IP1

Control and Optimization of Subsurface Flow

Controlling the flow of fluids (e.g. water, oil, natural gas or CO2) in subsurface porous media is a technical process with many mathematical challenges. The underlying physics can be described with coupled nearly-elliptic and nearly-hyperbolic nonlinear partial differential equations, which require the aid of large-scale numerical simulation. The strongly heterogeneous nature of subsurface rock leads to strong spatial variations in the coefficients. Moreover, the limited accessibility of the underground leads to very large uncertainties in those coefficients and severely limits the amount of control over the dynamic variables. In this talk I will address related system-theoretical aspects, reduced-order modeling techniques, and adjoint-based optimization methods.

Jan Dirk Jansen
TU Delft, The Netherlands
'j.d.jansen@tudelft.nl'

IP2

Analyzing and Generating BIG Networks

Graphs and networks are used to model interactions in a variety of contexts, and there is a growing need to accurately measure and model large-scale networks. We consider especially the role of triangles, useful for measuring social cohesion. This talk will focus on two topics: (1) Accurately estimating the number of triangles and clustering coefficients for BIG networks, and (2) Generating BIG artificial networks that capture the degree distribution and clustering coefficients of observed data. This is joint work with Ali Pinar, Todd Plantenga, and C. Seshadhri from Sandia National Labs and Christine Task from Purdue University.

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

IP3

Certified Reduced Models and Their Applications

Models of reduced computational complexity are used extensively throughout science and engineering to facilitate modeling of complex systems for control, design, multiscale analysis, uncertainty quantification etc. We shall discuss ongoing efforts to develop reduced methods endowed with rigorous error estimators to certify models, hence endowing it with predictive value. We outline the basic ideas behind certified models and discuss computational efficiency and efficient model construction. The performance of the certified reduced models will be illustrated through several examples and, time permitting, we conclude by discussing some ideas aiming to enable the development of certified reduced models for high-dimensional parametrized problems.

Jan S. Hesthaven
Brown University
Division of Applied Mathematic
Jan.Hesthaven@Brown.edu

IP4

Modeling Cardiac Function and Dysfunction

Simulating cardiac electrophysiological function is one of the most striking examples of a successful integrative multiscale modeling approach applied to a living system directly relevant to human disease. This presentation showcases specific examples of the state-of-the-art in cardiac integrative modeling, including 1) improving ventricular ablation procedure by using MRI reconstructed heart geometry and structure to investigate the reentrant circuits formed in the presence of an infarct scar; 2) developing a new out-of-the box high-frequency defibrillation methodology; 3) understanding the contributions of non-myocytes to cardiac function and dysfunction, and others.

Natalia A. Trayanova
John Hopkins University
Institute for Computational Medicine
ntrayanova@jhu.edu

IP5

Quantum Mechanics Without Wavefunctions

In principle, predictions of the electronic structure of atoms, molecules, and materials requires solving the many-body Schrdinger wave equation (SWE), whose eigenvalues and eigenfunctions delineate the distribution of electrons in energy and space, respectively. In fact, the SWE cannot be solved exactly for more than one electron but excellent approximations are available. However, such methods scale extremely poorly with system size and generally are not applicable to large numbers of atoms or to condensed matter. Alternatively, one can solve directly for the electron density distribution rather than the many-body wavefunction, within orbital-free density functional theory (OFDFT). By so doing, the problem greatly simplifies to 3 degrees of freedom rather 3N (N being the number of electrons). The state of the art of OFDFT, in terms of algorithms, physical models, and applications, will be discussed.

Emily A. Carter
Princeton University
eac@princeton.edu

IP6

Automated Astrophysics in the Big Data Era

Telescope projects are now routinely obtaining massive digital movies of the dynamic nights sky. But given the growing data volumes, coupled with the need to respond to transient events quickly with appropriate followup resources, it is no longer possible for people to be involved in the real-time loop. I discuss the development of a robotic telescopes, autonomous follow-up networks, and a machine-learning framework that act as a scalable, deterministic human surrogate for discovery and classification in astronomical imaging.

Joshua S. Bloom
University of California, Berkeley
jbloom@astro.berkeley.edu

IP7

PDE-Based Simulation Beyond Petascale

We explore fundamental computational complexity considerations that will drive algorithmic design choices for PDE-based simulation codes scaling to petascale and beyond. We argue that high-order methods using implicit or semi-implicit solvers are essential to efficient simulation of multiscale problems. These methods can be realized at per-point-costs equivalent low-order methods. We further show that multilevel solvers having bounded iteration counts can scale to billion-way concurrency. We analyze the scalability of (low- or high-order) domain decomposition approaches to predict parallel performance on exascale architectures. These predictions shed light on what exascale CFD computation might enable and provide insight to design requirements of exascale algorithms, codes, and architectures.

Paul F. Fischer
Argonne National Laboratory
fischer@mcs.anl.gov

IP8

Challenges for Algorithms and Software at Extreme Scale

Extreme scale systems face many challenges. The end of frequency scaling forces the use of extreme amounts of concurrency. Power constraints are forcing a reconsideration of the processor architecture, eliminating features that provide small performance benefit relative to the power consumed and making use of specialized processing elements such as GPUs. Future systems will need to combine these and other approaches to approach Exascale performance. This talk discusses how algorithms and software need to change to make effective use of extreme scale systems.

William D. Gropp
University of Illinois at Urbana-Champaign
Dept of Computer Science
wgropp@illinois.edu

IP9

Consistent Modelling of Interface Conditions for Multi-Physics Applications

In many multi-physics applications, information transport plays an important role for the efficiency and stability of the applied numerical algorithms. The global accuracy is quite often dominated by local effects at the interfaces, and local singularities can pollute the numerical solution. Here we discuss several issues such as hierarchical limiting techniques, local energy corrections and optimal estimates for the flux variables. The coupling is controlled by pairs of balance equations providing a very flexible framework. Different examples illustrate the abstract concepts, special focus is on surface based coupling techniques and highly non-linear dimension-reduced systems.

Barbara Wohlmuth
M2, Centre for Mathematical Sciences,
Technische Universität München, Germany
wohlmuth@ma.tum.de

CP1

Planning of MR-Guided Laser Induced Thermal Therapy Using UQ Methods

Magnetic resonance-guided laser induced heating of cancer metastases in brain is a minimally invasive alternative to surgery. MR thermal imaging provides thermal dose feedback to the surgeon during treatment. A generalized polynomial chaos implementation of the stochastic bioheat equation was critically evaluated for aiding in the planning and effective therapy delivery. High performance implementations are needed clinically. Statistical comparisons of the measured and predicted temperature field will be presented in several in-vivo datasets.

David Fuentes, Samuel Fahrenholtz, John Hazle, Jason Stafford
MD Anderson Cancer Center
fuentesdt@gmail.com, sjfahrenholtz@mdanderson.org, jhazle@mdanderson.org, jstafford@mdanderson.org

CP1

Renewables Integration in Power System Operations Modeling

The requirements for a smart, green, efficient and reliable grid are reshaping the US power grid. The integration of renewable resources in the power system operations introduces further complexity to utilities profit optimization problem. The Renewable Integration Model (RIM) a computer based system - offers insight on power systems management, and answers questions on the integration of renewable energies. The RIM's design, solution method, and initial results are studied in the following.

Cristina Marinovici
Pacific Northwest National Laboratory
cristina.marinovici@pnnl.gov

Harold Kirkham
Pacific Northwest National Laboratory
Richland WA
harold.kirkham@pnnl.gov

Leif Carlsen
PNNL
leif.carlsen@pnnl.gov

Kevin Glass
Pacific Northwest National Laboratory
Richland WA
kevin.glass@pnnl.gov

CP1

Heel Effect Adaptive Gain Correction of Flat Panel X-Ray Detectors

Anode heel effect causes large-scale variations in the X-ray beam intensity. Together with the localized gain variability of the detector, they contribute to the non-uniformities of digital radiographs. For mobile applications, conventional flat field gain correction techniques degrade substantially without recalibration after each source-to-image distance (SID) reset. We present a novel method to adaptively compensate for the heel effect with flat field techniques and minimize the effort for recalibration. The method is based on a mathematical-physical model to estimate the heel effect ratio at two different SIDs. The gain factor calibrated at the standard SID is applied to any setting with heel effect compensation. The model parameters are determined by fitting the model to directly measurable quantities. The effect of the proposed method is demonstrated on flat field images acquired at variable SIDs.

