optimized for a particular platform and unable to run efficiently on other platforms. The Xevolver framework is designed to separate such platform-specific optimizations from application codes. In Xevolver, a code portion that needs to be modified for system-specific code optimization is just annotated for user-defined code transformations to prevent messing up the original code. The usefulness and practicality of Xevolver are discussed on the basis of some case studies.

Hiroyuki Takizawa, Shoichi Hirasawa
Tohoku University/JST CREST
tacky@isc.tohoku.ac.jp, hirasawa@sc.isc.tohoku.ac.jp

Hiroaki Kobayashi
Tohoku University, Japan
koba@isc.tohoku.ac.jp

MS259

Microstructural Modeling of Fracture in Uranium dioxide Using a Phase-Field Based Approach

A phase-field model is implemented in MOOSE to investigate the effect of porosity and grain size on the intergranular brittle fracture in UO_2. The grain boundary fracture parameters are obtained from the results of molecular dynamics simulations. A numerical sensitivity study is then performed to obtain a microstructurally informed stress-based fracture model usable at the engineering scale.

Pritam Chakraborty, Michael Tonks
Idaho National Laboratory
pritam.chakraborty@inl.gov, michael.tonks@inl.gov

MS259

Computational Microstructure Science Using the Moose Framework

Mesoscale modeling and simulation provide a bridge between atomistic data and engineering scale computation. However, mesoscale modeling can be complicated due to

the multiphysics nature of the problems, and the typically high computation costs. The open source Multiphysics Object-Oriented Simulation Environment (MOOSE) provides a number of physics modules that are aimed specifically at mesoscale simulation. These tools model the coevolution of microstructure and properties due to applied load, temperature, and radiation damage. The Phase Field Module provides all the necessary tools to predict microstructure evolution using the phase field method and the Tensor Mechanics Module provides the tools for finite deformation mechanics simulations at the level of microstructure. Also, effective mechanical and thermal properties can be calculated as the microstructure evolves. Since these tools are based on MOOSE, they are massively parallel, work in 1D, 2D, and 3D and can use mesh and time step adaptivity. We have also developed the capability to reconstruct experimental microstructures directly into the simulations, to set the initial condition for simulations, to simplify validation and to facilitate virtual property measurement.

Bradley Fromm, Daniel Schwen, Michael Tonks
Idaho National Laboratory
bradley.fromm@inl.gov, daniel.schwen@inl.gov, michael.tonks@inl.gov

MS259

Fission Bubble Modeling in Uranium Carbide

Fission gas bubble behavior is a complex microscopic effect that leads to macroscopic changes in nuclear fuel. The interplay between thermal diffusion causing bubble growth, pitted against physical knock-outs from fission fragments causing bubble shrinkage, creates a non-linear problem that is not easily solved. MOOSE allows for characterization of this behavior within the code BISON by using lower-length scale effects to estimate the large bubble structure and swelling of irradiated uranium carbide.

Christopher Matthews, Andrew Klein
Oregon State University
matthchr@engr.orst.edu, andrew.klein@oregonstate.edu

MS259

Grizzly: A Simulation Tool for Nuclear Power Plant Component Aging

To determine the risk associated with extending the life of nuclear power plants, the effects of age-related degradation of critical components in those plants must be understood. Grizzly is a MOOSE-based tool being developed to simulate aging processes and understand the effects that age-related degradation will have on those components. An initial application of Grizzly is for coupled physics simulations to assess the susceptibility of aged reactor pressure vessels to fracture under accident conditions.

Benjamin Spencer
Idaho National Laboratory
benjamin.spencer@inl.gov

MS260

Iterative Solution Techniques in Reduced-order Modeling

Reduced-order models (ROMs) efficiently solve related linear systems arising in many-query applications. The full linear system is projected onto a smaller dimensional subspace resulting in a ROM whose solution approximates the solution of the original system. The reduced model size is problem dependent and could be smaller than the full problem, but large enough to make direct solution costly. In this scenario, the iterative solution of the ROM can be more effecient.

Virginia Forstall
University of Maryland at College Park
vhfors@gmail.com

Howard C. Elman
University of Maryland, College Park
elman@cs.umd.edu

MS260

Model Reduction in Physics-Based Sound Synthesis

Synthesizing physics-based sounds by time-stepping computational models of coupled solid and fluid systems is extremely time consuming due to long-time integration and the desire for real-time auditory feedback. This talk will highlight some of our recent work on reduced-order modeling of vibration and acoustic radiation: (1) basis compression to reduce the memory footprint of modal sound models, and (2) reduced-order modeling of bubble-based liquid sounds.

Doug L. James, Timothy Langlois
Cornell University
djames@cs.cornell.edu, langlois@cs.cornell.edu

MS260

Online Adaptive Model Reduction

Model reduction approximates the nonlinear manifold induced by full-order solutions with a (linear) reduced space. We present a nonlinear approximation of the manifold based on adapting the reduced space online. The adaptation relies on low-rank basis updates derived from sparse data of the full-order model. In particular in the presence of nonlinearities, our nonlinear approximation achieves a higher accuracy than the classical linear approximation. The sparsity of the data ensures computational efficiency.

Benjamin Peherstorfer
ACDL, Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
pehersto@mit.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@MIT.EDU

MS260

Reduced-order Models using Dynamic Mode Decomposition

Dynamic mode decomposition (DMD) is a technique that uses data to express "black-box" dynamics in terms of the eigenvalues, eigenfunctions, and modes of the corresponding Koopman operator. For a linear system, these modes are the eigenvectors, and they have a similar role in nonlinear settings. We describe Extended DMD, an extension that uses a richer set of functions to approximate the Koopman eigenfunctions, and demonstrate the method on

several examples, including flow past a cylinder.

Clarence Rowley
Princeton University
Department of Mechanical and Aerospace Engineering
cwrowley@princeton.edu

Matthew O. Williams
Applied Mathematics
University of Washington
mow2@Princeton.edu

Maziar S. Hemati, Scott Dawson
Princeton University
mhemati@princeton.edu, stdawson@princeton.edu

MS261

A Posteriori Analysis of the Parareal Algorithm: Efficient Resource Allocation and Convergence Criteria

We carry out a posteriori analysis of the parareal algorithm using adjoint based methods. The analysis is carried out relatively cheaply and allows for quantification of error in a quantity of interest. Furthermore, the analysis provides guidance in choosing discretization parameters so as to enhance load balancing in the parallel stage of the algorithm and by quantifying different sources of error, suggests appropriate strategies to accelerate convergence.

Jehanzeb H. Chaudhry
Florida State University
jehanzeb.hameed@gmail.com

Don Estep, Simon Tavener
Colorado State University
estep@stat.colostate.edu, tavener@math.colostate.edu

MS261

An Overview of the Multigrid Reduction in Time (MGRIT) Method

The multigrid-reduction-in-time (MGRIT) method is a scalable, truly multilevel approach to parallel time integration, derived based on multigrid reduction principles. In this talk, we present the algorithm and demonstrate that MGRIT offers excellent strong and weak parallel scaling up to thousands of processors. Complementary convergence analysis methodologies such as a semi-algebraic approach to mode analysis, which provides a predictive analysis tool for MGRIT and related algorithms on space-time grids, will also be discussed.

Stephanie Friedhoff
KU Lueven
stephanie.friedhoff@cs.kuleuven.be

MS261

Parareal Library for Time-Dependent PDEs in Medical Applications

The parallel-in-time integration method Parareal provides an additional direction for concurrency when solving time-dependent PDEs numerically. Thus, it can extend the strong scaling limit of a purely space-parallel approach. In this talk, we present a new C++ 'Library for the Parareal Method' (Lib4PrM). To illustrate its performance, a coupling with the ug4 package is used to solve a PDE modeling permeation through human skin.

Andreas Kreienbuehl
University of Lugano
andreas.kreienbuehl@usi.ch

Arne Naegel
Goethe University Frankfurt
arne.naegel@gcsc.uni-frankfurt.de

Daniel Ruprecht
Institute of Computational Science
Universita della Svizzera italiana
daniel.ruprecht@usi.ch

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich
r.speck@fz-juelich.de

Gabriel Wittum
Goethe Center for Scientific Computing
Goethe University Frankfurt
wittum@gcsc.uni-frankfurt.de

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

MS261

Multigrid Reduction in Time (MGRIT): A Flexible and Non-Intrusive Method

This talk highlights practical aspects of the parallel-in-time method, multigrid-reduction-in-time (MGRIT). The algorithm is non-intrusive and wraps existing time stepping codes. MGRIT is versatile by allowing for various time discretizations (e.g., Runge-Kutta and multistep) and for adaptive refinement and coarsening in time and space. Non-linear problems are handled through full approximation scheme (FAS) multigrid. Details of the software philosophy and practical experience and results (e.g., from a nonlinear compressible Navier-Stokes code), will also be given.

Jacob B. Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

Robert Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

Ulrike Meier Yang
Lawrence Livermore National Laboratory
yang11@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

Stephanie Friedhoff
KU Lueven
stephanie.friedhoff@cs.kuleuven.be

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

MS262

Demonstrating Improved Application Performance Using Dynamic Monitoring and Task Mapping

We present a framework enabling dynamic application mapping based on run-time analysis of system-wide network data, architecture-specific routing algorithms, and application communication patterns. We demonstrate dynamic remapping of MPI tasks based on route-length, bandwidth, and credit-stalls metrics for a parallel sparse matrix-vector multiplication kernel. Remapping based on dynamic network information in a congested environment in a shared network topology recovers up to 50% of the time lost to congestion, reducing matrix-vector multiplication time by 8%.

Ann Gentile, James Brandt, Karen D. Devine, Kevin Pedretti
Sandia National Laboratories
gentile@sandia.gov, brandt@sandia.gov, kddevin@sandia.gov, ktpedre@sandia.gov

MS262

Topology Aware Process Placement and Data Management

Programming multicore or manycore architectures efficiently is a challenge because numerous hardware characteristics have to be taken into account, especially the memory hierarchy. In this talk we will show how we can efficiently manage data and reduce communication cost by taking into account the topology of the machine and the affinity of the application processes in different contexts: process placement, load-balancing, resource allocation.

Emmanuel Jeannot
INRIA
emmanuel.jeannot@inria.fr

MS262

Local Search to Improve Geometric Task Mapping

We present a local search strategy to improve the mapping of a parallel job's tasks to the MPI ranks of its parallel allocation in order to reduce network congestion and the job's communication time. The goal is to reduce the number of network hops between communicating pairs of ranks. Using the miniGhost mini-app, which models the shock physics application CTH, we demonstrate that our strategy reduces application running time while also reducing the runtime variability.

Vitus Leung
Sandia National Laboratories
vjleung@sandia.gov

David Bunde
Knox College
dbunde@knox.edu

MS262

Process Mapping onto Complex Architectures and Partitions Thereof

The Scotch software computes process-processor mappings by assigning recursively parts of the process graphs to parts of the target graphs. To date, while regular target architectures can be described using pre-coded routines, irregular architectures or parts of regular ones require P^2 data structures, which makes them unpractical for very big machines. We will present a new, multilevel description of target architectures that alleviates this problem, trading-off memory for run time.

Francois Pellegrini
University of Bordeaux
francois.pellegrini@labri.fr

MS263

Fluctuating Hydrodynamics of Reactive Multispecies Mixtures

This paper describes the extension of the fluctuating Navier-Stokes (FNS) equations to multispecies, reactive mixtures. In FNS, hydrodynamic effects of thermal fluctuations are represented by adding stochastic flux terms, whose magnitude are set by fluctuation dissipation balance, to the Navier Stokes equations. Incorporating reactions into the system introduces a number of additional complexities. We discuss approaches to addressing these issues and present numerical results illustrating the impact of fluctuations in reacting systems.

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

MS263

Modeling Multi-Phase Flow Using Fluctuating Hydrodynamics

Incorporating thermal fluctuations in continuum Navier-Stokes equations requires the development of numerical methods that solve the complex stochastic partial differential equations of fluctuating hydrodynamics. The situation becomes more complex when more than one fluid phase is involved as in a liquid-vapor system. These stochastic PDE's require a special spatio-temporal discretization so that the correct Gibbs-Boltzmann equilibrium distribution is achieved at long times and the correct fluctuation-dissipation balance is preserved at each time step. We describe a stochastic method of lines discretization of the fully compressible Landau-Lifshitz-Navier-Stokes (fluctuating hydrodynamics) equations with the van der Waals equation of state. The diffuse interface method is used to model the order parameter (density) as a smooth variation across the interface. The surface tension effects give rise to Korteweg type stresses that show up as additional terms in the momentum and energy equations. The numerical scheme is validated by comparison of measured structure factors and capillary wave spectra with equilibrium theory. We also present several non-equilibrium examples to illustrate the capability of the algorithm to model multi-phase fluid phenomena in a neighborhood of the critical point. These examples include a study of the impact of fluctu-

ations on the spinodal decomposition following a rapid quench, as well as the piston effect in a cavity with supercooled walls. The conclusion in both cases is that thermal fluctuations affect the size and growth of the domains in off-critical quenches.

Anuj Chaudhri, John B. Bell
CCSE
Lawrence Berkeley Laboratory
achaudhri@lbl.gov, jbbell@lbl.gov

Alejandro Garcia
San Jose State University
algarcia@algarcia.org

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
aleks.donev@nyu.edu

MS263

The Long-time Tail of the Velocity Autocorrelation Function of a Particle in a Molecular Fluid

Through large-sized-ensemble runs of molecular dynamics simulation for a tracer particle suspended in a molecular fluid, its velocity autocorrelation function and diffusion coefficient are very accurately determined. The finite-system-size effects of molecular dynamics simulation and the effects of molecular character of the surrounding fluid are investigated. In addition, our results are compared with theoretical results obtained from the fluctuating hydrodynamics and computational results obtained from the smoothed-particle hydrodynamics.

Changho Kim, George E. Karniadakis
Brown University
Division of Applied Mathematics
changho_kim@brown.edu, george_karniadakis@brown.edu

MS264

A Low-Rank Approximation Method for High-Dimensional Uncertainty Quantification

This work introduces a model reduction technique that exploits the low-rank structure of the solution of interest, when exists, for fast propagation of high-dimensional uncertainties. To construct this low-rank approximation, the proposed method utilizes a hierarchy of models with lower fidelities, than the intended model, which can be simulated cheaply. Several numerical experiments will be provided to demonstrate the efficiency of the proposed approach.

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

Dongbin Xiu
University of Utah
dongbin.xu@utah.edu

Akil Narayan
University of Massachusetts Dartmouth
akil.narayan@umassd.edu

Hillary Fairbanks
Department of Applied Mathematics
University of Colorado, Boulder
hillary.fairbanks@colorado.edu

MS264

Bayesian Compressive Sensing Framework for High-Dimensional Surrogate Construction

Surrogate construction for high-dimensional models is challenged in two major ways: obtaining sufficient training model simulations becomes prohibitively expensive, and non-adaptive basis selection rules lead to excessively large basis sets. We enhanced select state-of-the-art tools from statistical learning to build efficient sparse surrogate representations, with quantified uncertainty, for high-dimensional complex models. Specifically, Bayesian compressive sensing techniques are supplemented by iterative basis growth and weighted regularization. Application to a 70-dimensional climate land model shows promising results.

Khachik Sargsyan, Cosmin Safta
Sandia National Laboratories
ksargsy@sandia.gov, csafta@sandia.gov

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

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

MS264

Inverse Subspace Iteration for Spectral Stochastic Finite Element Methods

We study random eigenvalue problems in the context of spectral stochastic finite elements. We assume that the matrix operator is given in the form of a polynomial chaos expansion, and we search for the coefficients of the polynomial chaos expansions of the eigenvectors and eigenvalues. We formulate a version of stochastic inverse subspace iteration, which is based on stochastic Galerkin finite element method, and we compare its accuracy with that of Monte Carlo and stochastic collocation methods.

Bedrich Sousedik
University of Maryland
sousedik@umbc.edu

Howard C. Elman
University of Maryland, College Park
elman@cs.umd.edu

MS264

Preconditioner for Parameter-dependent Linear Systems Based on an Empirical Interpolation of the Matrix Inverse

We consider parameter-dependent linear systems of equations, e.g. arising from the discretization of a parameter-dependent or stochastic PDE. We propose a parameter dependent preconditioner defined as an interpolation of the matrix inverse based on a Frobenius norm projection. We then show how such preconditioner can be used for projection based model reduction methods such as the reduced

basis method. We propose constructions of interpolation points that are dedicated either to the improvement of Galerkin projections or to the estimation of projection errors.

Loïc Giraldi
GeM, Ecole Centrale de Nantes
loic.giraldi@ec-nantes.fr

Anthony Nouy, Olivier Zahm
LUNAM Universite, Ecole Centrale Nantes, CNRS, GeM
anthony.nouy@ec-nantes.fr, olivier.zahm@ec-nantes.fr

MS265

Numerical Applications of Weak Galerkin Finite Element Methods

Weak Galerkin finite element methods are new numerical methods for solving partial differential equations that were first introduced by Wang and Ye for solving general second order elliptic partial differential equations (PDEs). The differential operators in PDEs are replaced by their weak forms through the integration by parts, which endows high flexibility for handling complex geometries, interface discontinuities, and solution singularities. This new method is a discontinuous finite element algorithm, which is parameter free, symmetric, and absolutely stable. Furthermore, through the Schur-complement technique, an effective implementation of the weak Galerkin is developed a linear system involving unknowns only associated with element boundaries. In this talk, several numerical applications of weak Galerkin methods will be discussed.

Lin Mu
Michigan State State University
nmu@math.msu.edu

MS265

BPX Preconditioner for Nonstandard Finite Element Methods for Diffusion Problems

We propose and analyze an optimal preconditioner for a general linear symmetric positive definite (SPD) system by following the basic idea of the well-known BPX framework. The SPD system arises from a large number of nonstandard finite element methods for diffusion problems, including the well-known hybridized Raviart-Thomas and Brezzi-Douglas-Marini mixed element methods, the hybridized discontinuous Galerkin method, the Weak Galerkin method, and the nonconforming Crouzeix-Raviart element method. We prove that the presented preconditioner is optimal, in the sense that the condition number of the preconditioned system is independent of the mesh size. Numerical experiments provided confirm the theoretical result.

Xiaoping Xie
Sichuan University, China
xpxie@scu.edu.cn

MS265

Recent Development of Weak Galerkin Methods

Newly developed weak Galerkin finite element methods will be introduced for solving partial differential equations. Weak Galerkin methods have the flexibility of employing discontinuous elements and share the simple formulations of continuous finite element methods at the same time. The Weak Galerkin method is an extension of the standard Galerkin finite element method where classical derivatives were substituted by weakly defined derivatives on functions with discontinuity. Recent development of weak Galerkin methods will be discussed in the presentation.

Xiu Ye
University of Arkansas, Little Rock
xxye@ualr.edu

MS265

A Divergence-free Weak Galerkin Finite Element

A weak Galerkin finite element is designed so that the computed velocity is divergence-free. The significance of such a method is shown by solving a low-viscosity Stokes problem. The traditional finite elements, weak Galerkin finite elements and discontinous Galerkin finite flements fail to produce a meaningful solution in solving such a test problem.

Shangyou Zhang
University of Delaware
szhang@udel.edu

MS266

Implementation of a Fast Multifrontal Solver Using Randomized HSS Compression

We present an efficient code for solving large sparse linear systems using the multifrontal method with hierarchically semi-separable (HSS) matrices. The low rank compression in HSS limits fill-in and reduces complexity of the solver. The HSS matrices are constructed using randomized sampling and rank-revealing QR. ULV decomposition replaces the traditional dense LU. The factorization acts as solver or preconditioner. Shared and distributed memory parallel results are presented for a range of applications.

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

Francois-Henry Rouet
Lawrence Berkeley National Laboratory
fhrouet@lbl.gov

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

MS266

Deterministic and Randomized CUR and Nystrom Approximations

The goal is to improve the accuracy and efficiency of CUR and Nyström approximations, by exploiting traditional matrix decompositions. To this end we establish connections between CUR approximations and LU decompositions, and derive conditions for the optimality of CUR decompositions. Furthermore, we express the standard Nyströ approximation as an incomplete Cholesky decomposition, the modified Nyström approximation as a CUR approximation, and derive residual bounds for deterministic approximations based on rank revealing QR decompositions. The results for the deterministic approximations are used to

guide and calibrate the randomized algorithms, and we derive new bounds for uniform and leverage score sampling.

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

MS266

Fast Generation of Random Orthogonal Matrices

Random orthogonal matrices have a wide variety of applications. They are used in the generation of various kinds of random matrices and random matrix polynomials. The random orthogonal matrix (ROM) simulation method uses random orthogonal matrices to generate multivariate random samples with the same mean and covariance as an observed sample. The natural distribution over the space of orthogonal matrices is the Haar distribution. One way to generate a random orthogonal matrix from the Haar distribution is to generate a random matrix A with elements from the standard normal distribution and compute its QR factorization $A = QR$, where R is chosen to have nonnegative diagonal elements; the orthogonal factor Q is then the required matrix. Stewart develops a more efficient algorithm that directly generates an $n \times n$ orthogonal matrix from the Haar distribution as a product of Householder transformations built from Householder vectors of dimensions $1, 2, \ldots, n-1$ chosen from the standard normal distribution. The goal of this work is to design an algorithm that reduces the computational cost of Stewart's algorithm by giving up the property that Q is exactly Haar distributed. We generate orthogonal random Q using k Householder transformations. We will argue that for the purpose of test matrix generation we can take k much less than the matrix dimension and still obtain an acceptable matrix. We will give performance results on a variety of architectures to show the benefits of the new algorithm.

Amal Khabou
NRIA Saclay, 4 Rue Jacques Monod 91893
amal.khabou@manchester.ac.uk

Francoise Tisseur
University of Manchester
Department of Mathematics
ftisseur@maths.manchester.ac.uk

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

MS266

Performance of Computing Low-Rank Approximation on Hybrid CPU/GPU Architectures

Low-rank matrix approximations play important roles in many statistical, scientific, and engineering applications. In this talk, we study the performance of various algorithms (including random projection/sampling) to compute such approximation of large sparse matrices on a hybrid CPU/GPU architecture. Ultimately, we would like to develop robust, efficient, and flexible software that computes such approximation for a wide range of applications. One objective for this talk is to seek inputs from such application experts.

Ichitaro Yamazaki
UTK
ic.yamazaki@gmail.com

Theo Mary
Universite de Toulouse, INPT ENSEEIHT
theo.mary@etu.enseeiht.fr

Jakub Kurzak, Stanimire Tomov
University of Tennessee, Knoxville
kurzak@icl.utk.edu, tomov@icl.utk.edu

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@cs.utk.edu

MS267

Optimized Reduced Order Modeling and Data Assimilation for Hydrodynamics with Large Time Step Observations

This presentation focuses on the theoretical development of numerical methods for problems originating from meteorology and application of two complementary methods of decomposing the flow field from experimental measurements into coherent structures namely: the Proper Orthogonal Decomposition (POD) and the Dynamic Mode Decomposition (DMD). Additionally, we aim to apply the four-dimensional variational approach of data assimilation (4D-Var) to seek for the model coefficients such that the derived ROM will inherit good dynamical properties.

Diana Bistrian
Dept. of Electrical Engineering and Industrial Information
University "Politehnica" of Timisoara
diana.bistrian@fih.upt.ro

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

MS267

Model Reduction and Sensor Placement in a Feedback Flow Control Problem

We consider the flow control problem of stabilizing the wake behind a circular cylinder through the cylinder's rotation. Model reduction methods are applied for computing the optimal linear feedback control as well as for choosing the location of a number of velocity measurements that can be used to practically implement the control law. Numerical results demonstrate the effectiveness of the controller reduction when compared to the full-state linear feedback control.

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

Serkan Gugercin
Virginia Tech.
Department of Mathematics

gugercin@math.vt.edu

Lizette Zietsman
Virginia Tech
Interdisciplinary Center for Applied Mathematics
lzietsma@vt.edu

MS267

Goal-Based Rom Adjoint for Optimal Sensor Locations and Data Assimilation

An goal-based reduced order modelling (ROM) adjoint approach is developed for optimising sensor locations. An adjoint (or sensitivity) based error measure is formulated which measures the error contribution of each solution variable to an overall goal (defined as a measure of what is deemed important in a problem). It provides importance maps for determining where best to place the monitoring devices. The expensive data resources will be efficiently used for data assimilation.

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

Zhizhao Che
Chemical Engineering, Imperial College London
z.che@imperial.ac.uk

Andrew G. Buchan, Pavlidis Dimitrios
Department of Earth Science & Engineering
Imperial College
andrew.buchan@imperial.ac.uk,
dimitrios.pavlidis@imperial.ac.uk

Dunhui Xiao
Department of Earth Science and Engineering
Imperial Colleg London
dh.xiao@imperial.ac.uk

MS267

Reduced Order Modelling (rom) of the Navier-Stokes Equations for 3D Free Surface Flows

This article presents a reduced order modelling method for Navier-Stokes Equations for 3D free surface flows. This work is the first to introduce ROM into 3D free surface flows. This newly ROM reduced CPU time for solving 3D free surface problems by orders of magnitude while keeping a high accuracy. This ROM has significant potential advantages in : interactive use, emergency response, control and uncertainty analysis.

Dunhui Xiao
Department of Earth Science and Engineering
Imperial Colleg London
dh.xiao@imperial.ac.uk

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

Andrew G. Buchan
Department of Earth Science & Engineering
Imperial College
andrew.buchan@imperial.ac.uk

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

MS268

Higher Order Finite Volume Approximations of the Inviscid Primitive Equations in a Complex Domain

We construct the cell-centered finite volume discretization of the two-dimensional inviscid primitive equations in a domain with topography. To compute the numerical fluxes, the so-called (first order) upwind scheme and the (second order) central-upwind scheme are introduced. Numerical simulations verify that the upwind and central-upwind are robust schemes regardless of the shape or size of the topography.

Gung-Min Gie
Indiana University
USA
gungmin.gie@louisville.edu

Arthur Bousquet, YoungJoon Hong
Indiana University
arthbous@indiana.edu, hongy@umail.iu.edu

Roger M. Temam
Inst. for Scientific Comput. and Appl. Math.
Indiana University
temam@indiana.edu

MS268

A New Atmospheric Dynamic Core using 4th Order Flux Reconstruction Method with WENO Limiting

In this talk, we present a dynamic core for compressible non-hydrostatic atmosphere by using a newly devised high-resolution scheme, so-called FR4-CD-WENO (4th order Flux-Reconstruction with Constrained Derivative using WENO limiting). Compared to the existing WENO-DG paradigm, the FR4-CD-WENO method shows superiorities in at least three aspects. 1) It makes use of the sub-cell information of the local reconstruction instead of the cell-averaged values, which results in a compact stencil for reconstruction, 2) It is more accurate than the WENO-DG scheme using the same degrees of freedom, and 3) It is algorithmically simpler and more computationally efficient. We have developed the FR4-CD-WENO dynamic core for atmospheric flows and shown it as a very promising framework for high-performance atmospheric models.

Xingliang Li
Center of Numerical Weather Predication,
China Meteorological Administration
lixl@ou.edu

Ziyao Sun
Department of Energy Sciences
Tokyo Institute of Technology
sun.z.ab@m.titech.ac.jp

Feng Xiao
Tokyo Institute of Technology
xiao@es.titech.ac.jp

Chungang Chen
Xi'an Jiaotong Unversity
cgchen@mail.xjtu.edu.cn

Xueshun Shen
Center of Numerical Weather Predication,
China Meteorological Administration
shenxs@cams.cma.gov.cn

Ming Xue
Center for Analysis and Prediction of Storms
University of Oklahoma
mxue@ou.edu

MS268

Numerical Weather Prediction in Two Dimensions with Topography, using a Finite Volume Method

We aim to study a finite volume scheme to solve the two dimensional inviscid primitive equations of the atmosphere with humidity and saturation, in presence of topography and subject to physically plausible boundary conditions to the system of equations. The equations are a nonlinear system of equations close to the Euler equations (the inviscid Primitive Equations), which are coupled with the equation for the content of water vapor. These equations form a nonlinear system of partial differential equations, with discontinuities due to the change of phase. A version of the projection method is introduced to enforce the compatibility condition on the horizontal velocity field, which comes from the boundary conditions. The resulting scheme allows for a significant reduction of the errors near the topography when compared to more standard finite volume schemes. In the numerical simulations, we first present the convergence results that are satisfied by the solutions simulated by our scheme when compared to particular analytic solutions. We then report on numerical experiments using realistic parameters. Finally, the effects of a random small-scale forcing on the velocity equation is numerically investigated. The numerical results show that such a forcing is responsible for recurrent large-scale patterns to emerge in the temperature and velocity fields.

Roger M. Temam
Inst. for Scientific Comput. and Appl. Math.
Indiana University
temam@indiana.edu

Arthur Bousquet
Indiana University
arthbous@indiana.edu

Mickael Chekroun
University of California, Los Angeles
Department of Atmospheric and Oceanic Sciences
mchekroun@atmos.ucla.edu

YoungJoon Hong
Indiana University
hongy@umail.iu.edu

Joseph J. Tribbia
Nat'l Center for Atmospheric
Research
tribbia@ncar.ucar.edu

MS268

Title Not Available at Time of Publication

Abstract not available at time of publication.

Yau Shu Wong
University of Alberta
yauwong@ualberta.ca

MS269

Towards Efficient N-x Contingency Selection Using Group Betweenness Centrality

The goal of N - x contingency selection is to pick a subset of critical cases to assess their potential to initiate a severe crippling of an electric power grid. Even for a moderate sized system there can be an overwhelmingly large number of contingency cases that need to be studied. The number grows exponentially with x. This combinatorial explosion renders any exhaustive search strategy computationally infeasible, even for small to medium sized systems. We propose a novel method for N - x contingency selection for x ¿ 1 using group betweenness centrality and show that computation can be relatively decoupled from the problem size. Thus, making contingency analysis feasible for large systems with x ¿ 1. Consequently, it may be that N - x (for x ¿ 1) contingency selection can be effectively deployed despite the combinatorial explosion of the number of potential N - x contingencies.

Mahantesh Halappanavar
Pacific Northwest National Laboratory
hala@pnnl.gov

Yousu Chen
Pacific Northwest National Lab
yousu.chen@pnnl.gov

Zhenyu Huang
Pacific Northwest National Laboratory
zhenyu.huang@pnnl.gov

MS269

Computational Study of Security-Constrained Economic Dispatch with Multi-Stage Rescheduling

Security-constrained economic dispatch with multiple stages of rescheduling gives rise to a linear program that is not solvable by traditional LP methods due to its large size. We devise a series of algorithmic enhancements based on the Benders' decomposition method to ameliorate the computational difficulty. We also propose a set of online measures to correct infeasibility encountered in the solution process. The approach, coded in GAMS, is able to

process large network cases.

Yanchao Liu
University of Wisconsin
yliu67@wisc.edu

Michael C. Ferris
University of Wisconsin
Department of Computer Science
ferris@cs.wisc.edu

Feng Zhao
ISO New England
fzhao@iso-ne.com

MS269

Moment-Based Relaxations of the Optimal Power Flow Problem

The optimal power flow (OPF) problem seeks an optimal operating point for an electric power system subject to constraints from both the network physics and engineering limits. We present a hierarchy of moment-based convex relaxations that globally solve many non-convex OPF problems for which existing relaxations fail. This talk demonstrates the capabilities of the moment relaxations using illustrations of the feasible spaces of small example OPF problems. We also describe recent work on computational improvements for the moment relaxations, including the exploitation of sparsity, which enable global solution of larger OPF problems.

Daniel Molzahn, Ian Hiskens
University of Michigan
dan.molzahn@gmail.com, hiskens@umich.edu

MS269

Decomposition Algorithms for Transmission and Generation Investment Planning

Stochastic transmission and generation expansion planning models are receiving increasing attention among researchers today. They are being used to explicitly model uncertainties that result from the increasing penetration of renewable energy technologies, as well as from long-term market and regulatory conditions. However, existing commercial planning tools still lack stochastic capabilities. We propose a two-stage investment-planning model that takes into account the aforementioned uncertainties, and we describe a scalable decomposition algorithm to solve real-sized problems. An application of our algorithm is illustrated using a 240-bus network representation of the Western Electricity Coordinating Council. We discuss its performance when implemented in both the Red Mesa/Sky supercomputer and a commodity multi-core workstation.

Francisco Munoz
Sandia National Laboratories
fdmunoz@sandia.gov

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

MS270

QuickSched - Using Tasks for Massively Parallel Hybrid Shared/distributed Memory Computing

Task-based parallelism is a well-known paradigm for efficient and scalable shared-memory parallel computing. In this talk, we will present QuickSched, a library for task-based parallelism which extends the usual model of tasks and dependencies by conflicts between tasks, and by specifying explicitly which resources are used by which task. We show how this approach allows us to provide good scaling for problems with complex task hierarchies. We also show how QuickSched extends task-based parallelism to CUDA-based GPUs. The task/resource decomposition is also useful for hybrid shared/distributed-memory computations, e.g. on clusters of multicores. The task graph can be used to equitably partition the work involved in a computation over a set of distributed-memory nodes, as opposed to simply partitioning the data. Furthermore, the task-based system can be used to implement fully asynchronous communication between distributed-memory nodes on top of MPI.

Pedro Gonnet
University of Oxford
gonnet@maths.ox.ac.uk

MS270

Parallelization Techniques for Tsunami Simulation Using Space-Fillig-Curve Orders

Samo(oa)^2 is a parallel code for element-oriented numerical solution of PDEs based on space-filling curve traversals. Heuristics and methods for adaptive, parallel grids are gained by sequential element orders induced by the Sierpinski curve, which provide an efficient approach that we use as basis for an application framework. We implemented a model that simulates tsunamis originating from time-dependent sea-floor displacements computed from an earthquake simulation (Galvez et al.). Parallel performance was studied with hybrid OpenMP/MPI parallelization and several load balancing techniques that exploit inherent properties of the space-filling curve.

Michael Bader, Oliver Meister
Technische Universität München
bader@in.tum.de, meistero@in.tum.de

MS270

A Patchwork Family - Task Distribution Patterns for Shallow Water Equations on Patch-structured AMR Grids

Spacetrees holding regular Cartesian patches are a powerful formalism to describe dynamically adaptive Cartesian grids that scale on manycores. In previous work, we successfully used this formalism to solve shallow water equations. The choice of a proper block size here is delicate. While large block sizes foster loop parallelism and vectorisation, they restrict the adaptivity's granularity. They increase the total memory footprint and lower the numerical accuracy per invested byte. Though small patches exhibit a high inter-block concurrency they are surprisingly outperformed by huge blocks processed by a plain parallel-for on Xeon Phis. This insight from our previous work has motivated us to introduce algorithms that automatically detect assemblies of patches that can be fused into one big patch and then to replace the assemblies by such a regular data structure. The present talk rolls back to a small-patch formulation and studies the patch performance, in particular the concurrency, en detail. We propose to formalise inter-block

parallelism as tasks each representing one patch update. These tasks then are fired into a scheduler to be deployed among the cores. We study the arising distribution patterns and try to put the best-case data and task assignment into relation to the data access patterns of the fused patches.

Tobias Weinzierl
Durham University
tobias.weinzierl@durham.ac.uk

MS270

Understanding Tsunami and Hurricane Deposits with a Mess-Scale Model for Sediment Dynamics

A meaningful event history for storms and tsunamis does not only include the number of events, but also contains information of the magnitude of each event. Conventional methods to simulate the sediment dynamics in tsunamis and storms have so far failed to reproduce the fine detail that real deposits feature. We employ a meso-scale approach to do this with some success. Our model runs on HPC platforms containing multi-CPU and multi-GPU achitectures.

Robert Weiss, Wei Cheng
Virginia Tech
Department of Geosciences
weiszr@vt.edu, chengwei@vt.edu

MS271

Block Adaptive MHD Simulations for Solar Coronal Dynamics

I will present the open source MPI-AMRVAC simulation toolkit [Keppens et al., 2012, JCP 231, 718-744], with a focus on solar physical applications modeled by its magnetohydrodynamic module. Spatial discretizations available cover standard shock-capturing finite volume algorithms, but also extensions to conservative high-order finite difference schemes, both employing many flavors of limited reconstruction strategies. Multi-step explicit time stepping includes strong stability preserving high order Runge-Kutta steppers to obtain stable evolutions in multi-dimensional applications realizing up to fourth order accuracy in space and time. The parallel scaling of the code is discussed and we obtain excellent weak scaling up to 30000 processors allowing to exploit modern peta-scale infrastructure. Solar physics applications target the formation of flux rope topologies through boundary-driven shearing of magnetic arcades, following the in situ condensation of prominences in radiatively controlled evolutions of arcades and fluxropes, and the enigmatic phenomenon of coronal rain, where small-scale condensations repeatedly form and rain down in thermodynamically structured magnetic arcades.

Rony Keppens
KU Leuven
Centre for Mathematical Plasma-Astrophysics
rony.keppens@wis.kuleuven.be

MS271

Globally Divergence-Free Projection Methods for Ideal Magnetohydrodynamics

It has been long established in the literature that controlling magnetic field divergence errors in compressible magnetohydrodynamic (MHD) simulations is necessary for numerical stability. Many approaches exist to achieve either approximately or exactly divergence-free discrete magnetic field values, including hyperbolic divergence-cleaning, constrained transport, and elliptic projection techniques. In this work we develop a high-order discontinuous Galerkin version of the elliptic projection method that produces a globally divergence-free magnetic field. After developing it, we show how to efficiently implement this method both on Cartesian and unstructured grids. The resulting scheme is applied to several standard test cases.

James A. Rossmanith
Iowa State University
Deparment of Mathematics
rossmani@iastate.edu

MS271

Multi-Fluid Plasma Modeling Through the Collisional Transition Regime

Plasmas consist of charged particles that interact through electromagnetic fields and collisions. Neutral particles when present interact only through collisions. For low collisionality, a multi-species kinetic model must be used. For high collisionality, a fluid model with lower dimensionality can be used. Multi-fluid plasma models approximate the velocity distribution function by a limited number of moments for each species. Models with higher moments extend the region of validity and relax the assumption of local thermodynamic equilibrium.

Uri Shumlak, Andrew Ho
Aerospace and Energetics Research Program
University of Washington
shumlak@uw.edu, andrewh0@uw.edu

Robert Lilly
University of Washington
rlilly@aa.washington.edu

Sean Miller
Aerospace and Energetics Research Program
University of Washington
smiller3@uw.edu

Noah F. Reddell, Eder Sousa
University of Washington
reddell@uw.edu, sousae@uw.edu

MS271

Positivity-Preserving Weno Schemes with Constrained Transport for Ideal MHD

We will present a novel positivity-preserving flux limiting technique developed for WENO schemes to solve idea magnetohydrodynamic (MHD) equations. There are two main steps in our MHD solver: first updating conservative quantities by WENO scheme with positivity-preserving limiter and then correcting the magnetic field and total energy by a high-order constrained transport approach. Several examples are presented to verify the order of accuracy and to demonstrate the efficiency of positivity-preserving limiter.

Qi Tang
Department of Mathematics
Michigan State University

tangqi@msu.edu

MS272

A High-Order Global Discontinuous Galerkin Non-Hydrostatic Atmospheric Model Using Hevi Time Integration Scheme

Availability of peta-scale computing resources facilitate development of high-resolution non-hydrostatic (NH) atmospheric model at a global scale. Discontinuous Galerkin (DG) method is an ideal candidate for such model development because of its inherent conservation property, geometric flexibility and excellent parallel efficiency. We consider a NH atmospheric model based on DG spatial discretization on cubed-sphere grid. The model uses the *horizontally-explicit and vertical-implicit* (HEVI) operator-split time integration method to overcome the CFL limitation resulting from the small vertical grid-spacing, and is constrained only by the minimum horizontal resolution. The performance of the DG-HEVI model will be evaluated through a suite of benchmark test cases.

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

Lei Bao
University of Colorado at Boulder
Dept. of Applied Mathematics
lei.bao@colorado.edu

MS272

Vertical Discretization of Geophysical Flows with the Hybrid Finite Element Method - Normal Mode and Wave Dispersion Properties

The Hybrid Finite Element Method is a new technique that takes advantage of the arbitrary order of accuracy of the discontinuous Galerkin and Spectral Element methods and provides a combined, natural setting for the implementation of grid staggering. We present a discrete normal mode analysis of this method for the vertical direction using the linearized Euler equations. 2-D simulations will be shown demonstrating atmospheric waves reproduced by the HFEM compared to traditional methods.

Jorge E. Guerra
Department of Land, Air and Water Resources
University of California, Davis
jeguerra@ucdavis.edu

Paul Ullrich
University of California, Davis
paullrich@ucdavis.edu

MS272

Toward Exa-Scale Computing in CAM-SE

Supercomputers are projected to achieve exa-flop scales within the next fire years, providing enough computational power to perform climate simulations with lateral resolutions approaching or exceeding the hydrostatic barrier. To perform accurate simulations at these resolutions, CAM-SE will require new technologies including: nonhydrostatic dynamical solvers, increased vertical accuracy, improved variable-resolution grids, and scale-aware parameterizations. We will discuss the need for these improvements and present our progress to date in their implementation.