Jue Wang
Union College
wangj@union.edu

Yongjian Yu
Varian Medical Systems
yyu6b@yahoo.com

CP1

A Parallel Two-Level Newton-Krylov-Schwarz Method For Three-Dimensional Blood Flow Simulations

We develop a parallel scalable domain decomposition method for numerical simulations of blood flows in three-dimensional compliant arteries. The corresponding fluid-structure interaction is modeled by a fully coupled system of linear elastic equation and the incompressible Navier-Stokes equations in an ALE framework. The resulting monolithic system is discretized with a fully-implicit finite element method on an unstructured moving mesh and solved by the Newton-Krylov method with a two-level Schwarz preconditioner. The investigation focuses on the efficiency and parallel scalability of the proposed two-level algorithm.

Yuqi Wu
Department of Applied Mathematics, University of Washington
yuqiwu@uw.edu

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

CP2

Nonlinear Scale Interactions and Energy Pathways in the Ocean

We refine a novel mathematical framework to analyze nonlinear scale-coupling and energy pathways in scale and in space from high-resolution ocean simulations. We test the validity of the traditional paradigm for such pathways at various locations such as in western boundary currents, near the equator, and in the deep ocean. We investigate the contribution of various nonlinear mechanisms to energy transfer across scales such as baroclinic and barotropic instabilities, barotropization, and Rossby wave generation.

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

Geoffrey Vallis, Kirk Bryan
Princeton/GFDL
gvk@princeton.edu, kbryan@princeton.edu

Robert Ecke
Center for Nonlinear Studies
Los Alamos National Laboratory
ecke@lanl.gov

Mathew Maltrud
LANL
maltrud@lanl.gov

Beth Wingate
Los Alamos National Laboratory
Computer And Computational Sciences Division
wingate@lanl.gov

CP2

A Non-Hydrostatic Spectral-Element Model of the Atmosphere

We are developing a non-hydrostatic version of the HOMME or CAM-SE spectral-element dynamical core employed by NCAR and the climate modeling community. We will discuss the design choices of this model including vertical coordinates and time stepping techniques methods, and demonstrate its performance on several standard hydrostatic and non-hydrostatic tests.

David M. Hall
National Center for Atmospheric Research
1850 Table Mesa Drive, Boulder CO 80305
david.hall@colorado.edu

CP2

Non-Dissipative Space Time Hp-Discontinuous Galerkin Method for the Time-Dependent Maxwell Equations

A space-time finite element method for the time-dependent Maxwell equations is presented. The method allows for local hp-refinement in space and time by employing a space-time Galerkin approach and is thus well suited for hp-adaptivity. Inspired by the so called cG methods for ODEs, nonequal test and trial spaces are employed. This allows for obtaining a non-dissipative method. For an efficient implementation, a hierarchical tensor product basis in space and time is proposed.

Martin Lilienthal
Graduate School of Computational Engineering
TU Darmstadt
lilienthal@gsc.tu-darmstadt.de

CP2

Discontinuous Collocation Method, Convergence and Applications

Collocation is a successful technique to simulate PDE in one spatial dimension. Recently, we have introduced discontinuous collocation method and applied it to PDE in several spatial variables. The technique employs triangular or tetrahedral partitions and piecewise polynomial interpolations. We proved that the method is L-infinity convergent. In the current work we prove L2 convergence under more general hypothesis. Our proof is similar to Bialeckis proof for rectangular partitions. We provide improved examples.

John A. Loustau
Hunter Collete CUNY
Dept of Math and Stat
jloustau@msn.com

Ariel Lindorff
Hunter College CUNY
Dept of Math Stat
alindorff@gmail.com

CP2

Automated Refinements for Discontinuous Collocation Method

Discontinuous collocation method is based on piecewise polynomial interpolation over a triangular domain partition. The method is discontinuous as the interpolation need not be continuous at element boundaries. If we expect the solution to be continuous then the size of the jump

discontinuities can be used to measure the error. In this work we use these values to guide a local refinement process which may be repeated until the implied error is within prescribed bounds.

Amy Wang
Hunter College CUNY - Department of Math and Stat
ms.amy.wang@gmail.com

John A. Loustau
Hunter Collete CUNY
Dept of Math and Stat
jloustau@msn.com

CP3

Wave Dynamics of Two-Phase Flow Models

Hyperbolic conservation laws with relaxation terms can be used to model two-phase flow in chemical non-equilibrium. Relaxation terms will introduce dispersion in the wave dynamics of the hyperbolic system. We use a relaxation parameter $\xi = k\varepsilon$ to study the transition between the equilibrium model and the frozen (non-equilibrium) model. Herein, we demonstrate the existence of a critical point for which the wave dynamics change from being equilibrium-like to being frozen-like.

Peder Aursand, Susanne Solem
Norwegian University of Science and Technology
peder.aursand@gmail.com, susanne.solem@ntnu.no

Tore Flåtten
SINTEF Energy Research
t.h.flatten@cma.uio.no

CP3

Dispersive Wave Propagation in Solids with Microstructure

The accuracy of dispersive wave models for solids with microstructure is analyzed through a series of numerical simulations. We compare numerical results obtained using various descriptions of the internal structure of a material: a micromorphic model for the microstructure, regular laminate structures, laminates with substructures, etc., for a large range of material parameters and wavelengths. Numerical computations are performed by means of the finite-volume numerical scheme, which belongs to the class of wave-propagation algorithms.

Arkadi Berezovski
Institute of Cybernetics at Tallinn University of Technology
arkadi.berezovski@cs.ioc.ee

Mihhail Berezovski
Department of Mathematical Sciences Worcester Polytechnic
Institute
mberezovski@wpi.edu

CP3

Efficient Boundary Element Methods for Molecular Electrostatics Using Python and Gpus

Many biomolecular processes are studied using electrostatics and continuum models. One such model for proteins in ionic solutions results in a coupled system of Poisson and linearized Poisson-Boltzmann equations. Using an integral formulation, this problem is solved with boundary element methods, BEM. We present a treecode-accelerated BEM that can solve biologically relevant problems using GPU hardware but with a friendly Python interface. We will present validation results and demonstrations with multiple proteins in solution.

Christopher Cooper
Boston University
cdcooper@bu.edu

Jaydeep Bardhan
Department of Electrical and Computer Engineering
Northeastern University
bardhan@alum.mit.edu

Lorena A. Barba
Department of Mechanical Engineering
Boston University
labarba@bu.edu

CP3

Enrichment by Exotic NURBS Geometrical Mappings in Isogeometric Analysis for fracture Mechanics

The mapping techniques are concerned with constructions of nonconventional exotic geometrical mappings such that the pushforward functions of B-spline functions defined on the parameter space into physical domain by the mappings generate singular functions that resemble the given point singularities. In this paper, for highly accurate stress analysis of elastic structures with cracks or corners, we mix NURBS basis functions and singular basis functions constructed by the design mapping and the exotic geometric mapping, respectively.

Hae-Soo Oh
The Department of Mathematics and Statistics
University of North Carolina at Charlotte
hso@uncc.edu

Hyunju Kim
University of North Carolina at Charlotte
The Department of Mathematics and Statistics
hkim22@uncc.edu

Jae Woo Jeong
Department of Mathematics
Miami University, Hamilton, OH 45011
jeongjw@muohio.edu

CP3

Computational Analysis of Non-Standard Wave Structures in Hydrodynamic and Magneto-Hydrodynamic Systems

We consider the Euler equations of hydrodynamics and the system of magneto-hydrodynamic equations both closed with a non-convex equation of state. We study the wave structure for both models and compute the approximate solution of several initial value problems. We analyze the formation of complex composite waves in both models. We investigate the influence of thermodynamical and magnetic field magnitudes in the formation of composite waves.

Susana Serna
Department of Mathematics

Universitat Autonoma de Barcelona
serna@mat.uab.es

CP3

Elimination of Oscillating Singularities at the Crack-Tips of An Interface Crack with a Help of a Curvature-Dependent Surface Tension

A new model of fracture mechanics incorporating a curvature-dependent surface tension acting on the boundaries of a crack is considered. The model is studied on the example of a single straight interface crack between two elastically dissimilar semi-planes. Linear elasticity is assumed for the behavior of the material of the plate in a bulk. A non-linear boundary condition with a consideration for a curvature-dependent surface tension is given on the crack boundary. It is well known from linear elastic fracture mechanics (LEFM) that oscillating singularities exist at the crack tips and lead to non-physical interpenetration and wrinkling of the crack boundaries. Using the methods of complex analysis, such as Dirichlet-to-Neumann mappings, the problem is reduced to a system of six singular integro-differential equations. It is proved that the introduction of the curvature-dependent surface tension eliminates both classical power singularities of the order 1/2 at the tips of the crack and oscillating singularities, thus resolving the classical contradictions of LEFM. Numerical computations are presented.

Anna Zemlyanova
Department of Mathematics
Texas A&M University
azem@math.tamu.edu

CP4

A Coupled Poroelasticity Numerical Model by Mixed Finite Elements

The numerical treatment of coupled poroelasticity is difficult because of the instabilities affecting the pressure solution. A coupled poroelasticity model based on Mixed Finite Elements has been developed so as to alleviate the numerical oscillations at the interface between different materials. Both a monolithic and sequential scheme are implemented, the former implying the use of ad-hoc preconditioned Krylov methods, the latter the separate solution of flow and deformation with an outer iteration to achieve convergence.

Massimiliano Ferronato
University of Padova
DMMMSA
ferronat@dmsa.unipd.it

Nicola Castelletto
DMMMSA
University of Padova
castel@dmsa.unipd.it

Carlo Janna
DMMMSA - University of Padova
janna@dmsa.unipd.it

CP4

Efficient Time-Stepping for Long-Time Simulations of Parabolic Reaction-Diffusion Equations

The flow of calcium ions in a heart cell is modeled by a system of time-dependent parabolic reaction-diffusion equations. The spatial discretization of this system by finite difference, finite element, and finite volume methods in a method of lines approach results in systems of ODEs that necessarily get stiffer and larger with increasing spatial mesh resolution. We will compare several state-of-the-art linear multi-step and Runge-Kutta methods, all specifically designed for problems of this type.

Xuan Huang
Department of Mathematics and Statistics
University of Maryland, Baltimore County
hu6@umbc.edu

Philipp Birken
Department 10, Mathematics and Natural Sciences
University of Kassel, Germany
birken@mathematik.uni-kassel.de

Matthias K. Gobbert
University of Maryland, Baltimore County
Department of Mathematics and Statistics
gobbert@umbc.edu

CP4

Adaptive Hp Finite Elements Using Cuda

Adaptive hp finite elements lead to exponential convergence to the exact solution for many PDEs. High polynomial order methods also offer the possibility of higher arithmetic intensity (ratio of computation performed to memory IO) on modern many-core architectures. We present a new high-performance implementation tuned for adaptive hp finite elements that hides the complexity of performing quadrature for multiple polynomial orders, arbitrary reference to physical mapping, and anisotropic and variable material properties. We also present performance results using NVIDIA CUDA.

Chetan Jhurani, Paul Mullowney
Tech-X Corporation
jhurani@txcorp.com, paulm@txcorp.com

CP4

A Robust Non-Negative Numerical Framework for Diffusion-Controlled Bimolecular-Reactive Systems

In this talk, we will present a novel non-negative computational framework for diffusive-reactive systems in heterogeneous anisotropic rigid porous media. The governing equations for the concentration of reactants and product will be written in terms of tensorial diffusion-reaction equations. We shall restrict ourselves to fast irreversible bimolecular reactions in which Damköhler number $Da >> 1$. We employ a linear transformation to rewrite the governing equations in terms of invariants, which are unaffected by the reaction. This results in two uncoupled tensorial diffusion equations in terms of these invariants, which are solved using a novel non-negative solver for tensorial diffusion-type equations. The concentrations of the reactants and product are then calculated from invariants using algebraic manipulations. Several representative numerical examples will be presented to illustrate the robustness, convergence, and the performance of the proposed computational framework. We will also compare the proposed formulation with other popular formulations. In particular, we will show that the standard single-field formulation does not produce reliable solutions, and the reason can be attributed to the fact that the single-field formulation does not guarantee non-negative solutions. We will also show that the clipping procedure (which produces non-negative solutions but is considered as a variational crime) over predicts the plume length.

Maruti K. Mudunuru, Kalyana Nakshatrala
University Of Houston - Main Campus
maruti.iitm@gmail.com, knakshatrala@uh.edu

Albert J. Valocchi
University of Illinois at Urbana-Champaign
valocchi@illinois.edu

CP4

Complexity-Friendly Algebraic Multigrid Preconditioners Based on Sparsified Coarsening

Traditional algebraic multigrid (AMG) methods use Galerkin coarsening where the sparsity of the multigrid hierarchy is dictated by its transfer operators. In many scenarios, the coarse operators tend to be much denser than the fine operator as the coarsening progresses. Such behavior is problematic in parallel computations where it imposes expensive communication overhead. We present a new technique for controlling the sparsity of the AMG operators. Our algorithm sparsifies the Galerking operators while preserving the energy of the algebraically smooth error modes. Numerical experiments for large-scale Graph-Laplacian problems demonstrate the potential of this approach.

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

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

Raanan Fattal
Computer Science and Engineering,
The Hebrew Unisersity of Jerusalem
raananf@cs.huji.ac.il

CP4

Localized Method of Approximate Particular Solutions for Solving Diffusion-Reaction Equa Tions in Two-Dimensional Space

An localized meshless method, the localized method of approximate particular solutions(LMAPS), was developed recently to solve elliptic partial differential equations. In this paper, the method is extended to solve diffusion-reaction equations in two-dimensional space, using positive definite radial basis function multiquadrics and non-positive definite radial basis function thin-plate splines. The comparison of method with different radial basis functions are tested on three examples. The numerical experiment suggested that the localized method with multiquadric is much more stable than the method with thin-plate splines.

Guangming Yao
Clarkson University
gyao@clarkson.edu

CP5

Spectral Methods on GPUs with Applications in Fluid Dynamics and Materials Science

GPGPU recently has drawn a big attention from the CSE community. This talk exploits how spectral methods can be benefited from GPGPU. Our focus is the design and implementation of spectral-collocation and spectral-Galerkin methods for elliptic systems. And we look at applications of the Navier-Stokes equation in incompressible fluids, the Cahn-Hilliard Equation in phase-field models, and nonlinear PDEs from the recently proposed phase-field-crystal model.

Feng Chen
Brown University
feng_chen_1@brown.edu

CP5

Comparative Study of Various Low Mach Number Preconditioners Applied to Steady and Unsteady Inviscid Flows

The performance of many existing compressible codes degrades as the Mach number of the flow tends to zero and leads to inaccurate computed results. To alleviate the problem, Roe-type based schemes have been developed for all speed flows, such as the preconditioned Roe and LM-Roe schemes. Most of these schemes tend to compute steady state solutions, but many fail to address unsteady solutions. This paper attempts to evaluate various preconditioners in terms of development, performance, accuracy and limitations in both steady and unsteady simulations at various Mach numbers by implementing them in a 3D Euler Solver using Roe Scheme on Unstructured meshes. While developing 3D Unstructured Euler Solver, mathematical models of each preconditioning is developed by using proper Non-Dimensionalization, Primitive variables and Characteristics based Boundary Conditions. Preconditioners proposed by W.R. Briley, Weiss-Smith, Turkel, Eriksson, Choi and Merkel, and LM-Roe are studied by computing flow over a NACA0012 airfoil and Cylinder at various Mach Numbers.