David M. Hall
University of Colorado at Boulder
Department of Computer Science
david.hall@colorado.edu

Henry Tufo
NCAR
University of Colorado at Boulder
tufo@cs.colorado.edu

MS272

Tempest: Efficient Computation of Atmospheric Flows Using High-Order Local Discretization Methods

The Tempest Framework composes several numerical methods to easily facilitate intercomparison of non-hydrostatic atmospheric flow calculations on the sphere. This framework includes the implementations of Spectral Elements, Discontinuous Galerkin, Flux Reconstruction, and Hybrid Finite Element methods with the goal of achieving optimal accuracy and efficiency in computing the solution of atmospheric problems. Several advantages of this approach are discussed such as: improved pressure gradient calculation, numerical stability by vertical/horizontal splitting, arbitrary order of accuracy, etc. The local numerical discretization allows for high performance parallel computation and efficient inclusion of parameterizations.

Paul Ullrich
University of California, Davis
paullrich@ucdavis.edu

MS273

Ongoing Developments in BigDFT towards the ab-initio Computation of Resonant States

Siegert's resonant states provide a convenient description of unoccupied electronic states of atoms and molecules, since few (discrete) resonant states carry as much information as many (in principle infinite) continuum states. The complex scaling method, and generalizations thereof, allows one to compute resonant states as particular eigenstates of complex-symmetric Hamiltonians. In my presentation, I will discuss the application of Polizzi's FEAST eigensolver to the iterative, matrix-free extraction of resonant states from complex-scaled Kohn-Sham Hamiltonians.

Alessandro Cerioni
Institute of Nanosciences and Cryogeny
CEA - Grenoble
alessandro.cerioni@gmail.com

Luigi Genovese
Institut Nanosciences et Cryogénie, CEA
luigi.genovese@cea.fr

Thierry Deutsch
Institute of Nanosciences and Cryogeny
CEA
thierry.deutsch@cea.fr

Ivan Duchemin
Institute of Nanosciences and Cryogeny
CEA France
ivan.duchemin@cea.fr

Maxime Moriniere
Institute of Nanosciences and Cryogeny
CEA - Grenoble
maxime.moriniere@cea.fr

MS273

Planning the Next Generation of Electronic Structure Codes

Pseudopotential density functional theory implemented in real space provides a means for computing the properties of materials in many different forms, including liquids, solids, and nanoscale structures. Current approaches yield self-consistent solutions for systems with thousands of atoms. However, there remain systems where quantum mechanical accuracy is desired, but scalability proves to be a hindrance. We will present an overview of our work on algorithms for this problem, which performs spectrum slicing in the eigensolver.

James R. Chelikowsky
Institute for Computational Engineering and Sciences
University of Texas at AUstin
jrc@utexas.edu

MS273

A Projected Preconditioned Conjugate Gradient Algorithm for Eigenvalue Calculation

We examine a projected gradient algorithm for computing a relatively large number of lowest eigenvalues of a Hermitian matrix. The algorithm performs fewer Rayleigh-Ritz calculations than some of the existing algorithms, thus has better parallel scalability. It is relatively easy to implement (for example in the Quantum Espresso package). We will discuss a number of practical issues for implementing this algorithm, and demonstrate its performance in Kohn-Sham density functional theory based electronic structure calculations.

Chao Yang, Eugene Vecharynski
Lawrence Berkeley National Lab
cyang@lbl.gov, evecharynski@lbl.gov

John Pask
Lawrence Livemore National Lab
pask1@llnl.gov

MS273

Updating Strategies for Efficient Large-Scale Electronic Structure Calculations

We give an introduction to our recently published Blocked Householder-CholeskyQR algorithm which has been developed for the dense symmetric eigensolver ELPA, whereat the QR-decomposition of tall and skinny matrices represents an important substep. Further, we show the benefits of this algorithm on today's HPC systems in terms of parallel efficiency. Finally, we give an outlook on modifying the classical ELPA solver,e.g. for updating the eigenspace information in case of small changes.

Roland Wittmann, Thomas K. Huckle
Institut fuer Informatik
Technische Universitaet Muenchen
wittmanr@in.tum.de, huckle@in.tum.de

MS274

Scalable Algorithms for Optimal Control of Systems Governed by PDEs under Uncertainty

We focus on optimal control of systems governed by PDEs with random coefficient functions. We seek controls that minimize an objective function that incorporates mean and variance of the control objective. To enable applications to problems with infinite-dimensional (high-dimensional when discretized) parameters, we consider linearization of the control objective at the mean of the uncertain parameters. As application, we consider optimal control of a porous medium flow model with a random permeability field.

Alen Alexanderian
University of Texas at Austin
alen@ices.utexas.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS274

A Scalable Compositional Approach to Uncertainty Quantification for the Optimization under Uncertainty of Multi-physics Systems

Complex multi-physics systems often exhibit a great deal of uncertainty due to unexpected interactions. To ensure a sufficiently reliable and robust system, rigorous design under uncertainty must be carried out with models capable of precisely representing multi-physics interactions. We propose a compositional approach to uncertainty quantification based on importance sampling that enables an offline/online approach to optimization under uncertainty with high-fidelity physics-based models. Our approach is demonstrated on the design of a gas turbine blade.

Doug Allaire
Texas A&M
dallaire@tamu.edu

Sergio Amaral
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
samaral@mit.edu

Kaiyu Li
Texas A&M
kaiyuli@tamu.edu

Karen E. Willcox
Massachusetts Institute of Technology

kwillcox@MIT.EDU

MS274

Stochastic Reduced-Order Models in Optimization and Inverse Problems

This work presents a novel approach for inverse problems in the presence of uncertainty using stochastic reduced order models (SROMs). Given the statistics of an uncertain observed quantity, the statistics of unknown system parameters are estimated through the solution of a stochastic optimization problem. The proposed framework is based on the representation of a random quantity using a SROM an optimal discrete approximation to a continuous random element that permits efficient and non-intrusive stochastic computations.

Wilkins Aquino
Associate Professor
Dept. of Civil and Environmental Engineering, Duke Univ.
wa20@duke.edu

James Warner
Nasa Langley Research Center
jw666@cornell.edu

Mircea Grigoriu
Cornell University
mdg12@cornell.edu

MS274

Integration of Approximate Schur Preconditioners and SQP Algorithms for Nonlinear PDE Optimization under Uncertainty

We study the integration of Schur preconditioners for optimality systems and matrix-free composite-step sequential quadratic programming (SQP) algorithms in the context of partial differential equation (PDE) constrained optimization under uncertainty. Our approach extends the recently introduced optimal solvers for PDE-constrained optimization to a wide range of problems in optimal engineering design, including those governed by PDEs with random inputs. We present numerical examples in thermal-fluid control, acoustic design and topology optimization.

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

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

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS275

Recursive Multilevel Approximate Inverse-Based Preconditioning

We present an algebraic recursive multilevel approximate inverse preconditioner based on a distributed Schur complement formulation, for solving systems of linear equations. The proposed solver uses recursive combinatorial algorithms to preprocess the matrix structure and to maximize sparsity in the approximate inverse factors. We describe the preordering, analysis and factorization phase of the construction, we assess their performance, consider strategies for automatic parameter selection, and we discuss applications to preconditioning least-squares problems.

Yiming Bu
University of Groningen
y.bu@rug.nl

Bruno Carpentieri
Insitute for Mathematics and Computation
University of Gronengin
bcarpentieri@gmail.com

MS275

A Recursive Multilevel Approximate Inverse-Based Preconditioner for Solving General Linear Systems

Abstract not available at time of publication.

Bruno Carpentieri
Insitute for Mathematics and Computation
University of Gronengin
bcarpentieri@gmail.com

MS275

Krylov Subspace Methods Preconditioned by Inner Iterations for Rank-Deficient Least Squares Problems

Inner-iteration preconditioning performed by several iterations of stationary iterative methods including the SOR-type method is applied to Krylov subspace methods such as the CG-, MINRES-, and GMRES-type methods for solving rank-deficient least squares problems. We present theoretical justifications for using these methods. Numerical experiments on large sparse problems show that the proposed methods are more robust and efficient compared to previous methods for some problems.

Keiichi Morikuni
Instiute of Computer Science
Academy of Sciences of the Czech Republic
morikuni@cs.cas.cz

Miroslav Rozloznik
Institute of Computer Science
Czech Academy of Sciences
miro@cs.cas.cz

Ken Hayami
National Institute of Informatics
hayami@nii.ac.jp

MS275

Global Adaptive Dropping in Incomplete Factorizations

Incomplete factorizations represent an important component in solving large sparse systems of equations and linear least squares problems. A new approach for SPD matrices is described. It is based on the factorized approximate inverse that links together the inverse and direct factors. The strategy is motivated by the floating-point behavior of the

decomposition that implies an algorithm where dropping of the incomplete decomposition mimicks propagation of the rounding errors. Numerical experiments demonstrate its efficiency.

Miroslav Tuma
Academy of Sciences of the Czech Republic
Institute of Computer Science
tuma@cs.cas.cz

Jiri Kopal
Technical University in Liberec
jiri.kopal@tul.cz

Miro Rozloznik
Czech Academy of Sciences
Prague, Czech Republic
miro@cs.cas.cz

MS276

Model-Based Sketching and Recovery with Expanders

Linear sketching and recovery of sparse vectors with randomly constructed sparse matrices has numerous applications in several areas, including compressive sensing, data stream computing, graph sketching, and combinatorial group testing. This work considers the same problem with the added twist that the sparse coefficients of the unknown vector exhibit further correlations as determined by a known sparsity model. We prove that exploiting model-based sparsity in recovery provably reduces the sketch size without sacrificing recovery quality. We also present the model-expander iterative hard thresholding algorithm for efficient recovery of model sparse signals from linear sketches.

Luca Baldassarre
École Polytechnique Fédérale de Lausanne
luca.baldassarre@epfl.ch

Volkan Cevher
EPFL
volkan.cevher@epfl.ch

MS276

Performance Limits of Ideal Decoders in Linear Inverse Problems

In this talk, we focus on the fundamental performance limits that can be expected from an ideal decoder given a general model in a linear inverse problem. We link the existence of an instance optimal decoder and the Null Space Property of the measurement operator in very general cases. We prove that in this general setting, the lower-RIP yields instance optimality with an operator-dependent norm called the M-norm while the upper-RIP allows to upper bound this M-norm by an operator-independent norm.

Anthony Bourrier
IRISA
anthony.bourrier@gipsa-lab.fr

MS276

Statistical Methods in Compressive Sensing: Theory and Experiment

This talk is concerned with compressive sensing of signals drawn from a Gaussian mixture model (GMM) with sparse precision matrices. Previous work has shown: (i) a signal drawn from a given GMM can be perfectly reconstructed from r noise-free measurements if the (dominant) rank of each covariance matrix is less than r; (ii) a sparse Gaussian graphical model can be efficiently estimated from fully-observed training signals using graphical lasso. This paper addresses a problem more challenging than both (i) and (ii), by assuming that the GMM is unknown and each signal is only partially observed through incomplete linear measurements. Under these challenging assumptions, we develop a hierarchical Bayesian method to simultaneously estimate the GMM and recover the signals using solely the incomplete measurements and a Bayesian shrinkage prior that promotes sparsity of the Gaussian precision matrices. In addition, we provide theoretical performance bounds to relate the reconstruction error to the number of signals for which measurements are available, the sparsity level of precision matrices, and the incompleteness of measurements. The proposed method is demonstrated extensively on compressive sensing of imagery and video, and the results with simulated and hardware-acquired real measurements show significant performance improvement over state-of-the-art methods.

Larry Carin
Electrical & Computer Engineering
Duke University
lcarin@duke.edu

MS276

Practical Compressed Sensing: On Asymptotic Structure

Compressed Sensing allows recovery from undersampled data, a desirable feat in many applications, e.g. MRI, microscopy, tomography, imaging, interferometry etc. However, a wide gap exists between practice and traditional CS theory as some of its principles such as sparsity and incoherence are unsuitable. This talk shows how new CS principles, namely asymptotic sparsity, asymptotic incoherence and multilevel sampling, provide a better fit for practical problems and help to better understand underlying phenomena and significantly improve results in real-world applications.

Bogdan Roman
University of Cambridge
abr28@cam.ac.uk

MS277

Optimal Energy Conserving Local Discontinuous Galerkin Methods for Second-Order Wave Equation in Heterogeneous Media

Abstract not available at time of publication.

Ching-Shan Chou
Department of Mathematics
Ohio State University
chou@math.ohio-state.edu

MS277

Performance Analysis of High-Order Discontinuous Galerkin Methods for First and Second Order For-

mulation of the Wave Equation

Abstract not available at time of publication.

Julien Diaz
Team-Project Magique-3D
INRIA Bordeaux Sud-Ouest
julien.diaz@inria.fr

MS277

The Double Absorbing Boundary Formulation of Complete Radiation Boundary Conditions

Complete radiation boundary conditions are local sequences implementing optimal rational approximations to the exact time-domain DtN map. Our original implementation of CRBCs was for first order systems and was not directly generalizable to second order formulations. The double absorbing boundary method provides an easy alternative formulation which can be directly applied to second order equations. We will outline the stability theory of double absorbing boundaries, and demonstrate implementations for a variety of volume discretizations.

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

MS277

Second-Order Wave Equation with Uncertain Parameters: Analysis and Computation

Abstract not available at time of publication.

Mohammed Motamed
Department of Mathematics and Statistics
The University of New Mexico
motamed@math.unm.edu

MS278

Adaptive Time Domain Boundary Element Methods (TD-BEM) for Scattering Problems

We investigate a generalized version of the MOT method for TD-BEM, which allows variable time-steps in an adaptive algorithm. The residual error estimator is based on new regularity results. We also discuss numerical quadrature schemes for the evaluation of the time-dependent Galerkin elements. Our theoretical results are underlined by several numerical examples.

Matthias Maischak, Matthias Gläfke
Brunel University
matthias.maischak@brunel.ac.uk, m_glaefke@hotmail.com

MS278

Convolution Quadrature Discretization of Volume Integral Equations

Time domain scattering from an inhomogeneous penetrable medium can be formulated as a time domain volume integral equation. This equation can be discretized in time using convolution quadrature. Using a Fourier basis in space, the resulting integral operators can be diagonalized to allow fast operator evaluation. The fully discrete system can then be solved by marching on in time and an iterative two grid procedure. We provide an analysis of this approach and some numerical results.

Peter B. Monk
Department of Mathematical Sciences
University of Delaware
monk@math.udel.edu

MS278

An Exponentially Convergent Convolution Quadrature Method for Time-Domain Boundary Integral Equations

The Convolution Quadrature method is a recent method to solve time-domain wave problems with Boundary Element Methods. We present new results on the exponentially convergence of the CQ solution to the exact solution of the underlying time-stepping scheme. These results rely on the analyticity of the frequency-domain solution and on the location of the resonant poles. We study the influence on the convergence of the scheme used for the time discretisation, the contour involved in the inverse transform, and the number of frequency problems solved.

Nicolas Salles, Timo Betcke
University College London
n.salles@ucl.ac.uk, t.betcke@ucl.ac.uk

MS278

Fast Galerkin Method for Parabolic Space-Time Boundary Integral Equations

Layer potentials of the heat equation are coercive in appropriate anisotropic Sobolev spaces. This fact implies stability and error estimates for Galerkin methods. The resulting linear systems are block-lower triangular and can be solved by block-forward elimination. To handle the cost of dense matrix calculation a space-time version of the fast multipole method can be used, which allows the computation a matrix vector product with nearly optimal cost. Further, if the space-time meshwidths satisfy $h_t \leq Ch_x^2$ the conditioning of the linear system in each time step does not grow with mesh refinements. We will also discuss the application of the methodology to three-dimensional transient Stokes flow. Perhaps surprisingly, this is not straight forward, because of different properties of the fundamental solutions of the heat and Stokes equations.

Johannes Tausch
Southern Methodist University
Department of Mathematics
tausch@mail.smu.edu

MS279

Overview Co-Design at the DOE NNSA Trilabs

Los Alamos, Lawrence Livermore, and Sandia National laboratories, collectively know as the Department of Energy NNSA Trilabs, have been conducting a variety of codesign studies. In this presentation we give an overview of a recently completed trilab milestone that focused on an exploration of the capabilities and characteristics of emerging and expected future architectures.

Richard Barrett
Sandia National Laboratories

rfbarre@sandia.gov

MS279

The ∇-Nabla Time-Composite Approach for Multi-Physics Applications Productivity

The numerical-analysis specific ∇ language provides a new approach for integrating next generation multi-physics large-scale scientific applications. Reflection and hierarchical-logical-time composition are combined to demonstrate productive practices and abilities for agile production at extreme-scale. The gain in abstraction with ∇ during design improves portability, allowing composable software integration for continually changing hardware architectures. This presentation will illustrate these new possibilities on several proxies to build a multi-physics application within the above-cited codesign space.

Jean-Sylvain Camier
CEA, DAM, DIF
Jean-Sylvain.Camier@cea.fr

MS279

Multi-Material ALE in the Blast Code

The goal of the BLAST project is to develop a new multi-physics ALE code based on higher-order finite elements. In this talk I will explain how we handle multi-material mixed cells during the Lagrangian and remap phases. The main topics will include the application of higher-order finite elements to multi-material closure models, remap approaches that preserve monotonicity, conservation and synchronization between different fields and materials, elimination of artificial mixing during remap.

Vladimir Tomov
Lawrence Livermore National Laboratory
tomov2@llnl.gov

MS279

Co-Design Studies Using Mini-Multifluid-Ppm

Abstract not available at time of publication.

Paul R. Woodward
Laboratory for Computational Science and Engineering
University of Minnesota
paul@lcse.umn.edu

MS280

Modeling Across the Undergraduate Curriculum

We report on the findings of the undergraduate group for the Modeling Across the Curriculum II Workshop. We discuss the curricular gaps between the status quo in academia today and what is needed to meet the challenges of a globally competitive workforce in the 21st Century. Along the lines of the two NRC reports Mathematical Sciences 2025 report and its companion piece Fueling Discovery and Innovation, we discuss several recommendations that came from our group.

Jeffrey Humpherys
Brigham Young University
jeffh@math.byu.edu

MS280

Mathematical Modeling in the Early Grades

The common core state standards for mathematics specifically name modeling with mathematics as a mathematical practice. Applied mathematicians, especially those engaged in mathematical modeling are in a great position to work with mathematics education specialists, teachers, curriculum developers and students to decide how students will enact this standard in the classroom. I will share the conclusions of such a group, describe a new three year project and solicit feedback from the audience about the potential for mathematical modeling in the early grades.

Rachel Levy
Harvey Mudd College
levy@hmc.edu

MS280

Applied and Computational Mathematics at the High School Level

Secondary schools in the US offer their students a varying level of access to rich mathematical modeling topics. Some schools offer entire courses centered on mathematical modeling practices while others offer very little. Our working group has developed a set of recommendations for questions to investigate and tasks to carry out that might help infuse models and modeling across the curriculum. Please join us to explore these ideas further in a constructive conversation.

Katherine Socha
Math for America
ksocha@parkschool.net

Kathleen Fowler
Clarkson University
Department of Mathematics
kfowler@clarkson.edu

MS280

Modeling Across the Curriculum: Introduction and Overview

This talk will begin with background information on the Modeling across the Curriculum, MaC, initiative, an NSF-sponsored SIAM program aimed at advancing modeling and computational applied mathematics throughout the educational spectrum. The main focus will be on the second MaC workshop held in January 2014 and the resulting report and recommendations. The subsequent talks will go more deeply into the three primary themes of the workshops.

Peter R. Turner
Clarkson University
School of Arts & Sciences
pturner@clarkson.edu

MS281

Singular Values and Convex Programming for Power System Synchrophasor Data Management

This talk centers on the efficient processing of the observations of phasor measurement units (PMUs) for power

system monitoring. Different from traditional PMU data analysis that typically analyzes individual PMU channels separately, we propose a spatial-temporal framework of collectively processing measurements of PMUs in electrically close regions. Leveraging the low-dimensionality of PMU data, various data management issues such as missing data recovery and data substitution detection can be addressed by solving computationally efficient convex programs.

Joe Chow, Meng Wang
Renssalaer Polytechnic Institute
chowj@rpi.edu, wangm7@rpi.edu

MS281

Policy-switching Schemes for Power System Protection

Current power system protection schemes have proven useful in order to mitigate several localized contingencies. However, there is still ample heterogeneity and uncertainty in the design of wide-area Remedial Action Schemes (RAS) for effective restraint of large cascading outages. We propose a policy-switching approach to select optimal protection schemes during and after a contingency. This talk describes the application of the method to a benchmarking test case in order to achieve optimal power system performance during emergency situations.

Rich Meier, Jesse Hostetler, Eduardo Cotilla-Sanchez, Alan Fern
Oregon State University
meierr@eecs.oregonstate.edu, hostetlj@onid.orst.edu, ecs@eecs.oregonstate.edu, afern@eecs.oregonstate.edu

MS281

Exploring State Estimation Techniques to Accommodate non-Gaussian Noises

Measurement data have always been a key element of power grid operation and planning. However, the use of measurement data is based on a key assumption of Gaussian noises. Driven by the investment from the American Recovery and Reinvestment Act of 2009, thousands more sensors are being put into the power grid. The rate and volume of the emerging measurement data are hundred times higher and larger. Power grid operation and planning functions are becoming more and more dependent on measurement data. More accurate understanding of the noise properties of the data is critical and can make large impact on the outcome of data applications. This talk will examine real phasor measurements and present latest findings in noise properties as well as potential impact on mathematical algorithms for power grid analysis.

Zhenyu Huang
Pacific Northwest National Laboratory
zhenyu.huang@pnnl.gov

Ning Zhou
Binghamton University
ningzhou@binghamton.edu

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

Shaobu Wang
Pacific Northwest National Laboratory
shaobu.wang@pnnl.gov

MS281

Efficient Algorithms for N-x Contingency Analysis for Power Grids

N-x contingency analysis on power flow evaluates the stability of a power system by simulating the failures of x transmission lines or generators. Currently each N - x contingency analysis is performed independently, but since the number of cases increases exponentially with x, this is computationally impractical. We propose new algorithms for the problem by observing that only a small portion of the system is changed when a component is removed.

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

Yu-Hong Yeung
Purdue University
yyeung@purdue.edu

Mahantesh Halappanavar
Pacific Northwest National Laboratory
hala@pnnl.gov

Zhengyu Huang
PNNL
zhenyu.huang@pnnl.gov

MS282

A Low-dimensional Approximation to the Stochastic Elliptic Interface Problem

In this presentation, we will discuss a numerical method of stochastic elliptic interface problem with random interface. An efficient finite element scheme is proposed to compute the covariance of the solutions by using a low-dimensional approximation of the random input. Error estimate is established and some numerical tests are performed.

Ju Ming
Beijing Computational Science Research Center
jming@csrc.ac.cn

MS282

A Finite Element Method for a Stokes Interface Problem

We present a higher-order finite element method for solving a class of interface problems in two dimensions. The method is based on correction terms added to the right-hand side of the natural method. We prove optimal error estimates of the method on general quasi-uniform meshes in the maximum norms. In addition, we apply the method to a Stokes interface problem obtaining optimal result.

Manuel A. Sanchez-Uribe
Brown University
manuel_sanchez_uribe@brown.edu

MS282

Surfactant Driven Tipstreaming in a Flow Focusing

Geometry

We model a surfactant-mediated tipstreaming in a microfluidic flow focusing geometry. That microfluidic method for production of submicrometer and potentially nanoscale droplets and particles uses the elongational flow along with dissolved surfactant in one of the liquid phases to create strong surfaces tension gradient. The concentration of bulk soluble surfactant was found to significantly effect the mode of formation and size of the emitted droplets. By carefully controlling the surfactant concentration and other flow quantities, droplets can be created that are an order of magnitude or more smaller than the scale of both the device and droplets produced in the absence of surfactant.

Jacek Wrobel
Department of Mathematics, Tulane University
Center for Computational Science
jwrobel@tulane.edu

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

Michael R. Booty
Department of Mathematical Sciences, NJIT
booty@njit.edu

MS282

Immersed Finite Element Methods with Enhanced Stability

In this talk, we present new immersed finite element (IFE) methods for the second-order elliptic interface problems. Comparing with classic IFE schemes using Galerkin formulation, these new IFE methods contain either partial stabilization terms on interface edges or full stabilization on all interior edges. *A priori* error estimates show that these new methods converge optimally in corresponding energy norms. Numerical experiments also indicate that these stabilized IFE methods outperform classic IFEs at vicinity of interfaces.

Xu Zhang
Purdue University
xuzhang@purdue.edu

Tao Lin
Department of Mathematics, Virginia Tech
tlin@math.vt.edu

MS283

Code Generation for Higher Level Spectral Methods with Spiral

The FFT is a ubiquitous tool in signal processing and for solving large-scale PDEs. Many use cases can be built on top of the 1D FFT, which has a very simple mathematical definition. However, extracting maximum performance on big parallel machines requires rethinking of the FFT as highest-level building block. We will investigate Spiral-based code generation for higher-level building blocks like convolution, interpolation, and Greens Function approaches and their implication of FFT-based performance optimization and show the performance potential of this approach.

Franz Franchetti
Department of Electrical and Computer Engineering
Carnegie Mellon University
franzf@ece.cmu.edu

MS283

Numerical Eigenvalue Engine towards Extreme-scale Computing Era

Towards Exa-scale computing Era, we have been studying next generation eigenvalue library, which must exploit high performance, highly parallelism and high portability. EigenExa has been developed and released in August 2013 as a prototype of exa-scale library. EigenExa performs on multi-peta scale system such as K computer and a BlueGene/Q. We investigate key points to be innovated in the heritage stage to 10 to 100 fold-scale complex systems. In this mini-symposium, project overview, our goal and key technologies in exa-scale library will be discussed.

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

Takeshi Fukaya
RIKEN, Japan
Japan
takeshi.fukaya@riken.jp

Yusuke Hirota
RIKEN AICS
yusuke.hirota@riken.jp

Susumu Yamada, Masahiko Machida
Japan Atomic Energy Agency
yamada.susumu@jaea.go.jp,
machida.masahiko@jaea.go.jp

MS283

Automatic Tuning for Parallel FFTs on GPU Clusters

In this talk, we propose an implementation of a parallel fast Fourier transform (FFT) with automatic performance tuning on GPU clusters. Because the parallel FFTs require all-to-all communications, one goal for parallel FFTs on GPU clusters is to minimize the PCI Express transfer time and the MPI communication time. Performance results of FFTs on a GPU cluster are reported.

Daisuke Takahashi
Faculty of Engineering, Information and Systems
University of Tsukuba
daisuke@cs.tsukuba.ac.jp

MS283

Statistical Performance Modeling and Autotuning for Dense QR Factorization in Hybrid CPU-GPU Systems

We investigate how to adaptively and automatically choose the block sizes of a dense QR factorization algorithm to maximize the use of CPU and GPU on the same computing node. The decision is based on statistical surrogate models of performance and an online monitor, which avoids unexpected occasional performance drops. Numerical results suggest that our approaches are efficient and can lead to

near-optimal block sizes.

Ray-Bing Chen
Department of Statistics
National Cheng Kung University
raybingc@gmail.com

Yaohung Tsai
Department of Matheamtics
National Taiwan University
ev.hunter@gmail.com

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

MS284
Multiphase Sub-Surface Flow Using Moose

Abstract not available at time of publication.

Chris Green
University of Melbourne, Australia
chris.green@csiro.au

Jonathan Ennis-King
CSIRO
jonathan.ennis-king@csiro.au

MS284
Low Mach and Two-Phase Flow Modeling with Moose Applications

A 7-equation two-phase flow fluid model has been implemented in the next-generation nuclear reactor safety code RELAP-7 built upon the MOOSE multiphysics framework. The entropy viscosity method, an artificial viscosity technique for hyperbolic conservation laws [Guermond et al., Entropy viscosity method for nonlinear conservation laws, Journal of Comput. Phys. 230 (2011) 4248–4267], has been extended to the low-Mach regimes for single and two-phase flows. The entropy viscosity method is a numerical stabilization technique that is spatial-discretization agnostic and satisfies the minimum entropy principle. The fluid flow governing equations are discretized using standard *continuous* finite elements and an implicit BDF2 time discretization. The resulting system of nonlinear equations are solved using the Jacobian-free Newton Krylov solver of MOOSE. Numerical results will be presented.

Jean C. Ragusa
Department of Nuclear Engineering
Texas A&M University
jean.ragusa@tamu.edu

Marco Delchini
Texas A&M University
delcmo@tamu.edu

Ray A. Berry
Idaho National Laboratory
ray.berry@inl.gov

MS284
Modeling Nuclear Fuel Behavior with BISON

BISON is a modern finite-element based nuclear fuel performance code that solves the coupled thermo-mechanics and unsteady species diffusion equations using the parallel Jacobian-Free Newton-Krylov framework in MOOSE. BISON can be used to investigate computationally large problems, e.g. a full stack of discrete pellets in a fuel rod. We provide a brief overview of the material and behavioral models and discuss successes with preconditioning and solving the large nonlinear system.

Shane Stafford
Idaho National Laboratory
shane.stafford@inl.gov

MS284
Stabilization Methods for High Peclet Number Flows in Heterogeneous Porous Media

A class of reconstructed Discontinuous Galerkin (rDG) methods are developed for fluid dynamics in heterogeneous porous media. Numerical examples demonstrate that the rDG methods are able to maintain stabilization of the solution in the case of high Peclet number flows, while rendering sharp resolution at the vicinity of large gradients or discontinuities, indicating a promising methodology for liquid convection in computational hydrogeology.

Yidong Xia
Idaho National Laboratory
yidong.xia@inl.gov

Hai Huang
Idaho National Laboratory
Idaho National Laboratory
hai.huang@inl.gov

Robert Podgorney
Idaho National Laboratory
robert.podgorney@inl.gov

MS285
Reassessing the Missing Point Estimation Model Order Reduction Method

When applied to nonlinear systems, projection-based model-order reduction requires a second-level approximation to achieve its full potential. This talk introduces a generalization of one such technique, the so-called Missing Point Estimation, which is characterized by its simplicity and generality. A greedy sampling algorithm based on a rigorous a priori error bound is proposed. The application to finite element discretizations will be considered.

Julien Cortial
Sandia National Laboratories (Livermore)
Quantitative Modeling and Analysis
julien.cortial@safran.fr

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

MS285
Reduced Basis Methods for Variational Inequalities

We consider parabolic variational inequalities with different trial and test spaces and a possibly non-coercive bilinear form. Fine discretizations that are needed for such problems resolve in high dimensional problems and in long

computing times. To reduce the dimensionality of these problems,we use the Reduced Basis Method. Error estimators could be obtained by combining Reduced Basis Methods with a space-time formulation of the variational inequality. We provide numerical results for a heat inequality model.

Silke Glas
University of Ulm, Germany
silke.glas@uni-ulm.de

MS285

Parsimonious Data Acquisition for Data-driven Model Reduction

Data-driven model reduction methods require input from an external source, such as snapshots for POD. Obtaining this input can be resource-intensive. This presentation will discuss methods that limit usage of memory and computation time in attempting to meet user-specified error tolerances or basis sizes.

Geoffrey M. Oxberry
Lawrence Livermore National Laboratory
oxberry1@llnl.gov

MS285

An Adaptive Parametrized-Background Data-Weak Approach to State Estimation; Application to Heat Transfer Companion Experiments

We present an adaptive Parametrized-Background Data-Weak (PBDW) formulation for data assimilation problems modeled by parametric PDEs. PBDW combines a prior space, which approximates the solution manifold associated with the parametric PDE, and M experimental observations to provide real-time, in-situ state estimation. The adaptive procedure exploits a novel a posteriori observation-based error estimator to refine the prior space using historical data-assimilation solutions. We illustrate our method through a synthetic example and a physical thermal patch problem.

Tommaso Taddei
MIT
ttaddei@mit.edu

Masayuki Yano, James Penn
Massachusetts Institute of Technology
myano@mit.edu, pennjam@mit.edu

Anthony T. Patera
Massachusetts Institute of Technology
Department of Mechanical Engineering
patera@mit.edu

MS286

Interweaving PFASST and Parallel Multigrid

The parallel full approximation scheme in space and time (PFASST) has been introduced by Emmett and Minion as an iterative method for the parallelization of ordinary differential equations or time-dependent PDEs. On each (time-)level the systems that arise have to be solved to the same accuracy. The usage of lower accuracy on levels with larger time steps is natural. In this talk different strategies for coupling PFASST iterations with multigrid methods are presented.

Matthias Bolten
University of Wuppertal
Department of Mathematics
bolten@math.uni-wuppertal.de

Michael Minion
Stanford University
mlminion@lbl.gov

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich
r.speck@fz-juelich.de

Matthew Emmett
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov

Daniel Ruprecht
Institute of Computational Science
Universita della Svizzera italiana
daniel.ruprecht@usi.ch

MS286

Towards a Multigrid Perspective of MLSDC

Spectral deferred corrections (SDC) are a class of iterative solvers for the collocation formulation of an ODE system. The multi-level extension MLSDC, which constitutes the basis for the time-parallel method PFASST, performs SDC sweeps on a full space-time hierarchy and employs an FAS technique well-known from non-linear multigrid methods. In this talk we analyze SDC's smoothing properties and investigate which concepts of standard multigrid theory translate to MLSDC and how these can be exploited further.

Dieter Moser, Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich
d.moser@fz-juelich.de, r.speck@fz-juelich.de

Matthias Bolten
University of Wuppertal
Department of Mathematics
bolten@math.uni-wuppertal.de

MS286

Parallel in Time Multigrid for Nonlinear Equations

The multigrid reduction in time method (MGRIT) creates a multilevel hierarchy of different temporal discretizations. For nonlinear equations each iteration of the parallel-in-time method requires expensive, nonlinear, spatial solves. Using a Picard method for the nonlinear solver, we investigate several methods for reducing the cost of these solves, including reducing solver accuracy on coarser temporal levels and introducing spatial coarsening on coarse temporal levels.

Ben O'Neill
University of Colorado, Boulder

Ben.Oneill@Colorado.edu

MS286

An Adaptive Spectral Deferred Time Integrator for Vesicle Suspensions

Vesicles are deformable and inextensible capsules, filled with and submerged in a viscous fluid. Their dynamics are governed by hydrodynamic and elastic forces which can be formulated as a system of integro-differential-algebraic equations. I will describe a spectral deferred correction algorithm that we use to construct high-order vesicle simulations. Then, I will show how we estimate the error with only one numerical solution and use this to construct an adaptive time stepping scheme.

Bryan D. Quaife
Institute for Computational Engineering and Sciences
University of Texas at Austin
quaife@ices.utexas.edu

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS287

Maximizing Throughput on a Dragonfly Network

In this talk, I will present our analysis of a 100+ Petaflop/s prototype machine with a dragonfly network, 92,160 high-radix routers and 8.8 million cores. We compare network throughput for various routing strategies, job placement policies, and application communication patterns. Our study is based on a novel model that predicts traffic on individual links for direct, indirect, and adaptive routing strategies. We analyze results for individual communication patterns and some common parallel job workloads.

Abhinav Bhatele
Lawrence Livermore National Laboratory
bhatele@llnl.gov

Nikhil Jain
University of Illinois
nikhil@illinois.edu

Xiang Ni, Laxmikant V Kale
University of Illinois at Urbana-Champaign
xiangni2@illinois.edu, kale@illinois.edu

MS287

Topology Aware Mapping using Graph Models for Exascale Systems

Communication time of parallel applications is limited by various features of the interconnection networks such as latency or bandwidths of the links. Topology aware task mapping methods that place application tasks on processors by exploiting information about the underlying network can help to avoid such limitations. In this work, by using the graph models for representing the network topology and applications communication requirements, we study topology aware task mapping methods to reduce bottlenecks in the applications' communication time.

Mehmet Deveci
The Ohio State University
mdeveci@bmi.osu.edu

Kamer Kaya
CERFACS, Toulouse, France
kamer@cerfacs.fr

Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
umit@bmi.osu.edu

MS287

Locality for Sparse Unstructured Communication Patterns

In high performance computing data centers, job allocation and mapping of the parallel tasks to the allocated nodes should comply with the job communication pattern to reduce the execution time. We discuss the potential of simultaneous allocation and mapping for jobs with sparse unstructured communication, and propose a graph-based algorithm that is based on breadth-first expansion both in node network and in job communication domains. Simulations show that our method outperforms other state-of-the-art approaches.

Ozan Tuncer
Boston University
otuncer@bu.edu

Vitus Leung
Sandia National Laboratories
vjleung@sandia.gov

Ayse Coskun
Boston University
Electrical and Computer Engineering Department
acoskun@bu.edu

MS288

Fluctuating Hydrodynamics Methods for Soft Materials

Many efficient implicit solvent coarse-grained (IS-CG) descriptions have been developed for equilibrium studies by removing the solvent degrees of freedom and treating their contributions implicitly in the free energy of interactions. To study many dynamic responses requires capturing momentum transfer and kinetic effects involving the solvent degrees of freedom. We present fluctuating hydrodynamic methods for extending IS-CG models. We present results for dynamic studies of polymeric materials and lipid bilayer membranes.

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

MS288

Temporal Integrators for Fluctuating Hydrodynamics

We develop temporal integrators for solving Langevin stochastic differential equations (SDEs) that arise in fluctuating hydrodynamics (FHD). These methods add fluctuations to standard second-order deterministic solvers in a way that maintains second-order weak accuracy for linearized FHD. We also construct integrators for integrating the overdamped limit of systems of equations with a fast and slow variable in the limit of inifinite separation of the

fast and slow timescales. We illustrate these integrators on applications involving particles suspended in a fluctuating fluid, as well as the development of giant nonequilibrium fluctuations in diffusively-mixing fluids.

Steven D. Delong
Courant Institute of Mathematical Sciences
New York University
steven.d.delong@gmail.com

Eric Vanden-Eijnden
Courant Institute
New York University
eve2@cims.nyu.edu

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
aleks.donev@nyu.edu

MS288

A Fluctuating Immersed Boundary Method for Brownian Suspensions of Rigid Particles

I will describe how to model Brownian suspensions of passive or active particles and rigid bodies using an immersed boundary (IB) approach. I will first discuss minimally-resolved models in which each suspended spherical particle is represented by a single IB marker [F. Balboa Usabiaga and R. Delgado-Buscalioni and B. E. Griffith and A. Donev, Computer Methods in Applied Mechanics and Engineering, 269:139-172, 2014; and S. Delong, F. Balboa Usabiaga, R. Delgado-Buscalioni, B. E. Griffith and A. Donev, J. Chem. Phys., 140, 134110, 2014]. More complex rigid bodies suspensed in fluid can be represented with different degrees of fidelity by enforcing a rigidity constraint for each partially- or fully-resolved body [B. Kallemov, A. Bhalla, A. Donev, and B. Griffith, in preparation]. Thermal fluctuations and thus Brownian motion can be consistently modeled by including a fluctuating (random) stress in the momentum equation, as dictated by fluctuating hydrodynamics.

Aleksandar Donev
Courant Institute
donev@courant.nyu.edu

Bakytzhan Kallemov
Courant Institute of Mathematical Sciences
NYU
kallemov@gmail.com

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

Steven D. Delong
Courant Institute of Mathematical Sciences
New York University
steven.d.delong@gmail.com

MS288

An Immersed Boundary Method for Rigid Bodies

We develop an immersed-boundary method for fluid-structure coupling at small and moderate Reynolds numbers. This is important in problems involving rigid and semi-rigid structures immersed in a fluid. We couple an immersed-boundary Lagrangian representation of rigid bodies to a finite-volume fluid solver. Our methods (1) Do not employ time splitting and are thus suitable for the steady Stokes (viscous-dominated or low Reynolds number) regime; (2) Strictly enforce the rigidity constraint; and, (3) Ensure fluctuation-dissipation balance in the Brownian regime in the overdamped limit even in the presence of nontrivial boundary conditions.

Bakytzhan Kallemov
Courant Institute of Mathematical Sciences
NYU
kallemov@gmail.com

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
aleks.donev@nyu.edu

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

Amneet Bhalla
Courant Institute of Mathematical Sciences
mail2amneet@gmail.com

MS289

Gaussian Processes in High-Dimensions

Gaussian process regression in high-dimensions is challenging task. Under special circumstances, active subspace methods can effectively deal with high-dimensions. Most of the times, the active subspace is defined in an intuitive, albeit ad hoc manner, using gradient information. Here, we show how traditional machine learning techniques can be used to infer the active subspace with no ad hoc assumptions and without gradient information.

Ilias Bilionis
Purdue University
ebilionis@gmail.com

Nicholas Zabaras
Cornell University
nzabaras@gmail.com

MS289

Numerical Solution for the High-Dimensional Joint Response-Excitation Pdf Evolution Equations

Evolution equations of the joint response-excitation probability density function (REPDF) generalize the existing PDF evolution equations and enable us to compute the PDF of the solution of stochastic systems with random initial condition, coefficient, forcing, involving colored noise. An efficient algorithm by using adaptive discontinuous Galerkin method and probabilistic collocation method has been developed for low dimensional systems. In this talk, we address the high-dimensionality of the REPDF system. We focus on two approximations, namely, the Proper Generalized Decomposition (PGD) involving the separated representation and the ANOVA approximation. Both approaches overcome the curse of dimensionality and can compute the PDFs of extremely high-dimensional systems. Here, we demonstrate the effectiveness of these methods to the Lorenz-96 model and the advection equation.

Heyrim Cho

Brown University
Providence, RI
heyrim_cho@brown.edu

Daniele Venturi
Division of Applied Mathematics
Brown University
daniele_venturi@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS289

High-Dimensional Hierarchical Uncertainty Quantification for Electronic Systems

We present a hierarchical uncertainty quantification approach for complex electronic systems with several high-dimensional subsystems. First, we obtain a sparse surrogate model for each subsystem, and utilize a tensor-train-based algorithm to obtain its orthonormal polynomials and Gauss quadrature points. Next, we treat each subsystem as a single parameter and perform the high-level simulation using stochastic spectral methods. The framework shows 90x speedup over hierarchical Monte Carlo on an MEMS/IC oscillator circuit with 184 random parameters.

Zheng Zhang
MIT
z_zhang@mit.edu

Xiu Yang
Pacific Northwest National Laboratory
xiu.yang@pnnl.gov

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

Ivan Oseledets
Institute of Numerical Mathematics
Russian Academy of Sciences, Moscow
ivan.oseledets@gmail.com

Luca Daniel
M.I.T.
Research Lab in Electronics
luca@mit.edu

MS289

Adaptive Multivariate Interpolation Algorithm on Nested Grids and Its Application to Stochastic Collocation

We propose an adaptive method for high dimensional polynomial interpolation for efficient stochastic collocation. The method utilizes least orthogonal interpolation on unstructured grids. It is based on a greedy algorithm to adaptively induces nested grids and can be highly flexible for practical stochastic computations. Along with numerical examples to demonstrate its effectiveness, we also provide convergence proof.

Xueyu Zhu
University of Utah
xzhu@sci.utah.edu

Yeonjong Shin
Scientific Computing and Imaging Institute
University of Utah
shin@math.utah.edu

Dongbin Xiu
University of Utah
dongbin.xu@utah.edu

MS290

BDDC Domain Decomposition for Weak Galerkin Methods

A Balancing domain decomposition by constraints (BDDC) algorithm is studied for solutions of large sparse linear algebraic systems arising from weak Galerkin discretization of second order elliptic boundary value problems. The condition number for the preconditioned system is estimated and numerical results are provided to confirm the results.

Xuemin Tu
University of Kansas
xuemin@ku,edu

MS290

Innovative Weak Galerkin Finite Element Methods with Application in Fluorescence Tomography

In this talk, I will discuss a new and efficient numerical algorithm by using weak Galerkin (WG) finite element methods for a fourth order elliptic problem arising from Fluorescence Tomography(FT) model. Fluorescence Tomography is an emerging, in vivo non-invasive 3-D imaging technique which reconstructs images that characterize the distribution of molecules that are tagged by fluorophores. Weak second order elliptic operator and its discrete version are introduced for a class of discontinuous functions defined on a finite element partition of the domain consisting of general polygons or polyhedra. An error estimate of optimal order is derived in an H^2-equivalent norm for the WG finite element solutions. Error estimates in the usual L^2 norm are established, yielding optimal order of convergence for all the WG finite element algorithms except the one corresponding to the lowest order (i.e., piecewise quadratic elements). Some numerical experiments are presented to illustrate the efficiency and accuracy of the numerical scheme.

Chunmei Wang
Georgia Institute of Technology
cwang462@math.gatech.edu

MS290

Weak Galerkin Mixed Finite Element Methods for Linear Elasticity Problems

In the talk, I will talk about solving linear elasticity problems by using Weak Galerkin mixed finite element method (WG-MFEM). It is shown that WG-MFEM provides an accurate approximation for both the stress tensor and the displacement field of linear elasticity problems. The numerical experiments will be provided to verify that WG-

MFEM is efficient and reliable in computing.

Yujie Zhang
Oklahoma State University, USA
yjzhang@mathdept.okstate.edu

MS290

A Weak Galerkin Finite Element Scheme for Solving the Brinkman Equations

A weak Galerkin (WG) finite element method for solving the Brinkman equation in two or three dimensional spaces by using polynomials is introduced and analyzed. The WG method is designed by using the generalized functions and their weak derivatives which are defined as distributions. The variational form we considered is based on two gradient operators which is different from the usual gradient-divergence operators. The WG method is highly flexible by allowing the use of discontinuous functions on arbitrary polygons or polyhedra with certain shape regularity. Optimal-order error estimates are established for the corresponding WG finite element solutions in various norms. Some computational results are presented to demonstrate the efficiency of the method.

Ran Zhang
Jilin University, China
zhangran@jlu.edu.cn

MS291

Preconditioning Stochastic Gradient Algorithms with Randomized Linear Algebra

We provide a method for combining SGD (Stochastic Gradient Descent) and RLA (Randomized Linear Algebra) algorithms. This involves reformulating a deterministic regression problem as a stochastic optimization problem that is of the form of an expectation over a nontrivial data-dependent probability distribution. This permits us to combine stochastic approximation (SA) methods and sample average approximation (SAA) methods from optimization theory to develop novel RLA-SGD-hybrid randomized algorithms for these deterministic regression problems.

Michael Mahoney
Stanford University
Applied and Computational Mathematics
mmahoney@stat.berkeley.edu

MS291

Randomized Methods for Accelerating Structured Matrix Computations

Methods based on randomized sampling have over the last several years proven to be powerful tools for computing low-rank approximations to matrices whose singular values exhibit appropriate decay. In this talk, we describe how such techniques can be extended to certain "rank-structured" matrices, for which only certain off-diagonal blocks (not the matrix itself) admit accurate low-rank approximations. Matrices of this type often arise in the construction of $O(N)$ direct solvers for elliptic PDEs.

Gunnar Martinsson
Univ. of Colorado at Boulder
martinss@colorado.edu

MS291

Using Random Butterfly Transformations to Avoid Pivoting in Sparse Direct Methods

Dynamic pivoting prevents parallel sparse direct solvers from achieving high performance and scalability. In this work, we investigate a statistical technique based on Random Butterfly Transformations to avoid pivoting. Previous works showed that this technique is successful in the dense case; here we investigate the sparse case. We will compare this method with the static pivoting and the partial pivoting approaches in various performance metrics, including robustness, sparsity, and runtime in a parallel environment.

Francois-Henry Rouet
Lawrence Berkeley National Laboratory
fhrouet@lbl.gov

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

Marc Baboulin
University of Paris-Sud/INRIA
marc.baboulin@inria.fr

MS292

A Pod Model for Resolving the Angular Dimension of the Boltzmann Transport Equation

A new method using POD in neutral particle transport problems is described. POD is used to represent the angular direction of particle travel when solving the Boltzmann equation for time independent problems. It is based on the method of snapshots which are of the angular flux distributions taken at different instances in space. The method substantially reduces the number of functions required to resolve angular direction, this reduces solving times whilst retaining solution accuracy.

Andrew G. Buchan
Department of Earth Science & Engineering
Imperial College
andrew.buchan@imperial.ac.uk

Atyab Calloo
Imperial College London
atyab.calloo13@imperial.ac.uk

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Steven Dargaville
Imperial College London
s.dargaville@imperial.ac.uk

Ionel M. Navon
Florida State University
Department of Scientific Computing

inavon@fsu.edu

MS292

Automated Adjoints for Mesh-Independent PDE-Constrained Optimisation

We will present a method for automatically deriving adjoints of finite element models with two major advantages over algorithmic differentiation: the adjoints enjoy approximately optimal efficiency (crucial for optimisation and inverse problems), and scale naturally in parallel without differentiating MPI calls or OpenMP directives. The method is used to derive adjoints of models employing the FEniCS framework. We will further discuss how discrete adjoints can be carefully used to achieve mesh independence in optimisation algorithms.

Patrick E. Farrell
Department of Mathematics
University of Oxford
patrick.farrell@maths.ox.ac.uk

MS292

Challenges in Assimilation of PM2.5 Observations for Air Pollution Forecast

Air pollution is a big issue in China. A good air quality forecast is very useful and important for warning and management of air pollution. More than 1,000 surface observation stations for PM2.5 monitoring were established recently. The big data from this observing network has the potential to improve the air quality forecast via advanced data assimilation techniques. Over recent years we have being worked on this problem. There are some progresses, and some failures too. This talk will focus on some challenges and discuss some possible solutions.

Jiang JZhu
Institute of Atmospheric Physics
Chinese Academy of Sciences
jzhu@mail.iap.ac.cn

MS292

Inversion of Geothermal Heat Flux in a Thermomechanically Coupled Nonlinear Stokes Ice Sheet Model

To project the contribution of polar ice sheets to future sea level rise, high-resolution numerical ice sheet models are critical. Yet, large uncertainties remain in the boundary conditions at the base of the ice sheet due to the lack of direct observations. Here we study mathematical and computational issues in inverse problems for basal boundary conditions, in particular, the geothermal heat flux, in a thermomechanically coupled nonlinear Stokes ice sheet model, using surface velocity observations.

Hongyu Zhu
Institute for Computational Engineering and Sciences
University of Texas at Austin
zhuhongyu@ices.utexas.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Toby Isaac
ICES
The University of Texas at Austin
tisaac@ices.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

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

MS293

Stability and Convergence of the Co-volume Scheme for the Stokes Problem

The co-volume scheme specifies the mass at cell centers and cell vertices, and both of the normal and tangential velocity components at cell edges. This scheme is extremely flexible, applicable to unstructured meshes, and avoids the need to reconstruct the tangential velocity component, as the classical C-grid scheme does. Even though the co-volume scheme has had some success in simulating geophysical flows, there has not been much progress in the theoretical analysis of the scheme. In this talk, we present stability and convergence results concerning the scheme applied to the classical Stokes problem on unstructured meshes. Both the linear and nonlinear cases will be discussed.

Qingshan Chen
Clemson University
qsc@clemson.edu

MS293

Semi-Analytical Time Differencing Methods for Stiff Problems

A semi-analytical method is developed based on conventional integrating factor (IF) and exponential time differencing (ETD) schemes for stiff problems. The latter means that there exists a thin layer with a large variation in their solutions. The occurrence of this stiff layer is due to the multiplication of a very small parameter ϵ with the transient term of the equation. Via singular perturbation analysis, an analytic approximation of the stiff layer, which is called a corrector, is sought for and embedded into the IF and ETD methods. These new schemes are then used to approximate the non-stiff part of the solution. Since the stiff part is resolved analytically by the corrector, the new method outperforms the conventional ones in terms of accuracy. In this paper, we apply our new method for both problems of ordinary differential equations and some partial differential equations.

Chang-Yeol Jung
Ulsan National Institute of Science and Technology
South Korea
changyeoljung@gmail.com

Thien Binh Nguyen
UNIST

ngthienbinh@gmail.com

MS293

A New Adaptive Weighted Essentially Non-oscillatory WENO-θ Scheme for Hypberbolic Conservation Laws

A new adaptive WENO-θ scheme is proposed. Depending on the smoothness of the large stencil used in the reconstruction procedure, a parameter θ is set adaptively to switch the scheme between a 5th-order upwind and 6th-order central approximation. A new set of smoothness indicators for both the sub-stencils and the large one is introduced. These are constructed symmetrically around x_j in Taylor expansions. Numerical results show that WENO-θ substantially improves other comparing WENO schemes.

Chang-Yeol Jung
Ulsan National Institute of Science and Technology
South Korea
changyeoljung@gmail.com

Thien Binh Nguyen
UNIST
ngthienbinh@gmail.com

MS293

The Effective Resolution of Advection Schemes

Typical classification of numerical schemes concerns the convergence order of the scheme and the relative accuracy in the limit of an infinitesimal grid (in either space or time). We define and analyze the effective resolution of numerical schemes designed for the advection equation by calculating the smallest spatial scale that is completely resolved by that scheme. This is done via dispersion relation analysis and numerical testing. We also briefly discuss the impact of these novel approaches to quantify the utility of a numerical algorithm at accurately representing the influence of waves on mean flow for nonlinear evolution equations in fluids that incorporate a three-wave interaction.

James Kent
University of Michigan
jdkent@umich.edu

Jared P. Whitehead
Brigham Young University
whitehead@mathematics.byu.edu

Christiane Jablonowski
University of Michigan
Ann Arbor MI 48109-2143
cjablono@umich.edu

Richard Rood
University of Michigan
Ann Arbor, MI 48109-2143
rbrood@umich.edu

MS294

Flooding with Equelle: A Domain Specific Language for Finite Volume Methods

Modern hardware is increasingly complex and difficult to utilize for researchers and scientists. Peak performance is only acquired through mastering a huge set of fields, including the physical problem at hand, the mathematical formulation, the numerical discretization, and the parallel implementation. As most simulator writing teams consist of one Ph.D. student, acquiring peak performance can be a daunting task. We address this challenge with a new open source domain specific language called Equelle. The language divides the different fields into different software components in order to maximize productivity of each researcher. For example, a parallelization expert would mostly contribute to the compiler back-end, and an expert in numerical methods would contribute with Equelle programs.

Andre R. Brodtkorb
SINTEF ICT, Department of Applied Mathematics
Andre.Brodtkorb@sintef.no

MS294

High Performance High Order Numerical Method for Tsunami Wave Propagation

We present a high order discontinuous Galerkin method for the accurate simulation of tsunami wave propagation. We discuss the acceleration of the method on modern many core hardware architectures such as GPUs, for the faster than real time predictions. The developed algorithms use a unified multi-threading approach OCCA. A computational kernel written in OCCA can be executed on several hardware architectures that support multi-threading approaches OpenCL, CUDA, OpenMP, Pthreads and Intel COI.

Rajesh Gandham
Rice University
Department of Computational and Applied Mathematics
rg20@rice.edu

Timothy Warburton
Department of Computational And Applied Mathematics
Rice University
timwar@rice.edu

David Medina
Rice University
dsm5@rice.edu

MS294

FEM Integration with Quadrature on the GPU

Efficient integration of low-order elements on a GPU has proven difficult. We have previously shown how to integrate a differential form (such as Laplace or elasticity) efficiently using algebraic simplification and exact integration. This, however, breaks down for multilinear forms or when using a coefficient. In this work, we show how to efficiently integrate an arbitrary form using quadrature. We introduce a technique we call "thread transposition" which matches the work done during evaluation at quadrature points to that done during basis coefficient evaluation. We are able to achieve more than 300GF/s for the variable-coefficient Laplacian, and provide a performance model to explain these results.

Matthew Knepley
University of Chicago
knepley@gmail.com

Karl Rupp
Institute for Analysis and Scientific Computing, TU Wien
Institute for Microelectronics, TU Wien
rupp@iue.tuwien.ac.at

Andy R. Terrel
Institute for Computational Engineering and Sciences
University of Texas at Austin
andy.terrel@gmail.com

MS294

Thermal Comfort Simulations on Massive Parallel Systems

Many engineering-based problems, which were deemed unsolvable a decade ago, can be simulated today using modern supercomputers such as SuperMUC installed at LRZ, Germany. Computational fluid dynamics play a dominant role in simulating urban floods or complex indoor air flow scenarios. We will present our multi-scale CFD approach based on hierarchic, block-structured Cartesian grids for solving a coupled thermal simulation together with a human manikin model, running on up to 140,000 cores in parallel.

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

Jérôme Frisch
Chair for Computation in Engineering
Technische Universität München
frisch@tum.de

MS295

Scalable Solvers for Extended MHD in the Low-β Regime

Extended MHD (XMHD) is a very challenging hyperbolic PDE system for implicit integration techniques due to the ill-conditioning introduced by fast dispersive waves. In this talk, we will describe our physics-based preconditioning approach for XMHD in the low-β regime, when a large guide field is present. The method exploits the nature of the hyperbolic couplings in XMHD to produce a block diagonally dominant PDE system, well-conditioned for multilevel techniques. Numerical experiments will demonstrate the scalability of the approach.

Luis Chacon
Los Alamos National Laboratory
chacon@lanl.gov

MS295

Multi-Fluid Magnetohydrodynamic Models for Partially-Ionized Non-Equilibrium Anisotropic Plasmas

A multi-fluid extended magnetohydrodynamic model and numerical solution method are proposed and described for the treatment of partially-ionized non-equilibrium anisotropic and strongly-magnetized plasmas. The Gaussian moment closure is used to describe the non-equilibrium transport of the neutral, ion, and electron components of the plasma in the multi-fluid model. The Gaussian closure is a maximum-entropy-based, strictly-hyperbolic closure providing approximate solutions to the Boltzmann equation which allow for strong pressure/temperature anisotropies. The treatment of reactive collisions is incorporated to allow for the effects of ionization-recombination and charge-exchange. The moment equations are coupled with Maxwells equations to complete the description for magnetized plasmas. A Godunov-type finite-volume method is proposed for the solution of the multi-fluid model by numerical means and results are described for both one- and two-dimensional problems.

Clinton P. Groth
University of Toronto Institute for Aerospace Studies
Canada
groth@utias.utoronto.ca

Ken Miura
University of Toronto Institute for Aerospace Studies
ken.miura@gmail.com

MS295

A High-order Block-adaptive Simulation Framework for Ideal and Resistive MHD Equations on Cubed-sphere Grids

Numerical simulations of large-scale space-physics problems typically require the computation of discrete solutions of complex multiphysics phenomena characterized by disparate spatial and temporal scales. Adequate numerical algorithms capable to resolve the wide range of dynamic scales of such simulations at reduced computational cost are highly desirable. This talk presents an adaptive, conservative, high-order CENO finite-volume approach for space-plasma phenomena described by the set of ideal and resistive MHD equations. Results for several benchmark problems are presented to illustrate the accuracy and computational performance of the fourth-order accurate method used in combination with a parallel, dynamically-adaptive, simulation framework on cubed-sphere grids.

Lucian Ivan, Hans De Sterck, Andree Susanto
University of Waterloo
Applied Mathematics
lucian.ivan@utoronto.ca, hdesterck@uwaterloo.ca,
asusanto@uwaterloo.ca

Clinton P. Groth
University of Toronto Institute for Aerospace Studies
Canada
groth@utias.utoronto.ca

MS295

Block Preconditioners for 3D Incompressible MHD

The scalable iterative solution of strongly coupled3D incompressible resistive magnetohydrodynamics (MHD) equations is very challenging. This study considers mixed FE integration for velocity/pressure (Q2/Q1) and edge-elements for magnetic induction and presents anew approximate block factorization (ABF) preconditioner for this system. The ABF preconditioner reduces the system toapproximate Schur complement systems that are solved by nodal and edge-element based AMG methods.

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Edward Phillips
Sandia National Laboratories
egphill@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

MS296

An HDG Method for Non-Hydrostatic Atmosphere

Abstract not available at time of publication.

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

MS296

Towards a Fully 3D Compressible Atmosphere Dynamical Core with Compatible Finite Elements

We have recently been developing compatible finite element methods (i.e., mixed finite element methods using families of spaces forming discrete de Rham complexes) as an extension of the C-grid staggered finite difference approach. These methods lead to higher-order discretisations on arbitrary meshes with flexibility to adjust the global DOF ratio between pressure and velocity in order to minimise the impact of spurious mode branches. These methods are being developed as part of the GungHo dynamical core project in the UK, a collaboration between the Met Office, STFC and several UK universities. In Phase I of the project, we successfully developed finite element discretisations for the nonlinear rotating shallow-water equations on the sphere, which serves as a useful testbed for the horizontal discretisation in dynamical cores. The compatible finite element structure allows for numerical schemes with stable advection of the implied diagnostic potential vorticity field, which is important for stable long-time integrations with minimal numerical dissipation. We are now developing these ideas in the context of a full 3D dynamical core, via vertical slice models. This talk will discuss our progress in addressing the following issues: * What is the finite element analogue of the Charney-Phillips vertical staggering (necessary for good representation of hydrostatic balance in the vertical)? What is the correct choice of finite element space for temperature? * How to obtain stable and accurate advection schemes for temperature in this case? * How to extend the velocity advection scheme to three dimensions? Can the 3D vorticity equation be incorporated? * How to discretise the nonlinear pressure gradient term in the theta-pi formulation? Proposed solutions to these questions will be illustrated with numerical experiments implemented using Firedrake, a high-performance finite element library targetting geophysical fluid dynamics applications.

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

MS296

A Higher-Order Finite Volume Nonhydrostatic Dynamical Core with Space-Time Refinement

We present an adaptive non-hydrostatic dynamical core based on a higher-order finite volume discretization on the cubed sphere. Adaptivity is both in space, using nested horizontal refinement; and in time, using subcycling in refined regions. The algorithm is able to maintain scalar conservation with careful flux construction at refinement boundaries, as well as conservative coarse-fine interpolation. We show results for simple tests as well as more challenging ones that highlight the benefits of refinement.

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

MS296

Optimization-based Spectral Element Semi-Lagrangian Tracer Transport

Transport algorithms are highly important in atmospheric modeling where hundreds of tracer species must be efficiently transported while maintaining conservation of tracer mass and preservation of physical bounds. We present a new transport algorithm that combines a high-order spectral element semi-Lagrangian scheme with an optimization algorithm to enforce mass conservation and bounds preservation. We evaluate the new method using several standard two- and three-dimensional transport problems on the sphere and compare to existing approaches.

Kara Peterson
Sandia Natl. Labs
kjpeter@sandia.gov

Mark A. Taylor
Sandia National Laboratories, Albuquerque, NM
mataylo@sandia.gov

MS297

Linear Response Eigenvalue Problem and Excited State Calculations

Linear response (LR) eigenvalue problems arise from excitation state calculations of collective motion of many particle systems. In this talk, we first present theoretical results for the LR eigenvalue problems such as minimization principles. Although the LR eigenvalue problems are nonsymmetric, these results mirror the well-known theoretical results for symmetric eigenproblems. Then we will discuss the best approximation of the few smallest positive eigenvalues via a structure-preserving projection, and describe conjugate gradient-like algorithms for simultaneously computing these eigenvalues.

Zhaojun Bai
Departments of Computer Science and Mathematics
University of California, Davis
bai@cs.ucdavis.edu

Ren-Cang Li
University of Texas - Arlington
rcli@uta.edu

MS297

Accelerating Quantum Transport Calculations Through the Feast Algorithm

In Quantum Transport the scattering region of a device is connected to semi-infinite leads, where the energy-dependent reflected modes must be determined by solving a polynomial eigenvalue problem. As fast decaying modes do not significantly contribute, it is sufficient to compute

only those eigenvalues with a small or vanishing imaginary part. The FEAST algorithm allows for drastically reducing the system size by specifically searching for the desired eigenvalues and for efficiently parallelizing the computational workload.

Sascha Brueck, Mauro Calderara
Integrated Systems Laboratory
ETH Zurich
bruecks@iis.ee.ethz.ch, mcalderara@iis.ee.ethz.ch

Hossein Bani-Hashemian, Joost VandeVondele
Nanoscale Simulations
ETH Zurich
hossein.banihashemian@mat.ethz.ch,
joost.vandevondele@mat.ethz.ch

Mathieu Luisier
Integrated Systems Laboratory
ETH Zurich
mluisier@iis.ee.ethz.ch

MS297

Parallel Solution of Eigenvalue Problems from Graphene Modeling with Solvers Based on Integration and Approximation

Graphene is a material that recently attracts a lot of interest among researchers. Studying its electronic properties involves solving (typically large scale) sparse eigenvalue problems. In this talk we will present FEAST-like and polynomial based solvers for the solution of these eigenvalue problems. We will also discuss approaches to the solution of the arising inner linear systems. We finally present numerical examples involving the solution of graphene eigenvalue problems.

Bruno Lang, Lukas Krämer
University of Wuppertal
lang@math.uni-wuppertal.de,
lukas.kraemer@math.uni-wuppertal.de

Martin Galgon
Yniversity of Wuppertal
galgon@math.uni-wuppertal.de

MS297

Polynomial Techniques and Primme for the Computation of Large Number of Eigenvalues

A computationally challenging task is the computation of a large number, M, of eigenvalues and their eigenvectors. Enforcing orthogonality of a Krylov subspace larger than M becomes a bottleneck, and computing eigenvectors a few at a time increases the iteration cost linearly with M. Polynomial filters have been used to trade orthogonalization for matrix-vector multiplications. We discuss polynomial filters both for single and multivectors on top of the PRIMME eigenvalue software.

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

MS298

Optimal Control Problems With Uncertain Model Parameters

I demonstrate solving stochastic optimal control problems by adding minimal modifications to their corresponding deterministic optimization solvers using Trilinos packages. The process of modifying the deterministic solver to a stochastic one can be done systematically and automatically given suitable code structure. I illustrate this process by solving optimal control problems with uncertain model parameters arising in heat-flux control and oil reservoir optimization.

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

MS298

On Risk-averse PDE-constrained Optimization using Convex Risk Measures Inspired by Conditional Value-at-risk

We consider a class of PDE-constrained optimization problems in which the underlying PDE-system contains uncertain parameters. In order to obtain robust controls that are both deterministic and risk-averse, we consider the minimization of the conditional value-at-risk (CVaR) of the reduced objective functional. In order to develop efficient numerical schemes, regularization approaches are suggested for the primal and the dual formulation of the minimization problem. The regularized CVaR is shown to be a convex risk measure. Sensitivity and consistency results of the regularized problems are derived. Finally, we present numerical results for several example problems.

Thomas M. Surowiec
Department of Mathematics
Humboldt University of Berlin
surowiec@math.hu-berlin.de

MS298

Optimization under Uncertainty: Application to Electrical Circuits

This presentation will demonstrate the application of stochastic optimization algorithms to the problem of model calibration and optimal control under uncertainty for simple electrical circuits. Our implementation uses the Rapid Optimization Library (ROL) from the Trilinos framework to solve optimization under uncertainty problems with algebraic and partial differential equation constraints. Our numerical experimentation considers the Shockley ideal diode equation and drift-diffusion model.

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

MS298

Maximizing AUC and Buffered AUC in Classification

We propose an alternative to the Area Under the Receiver Operating Characteristic Curve (AUC) performance metric called Buffered AUC (bAUC). We show that bAUC is an intuitive counterpart to AUC. We then show that bAUC, compared to AUC, can be a more informative measure of a classifiers ranking quality. In addition, we show

that bAUC is much easier to handle in optimization frameworks than AUC, specifically reducing to convex and linear programming. We use these friendly optimization properties to introduce the bAUC Efficiency Frontier, a concept that serves to partially resolve the incoherency that arises when misclassification costs need be considered. We conclude that bAUC avoids many of the numerically troublesome issues encountered by AUC and integrates much more smoothly into the general framework of model selection and evaluation.

Stan Uryasev
University of Florida
uryasev@ufl.edu

MS299

Facilitating Learners Cognitive Presence In A Self-Directed Online Course

The study will report our study on students self-directed learning and inquiry in a 15-week graduate-level online course. We will investigate learners cognitive presence and facilitating behaviors that help create higher-level cognitive presence. Online discussion transcripts, semi-structure interviews, and learning records (i.e. view frequencies, grades) will be used as the primary data sources. The purpose of this study is to reveal the effective facilitating strategies that trigger or promote learners critical thinking and deep reflection.

Ye Chen, Jing Lei
Syracuse University
ychen129@syr.edu, jlei@syr.edu

MS299

Towards Automating Analysis of Midterm Semester Feedback Surveys for Improving Course Effectiveness

This talk will report on the use of midterm semester feedback survey for improving teaching and learning effectiveness in college courses, including data from the Computational Modeling courses discussed during this session. Efforts to expand and partially automate analysis and feedback from these surveys will be explored, with some initial results.

Douglas Holton
Embry-Riddle Aeronautical University
douglas.holton@erau.edu

MS299

Tri-Located Course in Mathematical Modeling and Complementary Reu Summer Workshop

This talk will report the summative evaluation for the project funded by NSF TUES to create a cluster of collaborating institutions that combine students into common classes and use cyberlearning technologies to deliver and manage instruction. We will also share the course-based research experience that was complemented by innovative undergraduate research summer workshops. We will discuss ideas how to advance CSE Education through MOOC, project-based learning, and deep learning assessment technology.

Hong Liu
Department of Mathematics
Embry-Riddle Aeronautical University
liuho@erau.edu

Andrei Ludu
Embry-Riddle Aeronautical University
Mathematics Department
ludua@erau.edu

MS299

Deep Learning Assessment for Near Real-time, Formative Feedback during Complex Problem-solving Activities

A critical factor for progress in cyberlearning involves assessment, especially near real-time, formative feedback during complex problem-solving activities. Detecting misconceptions early is important to help learners gain competence and build confidence. This presentation will demonstrate one approach to assessing learning in complex domains that involves the analysis of student conceptualizations of the problem space over time and in comparison with those of experienced persons. These conceptualizations can be gathered explicitly in the form of annotated concept maps or through student models created using tools such as MATLAB.

Michael Spector
University of North Texas
Mike.Spector@unt.edu

MS300

Weighted Sparsity and Function Interpolation Via Infinite-dimensional Compressed Sensing

We introduce a framework for function interpolation using weighted l1 minimization. Two advantages of this framework over existing approaches are: (i) in the absence of noise it leads to interpolatory approximations, and (ii) it does not require a priori estimates on the expansion tail. We explain the critical role weighted sparsity plays in function interpolation; namely, that of regularizing the problem and removing aliased solutions. Finally, we present a number of near-optimal recovery guarantees.

Ben Adcock
Purdue University
ben_adcock@sfu.ca

MS300

Representation Using the Weyl Transform

The Weyl transform is introduced as a powerful framework for representing measurement data. Transform coefficients are connected to the Walsh-Hadamard transform of multiscale autocorrelations, and different forms of dyadic periodicity in a signal are shown to appear as different features in its Weyl coefficients. A large group of multiscale transformations is shown to support very fast pooling since the Weyl coefficients are unique up to permutation and phase changes when the original signal is transformed by any element of this group. The effectiveness of the Weyl transform is demonstrated through the example of textured image classification.

Robert Calderbank
Mathematics
Duke University

robert.calderbank@duke.edu

MS300

Compressive Parameter Estimation via Approximate Message Passing

The literature on compressive parameter estimation has been mostly focused on the use of sparsity dictionaries that encode a sampling of the parameter space; these dictionaries, however, suffer from coherence issues that must be controlled for successful estimation. We propose the use of statistical parameter estimation methods within the approximate message passing (AMP) algorithm for signal recovery. Our proposed work leverages the recently highlighted connection between statistical denoising methods and the thresholding step commonly used during recovery. As an example, we consider line spectral estimation by leveraging the well-known MUSIC algorithm. Numerical experiments show significant improvements in estimation performance.

Shermin Hamzehei
University of Massachusetts
shamzehei@engin.umass.edu

Marco F. Duarte
University of Massachusetts Amherst
mduarte@ecs.umass.edu

MS300

Fast and Robust Dictionary Learning, with Invariances and Multiresolution

We introduce a novel dictionary learning procedure for learning sparsifying representations for images. Sparse representations have been used successfully for a variety of imaging tasks, from denoting to imprinting to superresolution. Key features of our construction are (1) the ability of seamlessly construct a dictionary for patches of images at different scales, rather than on patches of a fixed size; (2) the ability to produce representations that are invariant under a known, given desired invariances (e.g. translations, rotations); (3) having fast algorithms for both the construction of the dictionary, and for computing the sparse representation of a signal onto the dictionary; (4) theoretical finite sample guarantees on the learning algorithm.

Mauro Maggioni, Samuel Gerber
Department of Mathematics
Duke University
mauro@math.duke.edu, sgerber@math.duke.edu

MS301

Upwind DG for Acoustic and Elastic Wave Equations

We develop and analyze a new strategy for the spatial discontinuous Galerkin discretization of wave equations in second order form. The method features a direct, parameter-free approach to defining interelement fluxes. Both energy-conserving and upwind discretizations can be devised. We derive a priori error estimates in the energy norm for certain fluxes and present numerical experiments showing that optimal convergence in L^2 is obtained.

Daniel Appelo
University of New Mexico
appelo@unm.edu

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

MS301

High-Order Upwind Methods for Second-Order Wave Equations on Curvilinear and Overlapping Grids

In this talk I will present preliminary results for a promising new technique to PDE discretization. The approach, called Galerkin finite differences, blends the attractive features from finite differences and finite elements to yield schemes with provable energy stability, high-order accuracy, and well-conditioned discrete systems even for very high order. Results are presented for the wave equation using both continuous and discontinuous representations.

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banksj3@rpi.edu

MS301

A Discontinuous Galerkin Method for the Spherically Reduced Einstein Field Equations with Second-Order Operators

A discontinuous Galerkin (dG) method for evolving a spherically reduced formulation of the Einstein field equations of general relativity is proposed. The system is discretized in its natural first-order in time second-order in space form. I will discuss in some detail our scheme as well as our treatment of the second-order spatial operators which appear in this particular system. By approximating the second-order spatial derivatives of the system, we avoid the need to introduce extra fields and equations which can lead to new challenges. We demonstrate stable, long-time evolutions achieved by our scheme, which also constitutes the first use of dG methods in computational general relativity. I will conclude with a brief overview of current status of dG methods within the computational relativity community.

Scott Field
Department of Astronomy
Cornell University
sfield@astro.cornell.edu

Jan Hesthaven
EPFL
jan.hesthaven@epfl.ch

Stephen Lau
Mathematics
The University of New Mexico
lau@math.unm.edu

Abdul Mroue
Canadian Institute for Theoretical Astrophysics
mroue@cita.utoronto.ca

MS301

Uncertainty Quantification for High Frequency

Waves

We consider the wave equation with highly oscillatory initial data, where there is uncertainty in the wave speed, initial phase and/or initial amplitude. To estimate quantities of interest related to the solution, and their uncertainty, we combine a high frequency method based on Gaussian beams with stochastic collocation. We show that this is an efficient approach for quadratic quantities of interest, since they vary smoothly with respect to the random variables, with derivatives bounded independently of the wave frequency.

Olof Runborg
KTH, Stockholm, Sweden
olofr@nada.kth.se

MS302

A Consistent and Robust Discrete Adjoint Solver for the SU2 Framework

In this talk we discuss the development of a discrete adjoint solver, which enables the computation of consistent gradients in a robust way based on the exploitation of the fixed-point structure of the flow solver. All occurring derivatives in this formulation can be calculated using advanced techniques of Algorithmic Differentiation so that the extension to arbitrary complex flow models can be performed along with the development of the primal code.

Tim Albring
TU Kaiserslautern
tim.albring@rhrk.uni-kl.de

Nicolas R. Gauger
RWTH Aachen University
nicolas.gauger@rhrk.uni-kl.de

MS302

High-Performance Optimizations of the Unstructured Open-Source SU2 Suite

Strategies and lessons learned from applying high-performance optimizations to an open-source computational fluid dynamics suite, SU2, are presented. We focus on performance optimizations of solver execution with emphasis on parallelization and on finding well-suited algorithms. The resulting code modifications are geared toward achieving high scalability with edge-based CFD solvers such as SU2, making efficient use of memory, and choosing appropriate algorithms for maximizing parallelism, especially for solving linear systems arising from implicit time discretizations.

Thomas Economon
Department of Aeronautics and Astronautics
Stanford University
economon@stanford.edu

Francisco Palacios
Stanford University
fpalacios@stanford.edu

Juan J. Alonso
Department of Aeronautics and Astronautics
Stanford University
jjalonso@stanford.edu

Gaurav Bansal
Software and Services Group
Intel Corporation, Hillsboro, OR, USA
gaurav2.bansal@intel.com

Anand Deshpande
Parallel Computing Laboratory
Intel Corporation, Bangalore, India
anand.m.deshpande@intel.com

Alexander Heinecke
Parallel Computing Laboratory
Intel Corporation, Santa Clara, CA, USA
alexander.heinecke@intel.com

Dheevatsa Mudigere
Parallel Computing Laboratory
Intel Corporation, Bangalore, India
dheevatsa.mudigere@intel.com

Mikhail Smelyanskiy
Intel Corporation
mikhail.smelyanskiy@intel.com

MS302

Large Scale Design Using Su2 and a Continuous Adjoint Rans Approach

Abstract not available at time of publication.

Francisco Palacios
Stanford University
fpalacios@stanford.edu

MS302

A Discrete Adjoint Framework for Lift-Constrained Noise Minimization Using SU2

This talk presents an discrete adjoint-based framework for aeroacoustic shape optimization. The unsteady adjoint solver is developed by applying algorithmic differentiation (AD) to the open source SU2 code. A practical lift-constrained noise minimization formulation is presented in which the acoustic signal is minimized at a far-field observation point while a lift constraint is imposed to ensure adequate aerodynamic performance in low speed operating conditions.

Beckett Zhou
TU Kaiserslautern
yuxiang.zhou@rhrk.uni-kl.de

Nicolas R. Gauger
RWTH Aachen University
nicolas.gauger@rhrk.uni-kl.de

Thomas Economon
Department of Aeronautics and Astronautics
Stanford University
economon@stanford.edu

Francisco Palacios
Stanford University
fpalacios@stanford.edu

Juan J. Alonso
Department of Aeronautics and Astronautics
Stanford University

jjalonso@stanford.edu

MS303

Time-Domain Simulation of Two Dimensional Elastic Scattering

We show how a combination of a simple Nystrom discretization in the Laplace domain and Convolution Quadrature can be used for simulation of transient elastic waves around a finite number of obstacles and cracks in two dimensions. The implementation of most of the integral operators is based on naive quadrature formulas and well chosen mixing parameters. Only the hypersingular integral operator requires the careful use of a regularization formula that is known in the literature. The Calderón Calculus thus defined can be used for exterior scattering problems, transmission problems in locally homogeneous materials, and even in problems in wave-structure interaction.

Victor Dominguez
Universidad Publica de Navarra
Spain
victor.dominguez@unavarra.es

Tonatiuh Sanchez-Vizuet
University of Delaware
tonatiuh@math.udel.edu

Francisco J. J. Sayas
Department of Mathematical Sciences
University of Delaware
fjsayas@math.udel.edu

MS303

Variable Order Fast Multipole Method for an Elastodynamic BEM

A fast time domain BEM based on the CQM is under study. Chebyschev polynomials are used as kernel expansion within the Fast Multipole Method. Furthermore, a directional clustering schema is applied to establish the hierarchical cluster tree. A variable order FMM is implemented to obtain optimal complexity while ensuring convergence. Some examples show the suitability of the proposed approach.

Thomas Traub
Graz Univeristy of Technology
Institute of Applied Mechanics
thomas.traub@tugraz.at

Martin Schanz
Technical University of Graz
m.schanz@tugraz.at

MS303

Accuracy of the Marching-on-in-Time Scheme for Td-Bie Methods

The marching-on-in-time (MOT) scheme is a popular discretization technique for Time Domain Integral Equation methods. The choice of temporal basis function has a profound impact on the computational characteristics. We will present an analysis of the accuracy in time of the MOT scheme based on the interpolation accuracy of the temporal basis functions. Surprisingly, a higher order of accuracy has been observed in numerical experiments with splines than for Lagrange basis functions with equal support.

Elwin Van 't Wout
University College London
Centre for Medical Image Computing
e.wout@ucl.ac.uk

MS303

Recent Advances in the Convolution Quadrature and Temporal Galerkin Approaches to Transient Electromagnetics

Time domain integral equation based methods for the solution of electromagnetics and scattering problems have been increasing in popularity in recent years. Two methods, convolution quadrature and temporal Galerkin, have become popular for the discretization of these equations, and have spawned many variants. This talk will compare and contrast these methods, especially in their most recent incarnations in electromagnetics, with respect to stability, accuracy, and efficiency.

Daniel Weile
University of Delaware
weile@ee.udel.edu

Balasubramaniam Shanker
Department of Electrical and Computer Engineering
Michigan State University
shanker@egr.msu.edu

MS304

Development of a Contact MiniApplication Using Kokkos

We present a new approach at the contact global search algorithm, targeted at providing high performance on many-core computer architectures. Sandias ACME contact library was chosen to serve as the reference implementation. A new global search algorithm was implemented, using the Kokkos performance portable programming model, that is based on a Morton-code linearized Bounding Volume Hierarchy (BVH) developed by Nvidia for execution on GPU co-processors. We conclude with results that compare the reference ACME search approach using MPI with the new Morton algorithm using MPI and multicore processing within each MPI rank.

Glen Hansen, Patrick Xavier, Sam Mish
Sandia National Laboratories
gahanse@sandia.gov, pgxavie@sandia.gov, smish@sandia.gov

MS304

Co-Designing Hierarchical Algorithms: Application to Vlasov-Maxwell Particle-in-Cell Methods

Often, multiscale mathematical descriptions admit a multilayered (hierarchical) description based on a systematic coarse-graining procedure. Recently, such hierarchical descriptions have been demonstrated to provide significant algorithmic acceleration in a number of challenging applications. Moreover, the layered nature of these algorithms has opened novel opportunities for co-design, which can be exploited to maximize FLOPs and minimize communication. In this presentation, we will discuss our co-design strategy for hierarchical algorithms, exemplified with a particle-

based implementation of the Vlasov-Maxwell system.