Ashish Gupta
University of Tennessee at Chattanooga
cxw842@mocs.utc.edu

CP5

Rankine-Hugoniot-Riemann Solver for Multi-Dimensional Balance Laws with Source Terms

We developed a *Rankine-Hugoniot-Riemann* (RHR) solver to solve systems of multi-dimensional balance laws with source terms. The solver incorporates the source terms and the cross fluxes as a singular term in each finite volume cell, yielding a jump in the solution satisfying a Rankine-Hugoniot condition. The Riemann problems at the cell boundaries can then be solved using a conventional Riemann solver. The solver is shown to be of second order for physically relevant equations.

Halvor Lund
Department of Energy and Process Engineering
Norwegian University of Science and Technology
halvor.lund@ntnu.no

Florian Müller
Institute of Fluid Dynamics

ETH Zürich
florian.mueller@sam.math.ethz.ch

Patrick Jenny
Institute of Fluid Dynamics
ETH Zurich
jenny@ifd.mavt.ethz.ch

CP5

Implementation of Multigrid Method for the Navier-Stokes Equations in Cuda

This contribution is concerned with implementation of the ultigrid method with Vanka type smoother on GPU. The method is used to solve problem of 2D flow over an urban canopy governed by Navier-Stokes equations and discretized by means of the mixed finite element method. GPU version of the algorithm achieved speed-up 5 compared to parallel code based on OpenMP and 26 compared to the sequential code.

Tomas Oberhuber
Faculty of Nuclear Sciences and Physical Engineering
Czech Technical University in Prague
tomas.oberhuber@fjfi.cvut.cz

Vladimir Klement, Vitezslav Zabka, Petr Bauer
Faculty of Nuclear Sciences and Physical Engineering
Czech T
wlada@post.cz, zabkav@gmail.com,
petr.bauer@fjfi.cvut.cz

CP5

Numerical Simulation of Two-Phase Incompressible Flows with Complex Interfaces

We consider a flow problem with two different immiscible incompressible newtonian phases (fluid-fluid or fluid-gas). A standard model for this consists of the Navier-Stokes equations with viscosity and density coefficients that are discontinuous across the interface and with a localized surface functional that models interface force. An example is the Boussinesq-Scriven model for viscous interfaces. We present finite element techniques for the the numerical treatment of such flow problems. Topics that are briefly addressed are the XFEM and space-time DG method. Results of numerical methods will be presented.

[1] S. Groß, A. Reusken, Numerical Methods for Two-phase Incompressible Flows, Springer 2011.

Arnold Reusken
Numerical Mathematics
RWTH Aachen University, Aachen, Germany
reusken@igpm.rwth-aachen.de

CP5

Mpp Limiter for Finite Difference Rk-Weno Scheme with Applications in Vlasov Simulations and Advection in Incompressible Flows

In this talk, we will give a maximum principle preserving limiter for WENO scheme with high order explicit Runge-Kutta time discretization with applications in Vlasov and advection of incompressible flow problems. We directly dealing with the numerical fluxes.Our approach is very easy, and can be applied to finite difference scheme, so that it could be straightforwardly extended to high dimensional problems. Numerical experiments demonstrate the high accuracy and high efficiency of our method.

Tao Xiong, Jingmei Qiu
University of Houston
xiongt001@gmail.com, jingqiu@math.uh.edu

Zhengfu Xu
Michigan Technological University
Dept. of Mathematical Sciences
zhengfux@mtu.edu

CP6

O(N) Parallel Algorithm for Computing Selected Elements of the Inverse of a Gram Matrix in Electronic Structure Calculations

Recently, there has been increased interest in developing O(N) algorithms for Density Functional First-Principles molecular dynamics at exascale. However, the energy functional formulation for general non-orthogonal orbitals requires the knowledge of selected elements of the inverse of an associated global Gram matrix. We present a scalable parallel algorithm, based on domain decomposition and approximate inverse techniques, for computing selected elements of the inverse of the Gram matrix in electronic structure calculations.

Daniel Osei-Kuffuor, Sebastien Hamel
Lawrence Livermore National Laboratory
dan.osei@gmail.com, hamel2@llnl.gov

Jean-Luc Fattebert
Lawrence Livermore National Lab.
fattebert1@llnl.gov

CP6

An Efficient State-Space Based Method for Direct Simulation of Particle-Laden Turbulent Flows

We present a new numerical method for the coupled simulation of particulate poly-disperse flow using the probabilistic description (Boltzmann like), within turbulent flows. The method is stated in a purely deterministic Eulerian framework using the particle density function of the velocity-position state space of the particles while the carrier phase obeys the incompressible Navier-Stokes equations. Application of the method will also be demonstrated by simulating a particle-laden turbulent flow using a parallel MPI-based solver implementation.

Carlos Pantano, Reetesh Ranjan
Mechanical Science and Engineering
University of Illinois at Urbana-Champaign
cpantano@illinois.edu, rranjan2@illinois.edu

CP6

On the Equivalence of P3M and NFFT-Based Fast Summation

The P3M method combines the Ewald sum with the fast Fourier transform in order to calculate long range Coulomb interactions for the classical N-body problem approximately with $\mathcal{O}(\mathcal{N} \log \mathcal{N})$ arithmetic operations. During this talk, we show that P3M belongs to a class of fast summation algorithms that are based on nonequispaced fast Fourier transforms (NFFT). This connection yields new theoretical and practical insights. Furthermore, we present performance results of our massively parallel implementation.

Michael Pippig
Chemnitz University of Technology
Department of Mathematics
michael.pippig@mathematik.tu-chemnitz.de

CP6

Fast Evaluating Matern Covariance Kernel by a CartesianTreecode

Evaluating sums of multivariate Matern kernels is a common computational task in statistical and machine learning community. The quadratic computational complexity of the summation is a signicant barrier to practical applications. We develop a Cartesian treecode algorithm to efficiently estimate sums of the Matern Kernel. The method uses a farfield Taylor expansion in Cartesian coordinates to compute particle-cluster interactions. The Taylor coefficients are obtained by recurrence relations which allows effifficient computation of high order approximations. In the serial code, for a given order of accuracy, the treecode CPU time scales as $O(N \log N)$ and the memory usage scales as $O(N)$, where N is the number of particles. Parallel code also gives promising scale.

Lei Wang
Argonne National Laboratory
lwang@mcs.anl.gov

CP6

Uncertainty Quantification of Molecular Systems

We propose a method to quantify the uncertainties in the molecular dynamics (MD) system to enhance the efficiency, we employ ℓ_1 minimization to recover the coefficients in general polynomial chaos (gPC) expansion given the prior knowledge that the coefficients are "sparse". We implement this method to quantify the dominant terms of gPC coefficients for the MD measurement (e.g.,viscosity, diffusivity, etc.) This method thoroughly exploits information from limited simulations, hence it is very efficient.

Huan Lei, Xiu Yang, George Karnidakis
Brown University
Huan_Lei@brown.edu, xiu_yang@brown.edu, george_karniadakis@brown.edu

CP6

A Faster Fft in the Mid-West

We describe in this paper an FFT library that was build while paying special attention to locality. The library achieves significantly better performance than FFTW, for long vectors. Unlike FFTW or Spiral, the recursive decomposition of the FFT is not created by a library generator; it is created by macro expansion that has a few selectable parameters. This provides an interface that can be more easily modified by users

Alexander Yee
University of Illinois Urbana-Champaign
a-yee@u.northwestern.edu

Marc Snir
University of Illinois at Urbana-Champaign
snir@illinois.edu

CP7

Efficient Computation of Eigenpairs for Large Scale-free Graphs

Large-scale network and data analysis is an area of great importance. The extreme eigenvalues and eigenvectors offer useful insights, but are difficult (expensive) to compute for large problems. We present a computational framework for large spectral graph computations based on the Trilinos toolkit. We study the effectiveness of several eigensolvers and preconditioners. Preliminary results indicate graph-based preconditioners can be highly effective.