Joshua Payne, Luis Chacon, Guangye Chen, Chris Newman, Dana Knoll, Allen McPherson
Los Alamos National Laboratory
payne@lanl.gov, chacon@lanl.gov, gchen@lanl.gov, cnewman@lanl.gov, nol@lanl.gov, mcpherson@lanl.gov

MS304

Uintah/Wasatch: Addressing Multiphsyics Complexity in a High-Performance Computing Environment

To address the coding and software challenges of modern hybrid architectures, we propose a modern approach to multiphysics code development. The approach is based on using a Domain Specific Language in tandem with runtime algorithm generation. When coupled with a large-scale parallel framework, the result is an architecture-proof code capable of executing on hybrid platforms. We share our experience developing such a code - an effort that spans an interdisciplinary team of engineers and computer scientists.

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

Christopher Earl
University of Utah
earl.christopher@gmail.com

Abhishek Bagusetty
Chemical Engineering
University of Utah
abhishek.bagusetty@utah.edu

Matthew Might
The University of Utah
might@cs.utah.edu

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

PP1

Lid Driven Cavity Simulations in 2d and 3d Using High Accurate Methods

in this poster, numerical simulations of two-dimensional and three-dimensional partial differential equations are presented by solving the steady navier-stokes equations in a lid driven cavity at high Reynolds numbers where it becomes difficult. in two dimensions, we use the streamfunction-vorticity formulation to solve the problem in a square domain. a numerical computational method is employed to discretize the problem in the x and y directions with a spectral collocation method. the problem is coded in the matlab programming environment. solutions at high reynolds numbers are obtained up to re=20000 on a fine grid. Also in this poster, the numerical computational simulation for the three-dimensional lid-driven cavity problem are obtained by solving the velocity-vorticity formulation of the navier-stokes equations (which is the first time that this has been simulated with special boundary conditions) for various Reynolds numbers.

Badr Alkahtani
King Saud University
alhaghog@gmail.com

PP1

Numerical Study of Thin Viscoelastic Films on Substrates

We numerically study the interfacial dynamics and instability of a thin viscoelastic film on a substrate. We use the long wave approximation to describe the non-linear evolution of the interface. We consider different regimes of slippage, and in each regime, we investigate the role of the liquid viscoelasticity and of the contact angle on the thin film break-up. Numerical solutions of the full non-linear equations are compared with the results of the linear stability analysis.

Valeria Barra
New Jersey Institute of Technology
vb82@njit.edu

Shahriar Afkhami
Department of Mathematical Sciences
New Jersey Institute of Technology
shahriar.afkhami@njit.edu

PP1

A Novel Modeling Approach for Multiscale, Multiphysics Flow

Turbulent combustion simulation remains a challenging problem due to the wide range of time and length scales associated with the governing physics. We propose a novel methodology, termed Lattice-Based Multiscale Simulation (LBMS) that creates a 3D network of widely spaced lines with fully resolved 1D physics along each line. Local turbulence along a line is modeled stochastically while lattice-scale advection is captured directly. LBMS is ideal for situations where multiscale coupling is key like wall-bounded flows, reacting flows, etc.

Derek A. Cline
University of Utah
Chemical Engineering Department
derek.cline@chemeng.utah.edu

PP1

Flusepa - a Navier-Stokes Solver for Unsteady Problems with Bodies in Relative Motion : Toward a Task-Based Parallel Version over a Runtime System for Large Simulations

FLUSEPA code is designed to handle unsteady problems with bodies in relative motion (stage separation) and strong shocks. A finite volume formulation is used to solve the RANS equations. Time integration consists on an explicit temporal adaptive solver. We present a task based parallel version of the aerodynamic solver designed from the previous MPI/OpenMP one and using a modern runtime system to schedule the tasks. An Ariane 5 booster separation computation will be presented.

Jean Marie Couteyen Carpaye
Airbus Defence and Space
Inria
jean-marie.couteyen@inria.fr

Jean Roman

INRIA
Jean.Roman@inria.fr

Pierre Brenner
Airbus Defence and Space
pierre.brenner@astrium.eads.net

PP1

Efficiency of an Adjoint Industrial CFD Code

We apply Algorithmic Differentiation on the commercial flow solver CFD-ACE+ to derive a discrete adjoint code. The objective is to assess its performance and to improve its efficiency in terms of memory consumption and runtime. To achieve that, we combine the flexibility of an operator overloading tool with the efficiency of an adjoint code generated by source transformation. In addition, we speed up the adjoint computation by exploiting the mathematical aspect of the involved fixed-point iteration.

Zahrasadat Dastouri
STCE,RWTH Aachen University
dastouri@stce.rwth-aachen.de

Johannes Lotz
STCE, RWTH Aachen University
Germany
lotz@stce.rwth-aachen.de

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

PP1

Production of Dissipative Vortices by Solid Bodies in the Inviscid Limit of Incompressible Fluid Flows: Comparison Between Prandtl, Navier-Stokes and Euler Solutions

We revisit the problem posed by Euler in 1748 that lead d'Alembert to formulate his paradox and we address the following question: does energy dissipate when boundary layer detaches from a solid body in the vanishing viscosity limit? To trigger detachment we consider a vortex dipole impinging onto a wall and we compare the numerical solutions of Euler, Prandtl, and Navier-Stokes equations. We observe the formation of two opposite-sign boundary layers whose thickness scales as predicted by Prandtl's 1904 theory. But after a certain time Prandtl's solution becomes singular, while the Navier-Stokes solution collapses down to much finer thickness for the boundary layers, in accordance with Kato's 1984 theorem. Then the boundary layers roll up and form vortices which detach from the wall and dissipate a finite amount of energy, even in the vanishing viscosity limit.

Marie Farge
LMD-CNRS
Ecole Normale Superieure
farge@lmd.ens.fr

Romain Nguyen van yen
Freie University, Berlin, Germany
romain.nguyenvanyen@gmail.com

Matthias Waidmann
Freie University, Berlin (Germany
matthias.waidman@math.fu-berlin.de

Kai Schneider
Universite de Provence, Aix-Marseille
Centre de Mathematiques et Informatique
kschneid@cmi.univ-mrs.fr

Rupert Klein
Freie Universität Berlin
rupert.klein@math.fu-berlin.de

PP1

A Numerical Study of Shock-Induced Cavity Collapse in a Solid Explosive

The shock-induced collapse of a gas-filled cavity in a solid explosive is examined. The system is modeled as a multimaterial compressible fluid with a mixture equation of state and an ignition-pressure reaction rate. The governing equations are solved numerically using a Godunov-type scheme designed to accommodate the large impedance mismatch at the fluid-solid interface. Results are described for ellipsoidal cavities of various shapes to determine whether the collapse initiates a detonation in the reactive material.

James R. Gambino
RPI
gambij@rpi.edu

Ashwani K. Kapila
Rensselaer Polytechnic Inst.
kapila@rpi.edu

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

William Henshaw
Rensselaer Polytechnic Institute
henshw@rpi.edu

PP1

Scalable Advection Algorithms for Multi-Tracers in Climate Codes

One of the important goals of ACES4BGC project (Applying Computationally Efficient Schemes for BioGeochemical Cycles) is to implement and optimize new computationally efficient advection algorithms for large number of tracer species. This work demonstrates a framework that uses intersection algorithms developed in MOAB (Mesh Oriented datABase), linked with a HOMME dynamical core driver.

Iulian Grindeanu
Argonne National Laboratory
iulian@mcs.anl.gov

Kara Peterson
Sandia Natl. Labs
kjpeter@sandia.gov

Vijay Mahadevan, Navamita Ray, Rajeev Jain
Argonne National Laboratory
vijay.m@gmail.com, navamitaray@gmail.com,

jain@mcs.anl.gov

PP1

Fast Ship Hydrodynamics Via Novel Methods

An existing high-order finite difference, potential flow solver for large-scale ocean wave modeling is extended to include interaction with floating bodies via a new immersed-boundary technique based on weighted least squares. High-order WENO schemes are used for the free-surface boundary conditions to obtain stable solutions for the linear seakeeping response of a ship at forward speed. The code is implemented on massively parallel GPU architectures using the CUDA API.

Stavros Kontos
DTU - Technical University of Denmark
stakon@mek.dtu.dk

Ole Lindberg
FORCE Technology
old@force.dk

Allan Engsig-Karup
The Technical University of Denmark
apek@dtu.dk

Harry Bingham
Technical University of Denmark
DTU MECHANICAL ENGINEERING
hbb@mek.dtu.dk

PP1

A Conservative, Positivity Preserving Scheme for Reactive Solute Transport Problems in Moving Domains

We study the mathematical models and numerical schemes for reactive transport of a soluble substance in deformable media. The problem is modeled by a convection-diffusion adsorption-desorption equation in moving domains. We present a conservative, positivity preserving, high resolution ALE-FCT scheme for this problem in the presence of dominant transport processes and wall reactions on the moving wall. A Patankar type time discretization is presented, which provides conservative treatment of nonlinear reactive terms. Consequences of this result are significant in the area of, e.g., nano-particle cancer drug delivery. Our result shows that periodic excitation of the cancerous tissue using, e.g., ultrasound, may enhance adsorption of cancer drugs carried by nano-particles via the human vasculature.

Sibusiso Mabuza
University of Houston
smabuza@math.uh.edu

PP1

Singly-Periodic Stokes Flow with a Wall

A closed formula for the 2D singly-periodic laminar velocity field near a solid straight boundary is derived using the method of images. The flow is induced by regularized periodic forces. This result provides a model, for instance, for the modeling of cilia.

Forest O. Mannan
Tulane University
fmannan@tulane.edu

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

PP1

Computational Hydrodynamics: How Portable and Scalable Are Heterogeneous Programming Paradigms?

Many-core era applications at the interface of mathematics and computer science adopt modern parallel programming paradigms and expose parallelism through proper algorithms. We present results for a novel massively parallel free surface wave model suitable for advanced experiments in Numerical Wave Tanks. Our application exhibits excellent performance portability and scalability using hybrid MPI-OpenCL/CUDA and running on arbitrary system sizes including desktops, superclusters and cloud utilizing heterogeneous devices like multi-core CPUs, GPUs, and Xeon Phi coprocessors.

Wojciech Pawlak
DTU Compute, Technical University of Denmark
Allan P. Engsig-Karup Associate Professor, Ph.D.
wojciechpawlak@gmail.com

Allan Engsig-Karup
The Technical University of Denmark
apek@dtu.dk

Stefan Glimberg
Lloyd's Register Consulting
stefan.glimberg@lr.org

PP1

A Fully Discrete Derivation of a Direct Ale Conservative Scheme for Compressible Hydrodynamics

Lagrange plus remap schemes have been used for years in industrial applications. However, this kind of scheme can lead to some difficulties with conservativity and computational cost. In this work, a direct ALE scheme is provided using a least action principle to a discretized action integral in both space and time. The internal energy equation follows from the conservation of the total energy. This mimetic procedure guarantees the physical compatibility between the thermodynamical variables.

Thibaud Vazquez-Gonzalez
CEA Bruyeres le Chatel
thibaud.vazquez-gonzalez@cea.fr

Antoine Llor, Christophe Fochesato
CEA
antoine.llor@cea.fr, christophe.fochesato@cea.fr

PP2

A Non Standard Scheme for Nagumo Type Differential Equations

In this work, we design explicit difference schemes for the Nagumo reaction-diffusion and the Allen-Cahn equations. The Nonstandard exact schemes that we design for the space-independent sub-equations help to reduce the computational cost common with difference schemes approx-

imation of singularly perturbed equations like the Allen-Cahn. A careful assemblage of these schemes with the energy preserving scheme of the steady state equation yields a uniformly convergent difference scheme for the equations.

Adebayo A. Aderogba
University of Pretoria
South Africa
aderogba_adebayo@yahoo.com

Michael Chapwanya, Pius Chin
Department of Mathematics and Appl Math.
University of Pretoria
michael.chapwanya@up.ac.za, pius.chin@up.ac.za

PP2

Reduced Basis Methods for Calibration and Option Pricing

We present reduced basis approximations for parametrized time-dependent variational (in-)equalities with the special focus on applications from finance. In particular, we consider the case of calibrating European and American options with the Heston model. With the use of an offline-online computational procedure, we significantly reduce the computational cost of the calibration phase. Numerical tests illustrate the approximation quality and convergence of the reduced basis methods and it's advantage in application to the calibration algorithms.

Olena Burkovska
Technical University of Munich
burkovsk@ma.tum.de

Kathrin Glau
Technical University Munich
kathrin.glau@tum.de

Mirco Mahlstedt
Technical University of Munich
Chair of Mathematical Finance
mirco.mahlstedt@tum.de

Barbara Wohlmuth
Technical University of Munich
barbara.wohlmuth@ma.tum.de

PP2

A Fast and Stable Explicit Operator Splitting Method for Phase-Field Models

We propose a combined finite difference and pseudo-spectral method for one- and two-dimensional nonlinear diffusion equations for thin film epitaxy with slope selection and Cahn-Hilliard equation. Equations are split into nonlinear part, solved using the method of lines approach together with an efficient large stability domain ODE solver; and linear part, solved by a pseudo-spectral method, which is based on the exact solution and thus has no stability restriction on the size of time step. Finally, an adaptive time-stepping strategy is introduced for long time simulation.

Yuanzhen Cheng
Tulane University
ycheng5@tulane.edu

Alexander Kurganov
Tulane University
Department of Mathematics
kurganov@math.tulane.edu

Zhuolin Qu
Tulane University
zqu1@tulane.edu

Tao Tang
Hong Kong Baptist University
ttang@math.hkbu.edu.hk

PP2

Method of Lines Transpose Schemes for Parabolic Problems

We present a novel numerical scheme suitable for solving parabolic differential equation model using the Method of Lines Transpose (MOL^T) combined with the successive convolution operators. The primary advantage is that the operators can be computed quickly in $O(N)$ work, to high precision; and a multi dimensional solution is formed by dimensional sweeps. We demonstrate our solver on the Allen-Cahn and Cahn-Hilliard equation.

Hana Cho
Michigan State University
chohana@math.msu.edu

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

Matt Causley
Kattering University
matt.causley@gmail.com

David Seal
Michigan State University
seal.david@gmail.com

PP2

Comparison of Nonlinear and Linear Stabilization Schemes for Advection-Diffusion Equations

Standard finite element discretizations of advection-diffusion equations introduce unphysical oscillations around steep gradients. Therefore, stabilization must be added to the discrete formulation to obtain correct solutions. The SUPG, dCG91, and Entropy Viscosity schemes are compared using stationary and non-stationary test equations. Differences in maximum overshoot and undershoot, smear, and convergence orders are compared using code written using deal.ii

Ryan R. Grove
SIAM Webmaster at Clemson University
rgrove@clemson.edu

Timo Heister
Clemson University
Mathematical Sciences

heister@clemson.edu

PP2

A Task-Parallel Approach for Solving PDEs on a Lattice

We solve PDE's on a lattice. This creates a unique set of communication patterns, at the intersections of the lines on the lattice, creating a coarse mesh, and between parallel lines, creatine overlapping fine meshes. We use a directed a-cyclical graph to order the execution and communication in parallel across the lattice, via sender-poller; allowing, in parallel, different pieces of the problem to be solved concurrently while communication is happening.

John T. Hutchins
Boise State University‘
Hutchinsjh@hotmail.com

Derek A. Cline
University of Utah
Chemical Engineering Department
derek.cline@chemeng.utah.edu

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

PP2

Finite Element Analysis of Free Material Optimization Problems

In Free Material Optimization, the design variable is the full material tensor of an elastic body. Written in matrix notation one obtains a control-in-the-coefficients problem for the material tensor. With this poster we present recent results in the finite element analysis in Free Material Optimization. We employ the variational discretization approach, where the control (i.e., material tensor) is only implicitly discretized. Numerical examples supplement our analytical findings.

Tobias Jordan
University of Hamburg
Department of Mathematics
tobias.jordan@uni-hamburg.de

Michael Hinze
Universität Hamburg
Department Mathematik
michael.hinze@uni-hamburg.de

PP2

Higher Order Numerical Schemes for Convection Diffusion Equation Based on B-Spline Quasi-Interpolation

In the present work, we propose B-Spline Quasi-Interpolation (BSQI) based higher order numerical scheme for convection-diffusion equation in one and two space dimensions. The linear stability of BSQI scheme is establish using the von-Neumann analysis. We find the CFL condition under which BSQI scheme is stable. Numerical experiments are performed for the non-linear problems like Burgers' equation, Buckley-Leverett, and the incompressible flow to measure the accuracy and rate of convergence of the BSQI scheme.

Rakesh Kumar
IIT Bombay
rakeshk@math.iitb.ac.in

Sambandam Baskar
Indian Institute of Technology Bombay
baskar@math.iitb.ac.in

PP2

NIST AMR Benchmarks

Adaptive mesh refinement (AMR) techniques for the numerical solution of PDEs have been under development for many years. Most research papers conclude with numerical results to demonstrate the effectiveness of a proposed method. Although there are some commonly used problems, like the L-domain problem, many disparate test problems have been used. NIST has mined the AMR literature to create a collection of standard test/benchmark problems for AMR. The resulting web resource will be demonstrated.

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

PP2

Finite Element Methods for the Evolution Problem in General Relativity

Gravitational waves can be understood as small ripples in the fabric of the Universe, caused by moving masses. To detect such weak waves, several new gravitational wave observatories are being built. Computer simulations are essential for determining expected signals and interpreting the data. This project consists of the design and implementation of a new mixed finite element method for the propagation of gravitational waves, by adapting the recently developed Finite Element Exterior Calculus framework.

Vincent Quenneville-Belair
School of Mathematics
University of Minnesota
vqb@umn.edu

PP2

Massively Parallel Radiation Transport Sweeps on Unstructured Grids

The massively parallel radiation transport code PDT has recently scaled to 432,000 processes with an efficiency greater than 70%. The discretization techniques employ a Discontinuous FEM method in space and a discrete-ordinate collocation in angle. The algorithm is based on a transport sweep, whereby the solution for a given angle is done one cell at a time. The parallel transport sweep algorithm is provably optimal for logically Cartesian meshes. However, complex geometries cannot be efficiently meshed with regular grids, even with point motion. In this work, we present an unstructured mesh generation capability that satisfies most of the requirements for optimal sweeps. Notably, the subdomain partitions remain convex and pipe-filling within a subdomain is preserved; however, a load imbalance in terms of spatial unknowns may exist between subdomains due to local geometrical features. Numerical results are presented and strategies to mitigate

spatial load imbalance are discussed.

Jean C. Ragusa, Tarek Ghaddar
Department of Nuclear Engineering
Texas A&M University
jean.ragusa@tamu.edu, tghaddar92@gmail.com

Michael Adams
Texas A&M University
i.am.michael.adams@gmail.com

PP2

The Discrete Maximum Principle in the Family of Mimetic Finite Difference Methods

The Maximum Principle is the important property of PDEs. To mimic this property in simulations is very desirable in wide range of applications. The mimetic finite difference method produces a family of schemes with equivalent properties such as the stencil, stability region, and convergence order. Each member is defined by parameters which can be chosen locally for every cell. We present a new adaptation methodology that identifies members which satisfies the discrete maximum principle.

Daniil Svyatskiy
Los Alamos National Laboratory
dasvyat@lanl.gov

Gianmarco Manzini
LANL
gmanzini@lanl.gov

Konstantin Lipnikov
Los Alamos National Laboratory
lipnikov@lanl.gov

PP2

A Locally Adaptive RBF-FD Method

Conventional Radial Basis Function (RBF) methods for numerically solving partial differential equations use global approximations resulting in dense matrices that grow in size if the algorithm refines. The Radial Basis Function – Finite Difference (RBF-FD) approach is a local approximation method that utilizes nearest neighbor nodes and yields a sparse implementation. Unfortunately RBF-FD matrices have fixed stencils and the approximations can lose accuracy. In this paper we propose using local approximations with locally adaptive stencils that take advantage of both global and local approximations. In this approach the stencil sizes stay fixed where the solution is smooth but grow in size only where refinement is needed. The advantage of this method is that it is computationally efficient and stable offering comparable accuracy to global approximations with significantly lower computational cost.

Wade Meyers
University of Wisconsin - Stout
meyersw3476@my.uwstout.edu

Talin Masihimirzakhanian
Cal Poly Pomona
talinm@csupomona.edu

Keith Wojciechowski
University of Wisconsin - Stout
wojciechowskik@uwstout.edu

PP2

A Second-Order Maximum Principle Preserving Lagrange Finite Element Technique for Nonlinear Scalar Conservation Equations

Based on one of our accepted paper (*J.-L. Guermond et al.*), this poster will show an explicit, (at least) second-order and maximum principle satisfying lagrange finite element method for solving nonlinear scalar conservation equations. The algorithm works for arbitrary meshes in any space dimension and for all Lipschitz fluxes. The formal second-order accuracy and the convergence properties are tested on a series of linear and nonlinear benchmark problems.

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

Murtazo Nazarov
TAMU
murtazo@math.tamu.edu

Bojan Popov
Department of Mathematics
Texas A&M University
popov@math.tamu.edu

Yong Yang
Texas A&M University
yytamu@math.tamu.edu

PP3

A New Test for Exclusion Algorithm to Find the Optimum Value of Function in R^n

The problem of finding the global minimum of a vector function is very common in science, economics and engineering. One of the most notable approaches to find the global minimum of a function is that based on interval analysis. In this area, the exclusion algorithms (EAs) are a well-known tool for finding the global minimum of a function over a compact domain. There are several choices for the minimization condition. In this paper, we introduce a new exclusion test and analyze the efficiency and computational complexity of exclusion algorithms based on this approach. We consider Lipschitz functions and give a new minimization condition for the exclusion algorithm. Then we study the convergence and complexity of the method.

Ibraheem Alolyan
King Saud University
ialolyan@ksu.edu.sa

PP3

Approximation and Error Estimation in High Dimensional Space for Stochastic Collocation Methods on Arbitrary Sparse Samples

We have develop a fast method that can capture piecewise smooth functions in high dimensions with high order and low computational cost. This method can be used for both approximation and error estimation of stochastic simulations where the computations can either be guided or

come from a legacy database. We demonstrate how this method compares to Gaussian processes.

Richard Archibald
Computational Mathematics Group
Oak Ridge National Labratory
archibaldrk@ornl.gov

PP3

Rational Least Squares Fitting using Krylov Spaces

For given square matrices A, F and a vector $\mathbf{v}$, we consider the problem of finding a rational function $R_m^{\min}$ of type (m, m) such that

$$\|F\mathbf{v} - R_m(A)\mathbf{v}\|_2^2 \to \min,$$

and propose an iterative algorithm for its solution. In the special case when $A = \mathrm{diag}(\lambda_j)$ and $F = \mathrm{diag}(\psi_j)$ are diagonal we have a weighted rational least squares fitting problem $\sum_{j=1}^{N} |v_j|^2 \cdot |\psi_j - R_m(\lambda_j)|^2 \to \min$, and compare our method to the popular vector fitting algorithm.

Mario Berljafa
School of Mathematics
The University of Manchester
m.berljafa@maths.man.ac.uk

Stefan Guettel
The University of Manchester
stefan.guettel@manchester.ac.uk

PP3

A Hybrid Openmp/mpi Cg Iterative Eigensolver for First-Principles Plane Wave Materials Science Codes

First-principles materials science codes based on density functional theory (DFT) and using plane waves (PW) have become the largest user (by method) of computer cycles at scientific computer centers around the world. We present a hybrid OpenMP/MPI Conjugate Gradient based iterative eigensolver that allows this approach to scale to tens of thousands of cores on modern many core parallel computers. Performance results will be presented for the Cray XE6 and XC30 architectures.

Andrew M. Canning
Lawrence Berkeley National Laboratory
acanning@lbl.gov

PP3

All Real Eigenvalues of Symmetric Tensors

This poster displays how to compute all real eigenvalues of a symmetric tensor. As is well known, the largest or smallest eigenvalue can be found by solving a polynomial optimization problem, while the other middle eigenvalues can not. We propose a new approach for computing all real eigenvalues sequentially, from the largest to the smallest. We show that each eigenvalue can be computed by solving a finite hierarchy of semidefinite relaxations. Numerical experiments are presented.

Cui Chunfeng, Dai Yu-Hong
Academy of Mathematics and Systems Science,
Chinese Academy of Sciences
cuichf@lsec.cc.ac.cn, dyh@lsec.cc.ac.cn

Nie Jiawang
Department of Mathematics,
University of California San Diego
njw@math.ucsd.edu

PP3

A Python Toolbox for Shape Optimization in Imaging and Data Analysis

Many data analysis and image processing problems are naturally expressed as shape optimization problems, e.g. image segmentation, surface reconstruction. The shape optimization approach is a great fit for such problems because of its intuitiveness and the ?exibility to easily incorporate data ?delity, geometric regularization and statistical prior terms. However, carrying out the actual minimization in an ef?cient and reliable manner requires overcoming many technical challenges. In this work, we introduce a Python toolbox that implements a diverse collection of shape energies for image processing, and state-of-the-art optimization methods to compute their solutions.

Gunay Dogan
Theiss Research, NIST
gunay.dogan@nist.gov

PP3

Discovering Block Structure in Graphs with Approximate Eigenvectors

Graphs and Networks are important in modeling structure across disciplines. We demonstrate that graph partitioning in a minimum cut sense can be improved with an ensemble of low-fidelity eigenvectors which can outperform a single high-fidelity eigenvector. These ensembles can be computed faster and be more helpful than one high-fidelity eigenvector. Since the individual ensemble members are independent they can be computed in parallel.

James Fairbanks
Georgia Institute of Technology
james.fairbanks@gatech.edu

Geoffrey D. Sanders
Center for Applied Scientific Computing
Lawrence Livermore National Lab
sanders29@llnl.gov

PP3

Efficient Multigrid Methods for Distributed Optimal Control Problems Constrained by Parabolic Equations

In this work we present numerical computations in support of our theoretical results regarding convergence properties of multigrid preconditioners for linear systems arising in the solution process of space-time distributed optimal control problems constrained by linear and semi-linear parabolic equations. As for elliptic-constrained problems, the number of preconditioned linear iterations per optimization iteration is shows to decrease with increasing resolution, but the rate of decrease is suboptimal by half an order.

Mona Hajghassem
Department of Mathematics and Statistics
University of Maryland, Baltimore County
mona4@umbc.edu

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

PP3

Computing the Heat Kernel of a Graph for a Local Clustering Algorithm

We present an efficient local clustering algorithm that finds cuts in large graphs by performing a sweep over a heat kernel pagerank vector. We show that for a subset S of Cheeger ratio ϕ, many vertices in S may serve as seeds for a heat kernel random walk which will find a cut of conductance $O(\sqrt{\phi})$. Further, the random walk process is performed in time sublinear in the size of the graph.

Olivia Simpson
University of California, San Diego
oliviamsimpson@gmail.com

PP4

Reduced Order Modelling for Optimal Cancer Treatment

We study reduced order modelling for optimal radiotherapy treatment plan. Boltzmann equation is used to model the interaction between radiative particles with tissue. At first, we solve optimization problems: minimizing the deviation from desired dose distribution. Then we consider a parameterized geometry. In offline stage we solve a problem for sampled parameter values. The online phase then consists of solving the reduced problem for the actual set of parameters. Numerical results are presented.

Bahodir Ahmedov
Aachen Institute for Advanced Study in
Computational Engineering Science (AICES)
ahmedov@aices.rwth-aachen.de

Michael Herty, Martin Grepl
RWTH Aachen University
herty@igpm.rwth-aachen.de, grepl@igpm.rwth-aachen.de

PP4

A Mesh Free Method for Numerical Simulation of Calcium Dynamics In Ventricular Myocytes

We consider a coupled system of non-linear reaction-diffusion equations that model the spatio-temporal variation of intracellular calcium concentration in ventricular myocytes. We introduce a modified mesh free method and utilize exponential time differencing to significantly reduce the simulation time. At the end we present numerical results demonstrating the stability of the method when used on uneven distribution of nodes.

Emmanuel O. Asante-Asamani
Department of Mathematical Sciences
University of Wisconsin-Milwaukee
eoa@uwm.edu

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

Zeyun Yu
University of Wisconsin-Milwaukee
yuz@uwm.edu

PP4

Computed Tear Film and Solute Dynamics on An Eye-Shaped Domain

The concentration of ions (osmolarity) in the tear film is a key variable to understanding its dynamics. We derive a system of nonlinear partial differential equations (PDEs) that couples solutes and fluid dynamics on a 2D eye-shaped domain. We solve these PDEs using the Overture computational framework with a hybrid BDF/RKC time-stepping scheme. Our results agree with existing 1D models and provide new insight into the osmolarity distribution and fluorescence imaging for in vivo experiments.

Richard Braun
University of Delaware
Department of Mathematical Sciences
braun@math.udel.edu

Longfei Li
Department of Mathematical Sciences
University of Delaware
longfei@udel.edu

Tobin Driscoll
University of Delaware
Mathematical Sciences
driscoll@udel.edu

William Henshaw, Jeffrey Banks
Rensselaer Polytechnic Institute
henshw@rpi.edu, banksj3@rpi.edu

P. Ewen King-Smith
Ohio State University
king-smith.1@osu.edu

PP4

The Transcriptomic Clock of Human Cerebral Cortex Development

Singular value decomposition and clustering analysis of stem cell RNA transcription data reveals the genes associated with the stages of human embryonic cerebral cortex development. The method discovered a previously unidentified stage between pluripotency and neural differentiation containing distinct transcriptional patterns for over two thousand differentially expressed genes. Enrichment analyses of genes associated with neurological diseases with respect to the resulting corticogenesis clock reveal distinct stages of development associated with root causes of the disease.

Elisabeth M. Brown
Rensselaer Polytechnic Institute
browne6@rpi.edu

Kristin Bennett
Mathematical Sciences
Rensselaer Polytechnic Institute
bennek@rpi.edu

Hannah De Los Santos, Joey Lea
Rensselaer Polytechnic Institute
delosh@rpi.edu, leaj2@rpi.edu

Nathan Boles, Thomas Kiehl, Sally Temple, Christopher Fasano
Neural Stem Cell Institute
nathanboles@neuralsci.org, tomkiehl@neuralsci.org, sallytemple@neuralsci.org, chrisfasano@neuralsci.org

PP4

Computational Methods to Study the Coordination of Mechanical Forces Involved in Amoeboid Cell Migration

We present a computational model to study the interplay of cellular mechanics, substrate mechanics and cell-matrix interaction and the resulting migration. Our mathematical framework considers a porous viscoelastic cytoplasm, adhesion dynamics, and substrate mechanics. Our model introduces a novel way of simulating a viscoelastic deforming network. Using our methodology we present insight into the 3D cell-substrate forces for cells migrating on flat substrates. The development and maintenance of multicellular organisms relies on cell migration. During these processes, cells encounter a wide range of extracellular environments and adapt their migration strategy in response to mechanical properties of their environment. One of the current challenges in cellular biology is to understand how the extracellular environment modulates the deployment of the cells' molecular machinery to produce different migratory behaviors.

Calina A. Copos
University of California Davis
ccopos@math.ucdavis.edu

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

PP4

Segmentation and Processing of Brain Images of Multiple Modalities in 2 and 3 Dimensions

Medical imaging of various modalities continues to advance, providing higher resolution and higher signal to noise ratios than ever before. Software must keep up by not only accounting for better quality images, but also with larger and larger datasets. We present work on segmenting and processing Magnetic resonance (MRI) and Electron microscopy (EM) images using advanced algorithms designed to scale in both 2 and 3 dimensions.

John Edwards, Brian Summa
Scientific Computing and Imaging Institute
University of Utah
jedwards@sci.utah.edu, bsumma@sci.utah.edu

Valerio Pascucci
SCI Institute - University of Utah
pascucci@sci.utah.edu

Christopher Johnson
University of Utah
Department of Computer Science
crj@sci.utah.edu

PP4

Improving Performance of Multi-Level Nonrigid Registration of Two Ct-Based Lung Images With Use of Gpu Computing

Graphics Processing Units (GPUs) have been successfully employed to accelerate scientific computing applications in several disciplines, including medical image processing. We demonstrate a novel computation- and memory-efficient diffeomorphic multi-level B-Spline transform composite method on GPUs for improving performance of non-rigid registration of two CT lung images. The GPU method is compared against its CPU counterpart, with GPU performance 112 times faster than the single-threaded CPU version occurring at highest resolutions while preserving accuracy.

Nathan Ellingwood, Youbing Yin
University of Iowa
nathan.ellingwood@gmail.com, youbing.yin@gmail.com

Matthew Smith
National Cheng Kung University
msmith@mail.ncku.edu.tw

Ching-Long Lin
University of Iowa
ching@engineering.uiowa.edu

PP4

Reaction of a Solid Tumor According to the Injection of Medical Supplies into Heart and Liver

A mathematical model is developed to describe the variation of a solid tumor cells density in response to medical supplies. Two factors, random motility and chemotaxis in response to TAF gradients are considered for the equation of tumor cells motion. The flow rate of medicines exerts influence on tumor cells density and tumor cells react sensitively to the medical supplies at the first second, and the density decreases to around 20% of the initial amount.

Jaegwi Go
Changwon National University
jggo@changwon.ac.kr

PP4

Shear Wave Filtering in Bouligand Structures

We propose that the presence of Bouligand-like structure in the dactyl club of the Stomatopod lead to wave filtering. The (nearly) periodicity of the microstructure suggest an interaction between the microstructure and propagating stress waves. The propose model use a combination of propagator matrix approach and the Floquet-Bloch theorem. From these combined analyses we compute the transmitted energy for different microstructural configurations.

Nicolas Guarin Zapata
Purdue University
nicoguarin@gmail.com

Juan Gomez
Universidad EAFIT
jgomezc1@eafit.edu.co

Nick Yaraghi, David Kisailus
University of California
nyaraghi@gmail.com, david@engr.ucr.edu

Pablo Zavattieri

Purdue University
zavattie@purdue.edu

PP4

Newtonian and Non-Newtonian Fluid Dynamics in Abdominal Aortic Aneurysms

Biomedical research has recently indicated that some specific dynamic characteristics, such as the blood wall shear stress and oscillatory shear index, of the blood flow inside arteries with aneurysms are risk factors for both the enlargement and rupture of the associated aneurysm. The primary objective of the project is to determine the influence that the geometry of an abdominal aortic aneurysm and the fluid constitutive model have on these specific characteristics.

Danielle D. Masse, Jason Howell
College of Charleston
massedd@g.cofc.edu, howelljs@cofc.edu

PP5

Adaptive Spectral Tensor-Train Decomposition for the Construction of Surrogate Models

We present a novel method for approximating high-dimensional functions, the cost of which scales linearly with the input parameter dimensionality. It hinges on the combination of the low-rank approximation of functions through the functional form of the tensor-train decomposition and on the theory of polynomial approximation [D. Bigoni, Y. M. Marzouk, and A. P. Engsig-Karup, Spectral tensor-train decomposition, arXiv preprint arXiv:1405.5713], using anisotropic adaptive strategies to meet the desired accuracy. The method is relevant for high-dimensional Uncertainty Quantification and inference. Synthetic and real applications will be presented.

Daniele Bigoni, Allan Engsig-Karup
The Technical University of Denmark
dabi@dtu.dk, apek@dtu.dk

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

PP5

A Posteriori Error Estimation for a Cut Cell Finite Volume Method in the Presence of Uncertainty

We consider a posteriori error estimates for interface problems. These are differential equations whose data is defined piecewise across a curve, or interface, partitioning the domain into two sections. Often, this interface is determined from a small number of measurements, each of which has uncertainty associated with it, giving a stochastic interface problem. We develop a method of computing the error for each sample interface that is independent of the number of samples.

James B. Collins, Simon Tavener, Don Estep
Colorado State University
jbcolli2@gmail.com, tavener@math.colostate.edu, estep@stat.colostate.edu

PP5

Probability Measures on Numerical Solutions of Odes for Uncertainty Quantification and Inference

Deterministic ODE solvers are widely used, but characterizing the error in numerical solutions within a coherent statistical framework is challenging. We successfully address this problem by constructing a probability measure over functions consistent with the ODE solution that provably contracts to a Dirac measure on the unique solution at rates determined by an underlying deterministic solver. The measure straightforwardly derives from important classes of numerical solvers and is illustrated on uncertainty quantification and inverse problems.

Patrick R. Conrad
Massachusetts Institute of Technology
p.conrad@warwick.ac.uk

Mark Girolami
University College London
m.girolami@warwick.ac.uk

Simo Sarkka
Aalto University
simo.sarkka@aalto.fi

Andrew Stuart
Mathematics Institute,
University of Warwick
a.m.stuart@warwick.ac.uk

Konstantinos Zygalakis
University of Southampton
k.zygalakis@soton.ac.uk

PP5

Adaptive Bayesian Selection, Calibration, and Validation of Coarse-Grained Models of Atomistic Systems

The predictive power of coarse-grained (CG) approximations of atomistic systems is explored. Bayesian methods for statistical calibration, validation, and model selection using model plausibilities are used to develop basic principles for developing CG and, eventually, macro-scale models. An adaptive algorithm for Bayesian calibration and selection of CG models is described and examples of application to polymer chains is presented.

Kathryn Farrell
Institute for Computational Engineering and Sciences
University of Texas at Austin
kfarrell@ices.utexas.edu

J. Tinsley Oden
The University of Texas at Austin
ICES
oden@ices.utexas.edu

Danial Faghihi
Institute for Computational Engineering and Sciences
University of Texas at Austin

danial@ices.utexas.edu

PP5

Emgr - Empirical Gramian Framework

Gramian-based model reduction is a well established method for linear state-space systems. Beyond linear systems, empirical gramians expand the scope of gramian-based methods to nonlinear systems. Furthermore, empirical gramians can also be used for parametric model order reduction, parameter identification and parameter reduction. The empirical gramian framework is a Matlab software toolbox enabling the computation of seven types of empirical gramians, which have applications in model reduction, system identification and uncertainty quantification.

Christian Himpe
Institute for Computational und Applied Mathematics
Universtiy of Muenster
christian.himpe@uni-muenster.de

Mario Ohlberger
Universität Münster
Institut für Numerische und Angewandte Mathematik
mario.ohlberger@uni-muenster.de

PP5

Kernel Density Estimation for Implicit Monte Carlo Radiation Transport

We use kernel density estimation in the Fleck-and-Cummings implicit Monte Carlo method to obtain smooth solution estimates in thermal radiative transfer problems. The kernel density estimators obtain conservative estimates when using reflective boundary corrections and resolve steep gradients with locally adaptive bandwidths. We show that solutions obtained using kernel density estimators are smoother and exhibit substantially less statistical noise than traditional histogram tallies.

Aaron M. Holgado
Texas A&M University
amholgado1@gmail.com

Robert Holladay
Virginia Polytechnic Institute and State University
rholla24@vt.edu

Allan Wollaber
LANL
wollaber@lanl.gov

Mathew Cleveland, Todd Urbatsch
Los Alamos National Laboratory
cleveland@lanl.gov, tmonster@lanl.gov

Ryan McClarren
Texas A&M University
USA
rgm@tamu.edu

PP5

Distribution Functions of Water Saturation for the Stochastic Buckley-Leverett Problem Via the Streamline Method

We give an analytical expression for the one-point distribution functions of the water saturation for the 1D stochastic Buckley-Leverett problem with uncertainty in porosity and in the total Darcy flux. These distribution functions particularly lead to any one-point statistics of the water saturation. Comparisons with both MC simulations and a low order approximation approach are provided. Finally, based on the streamline concept, a generalization to multiple spatial dimensions is outlined.

Fayadhoi Ibrahima
Stanford University
fibrahim@stanford.edu

Daniel W. Meyer
Institute of Fluid Dynamics
meyerda@ethz.ch

Hamdi Tchelepi
Stanford University
Energy Resources Engineering Department
tchelepi@stanford.edu

PP5

Hybridized Reduced Basis Method and Generalized Polynomial Chaos for Solving Partial Differential Equations

The generalized Polynomial Chaos(gPC) method is a popular method for solving partial differential equations (PDEs) with random parameters. However, when the probability space has high dimensionality, the solution ensemble size required for an accurate gPC approximation can be large.We show that this process can be made more efficient by closely hybridizing gPC with Reduced Basis Method(RBM).

Jiahua Jiang
University of Massachusetts, Dartmouth
jjiang@umassd.edu

PP5

Kriging and Spatial Design Accelerated by Orders of Magnitude: Combining Low-Rank Covariance Approximations with FFT-Techniques

Computational power poses heavy limitations to the achievable problem size for Kriging. In separate research lines, Kriging algorithms based on FFT and low-rank representations of covariance functions have been developed, both leading to drastic speedup factors. The current study combines these ideas, reducing the computational complexity of Kriging to $\mathcal{O}(k_q m d L^* \log L^*)$, where k_q is the rank of approximation, m the number of measurements, d dimension, L^* the number of lattice points along the longest edge of the regular d-dimensional lattice. These benefits can be fully exploited when leaving the final result in low-rank format, or when further low-rank operations follow. The current study assumes second-order stationarity and simple Kriging on a regular, equispaced lattice.

Alexander Litvinenko
SRI-UQ Center, King Abdullah University of Science and Techn
alexander.litvinenko@kaust.edu.sa

Wolfgang Nowak

Institute for Modelling Hydraulic and Environmental Systems
University of Stuttgart
wolfgang.nowak@iws.uni-stuttgart.de

PP5

A Nonlinear Non-Gaussian Smoother for Continuous Stochastic Dynamical Systems

We present a novel non-Gaussian smoothing methodology for high-dimensional stochastic fields governed by general nonlinear dynamics. The history of the system is quantitatively estimated using its stochastic dynamics and information gathered from noisy observations. Uncertainty is quantified using the reduced-order Dynamically-Orthogonal equations, and smoothing is performed by efficiently carrying out Bayesian-inference in an evolving low-dimensional dominant subspace. Various examples from computational-fluid-dynamics are provided, illustrating the superior performance of the smoother compared to its Gaussian/Monte-Carlo counterparts.

Tapovan Lolla, Pierre Lermusiaux
MIT
ltapovan@mit.edu, pierrel@mit.edu

PP5

Uncertainty Quantification in Incompressible Flow Using Sparse Grids

This work based on a bachelor thesis introduces a sparse grid stochastic collocation method for uncertainty quantification. We adapted the C++ sparse grid library SGpp for benchmark scenarios in Matlab. The method was tested on a simple ODE example as well as on an incompressible flow scenario in 1-D and 3-D random parameter space. The sparse grid collocation solutions show promising results with respect to accuracy and runtime compared to Monte Carlo methods.

Friedrich Menhorn
Department of Informatics, Technische Universität München
menhorn@in.tum.de

Tobias Neckel
Technische Universität München
neckel@in.tum.de

PP5

Matrix Splitting Techniques for Sampling a High-Dimensional Gaussian

Langevin and Hamiltonian proposals in the Metropolis-Hastings algorithm applied to Gaussian target distributions correspond to matrix splittings, similar to stationary iterative methods for solving systems of linear equations such as Gauss-Seidel and SSOR. We prove a result for how the efficiency of the Metropolis-Hastings algorithm depends on a general matrix splitting in high dimensions, revealing new efficient proposals for the Metropolis-Hastings algorithm.

Richard A. Norton, Colin Fox
University of Otago
Department of Physics
richard.norton@otago.ac.nz, colin.fox@otago.ac.nz

PP5

Time Series Estimation of a Stochastic Processes Coupled to Pdes for Multiscale Modeling

The well-studied molecular mechanics of atomic irradiation damage have long-standing simulation codes modeling binary collision cascades. Long range spatial correlations between damages are generally modeled using phase-field models. However, we propose a simulation scheme to model the meso-scale void accumulation damages using a stochastic cellular automata. This model allows us to bridge the gap between the micro and macro scales by modeling the stochastic cellular automata process using a time series. This time series can then be run at a much less computationally intensive simulation until reaching a steady state. The statistical properties of the steady state can then be coupled to the macro scale pde, ie continuum models, in a form that it can better understand.

Charlie Vollmer, Don Estep
Colorado State University
charlesv@rams.colostate.edu, estep@stat.colostate.edu

Anter A. El-Azab
Purdue University
aelazab@purdue.edu

PP5

Fast Stochastic Simulation of Non-Gaussian Correlated Process Variations

The recently developed stochastic spectral methods can estimate more efficiently the process variations effects compared with Monte-Carlo. However, existing publications mostly assume variation-parameters to be *independent* and *Gaussian*. In this paper, we develop an efficient simulation technique based on stochastic collocation for ***non-Gaussian*** and ***correlated*** random parameters. The technique is applied to silicon photonic process variations and shows **124**-times speedup compared with Monte-Carlo.

Tsui-Wei Weng
MIT
Research Lab of Electronics
twweng@mit.edu

Zheng Zhang
MIT
z_zhang@mit.edu

Luca Daniel
M.I.T.
Research Lab in Electronics
luca@mit.edu

PP5

Uncertainty Quantification for Integrated Circuits and MEMS

We present a simulator to quantify the uncertainties of nano-scale integrated circuits and microelectromechanical systems (MEMS). We show the stochastic spectral methods for stochastic static, transient and periodic steady-state (i.e., limit cycle) analysis, as well as some techniques to handle high parameter dimensionality and design hierar-

chy.

Zheng Zhang
MIT
z_zhang@mit.edu

Ibrahim Elfadel
Microsystem Engineering
Masdar Institute of Science and Engineering
ielfadel@masdar.ac.ae

Luca Daniel
M.I.T.
Research Lab in Electronics
luca@mit.edu

PP6

Reproducible Numerical Computing with HashDist

HashDist provides a critical component of the development workflow, enabling highly customizable, source-driven, and reproducible builds for scientific software stacks. HashDist features intelligent caching of sources and builds, parametrized build specifications, and the ability to interoperate with system compilers and packages. HashDist enables the easy specification of "software stacks", which allow both the novice user to install a default environment and the advanced user to configure every aspect of their build in a modular fashion. All HashDist builds are reproducible, with a unique build hash identifying exactly how each component of the software stack was installed.

Aron Ahmadia
US Army Engineer Research and Development Center
aron@ahmadia.net

Ondrej Certik
Department of Mathematics and Statistics
University of Nevada, Reno
ondrej@certik.cz

Christopher Kees
US Army Engineer Research and Development Center
christopher.e.kees@usace.army.mil

Dag Sverre Seljebotn
University of Oslo
d.s.seljebotn@astro.uio.no

Andy R. Terrel
Institute for Computational Engineering and Sciences
University of Texas at Austin
andy.terrel@gmail.com

PP6

A Scalable Fast Method for N-Body Problems Based on Exact Finite Element Basis Screen Functions

We introduce a fast method for computing N-body interactions based on the particle-particle–particle-mesh (P^3M) approach in which the calculation is split into rapidly decaying short-range interactions and mesh-resolvable long-range interactions. Our method employs screening functions designed to yield a long-range component that is found exactly by a finite element method, increasing the locality of the method compared to traditional FFT-based methods. Multigrid methods efficiently solve the resulting sparse matrix problem.

Natalie N. Beams
University of Illinois at Urbana-Champaign
beams2@illinois.edu

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

Jonathan B. Freund
University of Illinois at Urbana-Champaign
jbfreund@illinois.edu

PP6

Boltzmann Collision Operator for Cylindrically Symmetric Velocity Distributions in Plasmas

We develop a model for collision processes in industrially relevant plasmas. To reduce the computational cost of solving the Maxwell-Boltzmann equations, we assume that the velocity distribution function is cylindrically symmetric in velocity space and only axially dependent in physical space. We show that if the external force is only axially dependent, then the Boltzmann collision operator is also cylindrically symmetric, which reduces the Maxwell-Boltzmann system to two velocity and one spatial dimensions.

Yanping Chen, Yannan Shen, John Zweck, Matthew Goeckner
University of Texas at Dallas
yxc110030@utdallas.edu, yxs135630@utdallas.edu, jwz120030@utdallas.edu, goeckner@utdallas.edu

PP6

Analysis of a Heterogeneous Multiscale Method for Poroelasticity

In this paper, we develop a highly parallelizable numerical method to solve the heterogeneous linear poroelasticity equations in multiple dimensions via operator splitting and a finite-volume based heterogeneous multiscale method for the linear elasticity and reaction diffusion equations. We demonstrate convergence both analytically and numerically, and analyze its computational complexity on high performance computers.

Paul M. Delgado
UTEP
pmdelgado2@utep.edu

Vinod Kumar, Son Young Yi
University of Texas at El Paso
vkumar@utep.edu, syi@utep.edu

PP6

Cell List Algorithms for Nonequilibrium Molecular Dynamics

A common approach in the molecular simulation of homogeneous linear background flow is the use of boundary conditions that deform with the flow. Recent developments have been made in finding long-term compatible boundary conditions for a general class of three-dimensional flows. We present two modifications of the standard cell list algo-

rithm for nonequilibrium molecular dynamics. The modified algorithms handle the dynamic, deforming simulation geometry and reduce the computational complexity of force computations from $O(N^2)$ to $O(N)$ in the number of particles N.

Matthew Dobson
University of Massachusetts Amherst
dobson@math.umass.edu

PP6

Convex-Hull Classification of Molecular Data on a Cluster

There has been considerable interest in understanding variations in molecular signatures between normal and disease states using novel classification approaches. The proposed study will elucidate the choice of convex-hull ensemble classification for discerning distinct disease groups. The algorithms were implemented in R and parallelized across a high-performance computing cluster using readily available wrappers.

Sally R. Ellingson, Radha Nagarajan
University of Kentucky
sally@kcr.uky.edu, rnagarajan@uky.edu

PP6

Spectral Noise Filtering for Fourier Transform Profilometry

In Fourier transform profilometry, an optical system projects light with a sinusoidally-varying pattern on a 3-D surface and records the resulting image. The spectrum of the depth-modulated sinusoid is overlapped by other components of the spectrum, limiting accurate reconstruction of the surface. We present a new noise removal method to estimate the obscuring spectral components and filter them out. Surfaces reconstructed using this de-noising method feature increased precision and suffer from fewer aberrations.

Thomas Höft
University of St. Thomas
Department of Mathematics
hoft@stthomas.edu

PP6

Direct Evaluation of Unified Extended Splines

In this poster we demonstrate a method for direct evaluation of unified extended spline basis functions, which were first introduced by Wang and Fang in 2008. We also demonstrate the advantages of Bézier extraction of unified extended splines in Isogeometric analysis.

Ian D. Henriksen
Brigham Young University
insertinterestingnamehere@gmail.com

PP6

Numerical Modeling of Wave Propagation in Poroelastic Media Using Optimal Staggered Implicit Finite Differences

We apply a recently proposed optimal staggered implicit finite difference scheme to model wave propagation in poroelastic media. The advantage of this method is that it has high accuracy and only requires to solve tridiagonal linear systems of equations (to compute the spatial derivatives). The instabilities that arise in high contrast media are alleviated when using these schemes. We show examples for different media configurations.

Ursula Iturraran-Viveros, Reymundo Itza
Facultad de Ciencias
Universidad Nacional Autonoma de Mexico
ursula.iturraran@gmail.com, reymundo.itza@gmail.com

PP6

Graph-Based Analysis of Three-Dimensional, Large Scale Phase-Field Simulations

In large scale phase-field simulations of ternary eutectic solidification, multiple patterns emerge and evolve while growing. To study the transformation of patterns depending on thermodynamical properties or processing conditions and to analyze the respective microstructure characteristics, algorithms exploiting large datasets are discussed. Basic features of the arrangement of the three solid phases are derived from the voxel data via graph representations. The targeted reduction of the data allows to visualize single three-dimensional structures and evaluate the correlations of different patterns effectively. Thereby it is possible to augment current 2D statistics and analyze projections of the data.

Marcus Jainta
Karlsruhe Institute of Technology (KIT)
Institute of Applied Materials (IAM-ZBS)
marcus.jainta@kit.edu

Daniel Stubenvoll, Johannes Hötzer
Karlsruhe University of Applied Sciences (HSKA)
Institute of Materials and Processes (IMP)
stubenvoll.daniel@gmx.de,
johannes.htzer@hs-karlsruhe.de

Phillip Steinmetz, Britta Nestler
Karlsruhe Institute of Technology (KIT)
Institute of Applied Materials (IAM-ZBS)
philipp.steinmetz@kit.edu, britta.nestler@kit.edu

PP6

Uncertainty Quantification for the Estimation of the Diffusion Coefficient from Md Simulations

The diffusion coefficient D of a tracer particle suspended in a fluid is estimated either from the time integral of the velocity autocorrelation (VACF) or from the slope of the mean-squared displacement (MSD). We derive relations between the statistical errors present in VACF or MSD and in D and show that the statistical errors in D estimated by both methods have the same variance if VACF and MSD are calculated from the same MD trajectories.

Changho Kim, George E. Karniadakis
Brown University
Division of Applied Mathematics
changho_kim@brown.edu, george_karniadakis@brown.edu

PP6

Adaptive Model Order Reduction in Forward and Inverse Multi-Frequency Problem for Maxwell's

Equations

A model order reduction method is developed for solution of a forward and inverse multi-frequency problem for magnetotellurics. The Helmholtz decomposition allows to extend the MOR method to the case of a on-trivial operator's null space. We use Pade approximation of the forward response and the jacobian, which are analytic functions of frequency. An adaptive choice of interpolating frequencies, based on minmax optimization of several error estimates, uses a fast calculation of the residual across a frequency range.

Michal A. Kordy
Department of Mathematics, University of Utah
Energy & Geoscience Institute, University of Utah
kordy@math.utah.edu

Elena Cherkaev
University of Utah
Department of Mathematics
elena@math.utah.edu

Phil Wannamaker
University of Utah
pewanna@egi.utah.edu

PP6

Sparse Spectral Tau-Method for Binary Neutron Stars

We describe ongoing work toward construction of initial data for binary neutron stars via a multidomain modal tau-method. Sparse systems are achieved through the use of "integration preconditioning". We focus on (i) realization of the low-regularity interface between the stellar surface and exterior through tau conditions and (ii) necessary further preconditioning beyond the "integration preconditioning".

Stephen Lau
Mathematics
The University of New Mexico
lau@math.unm.edu

PP6

Time-Parallel Approaches for Complex Rotorcraft Calculations

Unsteady rotorcraft CFD calculations may require hundreds of thousands of timesteps to resolve the flow of spinning rotors. Spatial decomposition alone limits the number of processors that can be effectively utilized; exploiting parallelism in the temporal dimension could significantly increase the degree of scalability. We demonstrate several approaches to achieve temporal parallelism, including the Time-Spectral method for periodic and quasi-periodic cases and a Parallel in Time (PIT) approach for the general case of fully unsteady flow.

Joshua I. Leffell
United States Army Aeroflightdynamics Directorate
NASA Ames Research Center
joshua.i.leffell@nasa.gov

Jay Sitaraman
University of Wyoming
jsitaram@uwyo.edu

Andrew Wissink
United States Army Aeroflightdynamics Directorate
NASA Ames Research Center
andrew.m.wissink.civ@mail.mil

PP6

Matched Asymptotic Analysis to Solve the Narrow Escape Problem in a Domain with a Long Neck

In this study, we mainly consider the narrow escape problem in a two-dimensional domain Ω with a long neck, which is the two-dimensional analogue of a dendritic spine geometry. The narrow escape problem requires the computation of the mean escape time of a Brownian particle starting from the head until it exits from the end of the neck, where the particle is absorbed. We divide the domain into the neck part Ω_n and the head part Ω_h, with the common boundary Γ_ε. The escape time in Ω_h can be considered to be the time from the head to the end of the neck, while the escape time in Ω_n can be considered to be the time from the neck to the end of the neck. We compute the two exit times separately and match them by considering some boundary value problem with an impedance boundary condition on Γ_ε, which we refer to as the Neumann-Robin boundary model.

Xiaofei Li
Department of Mathematics, Inha University
xiaofeili@inha.edu

Hyundae Lee
Inha University
Department of Mathematics
hdlee18@gmail.com

PP6

Computational Homogenization for the Modeling of Soft Matter Materials

This contribution outlines the development of mathematically rigorous and computationally efficient homogenization frameworks for soft matter with intrinsic network microstructures. Those are commonly encountered in materials such as elastomers, hydrogels, soft biological tissues, non-woven fabrics, cellular foams, and muscles and are all microscopically composed of elongated 1D fibers. When these soft materials are subject to a macroscopic strain, the underlying microstructure undergoes a peculiar deformation and highly affect the macroscopic stress response of the material.

Christian Linder
Stanford University
linder@stanford.edu

PP6

Precice – Flexible Parallel Multi-Physics Coupling

Flexible and extensible partitioned multi-physics simulation environments require efficient and modular tools with a broad coupling functionality. preCICE is a library for flexible numerical coupling of single-physics solvers. It uses a partitioned black-box coupling scheme, thus requiring only minimal modifications to existing solvers. Codes currently coupled with preCICE comprise both commercial and academic solvers, with a particular focus on fluid-structure interaction. preCICE features a clean and modern software design with extensive unit and integration

testing while maintaining minimal external dependencies. Inter-solver parallelism, parallel communication and data mapping techniques will help to max out future exa-scale computers.

Florian Lindner, Miriam Mehl
Universität Stuttgart
florian.lindner@ipvs.uni-stuttgart.de,
miriam.mehl@ipvs.uni-stuttgart.de

Benjamin Uekermann
TU Munich
uekerman@in.tum.de

PP6

A Sparse Interpolation Algorithm for Dynamical Simulations in Computational Chemistry

We present an implementation of Smolyak's sparse grid interpolation algorithm designed for simulating reaction paths of photo-induced molecular transformations. Current reaction path methods are computationally burdensome, but Smolyak's algorithm yields a cheap surrogate model. Furthermore, our implementation of Smolyak's algorithm facilitates the computation of several thousand reactions paths simultaneously. We describe our new implementation of Smolyak's algorithm and compare its performance to MATLAB's Sparse Grid Interpolation Toolbox. We also present simulation results for 2-butene.

James Nance
North Carolina State Univ
jdnance2@ncsu.edu

ELENA Jakubikova
North Carolina State University
ejakubi@ncsu.edu

C.T. Kelley
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu

PP6

P^{2}NFFT - A Versatile Framework for Computing NFFT-based Fast Ewald Summation

The Particle-Particle–NFFT (P^{2}NFFT) is a framework for computing the long range Coulomb interactions of the classical N-body problem with $\mathcal{O}(N \log N)$ arithmetic operations on massively parallel architectures. Recently, it has been generalized to 2d- and 1d-periodic boundary conditions by a combination of P^{3}M and fast summation techniques based on nonequispaced fast Fourier transforms (NFFT). We review these new algorithms and present performance results of our publicly available implementation.

Michael Pippig, Franziska Nestler
Chemnitz University of Technology
Department of Mathematics
michael.pippig@mathematik.tu-chemnitz.de,
franziska.nestler@mathematik.tu-chemnitz.de

PP7

Fields That Cause Elastic Breakdown in Inhomogeneous Media

We study the inverse problem of an elastic two phase body of unknown interior geometry. We investigate which boundary fields will cause failure in the material, either fracture or plastic yielding for brittle and ductile materials respectively. Using knowledge of volume fractions and material properties, bounds are obtained on the fields which the material can safely support. From this we are able to determine which boundary fields may be sustained by the body.

Nathan C. Briggs, Graeme Milton, Zoe Last Koch, Andrew Boyles, Jonathan Boyle, Michael Primrose, Michael Zhao
University of Utah
nathancbriggs@gmail.com, milton2505@yahoo.com, kochzoe@gmail.com, eldrewomagnifico@gmail.com, jonathan.boyle@utah.edu, michael.s.primrose@gmail.com, mz84092@gmail.com

PP7

Charge Transfer Processes at Semiconductor-Electrolyte Interfaces in Solar Cell Modeling

This project discusses results from the numerical approximation and simulation of reactive semiconductor-electrolyte interfaces in solar cells using a drift-diffusion-Poisson model. A mixed finite element method is employed to approximate the potential and the electric field, while a local discontinuous Galerkin method is employed to compute the densities and currents. The non-linear reactive interface conditions are treated using a Schwarz domain decomposition method applied to the semiconductor and electrolyte regions.

Michael D. Harmon
Institute for Computational and Engineering Sciences
mharmon@ices.utexas.edu

PP7

Student Chapter Develops Future Professionals

SIAM at Embry-Riddle Aeronautical University has fostered an unrelenting passion for professional development of its student body. From K-12 outreach to robotics research, members of our organization have gained skills that translate to industry-level capabilities. The talk describes the components of the organization that creates an environment for growth and enriches each student and industry partner connection.

Stacey Joseph-Ellison
Embry-Riddle Aeronautical University
ellisos1@my.erau.edu

PP7

Optimal Control of Miscible Displacement Equations Using Discontinuous Galerkin Methods

In the energy industry, reservoir simulators enable oil companies to optimize oil and gas production. I analyze the accuracy of the discontinuous Galerkin method when solving an optimal control problem for the miscible displacement equations, which model a tertiary oil recovery process. The control variables are the flow rates at the injection wells and the state variables are the fluid mixture pressure and velocity, as well as the concentration of the injected fluid.

Brianna Lynn
Rice University

bkl3@rice.edu

PP7

Elastic Deformation Due to Dislocations in a Transversely Isotropic Viscoelastic Halfspace

Many materials in nature have transversely isotropic (TI) viscoelastic behavior and this time-dependent phenomenon can be addressed by advanced continuum mechanics and computational analysis. A model is developed to represent the viscoelastic behavior of a TI media due to polygonal dislocation loops. Applying the Laplace transform to the constitutive equations and adopting the algorithm proposed by Honig and Hirdes for the inverse transform, the deformation fields in the physical domain will be obtained.

Amirhossein Molavi Tabrizi, Ernian Pan, Ali Sangghaleh
University of Akron
am142@zips.uakron.edu, pan2@uakron.edu, as147@zips.uakron.edu

PP7

Identifying and Tracking Multiple Underwater Acoustic Sources Using Characteristic Signatures

When using passive sonar to locate sound sources, the general practice is to implement a particle filter to track the source using its previous locations. Our goal is to improve this method by including discrete characteristics to the filter and calculating the likelihood that the source falls under one or a set of these characteristics. This will improve our estimates of source location by narrowing down movement patterns, will help us distinguish between multiple sources, and can give us valuable information such as what the source type is.

Zoi-Heleni Michalopoulou, Jacob Moorman, Jake Brusca, Shan Fung
Department of Mathematical Sciences
New Jersey Institute of Technology
michalop@njit.edu, jdm47@njit.edu, jb327@njit.edu, ssf22@njit.edu

PP7

Simulation and Modeling of Unmanned Systems for Humanitarian Applications in Industry

Unmanned systems are a burgeoning technology in the robotics and autonomy space. Utilized for decades in the military and defense complex, unmanned systems are now poised to play a key role in solving some of the most pressing international capacitance issues. With notable applications across the agricultural space, we will demonstrate the importance and utility of simulation and modeling of complex systems to solve real-world problems.

Courtney E. Thurston
Commonwealth Connections Academy
thurscon@gmail.com

PP8

AWM Workshop: Sampling and Reconstruction in Inite-Dimensional Reproducing Kernel Subspace

We will introduce quasi-optimal signal reconstruction of admissible sampling schemes in Banach space setting. We will investigate the admissibility of sampling schemes for inite-dimensional reproducing kernel subspaces of Lp. At last, we will propose an iterative approximation-projection algorithm with boundary adjustment for signals to live in a inite-dimensional reproducing kernel subspaces.

Cheng Cheng
UNIVERSITY OF CENTRAL FLORIDA
cheng.cheng@knights.ucf.edu

PP8

AWM Workshop: A Lattice of Poincare Duality Algebras with Acyclic Annihilators and Finite Dimension Associated to a Manifold

One is motivated by the algebraic theory of surgery to study finite dimensional differential commutative algebras with an invariant nondegenerate pairing. Using Hodge decompositions, we show there is a lattice of such algebras, which share many familiar algebraic properties of differential forms on manifolds. We hope to use these ideas to solve differential equations on manifolds, which use d, *, and , and to give a straightforward combinatorial description of what it means to be a manifold.

Cameron Crowe
Stony Brook University
cameroncrowe1@gmail.com

PP8

AWM Workshop: Residual Based Aposteriori Error Estimation in a Fully Automatic Hp –fem for the Stokes Equations

Aposteriori error estimator as a computable quantity in terms of known quantities gives a tool to assess the approximation quality in order to improve the solution adaptively. In This work we present a fully automatic hp-adaptive refinement strategy using a residual based aposteriori error estimation which is based on the solution of local boundary value problems. The reliability and efficiency for estimator has been proved. Implementation for Stokes problem shows the convergence of our algorithm.

Arezou Ghesmati
Texas A&M University
aghesmati@math.tamu.edu

Markus Buerg
Texas A&M University
Visiting Professor
buerg@math.tamu.edu

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

PP8

AWM Workshop: Residual-Based A Posteriori Error Estimate for Interface Problems: Nonconforming Linear Elements

The residual-based a posteriori error estimation for the non-conforming linear finite element approximation to the interface problem is studied. We introduce a new and direct approach, without using the Helmholtz decomposition, to analyze the reliability of the estimator. It is proved that a slightly modified estimator is reliable with the constant independent of the jump of the interfaces, with- out the assumption that the diffusion coefficient is quasi-monotone.

Numerical results are also presented.

Cuiyu He
Purdue University
he75@math.purdue.edu

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

Shun Zhang
City University of Hong Kong
shun.zhang@city.edu.hk

PP8

AWM Workshop: Enhancements for Reduced Basis Methods: Reducing Offline Computational Costs

The reduced basis method (RBM) is a new technique for finding approximate solutions to parametric partial differential equations. It tries to reduce the total cost of computation, by approximating the solution space by a linear combination of pre-- computed solutions. However, the current algorithm (greedy algorithm) is still costly. So we are designing a more efficient greedy algorithm. The implementation of our new approach also opens the power of parallel computing into the world of RBM.

Jiahua Jiang
University of Massachusetts, Dartmouth
jjiang@umassd.edu

PP8

AWM Workshop: Combinatorial Navier-Stokes Equation

Form the cubical grid of three space. Replace differential forms, exterior d and wedge product by cochains, the coboundary operator and cup product. Replace the hodge star by the Poincare dual cell operation followed by translation. Now one can form the combinatorial analog of the continuum Navier-Stokes equation.

Aradhana Kumari
City University Of New York
aradidi87@gmail.com

PP8

AWM Workshop: An Adaptive Gmsfem for High-Contrast Flow Problems

In this paper, we derive an a-posteriori error indicator for the Generalized Multiscale Finite Element Method (GMsFEM) framework. This error indicator is further used to develop an adaptive enrichment algorithm for the linear elliptic equation with multiscale high-contrast coefficients. We consider two kinds of error indicators where one is based on the L2-norm of the local residual and the other is based on the weighted H-1-norm of the local residual where the weight is related to the coefficient of the elliptic equation. We show that the use of weighted H-1-norm residual gives a more robust error indicator which works well for cases with high contrast media. The convergence analysis of the method is given.

Guanglian Li
Texas A&M University
lotusli0707@gmail.com

PP8

AWM Workshop: Propagation Failure in Discrete Inhomogeneous Media Using a Caricature of the Cubic

We consider a bistable differential-difference equation with inhomogeneous diffusion across an interval lattice. Previous research uses a discontinuous nonlinearity to construct exact solutions. We employ a continuous piecewise linear nonlinearity, a caricature of the cubic, to derive exact steady-state front solutions. Diffusion coefficients are varied on a finite interval, representing the inhomogeneous media. The interval of propagation failure, dependent upon the diffusion coefficients and wave speed, determines when and where stationary front solutions occur.

Elizabeth Lydon
University Central Florida
elydon@knights.ucf.edu

PP8

AWM Workshop: Nontrivial Structure in Top Homology of a Space

The top homology of a reasonable compact topological space X is naturally sitting inside a vector space V with a canonical coordinate system. The position of this subspace relative to the axes is a non-trivial invariant of the homeomorphism type of X. The canonical coordinate axes of V are given by oriented components of the set of top dimensional manifold points.

Chandrika Sadanand
Stony Brook University
chandrika@math.sunysb.edu

PP9

A Synchronized Co-Volume Scheme for the Large-Scale Shallow Water Equations

The co-volume scheme specifies the mass at cell centers and cell vertices, and both of the normal and tangential velocity components at cell edges. This scheme is extremely flexible, applicable to unstructured meshes, and avoids the need to reconstruct the tangential velocity component, as the classical C-grid scheme does. But the co-volume has been shown to be a generalization of the Z-grid scheme, and therefore inherits all the known defects of the latter. From another point view, a recent study shows that, for the co-volume scheme on the f-plane, the mass-vorticity-divergence fields on the primary are completely decoupled from the mass-vorticity-divergence fields on the dual mesh. We propose to periodically synchronize the fields on the primary mesh and the fields on the dual mesh. In this presentation, we explore approaches for the synchronization, and the effects of them.

Qingshan Chen
Clemson University
qsc@clemson.edu

PP9

Physics-Based Preconditioning and Dual Timestepping for Stiff Combustion Problems

with Detailed Chemical Mechanisms

Time-accurate simulation of low-Mach combustion is challenging due in part to the wide range of tightly-coupled physical time scales. Traditional integration techniques are severely constrained when strongly nonlinear and stiff chemistry is introduced to the Navier-Stokes equations at low Mach number. Kinetic stiffness is exceptionally demanding when detailed, fine-grained reaction mechanisms are utilized. We are developing local physics-based preconditioning techniques for detailed chemical mechanisms and extending dual timestepping for high-fidelity combustion simulation on next-generation platforms.

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

PP9

Rods with Bend and Twist in a Brinkman Fluid

We develop a Lagrangian algorithm to model an elastic rod in a porous medium. The three dimensional fluid is governed by the incompressible Brinkman equation and the Kirchhoff rod model captures bend and twist of the rod. Regularized solutions are derived and we compare numerical results to asymptotics for swimming speeds in a Brinkman fluid. Numerical results showing bending and twisting energies in fluids of different porosity will be shown.

Nguyenho Ho
Mathematical Sciences Department
Worcester Polytechnic Institute
nho@wpi.edu

Sarah D. Olson
Worcester Polytechnic Institute
sdolson@wpi.edu

PP9

Discrete Exterior Calculus Solution of Incompressible Flows

Navier-Stokes equations are discretized using discrete exterior calculus (DEC) on simplicial meshes. The governing equations are first rewritten, using the smooth exterior calculus notation, in terms of velocity and pressure forms. The smooth forms and operators are then substituted with their corresponding discrete definitions based on the DEC framework. Several numerical test cases are presented to demonstrate numerical convergence, stability, and conservation properties of the developed numerical scheme.

Mamdouh S. Mohamed
Physical Sciences and Engineering Division
KAUST
mamdouh.mohamed@kaust.edu.sa

Sudantha Balage
Physical Sciences and Engineering Division
KAUST, Jeddah, KSA
sudantha.balage@kaust.edu.sa

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

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

PP9

Adaptive Wavelet Simulation for Weakly Compressible Flow Bounded by Solid Walls of Arbitrary Shape

We develop an adaptive simulation method for computing compressible flow bounded by solid walls of arbitrary shape. A finite volume approach is coupled with wavelet analysis for local grid refinement. A volume penalization method is employed to compute the flow in Cartesian coordinates. We assess the quality and efficiency of the method for a flow in a channel with an expanded section. The results are compared with a reference flow computed on a uniform grid.

Naoya Okamoto
Nagoya University
okamoto@ccs.engg.nagoya-u.ac.jp

Margarete Domingues
Instituto Nacional de Pesquisas Espaciais
margarete.domingues@inpe.br

Katsunori Yoshimatsu
Nagoya University
yosimatu@esi.nagoya-u.ac.jp

Kai Schneider
Universite de Provence, Aix-Marseille
Centre de Mathematiques et Informatique
kschneid@cmi.univ-mrs.fr

PP9

Multiple Solutions in Curved-Pipe Flow

The project deals with fluid flows in uniformly curved pipes driven by a steady axial pressure gradient, with focus on blood flow in arteries. The problem is known to have multiple solutions and the primary solution involving the formation of two-vortices(Dean vortices) was found and validated. Bifurcation was found for curved pipes with a square cross-section. In addition, wall shear stresses, stream function and vorticity were computed for different curvature ratios and Dean numbers.

Harsh Ranjan
Indian Institute of Technology Guwahati
harshranjan21@gmail.com

PP9

Cooperation and Efficiency in Sperm Motility Patterns

To fertilize the egg, sperm of certain species engage in cooperative swimming behaviors. These cooperative motility patterns result in differences in velocity and efficiency. We employ a preferred curvature flagellum model and the method of regularized Stokeslets to simulate the viscous fluid environment sperm encounter. Using these methods, we compare a single flagellum with two flagella systems to understand the empirical effects of cooperative swimming.

Owen Richfield
Tulane University

Center for Computational Science
orichfie@tulane.edu

Paul Cripe
Tulane University
pcripe@tulane.edu

Julie Simons
Tulane University
Department of Mathematics
jsimons@tulane.edu

PP9

Efficient Simulation of Fluid-Structure Interactions Modeled by Regularized Stokes Formulation Using Kernel-Independent Fast Multipole Method

Regularized Stokes formulation has been shown to be very effective at modeling fluid-structure interactions when the fluid is highly viscous. However, its computational cost grows quadratically with the number of particles immersed in the fluid. We demonstrate how kernel-independent fast multipole method can be applied to significantly improve the efficiency of this method, and present numerical results for simulating the dynamics of a large number of elastic rods immersed in 3D Stokes flows.

Minghao W. Rostami, Sarah D. Olson
Worcester Polytechnic Institute
mwu@wpi.edu, sdolson@wpi.edu

PP9

Scalable Parallel Solvers for Highly Heterogeneous Nonlinear Stokes Flow Discretized with Adaptive High-Order Finite Element

We present scalable parallel solvers for convection-driven flow in Earth's mantle with associated plate motions, which is governed by nonlinear Stokes equations. Crucial solver components are parallel geometric multigrid methods for preconditioning the linearized Stokes systems that arise upon discretization with high-order finite elements on adaptive meshes (resolution below 1km) and improved BFBT/LSC preconditioners for the Schur complement. We show robustness with respect to extreme viscosity variations and carry out global mantle flow simulations with real Earth data as we progress towards solving realistic global mantle flow inverse problems.

Johann Rudi, Toby Isaac
ICES
The University of Texas at Austin
johann@ices.utexas.edu, tisaac@ices.utexas.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Michael Gurnis
Caltech
gurnis@caltech.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

PP9

Discrete Adjoint Openfoam and Applications

A discrete adjoint version of OpenFOAM obtained using Algorithmic Differentiation[1] by operator overloading is applied to obtain derivatives based on external aerodynamics of a Volkswagen Polo Car. The results are validated qualitatively against a 'frozen turbulence' continuous adjoint implementation[2]. We demonstrate the robustness and flexibility of differentiating a turbulence model discretely to obtain exact derivatives. Strategies to tackle spatial and temporal complexities characteristic to discrete adjoint methods are also discussed. References: [1] Naumann, U., The Art of Differentiating Computer Programs, SIAM, 2012. [2] Othmer, C., A continuous adjoint formulation for the computation of topological and surface sensitivities of ducted flows,Int. J. Num. Meth. Fluids, 2006

Arindam Sen
Software and Tools for Computational Engineering, STCE
RWTH Aachen University
sen@stce.rwth-aachen.de

Markus Towara
STCE, RWTH Aachen University
towara@stce.rwth-aachen.de

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

PP9

Strategy for Efficiently Simulating Reactive Flows with Large Detailed Chemical Kinetics

A highly efficient numerical approach for simulating reactive flows with large detailed chemical kinetics is proposed. The present approach consists of a robust and fast explicit time integration method for chemical reaction equations and a species building technique for diffusion coefficient calculations. The computational results with a realistic combustion problem verify the high efficiency of the present approach.

Hiroshi Terashima
The University of Tokyo
htera@rocketlab.t.u-tokyo.ac.jp

Youhi Morii
Japan Aerospace Exploration Agency
morii.yuuhi@jaxa.jp

Mitsuo Koshi
Yokohama National University
koshim@jcom.home.ne.jp

PP9

An Exact and Consistent Adjoint Method for High-Fidelity Unsteady Compressible Flow Simulations

We consider high-resolution discretizations commonly used for compressible turbulent flows. A corresponding discrete-adjoint can produce gradients that include spurious numerical modes, restricting its utility. A continuous-adjoint ap-

proach does not, but provides an inaccurate gradient. We introduce a dual-consistent finite-difference discretization based on common workhorse methods for compressible turbulent flows, which enjoy the benefits of both approaches. Demonstrations with an unsteady mixing layer show that the gradient is both exact up to roundoff and consistent.

Ramanathan Vishnampet
Department of Mechanical Science and Engineering
University of Illinois at Urbana-Champaign
vishnam2@illinois.edu

Daniel J. Bodony
University of Illinois at Urbana-Champaign
Department of Aerospace Engineering
bodony@illinois.edu

Jonathan B. Freund
University of Illinois at Urbana-Champaign
jbfreund@illinois.edu

PP9

Periodic Stokes Flow in 2 Dimensional Space

We present a new boundary integral method to solve incompressible Stokes flow with low Reynolds number in a periodic background with high accuracy and efficiency.

Lin Zhao
Dartmouth College
lin.zhao.gr@dartmouth.edu

Alex H. Barnett
Department of Mathematics
Dartmouth College
ahb@math.dartmouth.edu

PP10

Etd Spectral Deferred Correction Methods

We introduce a new class of arbitrary-order exponential time differencing (ETD) methods based on spectral deferred correction. We study the stability and accuracy properties of these methods and conduct numerical experiments against existing ETD and implicit-explicit schemes. We find that our new high-order ETD spectral deferred correction schemes outperform state-of-the-art time integrators for computations requiring high accuracy, making them well-suited to work in conjunction with spatial spectral methods.

Tommaso Buvoli
Department of Applied Mathematics, University of Washington
Seattle, WA 98195 USA
buvoli@uw.edu

PP10

Adaptive Multigrid Methods for An Integrated Structural Health Monitoring (SHM) Systems for Composite Material with Fluid-Structure Interaction (FSI) Effect

To design a SHM system, it is important to understand phenomenologically and quantitatively the wave propagation in composite material and the influence of the geomaterial and mechanical properties of the structures. To accelerate the design of SHM systems, the FSI effect on the wave propagation has to be considered. Combining fluid dynamics with structural analysis traditionally poses a formidable challenge for even the most advanced numerical techniques due to the disconnected, domain-specific nature of analysis tools. We present the state-of-the-art in computational methods and techniques for wave propagation with FSI effect that go beyond the fundamentals of computational fluid and solid mechanics. Also this project aims to develop efficient numerical methods for wave propagation phenomena in composite material with FSI effect, which combine modern techniques from PDE-constrained optimization, adaptive and multigrid simulation methods.

Bhuiyan Shameem M. Ebna Hai
Numerical Methods for Computational Science and Engineering,
Mechanical Engineering Department, Helmut Schmidt University
ebnahaib@hsu-hh.de

Markus Bause
Helmut-Schmidt-University
University of the Federal Armed Forces Hamburg
bause@hsu-hh.de

PP10

Reducing the Impact of the Cfl Condition for Dispersive Wave Propagation Problems

To solve dispersive wave propagation problems in complex geometries efficiently and accurately, it is attractive to use explicit time-integration and high-order unstructured spectral element discretization in space. This combination may have a severe global conditional CFL time-step restriction. For specific models, we describe conditions placed on dispersive terms for making the CFL condition independent of discretisation method and meshsizes. In this sense the CFL condition impose a minimal constraint on the efficiency without compromising accuracy.

Allan P. Engsig-Karup
Technical University of Denmark
Department of Informatics and Mathematical Modeling
apek@imm.dtu.dk

Claes Eskilsson
Chalmars University of Technology
Gothenburg, Sweden.
claes.eskilsson@chalmers.se

PP10

Explicit Strong Strong Stability Preserving Multi Step Runge-Kutta Methods

High-order spatial discretizations of hyperbolic PDEs are often designed to have strong stability properties, such as monotonicity. We study explicit multistep RungeKutta strong stability preserving (SSP) time integration methods for use with such discretizations. We prove an upper bound on the SSP coefficient of explicit multistep RungeKutta methods of order two and above.Numerical optimization is used to find optimized explicit methods of up to five steps, eight stages, and tenth order. These methods are tested on the advection and Buckley-Leverett equations, and the results for the observed total variation diminishing and positivity preserving time-step are presented.

Zachary J. Grant

University of Massachusetts at Dartmouth
zgrant@umassd.edu

PP10

Robust Residual-Based A Posteriori Error Estimate for Interface Problems: Nonconforming Linear Elements

A robust residual-based a posteriori error estimation for the non-conforming linear finite element approximation to the interface problem is studied. We introduce a new and direct approach, without using the Helmholtz decomposition, to analyze the reliability of the estimator. It is proved that our estimator is reliable with the constant independent of the jump of the interfaces, without the assumption that the diffusion coefficient is quasi-monotone. Numerical results are also presented.

Cuiyu He
Purdue University
he75@math.purdue.edu

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

Shun Zhang
City University of Hong Kong
shun.zhang@city.edu.hk

PP10

Monolithic Multi-Time-Step Coupling Methods for First and Second-Order Transient Systems

New frameworks for coupling different integration methods (either temporal or spatial) for first- and second-order transient systems will be introduced. These methods allow different numerical time-integrators, time-steps, and numerical formulations in different regions of the computational domain. Unique features of the proposed methods will be demonstrated through several numerical examples.