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

Karen D. Devine, Richard B. Lehoucq
Sandia National Laboratories
kddevin@sandia.gov, rblehou@sandia.gov

Nicole Slattengren
Sandia National Labs
nlslatt@sandia.gov

Kevin Deweese
UC Santa Barbara
na

CP7

Versatile Batch QR Factorization on GPUs

This talk describes a versatile, GPU-based BatchQR library that performs the QR factorization of an array of moderately sized, dense matrices of variable dimensions. It addresses a growing need for batched matrix computations on the GPU as opposed to single, large matrix factorizations. The library exploits data as well as task-level parallelism by decomposing each dense QR into sets of independent tasks, which are performed entirely on the GPU.

Sharanyan Chetlur, Nuri Yeralan
University of Florida
schetlur@ufl.edu, nuriatuf@gmail.com

CP7

Utilizing Slater Matrix Properties to Design Better Preconditioners for Quantum Monte Carlo Methods

Quantum Monte Carlo (QMC) methods are often used to solve problems involving many body systems. In implementing QMC methods, sequences of Slater matrices are constructed. Solving linear systems for each matrix is a major part of the cost. Currently, preconditioned linear solves are utilized, reducing this cost from $O(n^3)$ to $O(n^2)$ power sweep (n updates). We analyze the properties of Slater matrices, attempting to further reduce the cost. Direct solvers will also be examined.

Arielle K. Grim Mcnally
Virginia Tech
Eric de Sturler
arielle5@vt.edu

Eric De Sturler
Virginia Tech

sturler@vt.edu

CP7

Augmenting and Shifting in a Restarted Lanczos Bidiagonalization Method

The Lanczos bidiagonalization method can be used to solve SVD and least squares problems. Generally, the method requires restarting. Two restarted methods for SVD problems are implicit shifting and augmenting the subspace after each restart. In this talk, we examine the Krylov subspaces that result from the bidiagonalization process applied in SVD and least squares schemes using harmonic Ritz values as shifts and harmonic Ritz vectors as augmenting vectors. Computed examples and results are presented.

Daniel J. Richmond, James Baglama
University of Rhode Island
dan@math.uri.edu, jbaglama@math.uri.edu

CP7

Multi-Preconditioning Gmres

Standard Krylov subspace methods for the solution of nonsymmetric linear systems only allow the user to choose a single preconditioner, although in many situations there may be a number of possibilities. Here we describe multi-preconditioned GMRES, which allows the use of more than one preconditioner. These multiple preconditioners can be combined in an optimal manner. We give some theoretical results, propose a practical algorithm, and present numerical results from problems in domain decomposition and PDE-constrained optimization. These numerical experiments illustrate the applicability and potential of the multi-preconditioned approach.

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

Chen Greif
Department of Computer Science
The University of British Columbia
greif@cs.ubc.ca

Tyrone Rees
Rutherford Appleton Lab, UK
tyrone.rees@stfc.ac.uk

CP7

Multifrontal Sparse QR Factorization on GPU

GPU devices are growing in popularity as heterogeneous architectures become prevalent in scientific computing. We discuss a multifrontal sparse QR factorization method implemented on GPU accelerators. Our method exploits task-based and data-based parallelism across the elimination tree with frontal matrices of differing sizes, within contribution block assembly, and within the numeric factorization. Finally, we show performance results for problems found in the UF Sparse Matrix collection.

Nuri Yeralan
University of Florida
nuriatuf@gmail.com

CP8

A Parallel Method for Computing Visibility Regions and Its Application to Dynamic Coverage Control

We present a parallel algorithm for computing visibility regions in complex terrain and explore its application to the problem of dynamic coverage control with a network of near-ground mobile sensor agents. The highly scalable method makes it tractable to compute the visibility region of each sensor and thus account for possible terrain obstructions. We show the scalability of this algorithm on Graphics Processing Units and provide coverage control simulations employing our novel sensor model.

Miles L. Detrixhe
University of California Santa Barbara
mileslucasd@gmail.com

CP8

State of the Art Hifoo: H_∞ Controller Optimization for Large and Sparse Systems

We present a new state of the art version of HIFOO, a MATLAB package for optimizing H_∞ and H_2 controllers for linear dynamical systems, supporting simultaneous multiple plant stabilization. Previous versions of HIFOO have generally been limited to smaller scale problems, due to the high asymptotic cost of computing the H_∞ norm of the transfer matrix. However, new sparse methods for computing the H_∞ norm allow HIFOO to extend efficiently to large and sparse systems.

Tim Mitchell
Courant Institute of Mathematical Sciences
New York University
tim.mitchell@cs.nyu.edu

Michael L. Overton
New York University
Courant Instit. of Math. Sci.
overton@cs.nyu.edu

CP8

Hybrid Functions Approach for Optimal Control Problems

In this work, we present a new direct computational method to solve optimal control problems. The approach is based on reducing the optimal control problems into a set of algebraic equations by first expanding the required solution as a hybrid function with unknown coefficient. These hybrid functions, which consist of block-pulse functions and Bernoulli polynomials are first introduced. Numerical examples are included to demonstrate the applicability and the accuracy of the proposed method and a comparison is made with the existing results.

Mohsen Razzaghi
Mississippi State University
razzaghi@math.msstate.edu

CP8

Multigrid Solution of a Distributed Optimal Control Problem Constrained by a Semilinear Elliptic Pde

We study a multigrid solution strategy for a distributed optimal control problem constrained by a semilinear elliptic PDE. Working in the discretize-then-optimize framework, we solve the reduced optimal control problem using Newtons method. Further, adjoint methods are used to compute matrix-vector multiplications for the reduced Hessian. In this work we introduce and analyze a matrix-free multigrid preconditioner for the reduced Hessian which proves to be of optimal order with respect the discretization.

Jyoti Saraswat
Department of Mathematics & Statistics
University of Maryland Baltimore County
jyo1@umbc.edu

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

CP8

Multigrid Preconditioners for Optimal Control Problems in Fluid Flow

We construct multigrid preconditioners to accelerate the solution process of optimal control problems constrained by the Stokes/Navier-Stokes equations. Our approach for the Stokes control problem is to eliminate the state and adjoint variables from the optimality system and to construct efficient multigrid preconditioners for the Schur-complement of the block associated with these variables. Similar preconditioners are constructed for the reduced Hessian in the Newton-PCG method for the optimal control of the stationary Navier-Stokes equations.

Ana Maria Soane
University Of Maryland, Baltimore County
Department of Mathematics and Statistics
asoane@math.umbc.edu

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

CP8

Thermoelastic Shape Optimization

Shape optimization is often devoted to systems with one physical discipline. This talk discusses specific problems, when dealing with shape optimization problems in multiphysics systems. In particular, the coupling of the heat equation with the elasticity equation is studied as it is the case in standard heating devices made from iron. Numerical approaches to the optimization of the iron body of hot plates under the coupled multi-physics problem are presented together with computational results.

Heinz Zorn
University of Trier
zorn4501@uni-trier.de

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

CP9

Uncertainty Quantification for Subsurface Flow Models Using Bayesian Nested Sampling Algorithm

We investigate an efficient sampling algorithm known as nested sampling, which can simultaneously perform the tasks of parameter estimation, sampling from the posterior distribution for uncertainty quantification, and estimating the Bayesian evidence for model comparison. Nested sampling has the advantage of generality of application, and computational feasibility. In this talk, we report the first successful application of nested sampling for calibration and model selection of several nonlinear subsurface flow problems.

Ahmed H. ElSheikh
Department of Earth Science and Engineering
Imperial College London, South Kensington Campus
aelsheikh@ices.utexas.edu

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

Ibrahim Hoteit
King Abdullah University of Science and Technology (KAUST)
ibrahim.hoteit@kaust.edu.sa

CP9

Variable Dimensional Bayesian Elastic Wave-Field Inversion of a 2D Heterogeneous Media

This paper introduces a methodology to infer the spatial variation of the elastic characteristics of a 2D heterogeneous earth model via Bayesian approach, given the probed medium's response to interrogating SH waves measured at the surface. A reduced dimension, self regularized treatment of the inverse problem using partition modeling is introduced, where the velocity field is discretized by a variable number of disjoint regions defined by Voronoi tessellations. A provided synthetic study indicates the strength of the technique.

Saba S. Esmailzadeh
Zachry Department of Civil Engineering,
Texas A&M University, College Station, TX, USA 77840
saba61@neo.tamu.edu

Zenon Medina-Cetina
Zachry Department of Civil Engineering
Texas A&M University
zenon@tamu.edu

Jun Won Kang
Department of Civil Engineering
New Mexico State University, NM 88003, USA.
jwkang@nmsu.edu

Loukas F. Kallivokas
The University of Texas at Austin
loukas@mail.utexas.edu

CP9

Optimization in the Presence of Uncertainty: Applications in Designing Random Heterogeneous

Media

The present paper discusses a sampling framework that enables the optimization/design of complex systems characterized by high-dimensional uncertainties and design variables. We are especially concerned with problems relating to random heterogeneous materials where uncertainties arise from the stochastic variability of their properties. In particular, we reformulate topology optimization problems to account for the design of truly random composites. The methodological advances proposed in this paper allows the analyst to identify several local maxima that provide important information with regards to the robustness of the design. We further propose a principled manner of introducing information from approximate models that can ultimately lead to further reductions in computational cost.

Phaedon S. Koutsourelakis
Heriot Watt University
p.s.koutsourelakis@tum.de

CP9

Adaptive Total Variation Regularization in Image Processing

We consider an adaptive version of the well-known total variation (TV) based regularization model in image processing. Adaptive TV based schemes try to avoid the staircasing artifacts associated with TV regularization based energy minimization models. An algorithm based on a modification of the split Bregman technique proposed by [Goldstein and Osher, The split Bregman algorithm for L^1 regularized problems, SIAM J. Imaging Sci. 2009], can be used for solving the adaptive case. Convergence analysis of such an alternating scheme is proved using the Fenchel duality and a recent result of [Svaiter, On weak convergence of the Douglas-Rachford method, SIAM J. Control Optim. 2011] on the weak convergence of Douglas-Rachford splitting method. We will provide detailed comparative experimental results using the modified split Bregman, dual minimization and additive operator splitting for the gradient descent scheme for the TV diffusion equation to highlight the efficiency of adaptive TV based schemes for image denoising, decomposition and segmentation problems.

Surya Prasath, Kannappan Palaniappan
University of Missouri-Columbia
prasaths@missouri.edu, palaniappank@missouri.edu

CP9

Transient Hydraulic Tomography Using the Geostatistical Approach

Transient Hydraulic Tomography is a well known method for characterizing aquifers that consists of measuring pressure (head) responses to estimate important parameters (eg. conductivity, specific storage). We discuss various computational issues regarding their large-scale implementation using the Geostatical approach and techniques for computing the uncertainty associated with the estimate.

Arvind Saibaba
Stanford University
arvindks@stanford.edu

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

CP9

Quantification of the Loss of Information in Source Attribution Problems in Global Atmospheric Chemical Transport Models

It is of crucial importance to be able to identify the location of atmospheric pollution sources in our planet in order to evaluate global emissions policies like the Kyoto protocol. It is shown that the ability to successfully use the adjoint method to reconstruct the location and magnitude of an instantaneous source in global chemical transport models is compromised in relevant global atmospheric time scales.

Mauricio Santillana
Harvard University, School of Engineering and Applied Scienc
msantill@fas.harvard.edu

CP10

Computing Solutions of Dynamic Sustainability Games

We introduce a dynamic Nash game among firms harvesting a renewable resource and propose a differential variational inequality (DVI) framework for modeling and solving such game. We consider the application of the proposed framework to a sustainable halibut fishery. The problem is modeled using the DVI, which in turn converted to a fixed-point problem that allows computationally efficient algorithms. A numerical example is presented to show the use of our proposed framework.

Sung Chung
New Mexico Institute of Mining and Technology
New Mexico Tech
schung@nmt.edu

CP10

How to Deal with Dynamic Contact Angles in Moving Contact-Line Simulations

We present an efficient algorithm within phase field framework on how to impose dynamic contact-angle boundary conditions for wall-bounded liquid-gas flows. Our strategy consists of two components: (i) we ignore the boundary conditions and re-formulate the 4-th order Cahn-Hilliard equation into two nominally de-coupled Helmholtz type equations; (ii) we treat the dynamic contat-angle boundary conditions in such a manner that the two Helmholtz type equations are truly de-coupled. We will also present a strategy on how to treat the variable-density Navier-Stokes equations such that only constant and time-independent coefficient matrices will result from the discretization, even though variable density and variable viscosity are involved in the governing equations. The overall method can deal with both dynamic and static contact angles for moving contact line problems involving large density ratios, and the computations for all flow variables have been completely de-coupled. Simulations of air-water two-phase flows involving solid walls of different hydrophobicities will be presented.

Suchuan Dong
Purdue University
sdong@math.purdue.edu

CP10

On An Implicit Time Integration Method for Wave Propagations

The Bathe method is known to be very effective for finite element solutions of linear and nonlinear problems in structural dynamics. In this presentation, we present dispersion properties of the scheme and show that its desired stability and accuracy characteristics for structural dynamics are also valuable for the solution of wave propagation problems. A dispersion analysis using the CFL number is presented and simple examples are solved to illustrate the capabilities of the scheme for wave propagations.

Gunwoo Noh, Seounghyun Ham
Dept. of Mechanical Engineering, M.I.T.
gunwoo@mit.edu, ham4@mit.edu

CP10

Numerical Simulation of the Damping Behavior of Particle-Filled Hollow Spheres

For many industrial applications, light-weight materials play an important role. A new material developed at the Fraunhofer Institute for Manufacturing Technology and Advanced Materials uses particle filled hollow spheres embedded into a sandwich structure to dampen vibration through friction. We present a numerical simulation of such a sphere using a discrete element model. For this purpose, techniques from molecular dynamics are adapted and extended correspondingly. First results illustrating the damping behavior are shown.

Tobias Steinle
Universität Paderborn
tobias.steinle@uni-paderborn.de

Jadran Vrabec
University of Paderborn
jadran.vrabec@uni-paderborn.de

Andrea Walther
Universität Paderborn
andrea.walther@uni-paderborn.de

CP10

Hybrid Parallel Transistor-Level Full-Chip Circuit Simulation

The emergence of multicore and manycore processors has provided an opportunity for application-specific simulation tools to take advantage of potential intra-node parallelism. CAD applications, being fundamental to the electronic design automation industry, need to harness the available hardware resources to be able to perform full-system simulation for modern technology nodes. We will present a hybrid (MPI+threads) approach for performing parallel transistor-level transient circuit simulation that achieves scalable performance for some challenging large-scale integrated circuits.

Heidi K. Thornquist, Siva Rajamanickam, Mike Heroux
Sandia National Laboratories
hkthorn@sandia.gov, srajama@sandia.gov,
maherou@sandia.gov

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

CP10

Numerical Optimization of Instrument Settings in Reference Material Certification

For metrological purposes, the National Institute of Standards and Technology (NIST) provides over 1200 Standard Reference Materials (SRM) of the highest quality. Developing SRMs to meet todays technological needs demands ever-increasing measurement accuracy and precision. Traditional methods of hand tuning chemical spectrometers by experienced users are no longer adequate. We have developed a new method for optimizing instrument settings that does not reply on having an expert instrument user. We demonstrate how a noisy objective function is optimized through a novel gradient approximation which employs an a priori estimate of noise levels.