Saeid Karimi
University of Houston
skarimi@central.uh.edu

Kalyana Nakshatrala
University Of Houston - Main Campus
knakshatrala@uh.edu

PP10

Total Order Function Space Spectral Collocation Methods Using the Padua Points

Multivariate Chebyshev spectral collocation methods typically are implemented using Lagrange interpolating polynomials corresponding to tensor products of Chebyshev nodes. The Padua points are an analogue of the Chebyshev nodes that interpolate the smaller bivariate total order function space $\mathbf{P}_n^2$. These points can be used to implement a spectral collocation scheme, however careful treatment of boundary conditions is required. This issue is explored for Poisson and Helmholtz problems with arbitrary, variable coefficient, Robin boundary conditions.

Scott Moe
University of Washington
smoe@uw.edu

PP10

Numerical Investigation of Influence of Node Alignment on Stable Calculation for Meshless Time Domain Method

The meshless time domain method (MTDM) is one of meshless methods. The shape function derived from the radial point interpolation method (RPIM) is adopted for discretization process instead of meshes in MTDM. In addition, the radial basis function (RBF) is employed in RPIM, and the multi quadratic function is adopted for the present study. The purpose of the present study is to investigate influence of node alignment and RBF on stable calculation.

Yoshiharu Ohi
RIKEN Advanced Institute for Computational Science
yoshiharu.ohi@riken.jp

Yoshihisa Fujita
Nagoya University
fujita.yoshihisa@nifs.ac.jp

Taku Itoh, Soichiro Ikuno
Tokyo University of Technology
taku@m.ieice.org, ikuno@cs.teu.ac.jp

PP10

Asymptotics of High-Frequency Scattering Problems

Various strategies for simulating wave scattering problems through the Helmholtz equation typically become computationally expensive for high frequencies. The solution of its integral representation exhibits asymptotic behaviour $q(\vec{y}) \sim k e^{ik\phi(\vec{y})} \sum_{j=0}^{\infty} a_j(\vec{y}) k^{-j}$ as the wave number k tends to ∞, which can reduce costs. Oscillatory integrals arising when solving for $q(\vec{y})$ on the boundary, can be approximated asymptotically using the (numerical) method of steepest descents. We explain this method, the case of multiple reflections and the extension to 3D.

Peter Opsomer
KU Leuven
peter.opsomer@cs.kuleuven.be

Daan Huybrechs
Department of Computer Science
K.U. Leuven
daan.huybrechs@cs.kuleuven.be

PP10

Accurate Derivative Computation for Finite Element Codes

It is common in physics to solve a PDE for some potential, while the physically relevant quantities are actually derivatives of said potential. Classical finite element methods, however, lose an order of accuracy for every derivative taken. We present an integral equation-based technique that, in the case of 2-D semilinear Poisson equations, enables the computation of arbitrary derivatives to the same order as a given FEM solution. We present particular applications to the Grad-Shafranov equation.

Lee F. Ricketson

UCLA
lfr224@cims.nyu.edu

Antoine Cerfon
NYU
cerfon@cims.nyu.edu

Manas Rachh
Courant Institute
NYU
manasrachh@gmail.com

PP10

Comparison of Weak Galerkin Finite Element Method with Dgfem and Mfem

This poster presents a comparative study on the newly introduced weak Galerkin finite element methods (WGFEMs) with the widely accepted discontinuous Galerkin finite element methods (DGFEMs) and the classical mixed finite element methods (MFEMs) for solving second-order elliptic boundary value problems. We examine the differences, similarities, and connection among these methods in scheme formulations, implementation strategies, accuracy, and computational cost. The comparison and numerical experiments demonstrate that WGFEMs are viable alternatives to MFEMs and hold some advantages over DGFEMs, due to their properties of local conservation, normal flux continuity, no need for penalty factor, and definiteness of discrete linear systems.

Farrah Sadre-Marandi
Colorado State University
sadre@math.colostate.edu

PP10

A Weighted Sequential Splitting Method for the 3D Maxwell's Equations

We present a Weighted Sequential Splitting (WSS) method for Maxwell's equations in 3D. The solution obtained by the WSS scheme in a given time step is a weighted average of solutions of several 1D Maxwell systems, discretized using a Crank Nicolson method. We prove convergence of the scheme for all weights $0 \leq \theta \leq 1$, and show that it is unconditionally stable of first order in time when $\theta \neq 0.5$, and second order when $\theta = 0.5$.

Puttha Sakkaplangkul
Oregon State University
sakkaplp@onid.oregonstate.edu

PP10

Asymptotically Compatible Schemes for Robust Discretization of Nonlocal Models

Nonlocality is ubiquitous in nature. While partial differential equations (PDE) have been used as effective models of many physical processes, nonlocal models and nonlocal balanced laws have become possible alternatives to treat anomalous process and singular behavior. In this talk, we use a recently developed nonlocal vector calculus and nonlocal calculus of variations to study a class of constrained value problems associated with nonlocal operators. In addition, we present asymptotically compatible discretizations that provide convergent approximations to both nonlocal models and their local limit. Such discretizations are useful for model validation and multiscale simulations.

Qiang Du, Xiaochuan Tian
Columbia University
Department of Applied Physics and Applied Mathematics
qd2125@columbia.edu, xt2156@columbia.edu

PP11

The Sparse Grid Combination Technique for Solving Eigenvalue Problems

Identifying microinstabilities in hot fusion plasmas can be done by solving the gyrokinetic eigenvalue problem which is a computationally demanding task. Our approach for solving this problem is the combination technique. It combines several solutions of the eigenvalue problem from different grid-sizes. That can diminish the curse of dimensionality, since it creates a sparse grid approximation. It also increases the scalability by introducing an additional layer of parallelism which is reusing the current parallelism.

Christoph Kowitz
TU Munich
kowitz@in.tum.de

Markus Hegland
Australian National Unversity
Markus.Hegland@anu.edu.au

Hans-Joachim Bungartz
Technische Universität München, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

PP11

A Tangential Interpolation Framework for MIMO Eigensystem Realization Algorithm

The Eigensystem Realization Algorithm (ERA) is a commonly used data- driven method for system identification and model reduction of dynamical systems. The main computational difficulty in ERA arises when the system under consideration has large number of inputs and outputs, thus requiring to compute a full SVD of a large-scale dense Hankel matrix. In this work, we present an algorithm that aims to resolve this computational bottleneck via tangential interpolation. A numerical example is presented.

Boris Kramer
Virginia Tech
bokr@vt.edu

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

PP11

Multi-Set Data Analysis and Simultaneous Matrix Block Diagonalization: Models and Algorithms

We present joint independent subspace analysis (JISA) of stationary sources as a statistical signal processing approach to multi-set data analysis. From an algebraic perspective, JISA may be viewed as a set of coupled block diagonalization (CBD) problems. We evoke some new results for JISA-CBD, and linkage to recent results on ISA of non-

stationary sources based on joint block diagonalization of covariance matrices. We focus on algorithms, performance, identifiability and the associated matrix factorizations.

Dana Lahat, Christian Jutten
GIPSA-Lab
Dana.Lahat@gipsa-lab.grenoble-inp.fr,
christian.jutten@gipsa-lab.grenoble-inp.fr

PP11

Overcoming the Gibbs Phenomenon : Fast Fourier Extension

Fourier series of smooth functions on an interval $[-1, 1]$ are known to exhibit the Gibbs phenomenon, and have overall slow convergence. The Fourier extension technique overcomes these problems by constructing a Fourier approximation on an extended interval $[-T, T]$, through least-squares optimization. We present FE algorithms for approximations from equidistant points that match the $O(N \log N)$ complexity of the discrete Fourier transform, and comment on the associated difficulties.

Roel Matthysen
KULeuven
roel.matthysen@cs.kuleuven.be

PP11

Big Graph Analytics of Human Connectome Networks

Analyzing the very dense human connectome network with billions of links has become a central and challenging topic in computational neuroscience. Yet, it offers neuroscientists great opportunities to understand the complex structure and functionality of the human brains. In this project we demonstrate how to use big data technology and tools, such as Spark GraphX and SAP HANA, to approximate various graph measurements from the brain connectivity graphs released by the Human Connectome Project.

Jürgen Ommen, Chih Lai, Yulin Yang
University of St. Thomas
Graduate Programs in Software
omme0003@stthomas.edu, clai@stthomas.edu,
yang9594@stthomas.edu

PP11

Componentwise Sensitivity of Matrix Functions and Applications

Matrix functions, such as the exponential, are used in a variety of applications. Previous sensitivity analyses involve mainly normwise condition numbers. Componentwise condition numbers and perturbation bounds are more appropriate when sensitivity with respect to particular components is of interest. We develop such analysis and apply it to a parameterized ODE problem arising in physics where the matrix components are flux coefficients describing chemical interaction.

Samuel Relton
University of Manchester, UK
samuel.relton@manchester.ac.uk

PP11

Inducing Approximately Optimal Flow Via Truthful Mediators

The classical approach to have selfish agents route optimally in network congestion games is to impose edge tolls that depend on agents' demands. However, this fails when demands are unknown to the mechanism designer. We design a weak mediator that can impose tolls such that it is an asymptotic ex-post Nash equilibrium for agents to truthfully report their demands to the mediator and faithfully follow its suggested route, which results in an approximately optimal flow.

Ryan M. Rogers
University of Pennsylvania
ryrogers@sas.upenn.edu

Aaron Roth
University of Pennsylvania, USA
aaroth@cis.upenn.edu

Jonathan Ullman
Columbia University
jullman@cs.columbia.edu

Steven Wu
University of Pennsylvania
wuzhiwei@cis.upenn.edu

PP11

Security in Data Mining Through Cloud Computing Using Expert System

This study presents an expert system to optimize the security of data mining in cloud computing.These systems still suffer from privacy and security concerns. The main security issues in cloud computing are identifying identity, access control, confidentiality, integrity, availability and non-repudiation. The main goal of this investigation is designing an expert system to increase the security of data mining in cloud computing systems. To do this, the efficiency of these techniques is examined for the same data set .Subsequently, an inference engine has been developed and this engine automatically selects the model which has best efficiency by using artificial intelligence techniques.

Ana Sadeghitohidi, Azadeh Roozbehi
Azad Tehran University
anasadeghii1@yahoo.com, azadeh.roozbehi@gmail.com

PP11

Big Data Analytics Application in Genomics Data Processing

Genomics datasets are large. In this project we are processing genomics data on a HPC cluster using Hadoop. Our goal is to identify from the genome data any disease pattern. Eventually we plan to have access to clinical data for minorities that will help identify risk factors when combined with genomics data. Our processing at this time is experimental and as such no patient privacy issues are involved because the clinical data will be anonymized.

S Srinivasan
Texas Southern University
JHJ School of Business
srinis@tsu.edu

Hector Miranda
Texas Southern University

mirandahc@tsu.edu

Daniel Vrinceanu
Texas SOuthern
vrinceanud@tsu.edu

Terence Vaughn
Texas Southern University
vaughntt@tsu.edu

PP11

Generalized Low Rank Models

Principal components analysis (PCA) is a well-known technique for approximating a data set represented by a matrix by a low rank matrix. Here, we extend the idea of PCA to handle arbitrary data sets consisting of numerical, Boolean, categorical, ordinal, and other data types. This framework encompasses many well known techniques in data analysis, such as nonnegative matrix factorization, matrix completion, sparse and robust PCA, k-means, k-SVD, and maximum margin matrix factorization. The method handles heterogeneous data sets, and leads to coherent schemes for compressing, denoising, and imputing missing entries across all data types simultaneously. It also admits a number of interesting interpretations of the low rank factors, which allow clustering of examples or of features. We propose several parallel algorithms for fitting generalized low rank models, and describe implementations and numerical results.

Madeleine R. Udell
Stanford University
Stanford University
madeleine.udell@gmail.com

Corinne Horn, Reza Zadeh
Stanford University
cehorn@stanford.edu, rezab@stanford.edu

Stephen Boyd
Stanford University
Department of Electrical Engineering
boyd@stanford.edu

PP11

A Structured Cholesky Factorization for Fock Matrix Construction

The bottleneck of Fock matrix construction is computation and storage of Electron Repulsion Integrals (ERIs). We present new high performance software implemented in GTFock which pre-computes a structured lazy evaluation pivoted Cholesky factorization of the matrix unfolding of the fourth-order ERI tensor. In addition to reducing ERI computation and storage by an order of magnitude, this new approach utilizes symmetry and sparsity structure.

Joseph Vokt
Cornell University
joseph.vokt@gatech.edu

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

PP11

Parallel Bayesian Global Optimization, With Application To Metrics Optimization at Yelp

We consider parallel global optimization of expensive-to-evaluate functions, and propose an efficient method based on stochastic approximation for implementing a conceptual Bayesian optimization algorithm proposed by Ginsbourger et al. (2010). We also introduce an open-source software implementation of this algorithm, called Metrics Optimization Engine, developed in collaboration with engineers at Yelp Inc. and used internally at Yelp to optimize prediction models and performance metrics.

Jialei Wang, Peter I. Frazier
Cornell University
jw865@cornell.edu, pf98@cornell.edu

Scott Clark, Eric Liu
Yelp Inc
sclark@yelp.com, eliu@yelp.com

PP12

The Modified Bidomain Model with Periodic Diffusive Inclusions

Bidomain equations are the standard way to model the electric potential in cardiac tissue. We propose the modification of this model for the case of the diseased heart, e.g. fibrosis of the heart tissue. On microscale, we assume to have periodic diffusive inclusions embedded in the healthy tissue modelled by the bidomain equations. We derive the macroscale model using the homogenisation technique. We recover a bidomain model with modified conductivities, that depend on the volume fraction of the diffusive inclusions but also on their geometries.

Andjela Davidovic
INRIA Bordeaux Sud-Ouest, Bordeaux, France
University Bordeaux, Bordeaux, France
andjela.davidovic@inria.fr

Yves Coudiere
Université Bordeaux 1
Inria & IMB (UMR 5251)
yves.coudiere@inria.fr

Clair Poignard
INRIA Bordeaux Sud-Ouest
University Bordeaux
clair.poignard@inria.fr

PP12

Simulation-Based Solute Transport in Kidney Cells

Based on solute and water conservation and electroneutrality constraints, we developed mathematical models of kidney cells. With the models, we computed intracellular Na^+ concentration and membrane hyperpolarization/depolarization behavior as a function of extracellular NaCl. Mathematical models like these can be adapted to formulate hypotheses of a system response after a perturbation. Furthermore, they can help in the development of new drugs prior to *in vivo* testing, thereby minimizing

research costs, and time.

Monica Nadal-Quiros
University of Puerto Rico, Rio Piedras Campus
Department of Biology
monica.nadal@upr.edu

Aniel Nieves-Gonzalez
University of Puerto Rico
Institute of Statistics and Information Systems
aniel.nieves@upr.edu

Leon Moore
SUNY HSC
Department of Physiology and Biophysics
leon.moore@stonybrook.edu

Mariano Marcano
University of Puerto Rico
mariano.marcano@upr.edu

PP12

An Adaptive Markov Chain Monte Carlo Method Applied to Simulation of a Tumor Growth Model

The inference of mechanistic descriptions of oncogenesis is often a challenge due to limited experimental data. Markov Chain Monte Carlo (MCMC) methods have been developed and widely used in Bayesian analysis. This project applies an adaptive MCMC technique based on the Metropolis-Hastings algorithm to calibrate the parameters of a tumor growth model to experimental data. The study is supported by the NIGMS of NIH grant as part of the West Virginia INBRE (P20GM103434).

Qing Wang, Zhijun Wang
Dept. of Computer Science, Math and Engineering
Shepherd University, Shepherdstown, WV
qwang@shepherd.edu, zwang@shepherd.edu

David Klinke
Dept. of Chemical Engineering
West Virginia University, Morgantown, WV
david.klinke@mail.wvu.edu

PP12

Numerical Simulation of a Tumor Cell Population Growth Dynamics Model Using Genetic Algorithm

Tumor cell growth models involve high-dimensional parameter spaces that require computationally tractable methods to solve. Genetic algorithm is applied to search for parameter values that fit experimental data from mice cancel cells. Dynamically variable crossover and mutation methods are used to produce new generations. Fitness functions, which measure the difference between calculated results and experimental data, are minimized. This study is supported by the NIGMS of NIH grant as part of the WV INBRE (P20GM103434).

Zhijun Wang, Qing Wang
Dept. of Computer Science, Math and Engineering
Shepherd University, Shepherdstown, WV
zwang@shepherd.edu, qwang@shepherd.edu

PP12

"Allostery": A Python Package for Network Analyses of Biomolecular Simulations

Allosteric interaction network is an intrinsic complex problem in computational bio-molecular researches and very scarce analysis tools are available. Here I present a python package named "Allostery" for calculation of biomolecular interaction networks using Lange-Grubmuller generalized correlation coefficient, based on Kraskov-Stogbauer-Grassberger Mutual Information estimator. It features: distance and k-d tree based implementation, memory efficient serial and cluster level parallel execution and user-friendly tools for data IO and visualization.

Yuhang Wang
University of Illinois, Urbana-Champaign
ywang148@illinois.edu

Emad Tajkhorshid
University of Illinois, Urbana-Chamnpaign
emad@life.illinois.edu

PP12

Numerical Methods for Protein Adsorption in Porous Membranes

Protein therapeutics are used as treatments for various illnesses (diabetes, cancer, hemophilia, infectious diseases). Most of the cost of protein production is associated with the downstream separation/purification. Developing more efficient purification methods would decrease the cost of protein therapeutics. In this presentation, we will discuss protein separation using multi-modal membranes developed in Clemson University's chemical engineering department. We will present numerical simulations of the advection-diffusion-reaction equation, comparing and contrasting solution methods and adsorption models.

Anastasia B. Wilson
Clemson University
Department of Mathematical Sciences
anastas@clemson.edu

PP12

Modeling Core Body Temperature during Exercise

Dynamics of core body temperature of rats running on treadmills with varied speeds at different ambient temperatures were modeled. The model was used to estimate heat generated in the body for thermoregulation and additional heat produced by the muscles during exercise. We found that the latter grows with exercise intensity in ambient temperature independent way. In contrast, the thermoregulatory component reduces at room temperature, while remaining unchanged at high temperature implying existence of compensatory mechanisms.

Yeonjoo Yoo
Indiana University-Purdue University Indianapolis
IUPUI
yjyoo@iupui.edu

PP12

Ensemble Kalman Filters for Dynamic Dipole Estimation from Magnetoencephalography

We discuss the dynamic imaging of electromagnetic cerebral activity based on magnetoencephalography (MEG) data. In this presentation, we develop two time evolution

models combined with the ensemble Kalman filter to localize the electric source currents and capture the rapid neural activity in the brain. The utilization of the Bayesian framework allows sequential updating of the preceding estimates of the activations. The effectiveness of the algorithms is demonstrated with simulated data.

Lijun Yu
Case Western Reserve University
lxy141@case.edu

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

PP13

Decomposition-Based Uncertainty Quantification with Application to Environmental Impacts of Aviation

We present a decomposition-based uncertainty quantification approach for feed-forward multicomponent systems. The aim is to decompose the uncertainty quantification task among the various components comprising a system, then synthesize these results to quantify uncertainty at the system level. We introduce concepts that extend our approach to systems containing a large number of component-to-component interface variables. We apply the proposed method to quantify how uncertainty in aviation technology affects uncertainty in environmental impacts.

Sergio Amaral
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
samaral@mit.edu

PP13

The Combined Block by Block - Monte Carlo Methods for Numerical Treatment of the Mixed Nonlinear Stochastic Integral Equation

In this paper, we study the random effects, W(s), on a mixed nonlinear integral equation of the second kind with time dependent. Two contributions are made: the first contribution discuses the existence and uniqueness of the solution of this equation and the second contribution presents an algorithm in which we combine both Block-by-Block method and Monte-Carlo method, to accurately and efficiently solve a mixed nonlinear stochastic integral equation. Numerical examples are given to illustrate the presented numerical method.

Abdallah A. Badr
Alexandria University
Department of Mathematics
a.badr@ubt.edu.sa

PP13

Computational Investigation of Quasi-Random Sequences for Error Estimation

Quasi-Monte Carlo methods has become a major vehicle for estimating high-dimensional integrations. However, the asymptotic error convergence order depends on the sequences used in estimation integration. In this paper, we investigate a family of related quasi-random sequences, known low-discrepancy sequences, and study their effects to convergence of error estimation. Our experimental results show that some low-discrepancy sequences significantly improve the performance of error estimation.

Hongmei Chi
Computer Science
Florida A&M University
hchi@cis.famu.edu

PP13

Optimal Source Encoding in Medium Parameter Reconstruction Problems

The computational cost of medium parameter reconstruction scales linearly with the number of sources used. To alleviate that bottleneck, several authors proposed to solve instead for a few linear combinations of the sources. In order to preserve consistency with the original problem, those combinations must be chosen randomly. Instead we propose to select the combinations that provide the most confidence in the reconstructed parameters, in the Bayesian sense. This leads us to minimize the nuclear norm of the variance of the posterior distribution evaluated at the MAP point.

Benjamin Crestel
University of Texas at Austin
Institute for Computational Engineering and Science
crestel@ices.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

PP13

Randomized Likelihood Method: A Scalable Approach to Big Data in Large-Scale Pde-Constrained Bayesian Inverse Problems

In this poster we present a scalable approach to tackle the big-data challenge in large-scale PDE-constrained Bayesian inverse problems. In particular, we propose a MAP approximation algorithm whose convergence is independent of the data dimension The idea is to randomize the likelihood to extract active/important/independent data which is in general much less compared to the original big data. The active data is then used to obtain approximate MAP point. We rigorously analyze the convergence of the proposed randomized likelihood method for a large class of inverse problems under mild conditions. One of the key results is that approximate MAPs converge to the big-data MAP in probability, independent of the amount of data. Several numerical results for large-scale PDE-constrained Bayesian inverse problems are presented to show the effectiveness of the proposed method and to verify the theoretical results.

Aaron Myers
UT Austin, ICES
aaron@ices.utexas.edu

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

Ellen Le
University of Texas at Austin - ICES
ellen@ices.utexas.edu

PP13

MUQ (MIT Uncertainty Quantification): Flexible Software for Connecting Algorithms and Applications

It is difficult to interface uncertainty quantification algorithms—e.g., Bayesian inference, robust optimization, surrogate building, sensitivity analysis—with computational models in a way that exposes model structure to the algorithms while accommodating easy development of new models and new algorithms. The MIT Uncertainty Quantification (MUQ) library addresses these issues with a flexible software framework for constructing complex multicomponent models and for developing new algorithms. We discuss our design goals and demonstrate MUQ by both implementing a new MCMC algorithm and constructing a coupled PDE model.

Matthew Parno, Patrick R. Conrad
Massachusetts Institute of Technology
mparno@mit.edu, p.conrad@warwick.ac.uk

Andrew Davis
MIT
davisad@mit.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

PP13

A Stochastic Dynamic Programming Method for Controlling a Combined Hydro/Wind Power Producer

We develop a computationally efficient nonlinear stochastic dynamic programming model to determine the optimal control of a combined hydro and wind power producer. The algorithm efficiency is improved through use of a radial basis function approximation within a corridor.

Kyle Perline
Cornell University
krp73@cornell.edu

PP13

Efficient Error Estimation for Elliptic PDEs with Random Data

When PDE models have uncertain inputs, efficient numerical methods for the forward problem are needed for uncertainty quantification. Stochastic Galerkin finite element methods have attractive approximation properties but are too expensive to implement on standard computers when the dimension of the random input space is high. To improve efficiency, we look to adaptive Galerkin schemes. We introduce a cheap a posteriori error estimator for elliptic problems that can be implemented in a non-intrusive way.

Catherine Powell
School of Mathematics
University of Manchester
c.powell@manchester.ac.uk

Alex Bespalov
University of Birmingham, United Kingdom
a.bespalov@bham.ac.uk

David Silvester
School of Mathematics
University of Manchester
david.silvester@manchester.ac.uk

PP13

Inference of Constitutive Parameters in a Nonlinear Stokes Mantle Flow Model

Motivated by inverse problems in mantle flow models with plate tectonics, we estimate constitutive parameters in a nonlinear Stokes system using a Bayesian inference approach. We compute the maximum a posteriori (MAP) point of the posterior parameter distribution, which amounts to the solution of a PDE-constrained optimization problem. To quantify the uncertainty in the parameters, we compare a local Gaussian approximation of the posterior distribution at the MAP point with results obtained using MCMC sampling.

Vishagan Ratnaswamy
California Institute of Technology
Seismological Laboratory
vratnasw@caltech.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Michael Gurnis
California Institute of Technology
gurnis@gps.caltech.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

PP13

Uncertainty Quantification for Thermally Driven Flow

The uncertainty in a scalar quantity of interest is quantified for a natural convective flow in a cavity with random boundary data. The performance of sparse grid stochastic collocation is compared to quasi Monte Carlo, depending on the variance of the random perturbation. The methods are accelerated by making use of simulation results from previously visited sampling points. In particular, improved starting conditions for the iterative non-linear solver and POD-Galerkin reduced-order modeling are considered.

Sebastian Ullmann, Lang Jens
TU Darmstadt
ullmann@gsc.tu-darmstadt.de, lang@mathematik.tu-

darmstadt.de

PP13

Computational and Statistical Tradeoffs: a Framework

The recent explosion in the quantity and dimension of data has brought up many new computational and statistical challenges that must be addressed simultaneously such that inference is both tractable and meaningful. We propose a framework that provides an explicit opportunity for the practitioner to specify how much statistical risk they are willing to accept for a given computational cost. We illustrate this with an example of estimation from a stream of iid normal variables.

Alexander Volfovsky
Statistics Department, Harvard University
Harvard University
volfovsky@fas.harvard.edu

Edoardo Airoldi, Daniel Sussman
Harvard University
airoldi@fas.harvard.edu, daniellsussman@fas.harvard.edu

PP13

Efficiency of the Girsanov Transformation Approach for Parametric Sensitivity Analysis of Stochastic Chemical Kinetics

Monte Carlo methods for sensitivity analysis of stochastic reaction networks can be classified into three categories, the pathwise derivative (PD) method, the finite difference (FD) method and the Girsanov transformation (GT) method. It has been numerically observed that when applicable, the PD method and FD method tend to be more efficient than the GT method. We provide a theoretical justification for this observation in terms of system size asymptotic analysis.

Ting Wang
Department of Mathematics and Statistics
University of Maryland Baltimore County
ting1@umbc.edu

Muruhan Rathinam
University of Maryland, Baltimore County
muruhan@umbc.edu

PP13

Convergence of the Robbins-Monro Algorithm in Infinite Dimensional Hilbert Spaces

Good proposal distributions can significantly improve the performance of random walk Metropolis algorithms. Recent work has suggested that finding the best Gaussian, fit with respect to the relative entropy metric against the target distribution, can provide such speedup. Minimizing relative entropy may be non-trivial, particularly in the case of distributions on infinite dimensional spaces. We demonstrate that the Robbins-Monro algorithm can be used to find the optimal Gaussian proposal distribution in infinite dimensional Hilbert spaces.

Daniel Watkins
Drexel University
danielmwatkins@gmail.com

Gideon Simpson
Department of Mathematics
Drexel University
simpson@math.drexel.edu

PP13

Utilizing Adjoint-Based Techniques to Effectively Perform Uq on Discontinuous Responses

Uncertainty is ubiquitous in predictive modeling and simulation due to unknown model parameters, initial conditions, etc. A number of recently developed methods for uncertainty quantification have focused on constructing surrogates of high-fidelity models using only a limited number of model evaluations. Unfortunately, many of these techniques perform poorly when the response surface is discontinuous. In this poster, we show how adjoint-based techniques can be used to efficiently construct discontinuous surrogate approximations.

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

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

PP13

Optimization of Modeled Land Surface Fluxes by Bayesian Parameter Calibration

The integration of observations and advanced modeling tools is critical for resolving issues of uncertainty in climate models. We present results from an isotopically-enabled land surface model, including experiments partitioning evapotranspiration into contributions from plant transpiration and surface evaporation. We demonstrate a model calibration approach in a Bayesian estimation framework, requiring Markov chain Monte Carlo sampling of the posterior distribution, which is shown to constrain uncertain parameters and inform relevant values for operational use.

Tony E. Wong
University of Colorado at Boulder
Department of Applied Mathematics
anthony.e.wong@colorado.edu

David Noone
University of Colorado at Boulder
Cooperative Institute for Research in Environmental Sciences
dcn@colorado.edu

William Kleiber
University of Colorado at Boulder
Department of Applied Mathematics

william.kleiber@colorado.edu

PP13

Efficient Monte Carlo Scheme for the Simulation of the Association of Lattice-model Proteins

A lattic Monte Carlo model with implicit membrane and water is used to model certain types of protein aggregation in biological membranes. We describe a sampling scheme that attempts to combine the merits of parallel tempering and multicanonical sampling. The key to this scheme is deriving best estimates for the log density of states from the MBAR estimator [Shirts, Michael R., and John D. Chodera. "Statistically optimal analysis of samples from multiple equilibrium states." The Journal of chemical physics 129.12 (2008): 124105.]

Yuanwei Xu, Mark Rodger
University of Warwick
Yuanwei.Xu@warwick.ac.uk, p.m.rodger@warwick.ac.uk

PP14

Visualising Protein Sequence Alignment

In bioinformatics, protein sequence alignment used strings of contiguous letters of amino acids, arranged in vertical register in order to highlight regions of similarity and difference. When sequences have different lengths, gap characters are inserted to denote insertions or deletions; however, gaps have no meaning in 3D structures. This project aims to revisit the protein sequence alignment problem by exploring possible ways to visualise the relationships between sequences in 3D space wile eliminating gap characters. In our work, an N-body Hamiltonian system in contact with a heat bath is built and it's configuration in 3D space forms the basis of our visualisation. The novel 3D visualisation method opens the possibility of analysing much larger alignments that wouldn't be possible with conventional visualisations produced by current alignment tools and editors.

Shaimaa M. Aljuhani
PhD student
The University of Manchester
smjuhani@gmail.com

Prof. Teresa Attwood
The University of Manchester
Faculty of Life Sciences
teresa.k.attwood@manchester.ac.uk

Dr. Tony Shardlow
University of Bath
Department of Mathematical Sciences
t.shardlow@bath.ac.uk

PP14

Reconstructing Physically Realistic Flow Fields from Sparse Experimental Data

Flow measurements from wind-tunnel experiments suffer from noise and low spatio-temporal resolution. To reconstruct complete flow fields we use physics-based interpolation, enforcing a vorticity formulation of momentum conservation using a vortex-based solver. Its mesh-free nature and desirable stability characteristics enable a flexible and efficient framework. We formulate the data-assimilation procedure as an adjoint-based optimization for efficiency. We apply the methodology to particle tracking velocimetry experiments and demonstrate reconstruction of smooth, physically-correct fields from limited data.

Iliass Azijli, Richard Dwight, Jan Schneiders
Delft University of Technology
i.azijli@tudelft.nl, r.p.dwight@tudelft.nl,
j.f.g.schneiders@tudelft.nl

Hester Bijl
Faculty Aerospace Engineering
Delft University of Technology, NL
h.bijl@tudelft.nl

PP14

Parallel-in-Time Integration with Pfasst++

Iterative parallel-in-time integration methods like the recently developed "parallel full approximation scheme in space and time' (PFASST) provide parallelism along the temporal axis by integrating multiple time-steps simultaneously. We present PFASST++, a new performance-oriented library enabling time-parallelization for existing codes. PFASST++ provides flexibility and ease of use while maintaining portability across many HPC architectures. We present preliminary results using a high-order Boris integrator for charged particle dynamics in external magnetic fields.

Torbjörn Klatt
Juelich Supercomputing Centre
t.klatt@fz-juelich.de

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich
r.speck@fz-juelich.de

Mathias Winkel, Daniel Ruprecht
Institute of Computational Science
Universita della Svizzera italiana
mathias.winkel@usi.ch, daniel.ruprecht@usi.ch

Matthew Emmett
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov

PP14

An Optimization-Based Approach Toward Elastoplasticity: Introducing a Projected Newton Algorithm

In this study, we explore an alternative treatment of classical plasticity by casting the theory into a mathematical program. We propose a precise development and implicit implementation of a projected Newton algorithm to directly solve the dual optimization problem of incremental state update in nonlinear models. In particular, implementation of multi-surface plasticity models turns out to be no different from that of single surface models.

Zahra S. Lotfian, Mettupalayam Sivaselvan
SUNY at Buffalo
zahrasad@buffalo.edu, mvs@buffalo.edu

PP14

Redesigning Laser-Plasma Simulations to Optimize

the Use of Limited Memory Bandwidth

Memory bandwidth limitations are a growing problem affecting pf3D, a large-scale multi-physics code simulating laser-plasma interactions. Implementing lossy hardware compression between DRAM and cache could improve available bandwidth use in the future. pf3D is resilient to errors from lossy compression at each time step. We predict pf3Ds hardware compression requirements, describe strategies to optimize data movement to utilize hardware compression for a variety of physics kernels, and present software experiments on several existing architectures.

Eileen R. Martin
Stanford University
ermartin@stanford.edu

Steve Langer
Lawrence Livermore National Laboratory
langer1@llnl.gov

PP14

Parallel Numerics for Partitioned Multiphysics Coupling

Massively parallel multi-physics simulations require efficient parallel coupling schemes. This holds in particular for partitioned simulations where singlephysics components are combined in a flexible and extensible way. Parallelism and more than two coupled fields pose new to a coupling scheme. Not all established coupling schemes can cope with these challenges. We propose a set of new implicit coupling schemes that meet all requirements by replacing the usual fixed-point equation by a different one. To achoeve fast convergence, the resulting fixed-point equation is combined with efficient quasi-Newton methods suitable for black-box solvers.

Miriam Mehl
Universität Stuttgart
miriam.mehl@ipvs.uni-stuttgart.de

Benjamin Uekermann
TU Munich
uekerman@in.tum.de

Florian Lindner
Universität Stuttgart
florian.lindner@ipvs.uni-stuttgart.de

PP14

Applied Math and CS R&D on Doe Leadership Computing Facilities

The DOE Leadership Computing Facilities at Argonne and Oak Ridge national laboratories welcome applied mathematicians and computer scientists to carry out their research on our largest supercomputers. These systems, among the most powerful in the world, provide a unique environment to explore algorithm scalability and implementation at full scale. Examples of current projects and ways to gain access to our systems will be described.

Paul C. Messina
Argonne National Laboratory
messina@anl.gov

PP14

Simple Yet Fast Integration Method Using Qss for Stiff Chemical Kinetic Odes

A simple yet robust and fast time integration method is proposed for efficiently solving stiff chemical kinetic ordinary differential equations. The proposed method is based on a general formula which preserves the conservation laws for any integration operators constructed using the Lagrange multiplier method. And a quasi-steady-state approximation is used as the integrator. The time step size is automatically controlled by keeping a Lagrange multiplier small. The proposed method, named ERENA, provides the good performances.

Youhi Morii
Japan Aerospace Exploration Agency
morii.yuuhi@jaxa.jp

Hiroshi Terashima
The University of Tokyo
htera@rocketlab.t.u-tokyo.ac.jp

Mitsuo Koshi
Yokohama National University
koshi@rocketlab.t.u-tokyo.ac.jp

Taro Shimizu, Eiji Shima
Japan Aerospace Exploration Agency
shimizu.taro@jaxa.jp, shima.eiji@jaxa.jp

PP14

PtychoLib: Parallel Ptychographic Reconstruction

Ptychography is a phase retrieval technique for reconstructing a specimen's image from its diffraction patterns in light, X-ray, and electron microscopes. Despite the potential for ptychography there is a lack of tools for data online analysis while it is being acquired. PtychoLib is a library for real-time ptychographic phase retrieval. It uses a hybrid parallel strategy to divide the computation between multiple graphics processing units then merge partial reconstructions into one coherent phase contrast image.

Youssef Nashed
Argonne National Laboratory
ynashed@anl.gov

David Vine
Advanced Photon Source
Argonne National Laboratory
dvine@anl.gov

Tom Peterka
Mathematics and Computer Science Division
Argonne National Laboratory
tpeterka@mcs.anl.gov

Junjing Deng
Applied Physics
Northwestern University
junjingdeng2011@u.northwestern.edu

Rob Ross
Argonne National Laboratory
rross@mcs.anl.gov

Chris Jacobsen
Advanced Photon Source
Argonne National Laboratory
cjacobsen@aps.anl.gov

PP14

Efficient Numerical Algorithm for Virtual Design in Nanoplasmonics

Nanomaterials have given rise to many devices, from high-density data storage to optical bio-sensors capable of detecting specific biochemicals. The design of new nanodevices relies increasingly on numerical simulations, driving a need for efficient numerical methods. In this work, integral equations are used to efficiently solve the electromagnetic transmission problem at the interface of a dielectric and a periodic metal nanostructure. Moreover, the same integral equations are used as the basis for an optimization method.

Alexandra Ortan
University of Minnesota
aortan@umn.edu

PP14

Model-Reduction for Closed-Loop Control of Unsteady Flows Using Plasma Actuators

Flow instabilities in aircraft aerodynamics increase noise emissions and drag. Active closed-loop flow control using plasma actuators is a promising way to reduce these effects. However, a rigorous model is required for the feedback control design. The spatially discretized actuator/flow dynamics is high-dimensional, nonlinear and, thus, computationally intractable for real-time applications. Therefore, a suitable control-oriented reduced-order model, which captures the essential physics dominating the system dynamic response, while accounting for the closed-loop dynamics, is required.

Laura Pasquale, Paul Houston, Pericle Zanchetta
University of Nottingham
laura.pasquale@nottingham.ac.uk,
pmzph1@exmail.nottingham.ac.uk,
eezpz@exmail.nottingham.ac.uk

PP14

Moving Pictures: Animating Still Images

We present a means of developing digital image transformations that allow a still image to be turned into a short and visually pleasing animation. Rather than manually altering successive frames to create the illusion of motion, the method presented here requires only the input of a few parameters for each transformation. We developed a mathematical framework wherein we defined animations as sequences of still images, and 'transformations' as composable functions on such sequences.
To implement this work, we have built a Matlab library of composable functions that streamline the process of turning still images into novel animations. Examples include manipulation of contrast, intensity, and colors of pixels, as well as warps of contours, positions, and size of select regions. The transformations allow for easy animation of regions of interest, giving some semblance of life to still images by turning them into animated GIFs.

Michael Pilosov
University of Colorado: Denver
michael.pilosov@ucdenver.edu

PP14

Bridging Multiple Structural Scales with a Generalized Finite Element Method

In thermo-mechanical problems subjected to intense heating effects, locally high solution fidelity is required near material-scale interfaces to capture strong gradients which might not be resolved using, for example, traditional homogenization approaches. Effective numerical methods must also bridge important localized, nonlinear behavior to the structural scale. This poster introduces a generalized finite element method (GFEM) involving the solution of interdependent coarse- and fine-scale problems for resolving localized material interfaces on the structural scale.

Julia A. Plews
Department of Civil and Environmental Engineering
University of Illinois at Urbana-Champaign
jplews2@illinois.edu

C. Armando Duarte
University of Illinois at Urbana-Champaign
caduarte@illinois.edu

PP14

A Parallelization Strategy for Large-Scale Vibronic Coupling Calculations

The vibronic coupling model of Köppel, Domcke, and Cederbaum is a powerful means to understand, predict, and analyze electronic spectra of molecules. We describe a new distributed-memory parallel algorithm for carrying out such calculations, based on a stencil representation of the required computational steps. Our algorithm utilizes coarse-grained parallelism to achieve near-linear scaling. The algorithm is demonstrated by simulating the low energy portion of the VUV spectrum of *trans*-1,3-butadiene.

Scott Rabidoux
Institute for Computational Engineering and Sciences
The University of Texas at Austin
scottrabidoux@gmail.com

Victor Eijkhout
The University of Texas at Austin
Texas Advanced Computing Center
eijkhout@tacc.utexas.edu

John Stanton
Department of Chemistry
The University of Texas at Austin
jfstanton@gmail.com

PP14

Some Numerical Methods for Modified Bessel Functions

The high-quality numerical methods for the computation of modified Bessel functions of the second kind with imaginary, complex order and real argument are elaborated. A Tau method computational scheme is applied for the constructive approximation of hypergeometric type differential equations and their systems. The numerical solution of some mixed boundary value problems for the Helmholtz

equation in wedge domains is developed.

Juri M. Rappoport
Russian Academy of Sciences
jmrap@landau.ac.ru

PP14

Oof: An Object-Oriented Finite-Element Solver for Materials Science

Materials scientists frequently have a requirement to build finite-element models of microstructural systems in order to explore structure-property relationships. The OOF software provides researchers with a modeling tool which starts from familiar image data formats, speaks the language of materials science, and provides push-button access to sophisticated tools for image manipulation and mesh construction. The recently-released 3D version allows for more complex and realistic virtual experiments to be performed.