William E. Wallace
National Institute of Standards and Technology
william.wallace@nist.gov

Anthony Kearsley
Mathematical and Computational Sciences Division
U.S.A. National Institute of Standards and Technology
ajk@nist.gov

CP11

Multi-Scale Modelling of Droplets Deformation

Deformation and fragmentation of droplets due to growth of hydrodynamical instabilities are studied using a finite volume-volume of fluid (FV-VOF) numerical technique for falling droplets and droplets in a stream. Problems are solved in both two and three dimensions. Results were compared with stability analyses that provided based on the linearized Navier-Stokes equations. Reasonable agreement is observed between the numerical and analytical solutions for the most amplified wave numbers, growth rate etc.

Maziyar Jalaal
Department of Mechanical Engineering, The University of British Columbia, Vancouver, Canada.
m_jalaal@yahoo.com

Kian Mehravaran
School of Engineering, The University of British
Columbia,
Kelowna, Canada.
kian.mehravaran@ubc.ca

CP11

Domain Decomposition Preconditioning for Variational Monte Carlo on Insulators

The main cost of the Variational Monte Carlo(VMC) methods is in constructing a sequence of Slater matrices and computing the ratios of determinants for successive Slater matrices. The scaling of constructing Slater matrices for insulators has been improved so that its cost is now linear with the number of particles. To improve computing the ratios of determinants, we employ preconditioned Krylov subspace iteration methods. The main work here is using domain decomposition preconditioner to solve the ratios of determinants iteratively. Iterative Krylov methods like GMRES are performed. Since there is one supposed

moving particle in every attempt of the VMC methods, two-subdomain decomposition is used most of the times. One subdomain is the local domain around the moving particle, the other includes the remainning particles and orbitals. To make the iteration robust and stable, an effective reordering of Slater matrices and incomplete LU decomposition are used .

Ming Li
virginia polytechnic institute and state university
liming@vt.edu

Eric De Sturler
Virginia Tech
sturler@vt.edu

Arielle Grim-McNally
virginia polytechnic institute and state university
arielle5@vt.edu

CP11

Coarse-Grained Stochastic Particle-Based Reaction-Diffusion Simulation Algorithm

The stochastic, diffusive motion of molecules influences the rate of molecular encounters and hence many biochemical processes in living cells. Tiny time steps are typically required to accurately resolve the probabilistic encounter events in a simulation of cellular signaling, creating the dilemma that only simplistic models can be studied at cell-biologically relevant timescales. Particular solutions of the diffusion equation, Green's functions, allow for larger time steps and more efficient simulations by detecting otherwise unnoticed encounter events. However, previous implementations of Green's functions based approaches were not generally applicable to reaction-diffusion problems. We generalized the method to reflect the physiologically relevant case of reversible, partially diffusion-controlled reactions. In my talk, I will discuss how the approach can be applied in computational cell biology, taking advantage of an efficient approximation technique that allows us to include the mathematically challenging case of 2D reaction-diffusion systems on membranes as well. The resulting algorithm is easily implementable, flexible and efficient and provides a coarse-grained but nevertheless detailed, stochastic representation of biochemical processes.

Thorsten Prustel, Martin Meier-Schellersheim
Laboratory of Systems Biology (LSB)
NIAID, NIH
prustelt@niaid.nih.gov, mms@niaid.nih.gov

CP11

A Three-Dimensional Domain Decomposition Method for Large-Scale DFT Electronic Structure Calculations

We present a 3D domain decomposition method for atoms and grids. A modified recursive bisection method is developed based on inertia tensor moment to reorder the atoms along a principal axis for atom decomposition. We define four data structures for grid decomposition and propose a method for solving the Poisson equation using FFT with communications minimized. Benchmark results show that the parallel efficiency at 131,072 cores is 67.7% compared to the baseline of 16,384 cores.

Truong Vinh Truong Duy
Research Center for Simulation Science
Japan Advanced Institute of Science and Technology
duytvt@jaist.ac.jp

Taisuke Ozaki
Japan Advanced Institute of Science & Technology
t-ozaki@jaist.ac.jp

CP11

Multiscale Smooth Dissipative Particle Simulation of Non-Isothermal Flows

Smooth Dissipative Particle Dynamics (SDPD) is a multiscale mesh-free method that provides a bridge between the continuum and molecular scales. SDPD is thermodynamically consistent, involves realistic transport coefficients, and includes fluctuation terms. The SDPD is implemented in our work for arbitrary 3D geometries with a methodology to model solid wall boundary conditions. The entropy equation is implemented with a velocity-entropy Verlet integration algorithm. Flows with heat transfer are simulated for verification of the SDPD. We present also the self-diffusion coefficient derived from SDPD simulations for gases and liquids. Results show the scale dependence on SDPD particle size.

Jun Yang, Nikolaos Gatsonis, Raffaele Potami
Worcester Polytechnic Institute
yangjun@wpi.edu, gatsonis@wpi.edu, raffaele@wpi.edu

CP11

Domain Decomposition Based Jacobi-Davidson Algorithm for Quantum Dot Simulation

A domain decomposition based Jacobi-Davidson algorithm is proposed to find the interior eigenpairs of large sparse polynomial eigenvalue problems arising from pyramidal quantum dot simulation. We apply domain decomposition method both to improve the convergence of correction equation that is the most expensive part in Jacobi-Davidson algorithm and to obtain initial guesses with rich information on target eigenpairs. The numerical results show that our algorithm is effective with good scalability on supercomputer.

Tao Zhao
Department of Computer Science
University of Colorado Boulder
tao.zhao@colorado.edu

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

Xiao-Chuan Cai
Department of Computer Science
University of Colorado Boulder
xiao-chuan.cai@colorado.edu

CP12

Mining Periodic Patterns in Digital Footprint 'Event' Data via Empirical Mode Decomposition (EMD)

Digital trace data can often be interpreted as a temporal series of events (such as the times at which emails are sent or posts on on-line social networks made). We present a novel technique for the categorisation of individuals according to the temporal characteristics inherent in personal 'event' data. Our method, exemplified on synthetic and real data sets, is underpinned by projecting individuals' data onto a basis of pseudo-periodic intrinsic mode functions obtained from EMD.

Duncan S. Barrack, Chienmin Chuang
Horizon Digital Economy Research Institute
University of Nottingham
duncan.barrack@nottingham.ac.uk,
chienmin.chuang@nottingham.ac.uk

Keith Hopcraft, Simon Preston
School of Mathematical Sciences
University of Nottingham
keith.hopcraf@nottingham.ac.uk,
simon.preston@nottingham.ac.uk

James Goulding, Gavin Smith
Horizon Digital Economy Research Institute
University of Nottingham
james.goulding@nottingham.ac.uk,
gavin.smith@nottingham.ac.uk

CP12

Recent Advances in Moment-Matching Model Order Reduction for Maxwell's Equations

We present some new results for the application of moment-matching methods in the field of model order reduction for Maxwell's equations. In detail, a different interpretation of existing heuristic error estimations in combination with an adaptive expansion point selection will be presented. Furthermore, we will show how to deal with the discrete, divergence conditions during the computation of the reduced order model.

André Bodendiek
Institut Computational Mathematics
Technical University Braunschweig
a.bodendiek@tu-bs.de

CP12

A Reduced Basis Kalman Filter for Parametrized Parabolic Partial Differential Equations

Realizing a Kalman filter for solution estimations of stochastically forced partial differential equations (SFPDEs) is often infeasible due to the high computational cost. Hence, we propose a reduced basis (RB) Kalman filter for parametrized linear SFPDEs giving solution estimations in real-time. Ingredients of our contribution are the RB-construction via KL-expansion and Duhamel's principle, certification by derivation of statistical information about error bounds for stochastically forced RB-models and application to a stochastically disturbed parametrized heat equation.

Markus Dihlmann, Bernard Haasdonk
University of Stuttgart
dihlmann@mathematik.uni-stuttgart.de,
haasdonk@mathematik.uni-stuttgart.de

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

CP12

Multi-Fidelity Modeling of Solar Irradiance

We describe our approach to enhancing the production-grade solar irradiance model, based on NASA/NOAA satellite imagery, by combining it with the real-time data from PV installations and weather sensors monitored by Locus Energy. The implications of leveraging ground data for thousands of monitored PV sites with the current short-term solar irradiance forecasting model, based on optical flow computations, is also discussed. The multi-fidelity forecasting model was calibrated and validated against SurfRad network's solar data.