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

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

PP14

Analyzing and Classifying "Two-Cycles" of Trigonometric Functions in Newton's Method

For this research we are analyzing trigonometric functions, specifically $f(x) = sin(x)$, in order to find initial values that when iterated upon, using Newton's Method, will give us a new value, which when iterated upon will give us the initial value again, This process then repeats. So far we have found reliable characterizing equations which give us the exact roots that correspond to these special initial values.

Morgan Rupard, Jennifer Switkes
Cal Poly Pomona
morganrupard@gmail.com, jmswitkes@csupomona.edu

PP14

Orbital Localization in Madness

The Multiresolution ADaptive Numerical Environment for Scientific Simulation (MADNESS) library provides a high-level environment for the solution of integral and differential equations in many dimensions using adaptive, fast methods with guaranteed precision based on multiresolution analysis and novel separated representations. These features provide a backdrop in which linear scaling electronic structure algorithms may be realized. Towards this goal, we will explore avenues such as fourth order molecular orbital localization to advance scalability with system size while maintaining accuracy and parallelism.

Bryan E. Sundahl
Institute for Advanced Computational Science
Stony Brook University
bryan.sundahl@stonybrook.edu

Robert Harrison
Brookhaven National Laboratory and Stony Brook University
rjharrison@gmail.com

Scott Thornton
Institute for Advanced Computational Science
wsttiger@gmail.com

PP14

Analysis of Anderson Acceleration for Coupled Neutronic and Thermal Hydraulic Calculations in a Light Water Reactor

In a coupling between Insilico (neutronics) and AMP (fuel performance, subchannel flow), Picard iteration shows instability due to oscillatory error modes arising at high enough power. We consider Anderson acceleration as an alternative solution method to Picard. We develop a simplified model which captures the relevant physics from the high fidelity simulation, and perform parametric studies which suggest that Anderson should be more robust and faster converging than Picard for the Insilico/AMP coupling.

Alexander R. Toth
North Carolina State University
artoth@ncsu.edu

C.T. Kelley
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu

Stuart Slattery
Computer Science and Mathematics Division
Oak Ridge National Laboratory
slatterysr@ornl.gov

Steven Hamilton
ORNL
hamiltonsp@ornl.gov

Kevin Clarno
Oak Ridge National Laboratory
clarnokt@ornl.gov

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

PP14

The Rapid Optimization Library (rol) in Trilinos

The Rapid Optimization Library (ROL) in Trilinos provides an object-oriented framework for large-scale, derivative-based optimization. The library is matrix-free and linear-algebra agnostic permitting easy interface with application code. ROL implements a suite of unconstrained and constrained optimization algorithms including: gradient descent, quasi-Newton, and inexact-Newton with line-search and trust-region globalization. A stochastic optimization subpackage (SOL) supplies default implementations of numerous risk measures and adaptive sampling capabilities. Examples demonstrate solutions of large op-

timization problems with model uncertainties.

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

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

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

PP14

Integrating Software Tools to Parallel Adaptive Simulations of Fusion Plasma in Tokamaks

Mesh-based methods are extensively applied in macroscopic studies of fusion plasmas in the tokamak. Software tools that include the parallel unstructured mesh management infrastructure (PUMI), the mesh adaption procedure (meshAdapt), Simmetrix meshing tool and PETSc equation solver are integrated to the physics simulation codes under development at Princeton Plasma Physics Lab (PPPL) such as M3D-C1.

Fan Zhang
Rensselaer Polytechnic Institute
fanzhangpku@gmail.com

Mark S. Shephard, E. Seegyoung Seol
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu, seols@rpi.edu

PP15

Using Radar Imagery Data to Invert for Maritime Environments

Predicting in situ maritime conditions is of great importance to radar detection in marine atmospheric boundary layers. To predict the maritime conditions, the radar wave propagation can be modeled using radar imagery data through proper orthogonal modes indexed on a specific environment. These modes live on the compact Stiefel manifold, and by exploiting the Riemannian structure of it, interpolation between the modes is possible and can be used to invert for the maritime environment.

Vasileios Fountoulakis, Christopher J. Earls
Cornell University
vf59@cornell.edu, cje23@cornell.edu

PP15

Thermal Imaging of Sub-Pixel Cracks Through Metal Plates

Active thermography has been studied as a non-destructive evaluation technique for structural components, using the well-understood theory of heat conduction to infer the presence, location, and characteristics of flaws in a solid. We here use both theory and simulation to determine optimal ranges of several operational parameters to be used in thermographic evaluation of thin metal plates, such as aircraft skins. We also implement a stochastic approach for characterizing cracks smaller than the imaging resolution.

Andrew Loeb, Christopher Earls
Cornell University
ael89@cornell.edu, earls@cornell.edu

PP15

Numerical Simulation of Ni Grain Growth in a Thermal Gradient

The Potts model is well developed to simulate normal, curvature-driven grain growth and has been used extensively to study various aspects of grain growth. In this work, we use this model, implemented into the SPPARKS code to simulate grain growth of Ni turbine blades that are heat treated with temperature gradient intentionally introduced to vary the grain size for varied mechanical properties across the extent of the blade. We will present the model, demonstrate its application and estimate mechanical properties as a function of position.

John A. Mitchell
Sandia National Laboratory
jamitch@sandia.gov

Veena Tikare
Sandia National Laboratories
vtikare@sandia.gov

PP15

Supply Chain Disruptions

Forecasting supply chain disruptions provides buyers with a tool to mitigate risk, avoid liability and increase revenue. A diffusion type algorithm is presented that can be used by a wholesale buyer to predict supply chain disruptions from multiple distributors. A client use case is presented with a discussion of the success of this algorithm via comparison to a logistic regression model designed for the same use case and comparison to chance.

Thomas Morrisey
Infosys
thomas_morrisey@infosys.com

Ravi Prasad
Infosys Limited
ravi_prasad02@infosys.com

PP15

Computational and Experimental Analysis of Dental Implants under Different Loading Conditions and Locations

The use of dental implants to solve different problems in dentistry has been growing rapidly. The success rates of dental implants are also very important for patients. This study investigates the stability of dental implants under different loading locations and conditions with the use of Finite Element Analysis. The CAD model is obtained from CT images. Designed dental implants were fabricated and experimentally tested (ISO 14801) to compare with computational(FEA) results.

Emre Ozyilmaz
Hitit University Engineering Faculty, Dep. of Mechanical Eng
emre.hitit@gmail.com

Eda Ozyilmaz, Halil Aykul
Hitit University Engineering Faculty,
Dept. of Mechanical Eng.
edaozyilmaz88@gmail.com, halilaykul@hitit.edu.tr

Mehmet Dalkiz
Mustafa Kemal University Dentistry Faculty
mdalkiz@yahoo.com

Ahmet Cini
Hitit University Engineering Faculty,
Dept. of Mechanical Eng.
ahmetcini@gmail.com

PP15

Math Projects with Tracker Video Analysis

Using Tracker freeware, students analyze own-recorded video clips and construct mathematical models in exploring real-world mathematical applications of concepts from algebra through calculus and differential equations.

Euguenia Peterson
Richard J. Daley College
epeterson@ccc.edu

PP15

Investigation of Numerical Models for New High Temperature Superconductors

Chad Sockwell, Florida State University

Dr. Janet Peterson and Dr. Max Gunzburger, Florida State University

Recent discoveries of new high temperature superconductors initiated investigations to harness these new material's properties. Unfortunately, many new high temperature superconductors come with odd properties such as multiband interactions and anisotropic behavior, as is the case for Magnesium Diboride. In this research Ginzburg Landau model variants are modified and simulated to try and simulate these materials more realistically. This extends to three dimensional domains, anisotropy, and microscopic corrections such as the Extended Ginzburg Landau model.

Chad Sockwell
Florida State University
Department of Scientific Computing
kcs12j@my.fsu.edu

PP15

Characterization of Ternary Eutectic Solidification Patterns from Phase-Field Simulations and Experimental Micrographs

Large scale three-dimensional phase-field simulations based on the grand potential are used to model the microstructure evolution from directional solidification of ternary eutectics. To compare the computational cross-sections and experimental micrographs, principal component analysis is applied. With this, the influence of single parameters from the different reported data sets on the evolving patterns can be quantified. This allows to determine both, the convergence of microstructure simulations and the necessary domain size for statistical reliable results.

Philipp Steinmetz, Johannes Hötzer, Marcus Jainta, Britta Nestler
Karlsruhe Institute of Technology (KIT)
Institute of Applied Materials (IAM-ZBS)
philipp.steinmetz@kit.edu, johannes.hoetzer@kit.edu, marcus.jainta@kit.edu, britta.nestler@kit.edu

Yuksel Yabansu, Surya Kalidindi
Georgia Institute of Technology
Materials Informatics for Engineering Design
yabansu@gatech.edu, surya.kalidindi@me.gatech.edu

PP15

Multidisciplinary Development of An Autonomous Underwater Vehicle: Cooperative Fleet for Surveillance Mission

A fleet of cooperative autonomous underwater vehicles (AUV) called Eco-Dolphin was designed to collect coastal science data under saltwater condition. One AUV equipped both long-range wireless and Wireless sonar communication and the others equipped short-range acoustic and wireless sonar communication. The fleet position can be tracked via GPS and Wi-Fi and individual positions can be calculated. Successful achieved wireless sonar communication on prototype one this summer.

Ci Wen, Stacey Joseph-Ellison, Junzhen Shao, Qi Zhou, Jonathan Jaworski, Zakaria Daud
Embry-Riddle Aeronautical University
wenc1@my.erau.edu, ellisos1@my.erau.edu, shaoj@my.erau.edu, zhouq@my.erau.edu, jaworskj@my.erau.edu, daudz@my.erau.edu

PP15

Towards Real-Time Blob-Filaments Detection in Fusion Plasma

Magnetic fusion is a viable energy source for the future. The success of magnetically confined fusion reactors demand steady-state plasma confinement which is challenged by the edge turbulence such as the blob-filaments. We present a real-time outlier detection algorithm to efficiently find blob-filaments in fusion simulations and experiments. We have implemented this algorithm with hybrid MPI/OpenMP and demonstrated the accuracy and efficiency with a set of data from the XGC1 fusion simulations.

Lingfei Wu
Department of Computer Science
College of William & Mary
lfwu@cs.wm.edu

Kesheng Wu
NERSC, Lawrence Berkeley National Lab
kwu@lbl.gov

Alex Sim
Lawrence Berkeley National Lab
asim@lbl.gov

Andreas Stathopoulos
College of William & Mary
Department of Computer Science

andreas@cs.wm.edu

PP101

Bet: Algorithmic and Error Analyses

We review the basic theory behind the non-intrusive measure-theoretic algorithm for solving stochastic inverse problems for deterministic models. We define sigma-algebra constraints on sampling methods to achieve certain approximation properties of events. We provide convergence results and a priori and a posteriori error analyses producing fully computable a posteriori error estimates. The fully computable a posteriori error estimates provide lower and upper bounds for computed probabilities of events. We demonstrate the use of this analysis on numerical results utilizing a posteriori error estimates to improve the accuracy of computed probability measures.

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

PP101

BET: Applications for an Open Source Inverse Problems Package

We present example applications of the BET Python Package to solve the stochastic inverse problem within a measure theoretic framework. We demonstrate various features of the BET package for approximating inverse probability measures including goal-oriented adaptive sampling algorithms, parallel capabilities, differing approximation options, and visualization capabilities. BET can interface with a variety of physics-based computational models.

Lindley C. Graham
University of Austin at Texas
Institute for Computational Engineering and Sciences (ICES)
lgraham@ices.utexas.edu

Steven Mattis
University of Texas at Austin
steve.a.mattis@gmail.com

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

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

PP101

Bet: Modifications and Analysis for Model Discrepancies

We consider the selection of a model as an uncertain input parameter and we quantify this uncertainty in the measure-theoretic framework using the BET Python Package. For example, in multi-scale models, optimal mesoscale models that upscale fine-scale parameters for use in the macro-scale model are often uncertain. We describe and analyze modifications to the BET package to optimize load balancing taking into account the widely different computational demands of different models with different resolutions

Nishant Panda
Colorado State University
nishant.panda@gmail.com

PP102

Tsunami Modeling In North Africa Using Geoclaw Software: a Tool for the Tsunami Scenario Database in the West Mediterranean

North Africa exhibits an active plate boundary with Europe, highlighted by a moderate seismicity.The triggering of regional tsunamis is due to (1) earthquakes generated in North Algeria and (2) landslides. Modeling carried out using tsunami sources located along the Algerian coast is presented hereby. Computing scenarios using the Riemanns algorithm implemented within the GEOCLAW software to solve the non linear shallow water equations provides accurate results in comparison with other tsunami software.

Lubna Amir
USTHB University
lamir@usthb.dz

Walter Dudley
University of Hawaii at Hilo
200 W. Kawili Street, Hilo, Hawaii, USA
dudley@hawaii.edu

Jean Roger
G-Mer Etudes Marines
Av. de l'Europe - Saint Francois Guadeloupe, France
jeanrog@hotmail.fr

PP102

ForestClaw : Parallel, Adaptive, Multiblock Simulations for Clawpack

We describe our recent progress in developing ForestClaw, an adaptive mesh refinement code based on the dynamic octree library p4est (C. Burstedde, Univ. of Bonn) and finite volume solvers, including ClawPack (R. J. LeVeque). One particular feature of ForestClaw is the ease with which it can handle multi-block domains such as the cubed-sphere, or other geometries not easily described by a single logically rectangular domain. ForestClaw preserves all of the scalability and performance of the underlying grid management library p4est.

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

PP102

Adjoint Methods for Guiding Adaptive Mesh Refinement in Wave Propagation Problems

AMRClaw and GeoClaw software uses block-structured adaptive mesh refinement to selectively refine around propagating waves. The refinement criterion is often based on local estimates of error via Richardson extrapolation, or some measure of solution smoothness such as the gradient. For problems where a small region of the solution is of primary interest, solving the time-dependent adjoint equation and using a suitable inner product with the forward solu-

tion allows more precise refinement of the relevant waves.

Brisa Davis
University of Washington
bndavis@uw.edu

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

PP102

A Community-Driven Collection of Approximate Riemann Solvers for Hyperbolic Problems

The key ingredient in modern numerical methods for hyperbolic problems is the Riemann solver. The same algorithms can then be used to solve a wide range of hyperbolic problems – water waves, fluid dynamics, elasticity, electromagnetics, etc. We present an effort to make available a large library of Riemann solvers for general use, and the development of a set of IPython notebooks that describe and interactively explain the solvers for important systems.

Mauricio J. Del Razo
University of Washington
mauricio_delrazo@yahoo.com

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

PP102

High Resolution Tsunami Modeling at the Mediterranean Coast of Israel Towards An Early Warning Tsunami Scenarios Data Bank

A tsunami modelling investigation using the state of the art, open source tsunami model (GeoClaw a subset of the software package Clawpack), its adaptation to investigate the impact of tsunami wave generation, propagation and inundation at the Mediterranean coast of Israel using high resolution bathymetric and topographic grid, aided by additional tsunami generation modelling tools simulating the initial stages of tsunami generation by earthquake induced tectonic plates rupture and movement or by landslide on the coastal shelf.

Barak Galanti
Inter-University Computational Center
Tel Aviv University
galanti@mail.iucc.ac.il

PP102

PyClaw: Accurate, Scalable Solution of Hyperbolic PDEs in Python

PyClaw is a Python-based interface to fast and accurate Godunov-type numerical algorithms for modeling nonlinear waves. It is massively scalable and includes both 2nd-order TVD and higher-order WENO discretizations. It is built on existing codes, including Clawpack, SharpClaw, and PETSc. We describe PyClaw's design, development, and some applications.

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Aron Ahmadia
US Army Engineer Research and Development Center
aron@ahmadia.net

Kyle T. Mandli
University of Texas at Austin
ICES
ktm2132@columbia.edu

PP102

Practical Applications of GeoClaw to Tsunami Hazard Assessment

The GeoClaw software has recently been used for several tsunami hazard assessment projects and two examples will be presented in this poster. The first is the analysis of an elementary school site in Ocosta, WA that has since been selected for the first "vertical evacuation structure" to be built in the United States. The second is the development of improved probabilistic tsunami hazard assessment techniques.

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

Loyce Adams
University of Washington
lma3@uw.edu

Frank I. Gonzalez
U. Washington, Dept. of Earth and Space Sciences
figonzal@u.washington.edu

PP102

CUDACLAW: A GPU Framework for the Solution of Hyperbolic Pdes

CUDACLAW is a data-parallel framework that allows users to take advantage of GPU accelerators to solve hyperbolic PDEs, without being burdened by the need to write CUDA code, worry about thread and block details, data layout, and data movement between the different levels of the memory hierarchy. The user defines the set of PDEs via a serial Riemann solver and the framework takes care of orchestrating the computations and data transfers to maximize arithmetic throughput. A prototype implementation shows that, on 2D and 3D acoustics wave equations and shallow water equations, the framework can achieve performance comparable to that of manually-tuned code.

George M. Turkiyyah
American University of Beirut
gt02@aub.edu.lb

H. Gorune Ohannessian
University of Wisconsin
gorune@cs.wisc.edu

Aron Ahmadia
US Army Engineer Research and Development Center
aron@ahmadia.net

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

PP103
An Overview of PETSc

The changing landscape of scientific application needs and high-performance architectures requires continued innovation in mathematical algorithms for the robust solution of large-scale numerical problems. This poster provides an overview of the Portable, Extensible Toolkit for Scientific Computing (PETSc), a suite of data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations (PDEs). We also will provide demos of a variety of PDE-based examples, including advances in composable linear, nonlinear, and timestepping solvers, which incorporate multiple levels of nested algorithms and data models to exploit architectural features and/or problem-specific structure.

Satish Balay
Argonne National Laboratory
balay@anl.gov

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory and CU Boulder
jed@jedbrown.org

William D. Gropp
University of Illinois at Urbana-Champaign
Dept of Computer Science
wgropp@illinois.edu

Matthew Knepley
University of Chicago
knepley@gmail.com

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Hong Zhang
MCS, Argonne National Laboratory
hzhang@mcs.anl.gov

PP103
Construction of Parallel Adaptive Simulation Loops

The automation of reliable large-scale simulations requires the integration of a number of operations that include mesh generation, equation discretization, equation solution, error estimation and discretization improvement within a loop that requires regular dynamic load balancing to maintain scalability. This poster presents a set of components to support adaptive unstructured simulations on massively parallel computers. The in-memory integration of these components with multiple unstructured mesh finite element and finite volume codes will be included.

Brian Granzow
Rensselaer Polytechnic Institute
granzb@rpi.edu

PP103
MueLu: Multigrid Framework for Advanced Architectures

MueLu multigrid library is a part of Sandia's Trilinos project. MueLu is designed to be flexible, easily extensible, and efficient on emerging architectures. MueLu allows users to specify preferred data (ordinal or scalar) and shared memory ("node") types by leveraging Trilinos' sparse linear algebra libraries Kokkos and Tpetra. Current algorithms include smoothed aggregation (SA), nonsymmetric SA, Maxwell solvers, and energy-minimization. In this poster, we highlight important design features and large-scale results on DOE supercomputers.

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

Andrey Prokopenko
Sandia National Laboratories
aprokop@sandia.gov

PP103
Parallel Unstructured Mesh Infrastructure

The Parallel Unstructured Mesh Infrastructure (PUMI) supports the needs of parallel mesh-based analysis codes to develop complete simulation workflows. This poster will focus on recent developments including array-based mesh data structures and field abstractions that provide substantial reductions in memory usage. Component reorganization via dependency inversion provides a substantial decrease in code size. The integration of PUMI with dynamic load balancing and its use in simulation workflows will be included.

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

E. Seegyoung Seol
Rensselaer Polytechnic Institute
Scientific Computation Research Center
seols@rpi.edu

Gerrett Diamond
Rensselaer Polytechnic Institute
diamog@rpi.edu

PP103
Massively Parallel Adaptive Simulations Using Petsc for Turbulent Boundary Layer Flows

A set of tools and techniques are presented on adaptive methods for boundary layer meshes. Such meshes are useful in wall-bounded turbulent flows. An adaptive approach

for such meshes must maintain highly anisotropic, graded, and layered elements near the walls. Parallel procedures must account for mixed elements, i.e., in mesh modifications and load balancing. Additionally, we study the efficiency and scalability of the linear solvers and preconditioners available via PETSc for this class of problems.

Michel Rasquin
University of Colorado Boulder
michel.rasquin@colorado.edu

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

Benjamin Matthews
University of Colorado Boulder
benjamin.a.matthews@colorado.edu

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

Onkar Sahni
Rensselaer Polytechnic Institute
sahni@rpi.edu

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

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

PP103

Sparse Direct Solvers and Preconditioners on Manycore Systems

We develop scalable sparse direct linear solvers and effective preconditioners for the most challenging linear systems on manycore parallel machines. Our focal efforts are the developments of two types of solvers: The first is a pure direct solver, encapsulated in SuperLU software. The second is the nearly-optimal preconditioners using the HSS low-rank approximate factorization of the dense submatrices, encapsulated in STRUMPACK software.

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

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

Francois-Henry Rouet
Lawrence Berkeley National Laboratory
fhrouet@lbl.gov

Piyush Sao, Richard Vuduc
Georgia Institute of Technology
piyush3@gatech.edu, richie@cc.gatech.edu

PP103

Dynamic Partitioning Using Mesh Adjacencies

Parallel unstructured mesh-based applications running on the latest petascale systems require partitions optimizing specific balance metrics. Methods combining the most powerful graph based and geometric methods with diffusive methods directly operating on the unstructured mesh are discussed. Implementation details will be provided for the components comprising the direct unstructured mesh based methods as will initial results for applications with meshes of several billion elements on over 500k parts.

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

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
dibanez@scorec.rpi.edu

Gerrett Diamond
Rensselaer Polytechnic Institute
diamog@rpi.edu

PP103

Fastmath Structured Mesh and Particle Technologies

Scientific phenomenon happen over a wide range of length and time scales. Chombo and BoxLib are libraries that are developing and deploying state-of-the-art structured adaptive mesh and particle-in-cell technologies under the aegis of FASTMath. The evolving algorithms and computational frameworks provided by these libraries allow scientific application codes to capture these scales efficiently and enable scientific discovery at scale currently, as well as prepare for the future architectures.

Anshu Dubey, Phillip Colella
Lawrence Berkeley National Laboratory
adubey@lbl.gov, PColella@lbl.gov

Mark Adams
Lawrence Berkeley Laboratory
mfadams@lbl.gov

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

Dan Graves, Terry J. Ligocki
Lawrence Berkeley Laboratory
dtgraves@lbl.gov, TJLigocki@lbl.gov

Brian Van Straalen
Lawrence Berkeley National Laboratory
Compuational Research Division
bvstraalen@lbl.gov

Milo Dorr
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

dorr1@llnl.gov

PP103

Sundials: Suite of Nonlinear and Differential/algebraic Solvers

In this poster we present the SUNDIALS library suite, consisting of the solver libraries CVODE, CVODES, ARKode, IDA, IDAS, and KINSOL. This DOE-funded solver library focuses on robust and scalable solvers for systems of ordinary differential equations, systems of differential-algebraic equations and systems of nonlinear equations. Sharing a common set of vector libraries and linear solvers (both direct and iterative), these solvers are designed for algorithmic flexibility, with interfaces in C, C++ and Fortran (some even in Matlab), and may be easily adapted to application-specific and user-defined data structures.

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

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

PP103

Hypre: High Performance Preconditioners

The hypre software library provides high performance preconditioners and solvers for the solution of large sparse linear systems on massively parallel computers. One of its attractive features is the provision of conceptual interfaces, which include a structured, a semi-structured, and a traditional linear-algebra based interface. These interfaces give application users a more natural means for describing their linear systems, and provide access to hypre's solvers, which include structured and unstructured algebraic multigrid solvers.

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

Jacob B. Schroder, Ulrike M. Yang
Lawrence Livermore National Laboratory
schroder2@llnl.gov, umyang@llnl.gov

PP104

Stochastic Methods in Mesoscopic Materials Modeling

For soft materials, a central challenge is to link bulk properties with the behaviours of the material mesostructures. Such interactions/kinetics often involve a subtle interplay of enthalpic and entropic effects extending over a wide range of scales presenting well-known challenges for modelling and simulation. We present our recent progress on dynamic implicit-solvent coarse-grained approaches that utilize a fluctuating hydrodynamics thermostat. We investigate the properties of polymeric materials and biophysical questions for lipid bilayer membranes.

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

PP104

Concurrent Coupling Methods in Mesoscopic Materials Modeling

Mesoscopic methods with thermodynamic fluctuations, such as dissipative particle dynamics, are powerful tools to study material properties at relevant spatial-temporal scales. Concurrently coupling a mesoscopic model with a microscopic description, such as molecular dynamics, can benefit further from microscopic description with great details. Concurrently coupling a mesoscopic model with a continuum discretization, such as smoothed particle hydrodynamics, enables to simulate real applications at large scale with great computational efficiency.

Xin Bian
Brown University
Xin_Bian@brown.edu

PP104

Overview of Mathematics for Mesoscopic Modeling of Materials

The Collaboratory on Mathematics for Mesoscopic Modeling of Materials, or CM4, develops systematic mathematical foundations for understanding and controlling the fundamental mechanisms in mesoscale processes. CM4 focuses on the development and integration of particle-based, continuum-based, and stochastic methods, as well as the concurrent coupling between them. As science drivers, a set of demonstration problems will facilitate integration of CM4 mathematical models and numerical approaches in the context of important materials modeling challenges.

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

PP104

Hierarchical Coarse-graining and Parallelization Methods for Mesoscale material Models

We discuss information theory-based methods for the parameterized coarse-graining of non-equilibrium extended systems, typically associated with coupled physicochemical mechanisms or driven systems, and where steady states are unknown altogether, e.g. do not have a Gibbs structure. Our pathwise information theory tools provide reliable molecular model parameterizations and rational model selection through suitable dynamics-based information criteria. We also discuss connections of the developed tools to force-matching in a dynamics context and the transferability of coarse-graining maps.

Markos A. Katsoulakis
University of Massachusetts, Amherst
Dept of Mathematics and Statistics
markos@math.umass.edu

PP104

Particle-Based Methods in Mesoscopic Materials Modeling

Coarse-grained particle-based method is capable of simulating mesoscopic phenomena with a discrete or discon-

tinuous nature. Its governing equations can be obtained by applying the Mori-Zwanzig projection on a microscopic dynamics. We present a bottom-up approach to quantify the coarse-grained model via the Mori-Zwanzig formulation with microscopic data provided by molecular dynamics (MD) simulations. Results show that such MD-informed coarse-grained model is able to accurately reproduce its underlying MD system.

Zhen Li
Brown University
Brown University
Zhen_Li@brown.edu

PP104

Grid-Based Methods in Mesoscopic Materials Modeling

Our focus is on developing continuum-level simulations to describe the dynamics of colloidal systems, self-assembly and flow transport in micro-channels. Generally the approach is to use grid-based methods but we also work with Lagrangian particle-based methods. Here we evaluate the electrostatic or magnetostatic moments of particles in suspension, subject to an external field, and the resulting force interactions. We use both immersive methods (smoothed profile method) and finite multipoles, and compare accuracy and efficiency.

Martin Maxey
Division of Applied Mathematics,
Brown University
maxey@cfm.brown.edu

PP104

Fast Solvers for Mesoscopic Materials Modeling

The Collaboratory on Mathematics for Mesoscopic Modeling of Materials is developing numerical solvers designed to efficiently solve systems modeling charge transport. These phenomena are commonly encountered in engineering applications and are also of growing importance to biological contexts as well. We conduct a numerical analysis of our proposed schemes to understand important aspects of our solvers, such as stability, consistency, and well-posedness. We consider fast solvers for the Poisson-Nernst-Planck (PNP) system using a finite element discretization. We leverage qualitative mathematical structures and physical properties of this system to improve the quality and efficiency of computed solutions. Our proposed discretization scheme readily leads to an energy estimate for the computed solution that is characteristic of the continuous PNP system. Further analysis for the efficiency and robustness of the solvers is considered as a separate component of the analysis.

Jinchao Xu
Pennsylvania State University
xu@math.psu.edu

Chun Liu
Department of Mathematics, Penn State University
University Park, PA 16802
liu@math.psu.edu

Xiaozhe Hu
The Pennsylvania State University
xiaozhe.hu@tufts.edu

Maximilian S. Metti
Penn State University
msm37@psu.edu

PP104

Applications in Mesoscopic Modeling of Materials

Recent applications in micro-/nano-technology, material assembly and biological systems demand robust and accurate computational modeling of multiphysical processes at the mesoscale. In this poster we focus on mathematical models and numerical schemes that can effectively capture mesoscale multiphysics (hydrodynamics, transport, electrostatics and chemical reaction). Specifically, we show simulation results on mixing and separation processes in micro-/nano-channel, electrokinetic flow through semi-permeable membranes, and diffusive reaction on biomolecules.

Wenxiao Pan
Pacific Northwest National Laboratory
wenxiao.pan@pnnl.gov

Mauro Perego
CSRI Sandia National Laboratories
mperego@fsu.edu

PP104

Coarse-Graining in Mesoscopic Materials Modeling

The mesoscopic modeling of materials is a flourishing subject of research due to its importance for many technological and societally relevant applications. By definition, mesoscopic models entail the process of coarse-graining of microscopic (molecular) models. The resulting models may exhibit space/time symmetry breaking. In our recent work as part of CM4 we have focused on the use of renormalization in order to construct coarse-grained models which are adaptive, in the sense that the relevant coefficients of the model are estimated on the fly as the system evolves. We believe that renormalization (perturbative or non-perturbative) is a viable method for constructing coarse-grained models of controlled accuracy for complex systems. The purpose of the current poster is to present our recent results on the subject as well as outline future directions and open problems.

Panos Stinis
University of Minnesota, Twin Cities
stinis@umn.edu

PP105

Realizability Limiting in High-Order Numerical Solutions of Entropy-Based Moment Closures

Entropy moment closures for kinetic equations (colloquially known as M_N models) have attractive theoretical properties (hyperbolicity, entropy dissipation, and positivity) but are only defined in the set of realizable moment vectors, that is those which are consistent with a positive distribution. High-order numerical solutions do not always stay in this set, so we investigate the use of a limiter to handle nonrealizable moments in the implementation of a high-order discontinuous Galerkin method.

Graham Alldredge
RWTH AACHEN University
alldredge@mathcces.rwth-aachen.de

Florian Schneider
TU Kaiserslautern
fschneid@mathematik.uni-kl.de

PP105

Exploration and Validation of Full-Domain Massively Parallel Transport Sweep Algorithms

Scaling of discrete-ordinates transport sweep-based algorithms is currently of great general interest. Recent research on volume-based spatial decomposition approaches shows great potential. Schedule conflict resolution and spatial domain overloading are essential algorithm characteristics for obtaining good scaling. We describe computational experiments performed on the IBM BG/Q machine (Sequoia) at LLNL that indicate MPI based algorithms scale to the full machine with good efficiencies predicted by parallel performance models. We also present simple studies of single node performance using task, thread and instruction level parallelism for these algorithms.

Teresa S. Bailey
LLNL
bailey42@llnl.gov

Peter Brown, Adam Kunen
Lawrence Livermore National Laboratory
pnbrown@llnl.gov, kunen1@llnl.gov

PP105

A New Moment Method in the Kinetic Theory of Gases Based on the L^2 Function Space

Recently, a framework for hyperbolic model reduction in the kinetic theory was proposed. Based on this framework, we reconsider the convergence of Grad's expansion, and find the function space may be insufficient in some cases. To make improvements, we replace the weighted L^2 space used in the classic moment method by the plain L^2 space, and special projections are used in the model recduction to maintain both hyperbolicity and conservation.

Zhenning Cai
Center for Computational Engineering Science
RWTH Aachen University
cai@mathcces.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen University
mt@mathcces.rwth-aachen.de

PP105

Markov Chain Formalism for Radiative Transfer in Planetary Atmospheres: Forward Modeling, Including Linearization

The 1D vector radiative transfer problem is solved with a Markov chain approach. The matrix structure of that method enables efficient numerical computation of polarized radiation fields throughout the atmosphere, as well as straightforward linearization to derive Jacobian matrixes exactly. We applied this computational forward modeling framework for remote sensing signals to aerosol/surface property retrievals from air- and spaceborne observations on Earth and to analyses of polarized imagery of Titan from the Cassini probe.

Anthony B. Davis, Feng Xu, Robert West, David Diner
Jet Propulsion Laboratory
California Institute of Technology
anthony.b.davis@jpl.nasa.gov, feng.xu@jpl.nasa.gov,
robert.a.west@jpl.nasa.gov, david.j.diner@jpl.nasa.gov

PP105

On the Hyperbolicity of Grads 13 Moment System

In gas kinetic theory, lack of global hyperbolicity is a major critique for Grad's moment method. For the well-known Grad's 13 moment system, I. Muller et. al.(1998) pointed out it is not globally hyperbolic for 1D flow. In this presentation, we will point out for 3D case, for each equilibrium state, none of its neighbourhoods is contained in the hyperbolicity region. And a globally hyperbolic regularization for it will be proposed.

Yuwei Fan
School of Mathematical Sciences
Peking University, Beijing, P.R.China
ywfan@pku.edu.cn

Zhenning Cai
Center for Computational Engineering Science
RWTH Aachen University
cai@mathcces.rwth-aachen.de

Ruo Li
School of Mathematical Science
Peking University
rli@math.pku.edu.cn

PP105

Convergence of Filtered Spherical Harmonic Equations for Radiation Transport

We analyze the global convergence properties of the filtered spherical harmonic (FP_N) equations for radiation transport. The well-known spherical harmonic (P_N) equations are a spectral method (in angle) for the radiation transport equation and are known to suffer from Gibbs phenomena around discontinuities. The filtered equations include additional terms to address this issue that are derived via a spectral filtering procedure. We show explicitly how the global L^2 convergence rate (in space and angle) of the spectral method to the solution of the transport equation depends on the smoothness of the solution (in angle only) and on the order of the filter. The results are confirmed by numerical experiments. Numerical tests have been implemented in MATLAB and are available online.

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

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

Kerstin Kuepper
RWTH Aachen University
Center for Computational Engineering Science
kuepper@mathcces.rwth-aachen.de

PP105

A High Order, Implicit, Hybrid Solver for Linear

Kinetic Equations

Recently, an implicit, hybrid solver for linear kinetic equations has been investigated using a backward Euler solver in time. We present a high order time discretization based on deferred corrections with respect to the backward Euler method. Also, several numerical results are given to show the efficacy of the hybrid method for various orders of accuracy in time as well as streaming versus highly collisional regimes.

Michael Crockat
Michigan State University
crockat1@msu.edu

Charles K. Garrett, Cory Hauck
Oak Ridge National Laboratory
garrettck@ornl.gov, hauckc@ornl.gov

PP105

A High Order Time Splitting Method Based on Integral Deferred Correction for Semi-Lagrangian Vlasov Simulations

The dimensional splitting semi-Lagrangian methods with different reconstruction/interpolation strategies have been applied to kinetic simulations in various settings. However, the temporal error is dominated by the splitting error. In order to have numerical algorithms that are high order both in space and in time, we propose to use the integral deferred correction (IDC) framework to reduce the splitting error. The proposed algorithm is applied to the Vlasov-Poisson system, the guiding center model and incompressible flows.We show numerically that the IDC procedure can automatically increase the order of accuracy in time.

Wei Guo
Michigan State University
wguo126@gmail.com

PP105

Analysis of Discontinuous Galerkin Algorithms for Diffusion and for Energy-Conserving Hamiltonian Dynamics

We present mixed continuous/discontinuous Galerkin schemes for solution of a class of kinetic Vlasov-Boltzmann problems in Hamiltonian Poisson-bracket form (plus collisions). These schemes conserve energy, and optionally the L_2 norm, exactly. Application to electromagnetic gyrokinetic problems requires novel extension to avoid the Ampere cancellation problem and significant time step limitations. There are subtle properties of commonly used DG schemes for second order derivatives, such as from diffusion, which we compare with a recovery-based approach.

Greg Hammett
Princeton Plasma Physics Laboratory
Princeton University
hammett@princeton.edu

Ammar Hakim
Princeton Plasma Physics Laboratory
ahakim@pppl.gov

Eric Shi, Ian Abel, Tim Stoltzfus-Dueck
Princeton University
eshi@pppl.gov, iabel@princeton.edu, tstoltzf@princeton.edu

PP105

Massively Parallel Calculations of Neutronics Experiments Using Pdt

We demonstrate efficient parallel solution algorithms for high-fidelity neutron transport calculations and thermal radiative-transfer calculations. Calculations employ spatial grids that are structured at a coarse level but unstructured at a fine level to accurately represent geometries of simulated experiments. Calculations employ optimal parallel transport sweeps and physics-based preconditioners.

Marvin L. Adams, Aaron Holzaepfel, W. Daryl Hawkins, Michael Adams, Anthony Barbu, Timmie Smith
Texas A&M University
mladams@tamu.edu, aholzaepfel@neo.tamu.edu, dhawkins@tamu.edu, i.am.michael.adams@gmail.com, apbarbu@gmail.com, timmie@tamu.edu

PP105

Numerical Solution of the Boltzmann Equation Using Quadrature-Based Projection Methods

The lack of hyperbolicity has posed many problems for existing moment models based on the Boltzmann equation. We derive globally hyperbolic and rotationally invariant equations using the Quadrature-Based Moment Method. The method is explained by a newly developed framework of projection operators. In order to investigate the approximation quality of the emerging non-conservative equations, we apply dedicated numerical methods on unstructured two-dimensional quad grids and the results are compared to reference solutions.

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

PP105

Positive Filtered P_N Closures for Linear Kinetic Transport Equations, with some Convergence Results

We propose a modification to the standard spherical harmonic closure, known as P_N closure, for kinetic equations. The modification produces smooth, nonnegative polynomial ansatzes by applying two-step filtering on the oscillatory, partially negative P_N solutions, and integrates the closed moment system with filtered moments. We formulate the filtering process as a quadratic program, prove convergence of the proposed closure when the moment order goes to infinity, and report numerical results on the line source benchmark.

Ming Tse P. Laiu
Department of Electrical and Computer Engineering
University of Maryland - College Park
mtlaiu@umd.edu

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

Dianne P. O'Leary
University of Maryland, College Park
Department of Computer Science
oleary@cs.umd.edu

André Tits
University of Maryland at College Park
andre@umd.edu

PP105

Implicit, Filtered Pn Methods for Radiation Transport

The solution of thermal radiation transport as part of radiation-hydrodynamics calculations is important in the simulation of astrophysical phenomenon as well as high-energy density physics applications such as inertial confinement fusion. In this work we present an implicit method for solving the spherical harmonics (PN) equations of radiation transport using filtered expansions. Since the introduction of such filtered PN methods by McClarren and Hauck, these approaches have been successful in producing high fidelity solutions to difficult transport problems. . In this poster we present results of implicit filtered PN radiation transport simulations and discuss preconditioning strategies as well as the effect of implicit time integration on the necessary filter strength. We compare the results to reference Monte-Carlo calculations for several standard test problems, including radiation transport in a laser-driven shock tube experiment.

Ryan G. McClarren
Texas A & M
rmcclarren@ne.tamu.edu

Vincent Laboure
Texas A&M University
vincent.laboure@neo.tamu.edu

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

PP105

Massively Parallel Nuclear Reactor Analysis Using Pdt

We demonstrate efficient parallel solution algorithms for high-fidelity neutron and gamma transport calculations in nuclear reactors. Calculations employ spatial grids that are structured at a coarse level but unstructured at a fine level to accurately represent fuel and structural geometries. Calculations employ optimal parallel transport sweeps and sophisticated physics-based preconditioners.

Marvin L. Adams, Carolyn McGraw, W. Daryl Hawkins, Michael Adams, Timmie Smith
Texas A&M University
mladams@tamu.edu, car3262@tamu.edu, dhawkins@tamu.edu, i.am.michael.adams@gmail.com, timmie@tamu.edu

PP105

High Order Asymptotic Preserving Nodal Discontinuous Galerkin Imex Schemes for the Bgk Equation

We develop high-order asymptotic preserving schemes for the BGK equation in a hyperbolic scaling. Our approaches are based on the micro-macro formulation of the kinetic equation which involves a natural decomposition of the problem to the equilibrium and the non-equilibrium parts. The new ingredients for the proposed methods to achieve high order accuracy are the following: we introduce discontinuous Galerkin discretization of arbitrary order of accuracy with nodal Lagrangian basis functions in space; we employ a high order globally stiffly accurate implicit-explicit Runge-Kutta scheme as time discretization. It is demonstrated that the proposed scheme becomes a local DG discretization with an explicit RK method for the macroscopic compressible Navier-Stokes equations, a method in a similar spirit to the ones in [Bassi & Rabey 1997, Cockburn & Shu 1998]. Numerical results are presented for a wide range of Knudsen number to illustrate the effectiveness and high order accuracy.

Jingmei Qiu
University of Houston
jingqiu@math.uh.edu

Juhi Jang
Math department
UC Riverside
juhi.jang@ucr.edu

Fengyan Li
Rensselaer Polytechnic Institute
lif@rpi.edu

Tao Xiong
Math Department
University of Houston
txiong@math.uh.edu

PP105

Towards Hyperbolic Moment Approximations of Multicomponent Plasmas

Travelling shock waves through initially neutral gases can initiate ionisation reactions leading to a multicomponent plasma. The description of such shock waves poses a great modelling challenge since not only collisional processes including reactive collisions, but also the interactions of the gas with the macroscopic force fields have to be taken into account. In the most general case a kinetic description of the plasma would require the numerical solution of the Boltzmann equation for each species coupled to the Maxwell equations. We consider hyperbolic moment approximations of the Boltzmann equation, which promise an efficient description of the plasma components for Knudsen numbers in the continuum and transition regime.

Roman P. Schärer, Manuel Torrilhon
RWTH Aachen University
schaerer@mathcces.rwth-aachen.de, mt@mathcces.rwth-aachen.de

PP105

Ifp: An Optimal, Fully Conservative, Fully Implicit, Vlasov-Fokker-Planck Solver: Poster

We introduce a new, exactly conservative (mass, momentum, and energy), and fully nonlinearly implicit solver for a multi-species 1D-2V Vlasov-Rosenbluth-Fokker-Planck system. The new solver optimizes mesh resolution requirements by 1) adapting the velocity-space mesh based on the species' local thermal-velocity, and 2) treating the cross-

species collisions exhibiting disparate thermal velocities by an asymptotic formulation. We demonstrate the efficiency and accuracy properties of the approach with challenging numerical examples.

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

PP105

Residual Monte Carlo Methods Within the Moment-Based Acceleration Framework

Over the past several years, Moment-Based Acceleration, or High-Order/Low-Order methods, have been used to solve kinetic equations efficiently and accurately. Recent work has adapted these methods to include Monte Carlo transport solvers. These Monte Carlo simulations often contain a significant level of stochastic noise which can corrupt the low-order solver, resulting in a breakdown of the method or inaccurate results. In this poster we demonstrate that a residual form of the Monte Carlo method can be used, producing significantly more accurate results. We display results for both thermal radiation transport and neutronics applications.

Jeffrey A. Willert
Los Alamos National Laboratory
jaw@lanl.gov

PP105

Asymptotic Preserving Discontinuous Galerkin Method for the Radiative Transfer Equation

Many kinetic equations converge to macroscopic models, known as the asymptotic limit of the kinetic equations, when ϵ (the ratio of mean free path over macroscopic size) $\to 0$. Asymptotic preserving (AP) methods are designed to preserve the asymptotic limits from microscopic to macroscopic models in the discrete setting. In this poster, we consider the radiative transfer equation which models isotropic particle scattering in a medium. Discontinuous Galerkin (DG) (with upwind flux and at least piecewise linear polynomial) methods are known to have such property, but require much more degree of freedom, especially in high dimensional applications. We will present an AP mixed DG-FV method which has comparable computational cost and memory as FV method. Rigorous analysis will be provided to show that the proposed methods are consistent with the limit equation in the asymptotic regimes. Some numerical examples are presented to demonstrate the performance of our methods.

Yulong Xing
Department of Mathematics
Univeristy of Tennessee / Oak Ridge National Lab
xingy@math.utk.edu

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

PP106

Architecture Portable Assembly for Maxwell's Equations

This poster presents an architecture portable implementation of Maxwell's equations. To achieve portability the Kokkos library is employed to handle memory layout and parallelization over a large number of threads in the assembly kernel. The PDE is discretized using edge-basis functions for the eletric field and face-basis functions for magnetics. We also explore the memory to computation tradeoffs of using these vector basis functions. Results illustrating the thread scalability of our approach will be presented.

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Irina Demeshko
Sandia National Laboratories
ipdemes@sandia.gov

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Matthew Bettencourt
Sandia National Laboratories
mbetten@sandia.gov

PP106

Towards Exascale Implementation of the Finite Element Based Application Development Environment

It is not clear today which architecture Exascale supercomputers will have, but we believe that multicore-CPUs and manycore-accelerator devices, that promise improvements in both runtime performance and energy consumption, are major candidates to be used in the future HPC systems. Nowadays, the Finite Element Method is actively used in diverse variety of scientific and industry simulations, that require High Performance Computing (HPC) technology. In this poster we present our strategy for adapting Finite Element-based codes for future architecture, based on using a new Phalanx-Kokkoslocal field evaluation kernel from Trilinos, specifically designed for general partial differential equation solvers and portable acrossdiverse multicore/ manycore architectures.We present performance evaluation results for our strategy on the example of Finite Element Assembly in Greenland Ice-Sheet simulations code.

Irina Demeshko, H. Carter Edwards, Michael Heroux, ROGER P. Pawkowski
Sandia National Laboratories
ipdemes@sandia.gov, hcedwar@sandia.gov, maherou@sandia.gov, rppawlo@sandia.gov

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

Andrew Salinger
CSRI
Sandia National Labs

agsalin@sandia.gov

PP106

Multicore Finite Element Assembly Via Scans

Minisymposterium placeholder abstract to follow

Robert C. Kirby
Baylor University
Robert_Kirby@baylor.edu

PP106

Operator Transformation and Code Generation for Scientific Computing

The present poster discusses three software components for scientific computing. 'Pymbolic' is an expression tree with traversal and rewriting capabilities. Both its mathematical vocabulary and its traversal operations are easily extended for use in mathematical abstractions. 'PyOpenCL' is a toolkit to facilitate code generation, access to massively parallel hardware, and basic parallel algorithms. 'Loopy' is a code generator targeting shared-memory, massively parallel machines, that cleanly separates task specification and transformation-based optimization.

Andreas Kloeckner
Department of Computer Science
University of Illinois at Urbana-Champaign
andreask@illinois.edu

PP106

OCCA: A Unified Approach to Multi-Threading Languages

The inability to predict lasting languages and architectures led us to develop OCCA, a library focused on host-device interactions. Using run-time compilation and macro expansions, the result is a novel single kernel language that expands to multiple threading languages. Currently, OCCA supports device kernel expansions for the Pthreads, OpenMP, OpenCL, CUDA and COI platforms. OCCA can be used through the front-ends provided for C++, C, Matlab, Python, Julia and C#.

David Medina
Rice University
dmedina@rice.edu

Tim Warburton
Rice University
Department of Computational and Applied
timwar@rice.edu

Amik St-Cyr
Computation and Modeling,
Shell International E&P, Inc.
amik.st-cyr@shell.com

PP106

Assembly Algorithms for Pdes with Uncertain Input Data on Emerging Multicore Architectures

In this work we explore assembly algorithms for PDEs discretized with stochastic Galerkin and sampling-based uncertainty quantification algorithms. We demonstrate that rearrangements of these classical uncertainty propagation algorithms lead to improved PDE assembly performance by amortizing communication latency, improving memory access patterns, and exposing new dimensions of fine-grained parallelism. We also describe simple approaches for incorporating these ideas in complex simulation codes.

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

H. Carter Edwards
Sandia National Laboratories
hcedwar@sandia.gov

PP106

Using Multicore Parallelism for Common Finite Element Operations

Whereas finite element assembly with MPI is well understood, assembly using a multithreading model has been less explored. This later model becomes more important as the number of cores on a chip, and thus in a node, increases. Using graph coloring and MapReduce on cells or sets of cells, not only good scalability but bitwise reproducibility can be achieved for matrix assembly and other common operations such as post-processing, estimating discretization errors, etc.

Bruno Turcksin
Texas A&M University
turcksin@math.tamu.edu

Martin Kronbichler
Technische Universitat Munchen
kronbichler@lnm.mw.tum.de

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

PP201

Forward Backward Doubly Stochastic Differential Equations and Applications to The Optimal Filtering Problem

We consider the filtering problem where a signal process is modeled by an SDE and the observation is perturbed by a white noise. The goal is to find the best estimation of the signal process based on the observation. Kalman Filter, Particle Filter, Zakai Filter are some well-known approaches to solve optimal filter problems. In this effort, we shall show the optimal filter problem can also be solved using forward backward doubly stochastic differential equations.

Feng Bao
Oak Ridge National Laboratory
baof@ornlgov

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

Clayton G. Webster, Guannan Zhang
Oak Ridge National Laboratory

webstercg@ornl.gov, zhangg@ornl.gov

PP201

Embedded Sampling-Based Uncertainty Quantification Approaches for Emerging Computer Architectures

In this work we explore embedded sampling-based uncertainty quantification approaches geared towards emerging computer architectures. These approaches improve performance by leveraging reuse of simulation information between samples, improve memory access patterns, and expose new dimensions of fine-grained parallelism. We investigate the applicability of these ideas to relevant partial differential equations with uncertain input data.

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

Marta D'Elia, H. Carter Edwards
Sandia National Laboratories
mdelia@sandia.gov, hcedwar@sandia.gov

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

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

PP201

A Unified Framework for Uncertainty and Sensitivity Analysis of Computational Models with Many Input Parameters

Computational models have found wide applications in simulating physical systems. Uncertainties in input parameters of the system can greatly influence the outputs, which are studied by Uncertainty Analysis (UA) and Sensitivity Analysis (SA). As the system becomes more complex, the number of input parameters can be large and existing methods for UA and SA are computationally intensive or prohibitive. We propose a unified framework by using a hierarchical variable selection approach to connect UA and SA with one design. By incorporating the effect hierarchy principle and the effect heredity principle, the approach works well especially when the number of input parameters is large. Since the whole procedure requires only one design, it is economical in run size and computationally efficient.

Li Gu, C. F. Jeff Wu
Georgia Institute of Technology
lgu@gatech.edu, jeffwu@isye.gatech.edu

PP201

Florida State University Efforts Withing the Equinox Project

We present the results of projects funded by the EQUINOX grant from the Department of Energy Advanced Simulation Computing Research program of the US Department of Energy. Included in the presentation are extensions of $1/f^{\alpha}$ noise to multi-dimensions, the construction of total degree interpolation points that minimize the Lebesgue constant, and the treatment of uncertainty quantification for nonlocal diffusion equations, including fractional Laplacian models.

Max Gunzburger
Florida State University
School for Computational Sciences
mgunzburger@fsu.edu

PP201

A Mathematical Environment for Quantifying Uncertainty: Integrated and Optimized at the EXtreme Scale (equinox)

This poster describes a modern mathematical and statistical foundation that will enable next-generation, complex, stochastic predictive simulations. This is a collaborative effort funded by the Advanced Scientific Computing Research at the US Department of Energy, focused on combing novel paradigms in applied mathematics, statistics and computational science into a unified framework, which we call an Environment for Quantifying Uncertainty: Integrated aNd Optimized at the eXtremescale (EQUINOX). This rigorous UQ methodology includes adaptive and quasi-optimal approximations, hierarchical and multilevel methods, architecture aware UQ paradigms, and robust experimental design strategies.

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

PP201

Hierarchical Acceleration of Multilevel Methods for Pdes with Random Input Data

Multilevel methods seek to decrease computational complexity by balancing spatial and stochastic discretization errors. As a form of variance reduction, multilevel Monte Carlo (MLMC) and multilevel stochastic collocation have been developed. We present an approach to adaptively accelerate a sequence of hierarchical interpolants required by a multilevel method. Taking advantage of the hierarchical structure, we build new iterates and improved preconditioners, at each level, by using the interpolant from the previous level. We provide rigorous complexity analysis of the fully discrete problem and demonstrate the increased computational efficiency.

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

PP202

Compatible Discrete Operator Schemes for Advection-Diffusion Equations

Compatible Discrete Operator (CDO) schemes belong to the class of mimetic (or structure-preserving) schemes and can be deployed on polyhedral meshes. We extend the analysis of CDO schemes, introduced by J. Bonelle and A. Ern for elliptic and Stokes equations, to advection-diffusion equations. The novelties are a discrete contraction operator for the advective derivative designed using ideas from Friedrichs' systems and boundary Hodge operators to weakly enforce Dirichlet conditions. We prove stability and error bounds for the discrete problems. Two salient aspects are Péclet-robust error estimates and the treatment

of divergence-free advection fields in the absence of reactive dissipation. Finally, we present numerical results on three-dimensional polyhedral meshes.

Pierre Cantin
Universite Paris-Est, Ecole des Ponts, Paris Tech,
CERMICS,
77455 Marne La Vallee, Cedex 2, France
pircantin@gmail.com

Alexandre Ern
Universite Paris-Est
CERMICS, Ecole des Ponts
ern@cermics.enpc.fr

Jérome Bonelle
EDF R&D
jerome.bonelle@edf.fr

PP202

Upwinding in the Mimetic Finite Difference Method for Richards' Equation

In porous media applications, harmonic averaging of material properties is the intrinsic property of FV, mimetic FD, and mixed FE methods reflecting continuity of the Darcy flux. In some cases, as as infiltration in a dry soil, it may lead to a strong nonlinear stability condition. To relax this condition, upwinding of a relative permeability is often used. To preserve fundamental properties of the mimetic methods such as symmetry and positive definiteness of discrete operators, a special upwind technique is required that we present and analyze in this poster.

Konstantin Lipnikov
Los Alamos National Laboratory
lipnikov@lanl.gov

PP202

The Virtual Element Method

We present a generalization of the finite element methods to unstructured polygonal and polyhedral meshes. The development is based on a new paradigm, dubbed virtual element method, or VEM. In this new framework only a part of the approximation space is constructed explicitly, while the remaining (virtual) part of the space is given only in terms of some algebraic conditions. This new family of numerical methods is suitable to polygonal and polyhedral unstructured meshes and is applied to the numerical resolution of diffusion, reaction-diffusion and convection-diffusion problems, the Stokes equations and compressible and incompressible elasticity equations.

Gianmarco Manzini
Los Alamos National Laboratory
gm.manzini@gmail.com

PP203

Optimal Control for Mass Conservative Level Set Methods

We present a conservative level set method for numerical simulation of evolving interfaces. A PDE-constrained optimization problem is formulated and solved in an iterative fashion. The proposed optimal control procedure constrains the level set function to satisfy a conservation law for the corresponding Heaviside function. The control corrects the convective flux in the nonlinear state equation so as to enforce mass conservation while minimizing deviations from the target state. The methodology is evaluated numerically.

Christopher Basting, Dmitri Kuzmin
Dortmund University of Technology
christopher.basting@math.tu-dortmund.de,
kuzmin@math.uni-dortmund.de

PP203

A New Partitioned Algorithm for Explicit Elastodynamics Based on Variational Flux Recovery

We present a new partitioned algorithm for explicit elastodynamics, based on variational flux recovery. This method works by the exchange of tractions which depend on both the known state and the unknown state at the future time step. The exchange of recovered tractions between the bodies formally involves modification of both the forcing term and the mass matrix on each subdomain. Reliance on the future time step distinguishes our approach from traditional partitioned solution algorithms, which use interface conditions between subdomains and the current time step solution to define boundary conditions for the models. Our approach offers some distinct numerical and theoretical advantages. It passes a linear patch test and is second order accurate in space. Furthermore, if interface grids match, our new partitioned method recovers the solution of a monolithic coupling scheme for the solid-solid interaction problem.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Paul Kuberry
Sandia National Laboratories
pakuber@sandia.gov

PP203

Analysis of a Fluid-Structure Interaction Problem Decoupled by Optimal Control

A new method is presented which allows fluid-structure interaction problems to be stably decoupled. This method permits the use of existing solvers for the fluid and structure subsystems. With existence of an optimal solution and Lagrange multipliers proven, the optimization problem constrained by PDEs can be written unconstrained using the Lagrange multiplier rule. The first order optimality system is derived and and a gradient based algorithm is given along with numerical results.

Paul A. Kuberry
Clemson University
pkuberr@clemson.edu

Hyesuk Lee
Clemson University
Dep. of Mathematical Sciences
hklee@clemson.edu

PP203

Feature-Preserving Finite Element Transport

Across Interfaces: Part 2, Direct Flux Recovery

We present an optimization-driven approach for coupling transport codes across non-coincident interfaces. Our approach relies on the recently introduced methods for optimization-based finite element transport (see Part 1). It uses constrained interpolation to recover and exchange flux variables on the interfaces, and maintains key physical properties, such as mass conservation and monotonicity.

Kara Peterson
Sandia Natl. Labs
kjpeter@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

PP203

Feature-Preserving Finite Element Transport Across Interfaces: Part 1, Optimization-Based Transport

We discuss an optimization framework for the design of robust feature-preserving schemes for finite element transport. Our optimization models and algorithms preserve key physical features such as monotonicity and mass conservation, and can be used to efficiently couple simulation codes for transport across non-coincident interfaces (see Part 2).

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

Kara Peterson
Sandia Natl. Labs
kjpeter@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

PP203

Higher Order Finite Element Methods for Interface Problems

We present a higher-order finite element method for solving a class of interface problems in two dimensions. The method is based on correction terms added to the right-hand side of the natural method. We prove optimal error estimates of the method on general quasi-uniform meshes in the maximum norms. In addition, we apply the method to a Stokes interface problem obtaining optimal result.

Manuel A. Sanchez-Uribe
Brown University
manuel_sanchez_uribe@brown.edu

PP203

Extended and Conformal Decomposition Finite Elements for 3D Compatible Discretizations

Current tends to concentrate near material interfaces, ergo accurate simulation of electromagnetics requires the resolution of said interfaces. Large deformations make body-fit meshes impractical, but the continuity requirements of physics-compatible discretizations are not satisfied without body-fitting. We propose using the existing non-conforming node and edge element bases to dynamically construct an interface-conforming basis in 3D. We also develop a conformal decomposition finite element method for edges and demonstrate their equivalence for tetrahedral elements.

Christopher Siefert, Richard Kramer
Sandia National Laboratories
csiefer@sandia.gov, rmkrame@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Thomas Voth
Sandia National Laboratory
tevoth@sandia.gov

PP204

Chebfun

Chebfun is an open-source MATLAB package for computing with functions instead of numbers. By overloading common MATLAB commands to operate on function objects called chebfuns, the package provides a computing experience which feels symbolic but is actually driven by the numerical tools of Chebyshev polynomial approximation technology. Users can do things like plot and evaluate functions, compute roots and extrema, and solve differential equations using a simple syntax intuitive to anyone familiar with MATLAB. This poster will provide a tour of some of Chebfun's major features. A member of the Chebfun development team will be on site to present demos and discuss the software with anyone interested.

Anthony Austin
University of Oxford
austin@maths.ox.ac.uk

PP204

Feel++: A Versatile High Performance Finite Element Embedded Library into C++

Feel++ (Finite Element method Embedded Language in C++) offers a domain specific language to partial differential equations embedded in C++. It allows to use a very wide range of possibly, high order, Galerkin methods, and advanced numerical methods such as domain decomposition methods including h-p mortar methods, levelset methods, fictitious domain methods or certified reduced basis. We illustrate on some applications the numerical behavor up to millions and even billions of unknowns.

Vincent Chabannes
Laboratoire Jean Kuntzmann
Université de Grenoble

vincent.chabannes@feelpp.org

PP204

Dolfin-Adjoint

dolfin-adjoint automatically derives parallel, efficient adjoint and tangent linear models from finite-element models written in the FEniCS environment. It also provides high-level tools to solve PDE-constrained optimisation and generalised stability problems. This poster presents an overview of dolfin-adjoint, including examples and recent developments. The dolfin-adjoint developers will live-demo the software and answer questions.

Simon W. Funke
Center for Biomedical Computing
Simula Research Laboratory
simon@simula.no

Marie E. Rognes
Simula Research Laboratory
meg@simula.no

Patrick E. Farrell
Department of Mathematics
University of Oxford
patrick.farrell@maths.ox.ac.uk

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

PP204

Building Performance Transportable Codes for Extreme Scale

The goal of the Center for Exascale Simulation of Plasma-Coupled Combustion is to explore and understand a new approach to controlling combustion though the use of plasmas. The computation is multiscale and multiphysics, and is a challenge even with extreme scale computers. This poster highlights the approach to transportable performance being being taken by the center, including the use of a golden copy for development and source transformations for data structure and loop optimization.

William D. Gropp
University of Illinois at Urbana-Champaign
Dept of Computer Science
wgropp@illinois.edu

PP204

Firedrake: Automating Finite Element by Composing Abstractions

Firedrake automates the portable solution of partial differential equations using the finite element method. Firedrake takes separation of concerns in automated FEM to a new level. In addition to the Unified Form Language from the FEniCS project, and PETSc's linear algebra abstraction, Firedrake introduces the PyOP2 abstraction for mesh iteration and the COFFEE abstraction for kernel vectorisation and optimisation. The result is faster, more flexible and more capable automated simulation.

David Ham, Florian Rathgeber
Imperial College London
david.ham@imperial.ac.uk, f.rathgeber10@imperial.ac.uk

Lawrence Mitchell
Department of Computing
Imperial College London
lawrence.mitchell@imperial.ac.uk

Michael Lange
Department of Earth Science and Engineering
Imperial College London
michael.lange@imperial.ac.uk

Andrew McRae
Department of Mathematics
Imperial College London
a.mcrae12@imperial.ac.uk

Gheorghe-Teodor Bercea, Fabio Luporini
Department of Computing
Imperial College London
gheorghe-teodor.bercea08@imperial.ac.uk,
f.luporini12@imperial.ac.uk

Paul Kelly
Imperial College London
p.kelly@imperial.ac.uk

PP204

The DEAL.II Finite Element Library

deal.II is a C++ software library supporting the creation of finite element codes and an open community of users and developers. In this poster session, we present an overview of deal.II, highlight some applications, and present future plans.

Timo Heister
Clemson University
Mathematical Sciences
heister@clemson.edu

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

Guido Kanschat, Matthias Maier
Universität Heidelberg, Germany
kanschat@uni-heidelberg.de, matthias.maier@iwr.uni-heidelberg.de

PP204

An Overview of the Trilinos Project

Trilinos is a large collection of interoperable packages for formulation, solution and analysis of large scale modeling and simulation problems. Trilinos provides libraries for (i) portable architecture-aware data containers and algorithms; (ii) geometry, meshing, discretization, load balancing and scalable solution of linear, nonlinear and transient problems; (iii) optimization, uncertainty quantification and analysis and (iv) scalable IO. Trilinos also provides software tools and processes for developing scientific software from inception to post-delivery maintenance, supporting components-based application development strategies intended to improve software quality, and reduce time to completion and cost.

Michael Heroux
Sandia National Laboratories

maherou@sandia.gov

PP204

FEniCS High Performance Computing with Applications in Aerodynamics, Environmental Science and Biomedicine

Adaptive finite element methods (AFEM) using unstructured meshes pose serious challenges for the development of efficient algorithms and software implementations for massively parallel architectures. We describe the open source software DOLFIN-- HPC/Unicorn that implements AFEM for a general class of problems in computational mechanics, as part of the software project FEniCS [1,2]. DOLFIN--HPC is based on a hybrid MPI+OpenMP/MPI+PGAS programming model, and shows near optimal scaling up to tens of thousands of processing elements for applications in computational fluid dynamics (CFD) [3]. The method is based on the computation of an adjoint (or dual) problem to derive a posteriori estimates of the error in a certain output of interest, such as the drag or lift of an airplane, which guides the mesh refinement algorithms. Since the problem is nonlinear, the primal solution appears as data in the adjoint problem, and since the adjoint problem is solved backwards in time, the time series of the primal problem must be stored which is a challenge for large problems. In addition, algorithms for local mesh refinement pose challenges in terms of dynamic load balancing and efficient communication over the unstructured mesh data. In this poster presentation we describe the methods and software implementation, and highlight application projects based on the computational technology, including simulation of the aerodynamics of airplanes, blood flow in the human heart, and human phonation.

Johan Hoffman
Royal Institute of Technology KTH
jhoffman@kth.se

PP204

EMatter: A Materials Simulation Framework As a Service

eMatter is a fledging new computational service that aims to simplify life for computational materials scientists. Built on top of a stack of scalable high-performance libraries, including PETSc, libMesh and MOOSE, eMatter presents their capabilities as a service, or as a platform. Library maintenance, configuration and building, as well as job submission, data handling and basic visualization capabilities are taken care of by the framework. eMatter demos will be provided and developers will be available on site for questions and hands-on exercises.

Dmitry A. Karpeyev
University of Chicago
karpeev@uchicago.edu

PP204

Composability in Petsc

Composability of lower-level abstractions is crucial for building large software systems while controlling complexity. This fact is often overlooked in computational science since it is not emphasized in the numerical analysis. We present two examples of its role in PETSc. In the design of nonlinear solvers, we allow preconditioning of one solve by another, in analogy with the linear case, giving us much richer and more robust alternatives for nonlinear solves. Likewise, for small scale parallelism, we propose a system that allows different threading models to interact, avoiding resource contention and over-provisioning. We attempt to minimize assumptions about our environment so as to maximize the ability to compose with other packages and users that chose other models.

Matthew G. Knepley
University of Chicago
knepley@ci.uchicago.edu

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory and CU Boulder
jed@jedbrown.org

PP204

Fenics

FEniCS is a high-level, high-performance software environment for automated, efficient, and intuitive solution of partial differential equations by the finite element method. In this poster session, an overview of FEniCS is presented, together with explained examples and a summary of recently added new features. Live demos will be presented and FEniCS developers will be on spot to chat with users.

Anders Logg
Chalmers University of Technology
logg@chalmers.se

PP204

Sigma: Scalable Interface for Geometry and Mesh Based Applications

SIGMA provides components to define/query geometric entities (CGM), and apply efficient mesh generation algorithms (MeshKit) based on relational definitions (Lasso) to create high quality unstructured mesh databases (MOAB), which serves as a data backplane for applications (Nuclear, CFD, Climate, FEA) running on desktop to petascale architectures. SIGMA tools handle several I/O formats with interfaces for visualization, solvers (PETSc) and scalable conservative solution transfers (enabling multi-physics solvers). We will present interactive workflow demos to encourage further discussions with audience.

Vijay Mahadevan, Iulian Grindeanu, Rajeev Jain
Argonne National Laboratory
mahadevan@anl.gov, iulian@mcs.anl.gov,
jain@mcs.anl.gov

Navamita Ray
Argonne National Labratory
nray@mcs.anl.gov

Danqing Wu
Argonne National Laboratory
wuda@mcs.anl.gov

Paul Wilson
University of Wisconsin - Madison
wilsonp@engr.wisc.edu

PP204

Dune - The Distributed and Unified Numerics En-

vironment

DUNE is a modular open-source framework for the solution of partial differential equations using grid-based methods, written in modern C++. It offers clear abstractions at all levels of a PDE simulation, providing the user with high-level abstractions for productive development, while also allowing access to lower-level functionality, enabling scalability from notebooks to HPC applications. Due to its modularity, DUNE easily integrates with legacy codes like existing grid managers and LA libraries.

Steffen Müthing
Heidelberg University
steffen.muething@iwr.uni-heidelberg.de

PP204

Elemental

Elemental is a modern C++11 implementation of distributed-memory dense linear algebra that has been extended to various sparse-direct solvers, preconditioners, optimization, and medical imaging applications. In this poster session, an overview of Elemental is presented, with a focus on the idioms that the project has converged upon for manipulating (distributed) matrices and exposing abstract interfaces without incurring performance penalties.

Jack L. Poulson
Computational Science and Engineering
Georgia Institute of Technology
jack.poulson@gmail.com

PP204

Jupyter Widgets: Interactive Computing Through the Browser in Any Programming Language

The Jupyter project, evolved from IPython, provides a generic architecture for interactive computing with a language-independent protocol for controlling code execution over the network. In particular, it can be used from web browsers to provide interactive graphical exploration of complex computations by combining Javascript widgets coupled to computational kernels written in Python, Julia, R or any other language. This facilitates both rapid exploration and illustration of computational concepts for education.

Min Ragan-Kelley
UC Berkeley AS&T
MinRK@Berkeley.edu

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

PP204

Camellia: A Software Framework for Discontinuous Petrov-Galerkin Methods

The discontinuous Petrov-Galerkin (DPG) methodology of Demkowicz and Gopalakrishnan minimizes the solution residual in a user-determinable energy norm and offers a built-in mechanism for evaluating error in the energy norm, among other desirable features. However, the methodology brings with it some additional complexity for researchers who wish to experiment with DPG in their computations. In this poster, we introduce *Camellia*, a software framework whose central design goal is to enable developers to create efficient *hp*-adaptive DPG solvers with minimal effort.

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

PP204

ViennaCL - Fast Linear Algebra for Multi and Many-Core Architectures

An overview of the linear algebra functionality and solvers available in the Vienna Computing Library (ViennaCL) for multi-core CPUs as well as GPUs and Intel's MIC architecture is given. Compute backends using OpenMP, CUDA, and OpenCL allow for maximum flexibility and best performance on the underlying hardware from all major vendors.

Karl Rupp
Institute for Analysis and Scientific Computing, TU Wien
Institute for Microelectronics, TU Wien
rupp@iue.tuwien.ac.at

Philippe Tillet, Toby St Clere Smithe, Namik Karovic, Josef Weinbub, Florian Rudolf
Institute for Microelectronics
Vienna University of Technology
phil.tillet@gmail.com, mail@tsmithe.net, namik.karovic@gmail.com, weinbub@iue.tuwien.ac.at, rudolf@iue.tuwien.ac.at

PP205

Radical Optimization Techniques for Asynchronous Iterative Algorithms on Gpus

Asynchronous algorithms, with their ability to tolerate memory latency, form an important class of algorithms for modern computer architectures. For the recently proposed asynchronous iterative algorithm for computing incomplete factorizations we present non-traditional optimizations to leverage the computing power of GPUs. These include controlling the order in which variables are updated, taking advantage of cache reuse between thread blocks, and optimizing the trade-off between parallelism and algorithm convergence for time-to-solution performance.

Hartwig Anzt
Innovate Computing Lab
University of Tennessee
hanzt@icl.utk.edu

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

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@cs.utk.edu

PP205

Experiences in Autotuning Linear Algebra Operations for Energy Minimization on Gpus

Approaching the Exascale computing era requires a

paradigm shift from pure runtime performance to metrics accounting also for resource efficiency. In this line, we present experiences about using the BEAST autotuning framework for optimizing the energy efficiency of GPU implementations of basic linear algebra building blocks. We analyze the kernel metrics to correlate the pruning parameters to performance and resource usage.

Hartwig Anzt
Innovate Computing Lab
University of Tennessee
hanzt@icl.utk.edu

Blake Haugen
Innovative Computing Lab
University of Tennessee
bhaugen@utk.edu

Jakub Kurzak
University of Tennessee, Knoxville
kurzak@icl.utk.edu

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@cs.utk.edu

PP205

CUSP: A Parallel Sparse Matrix Package for Gpus

CUSP, an open-source C++ templated library for sparse matrix and graph operations, enables collection oriented parallel processing on modern architec- tures, such as GPUs. In this presentation we will discuss the Thrust-based ab- stract programming model to express highly irregular computations involving sparse matrices, such as addition and multiplication, and the integration of GPU oriented programming systems, such as CUB, to implement high-performance routines to perform reordering, coloring, and maximum-ow computations on GPU architectures.

Steven Dalton
University of Illinois at Urbana-Champaign
dalton6@illinois.edu

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

PP205

Sparse Matrix-Matrix Multiplication on High-Throughput Architectures

Many computations in the physical and data sciences rely on sparse matrix operations such as the sparse matrix-matrix multiplication (SpMM) operation. Yet an efficient SpMM on high-throughput architectures requires exposing fine-grained parallelism in the operation while limiting the use of the off-chip memory bandwidth. In this poster we highlight an approach that decomposes the SpMM into three phases: expansion, sorting, and contraction, thereby allowing the use of optimized throughput-oriented kernels. The result is a SpMM algorithm that requires low a priori analysis of the matrix structures, yet yields substantial savings for general classes of unstructured matrices.

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

Steven Dalton
University of Illinois at Urbana-Champaign
dalton6@illinois.edu

PP206

Parallel Petascale Modeling of Transportation Accidents Involving High Explosives

Simulations of multiple explosive devices undergoing combustion requires algorithms and models capable of solving fluid-structure interactions, compressible flow, and fast solid-¿gas reactions. We have been developing these tools in the Uintah:MPMICE, which is capable of scaling up to 512k. This research will give insight in to the physical mechanism of Deflagration to Detonation Transition (DDT) for large-scale explosions and help determine safe packaging configurations to eliminate DDT in transportation accidents.

Jacqueline Beckvermit
Department of Chemsitry
University of Utah
beckvermit@gmail.com

Andrew Bezdjian
University of Utah
andrewbezdjian@gmail.com

Todd Harman
Department of Mechanical Engineering
University of Utah
t.harman@utah.edu

John A. Schmidt
SCI Institute
University of Utah
jas@sci.utah.edu

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

Chuck Wight
Weber State University
cwight@weber.edu

PP206

Radiation Modeling Using Reverse Monte Carlo Ray Tracing Within the Uintah Framework

The Uintah ARCHES component has previously used the Discrete Ordinates algorithm for solving the Radiative Transport Equation (RTE). This approach is expensive due to the linear solves involved. We are exploring the use of the Reverse Monte Carlo Ray Tracing (RMCRT) technique for solving the RTE on both CPUs and GPUs to reduce this computational cost. Here we we discuss the challenges involved in moving to an RMCRT approach for solving the RTE at scale.

Alan Humphrey
Scientific Computing and Imaging Institute
University of Utah

ahumphrey@sci.utah.edu

PP206

Using Uintah:mpmice for High Resolution Urban Flow Studies

Urban flow transport plays a critical role in investigations of the impact of local and global climate on urban inhabitants. We introduce Uintah:MPMICE for the simulation of fluid-structure interactions in urban flows. Uintah:MPMICE has been developed in a massively parallel computational infrastructure, uses material points to represent buildings, and LES technique to represent momentum and scalar transport. Using Uintah:MPMICE for highly turbulent flows in urban areas introduce challenges on the implementation of appropriate boundary conditions.

Arash Nemati Hayati, Rob Stoll
Department of Mechanical Engineering University of Utah
a.nematihayati@utah.edu, rstoll2@gmail.com

Todd Harman
Department of Mechanical Engineering
University of Utah
t.harman@utah.edu

Eric Pardyjak
Department of Mechanical Engineering University of Utah
pardyjak@gmail.com

PP206

Wasatch: A CPU/GPU-Ready Multiphysics Code Using a Domain Specific Language

Hybrid computing poses a challenge for computational science. Legacy code is unable to leverage the benefits of CPU/GPU platforms without significant refactor. Here, we present a promising approach to making computational codes architecture-proof. Our method consists of using a Domain Specific Language to separate implementation and usage. At the outset, one can write a single code that may be executed on multiple backends. This is shown to work with our flagship multiphysics code, Wasatch.

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

Abhishek Bagusetty
Chemical Engineering
University of Utah
abhishek.bagusetty@utah.edu

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

PP206

Applied Large Eddy Simulation: Validation and Uncertainty Quantification of Lab and Pilot-Scale, Oxy-Coal Boiler Simulations

This work seeks to advance simulation capabilitiesrelatedto oxy-coal combustion. The approach combines experimental and simulation data from lab and pilot-scale oxy-coal systems with Large Eddy Simulation (LES) in a rigorous validation and uncertainty quantification (V/UQ) Bayesian technique. The method identifies a subset of parameters that control the heat flux, temperature, and species concentrations within the boilers. Given the uncertain ranges in these variables, the V/UQ method discovers a range of simultaneously consistent values constrained by the experimental data with uncertainties. The experimental systems of interest are the 15MW Alstom Boiler Simulation Facility and the 150 KW Oxy-Fired Combustor located at the University of Utah. The simulation work was performedusing Arches, a large-eddy simulation component with DQMOM multiphase capabilities in the Uintah framework, at large scale.

Jeremy Thornock
The University of Utah
jeremy.thornock@utah.edu

Wu Yuxin
Institute for Clean and Secure Energy
University of Utah
wuyx02@gmail.com

Ben Isaac
University of Utah & Université Libre de Bruxelles
bisaac2@gmail.com

Sean Smith
The University of Utah
sean.t.smith@utah.edu

Philip J. Smith
University of Utah
Philip.Smith@utah.edu

PP207

The Periodic Table of the Finite Elements

Since the dawn of the finite element method half a century ago, increasingly sophisticated finite element spaces have been developed and applied. Even to specialists, the resulting collection of finite elements can seem a disorganized zoo of possibilities. By taking the viewpoint of finite element exterior calculus these spaces can be arranged in a sort of periodic table, which explains their shape functions, degrees of freedom, and interrelationships, and guides implementation in advanced software environments.

Douglas N. Arnold
School of Mathematics
University of Minnesota
arnold@umn.edu

Anders Logg
Chalmers University of Technology
logg@chalmers.se

PP207

Convolution-Translation and Bounded Cochain Projections for the Elasticity Complex

Uniformly bounded projections in L^2 which commute with the exterior derivative play an essential role in the finite element exterior calculus. The extension part of the process limits the construction of bounded cochain projections for the elasticity complex to star shaped domains. In this work, we avoid the extension operator by using

convolution-translation. The construction is revisited from that point of view and cochain projections without the restriction of star-shaped domains are constructed for the elasticity complex.

Gerard Awanou
Chicago
awanou@uic.edu

PP207

What is a Good Linear Finite Element... On a Generic Polytope?

The notion of what constitutes a "good" linear finite element on geometries other than simplices and cubes remains largely unexplored. We use harmonic coordinates as a means to investigate this question, arriving at a few key conclusions. On convex polygons, harmonic coordinates in general provide no improvement over standard interpolation on the constrained Delaunay triangulation of the polygon. On non-convex polygons, however, harmonic coordinates can provide optimal interpolation estimates even when all triangulations fail to do so. We also present the extension and implication of these results to finite elements on non-convex polyhedra in 3D.

Andrew Gillette
University of Arizona
Department of Mathematics
agillette@math.arizona.edu

Alexander Rand
CD-adapco
Austin, Texas
alexander.rand@cd-adapco.com

PP207

Stokes Elements on Cubic Meshes Yielding Divergence-Free Approximations

Using a finite element exterior calculus framework, conforming piecewise polynomial spaces with respect to cubic meshes are constructed for the Stokes problem in arbitrary dimensions yielding exactly divergence–free velocity approximations. We first present the construction of the lowest order case, its implementation, and convergence analysis. We then introduce finite element spaces with continuous pressure approximations leading to a system of less unknowns. Finally, numerical experiments are shown verifying the theoretical results.

Michael J. Neilan
University of Pittsburgh
Department of Mathematics
neilan@pitt.edu

Duygu Sap
University of Pittsburgh
duygusap@googlemail.com

PP207

Weak Galerkin Finite Element Methods

Weak Galerkin finite element method (WG-FEM)is a new approximating technique for partial differential equations. The main idea behind the WG method is the use of discrete weak differential operators in the conventional variational formulations for the corresponding PDEs. This presentation shall introduce the WG-FEM through some model problems: (1) mixed formulation for second order elliptic equations, (2) the Stokes equations, and (3) div-curl problems or Maxwell equations.

Junping Wang
National Science Foundation
jwang@nsf.gov

Chunmei Wang
Georgia Institute of Technology
cwang462@math.gatech.edu

Xiu Ye
University of Arkansas, Little Rock
xxye@ualr.edu

PP208

Approximate Active Bayesian Inference of Nonlinear Dynamical Systems

This poster introduces a computationally efficient approximation method to design experiments for Bayesian model selection of nonlinear dynamical systems. The methods employs surrogate functions to reconstruct the total ordering of the candidate experiments induced by the mutual information between the data and the models. Under mild assumptions, it can be shown that such approximation achieves almost indistinguishable design results with a significant computational speedup.

AlbertoGiovanni Busetto
UCSB
agb@ece.ucsb.edu

PP208

Variational Reformulation of Bayesian Inverse Problems

Despite of its computational appeal, the classical approach to inverse problems suffers from many shortcomings. The Bayesian formalism to inverse problems avoids most of these difficulties, albeit at an increased computational cost. In this work, we use information theoretic arguments in order to cast the Bayesian inference problem in terms of an optimization problem. The resulting scheme combines the theoretical soundness of fully Bayesian inference with the computational efficiency of a simple optimization.

Ilias Bilionis
Purdue University
ebilionis@gmail.com

Panagiotis Tsilifis
University of Southern California
pantsili@gmail.com

Nicholas Zabaras
Cornell University
nzabaras@gmail.com

PP208

Bayesian Model Selection for Exploring Mechanisms Contributing to Differential Signaling

An important phenomenon in cytokine signaling, termed "differential signaling," is the ability of structurally similar ligands with different binding affinities to elicit diverse

biological actions through the same receptor. We model signaling networks as linear steady state models for which we derive explicit signaling differentials for single cells and cell populations. We explore the capacity of candidate signaling mechanisms to reproduce observed differential signaling dose response data by employing Bayesian model selection.

Pencho Yordanov
ETH Zurich
pencho.yordanov@bsse.ethz.ch

Joerg Stelling
Department of Biosystems Science and Engineering, ETH Zurich
joerg.stelling@bsse.ethz.ch