Sergey Koltakov
Stanford University
koltakov@stanford.edu

CP12

Efficient Reduced Order Models for Solving Problems in Fluid Mechanics Using Stochastic Collocation

In this work, we propose a novel computational algorithm that employs proper orthogonal decomposition (POD) and sparse grid stochastic collocation method to approximate the solution to stochastic partial differential equations describing fluids. In particular, we will describe the stability and accuracy of the proposed parallel algorithm leading to a reduction of the overall computational cost. Numerical results validating the performance of the proposed method will be presented for benchmark problems.

Maziar Raissi
Mathematical Sciences
George Mason University
mraissi@gmu.edu

Padmanabhan Seshaiyer
George Mason University
pseshaiy@gmu.edu

CP12

Efficient New Greedy Algorithms for Reduced Basis Methods

We propose two new algorithms to improve greedy sampling of high-dimensional functions. While the techniques have a substantial degree of generality, we frame the discussion in the context of reduced basis methods for high-dimensional parametrized functions. The first algorithm, based on a saturation assumption of the error in the greedy algorithm, is shown to result in a significant reduction of the workload over the standard greedy algorithm. In a further improved approach, this is combined with an algorithm in which the train set for the greedy approach is adaptively sparsefied and enriched. A safety check step is added at the end of the algorithm to certify the quality of the sampling.

Shun Zhang, Jan Hesthaven
Division of Applied Mathematics
Brown University
Shun_Zhang@brown.edu, jan_hesthaven@brown.edu

Benjamin Stamm
University of California, Berkeley, US
stamm@math.berkeley.edu

MS1

Discontinuous Galerkin Methods in Convection Dominated Application Models

We present discontinuous Galerkin (DG) methods arising in convection dominated application models. The first model is the DG Advanced Circulation model for hurricane storm surge, where recent additions to the code include layered dynamic sediment, local time-stepping algorithms, and non-hydrostatic pressure formalisms. The second model we address, is a reduced magnetohydrodynamic code for solving convective scrape-off layer dynamics in tokamak edge plasmas.

Craig Michoski
UT Austin
michoski@ices.utexas.edu

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

Kyle T. Mandli
University of Texas at Austin
ICES
kyle@ices.utexas.edu

Francois Waelbroeck
UT IFS
Austing
flw@mail.utexas.edu

MS1

Numerical Simulation of Cylindrical Solitary Waves in Periodic Media

We study the behavior of nonlinear waves in a two-dimensional medium with material properties that vary periodically in space. The system considered has no dispersive terms; however, the material heterogeneity leads to reflections and effective dispersion. This dispersion combined with the nonlinearity of the system breaks the initial profile into solitary waves. Interaction of two solitary waves is studied. The results are typical of solitary waves; the two pulses retain their identity after the interaction.

Manuel Quezada De Luna
KAUST
mquezada@math.tamu.edu

MS1

Numerical Inverse Scattering: Uniformly Accurate Resolution of Dispersion

The Riemann–Hilbert formulations of the Korteweg-de Vries and Nonlinear Schrödinger equations have proved to be numerically valuable. Borrowing ideas from the method of nonlinear steepest descent, the resulting numerical schemes are seen to be asymptotically reliable. We derive sufficient conditions for a uniformly accurate numerical method. We resolve highly-oscillatory solutions of dispersive equations on unbounded domains with uniform accuracy.

Thomas D. Trogdon
Department of Applied Mathematics
University of Washington
trogdon@amath.washington.edu

Sheehan Olver
The University of Sydney
Sheehan.Olver@sydney.edu.au

MS1

Grid Resolution Requirements in Nonhydrostatic Internal Wave Modeling

Abstract not available at time of publication.

Sean Vitousek
Environmental Fluid Mechanics Laboratory
seanv@stanford.edu

MS2

Kalman Smoothing Approach for 4D Imaging

We present a formulation for modeling of large scale images that change dynamically in time. We use an optimization formulation of Kalman smoothing, combining a slowly varying dynamic prior with PDE constrained optimization used in 3D imaging. We discuss computational aspects and show how matrix free algorithms can be used to solve the problem.

Aleksandr Aravkin
the University of British Columbia
saravkin@eos.ubc.ca

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

MS2

4D Biomedical Imaging and Compressed Sensing

In state-of-the-art biomedicine, 4D time-evolving processes arise at different scales in space and time, e.g. researchers study cell migration in high-resolution/-speed microscopy or motion of organs (heart, lung,) in tomography. These problems offer interesting challenges for modeling and nonlinear large-scale PDE optimization. The goal of this talk is to present recent modeling and numerics for 4D imaging subject to transport equations. Moreover, we will focus on the design question of compressed sensing in space and time.

Christoph Brune
Department of Mathematics
University of California Los Angeles
christoph.brune@uni-muenster.de

MS2

Electromagnetic Imaging of Subsurface Flow

Abstract not available at time of publication.

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

MS2

Reliability of 4D Image Based Simulations of Blood

Flow in Arteries

Integration of 4D images and computational fluid dynamics is an effective approach for the numerical simulation of blood in moving arteries. The reliability of this approach depends on the quality of the images and the registration process tracking the wall motion in time. Here we discuss the accuracy of the hemodynamics simulations in a patient specific aorta for different sampling frequencies of the images. The relevance of appropriate selection of boundary conditions will be addressed.

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

Tiziano Passerini
Department of Math & CS
Emory University
tiziano@mathcs.emory.edu

Marina Piccinelli
Department of Math&CS
Emory University
marina@mathcs.emory.edu

MS3

An Overview of Silent Data Corruption and its Effects on Linear Solvers

Exascale machines are expected to challenge current HPC computing methods. Fault rates due to high component count will likely be exacerbated by decreasing component reliability. To help algorithm developers prepare for a future that may include a constant rate of faults, I will discuss current hardware trends, how they relate to existing failure modes, and how simulations of faults affect existing codes. I will end with a few glimmers of hope for future architectures.

Sean Blanchard
Los Alamos National Laboratory
seanb@lanl.gov

Nathan DeBardeleben
Los Alamos National Lab
ndebard@lanl.gov

MS3

Fault-tolerant Iterative Linear Solvers via Selective Reliability

Protecting computations and data from corruption due to hardware errors costs energy. However, energy increasingly constrains modern computers. As processor counts continue to grow, it will become too expensive to correct errors at system levels, before they reach user code. We will show instead that if the system provides a programming model that lets applications apply reliability only when and where needed, we can develop algorithms that compute the right answer despite hardware faults.

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

MS3

Fault Tolerant Qmr

As architectures are developed with smaller components operating a lower voltages, soft errors (i.e., bit flips) become potentially worrying, and in fact may not even be detected. Previous work has investigated the sensitivity several Krylov methods to soft errors. Here, we explore fault tolerance issues in QMR, compare its performance in the presence of bit flips, and introduce a fault tolerant variant of QMR.

Victoria Howle, Ashley Meek
Texas Tech
victoria.howle@ttu.edu, ashley.meek@ttu.edu

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

MS3

Algorithmic Strategies for Soft-error Resilience in Sparse Linear Solvers

Increasing rates of soft errors in many-core and multi-core architectures can adversely affect the accuracy and convergence of a wide variety of parallel numerical applications, such as, implicit and explicit methods for the solution of partial differential equations, the latter being particularly susceptible to instabilities arising from soft errors. In this talk we present methods to characterize the growth and propagation of soft errors in sparse iterative kernels and discuss light-weight strategies to recover from the resulting error states.

Radhika S. Saksena
Pennsylvania State University
rss29@psu.edu

Manu Shantharam
The University of Utah
shantharam.manu@gmail.com

Padma Raghavan
The Pennsylvania State Univ.
Dept of Computer Science Engr.
raghavan@cse.psu.edu

MS5

Fast Implicit Maxwell Solver for Linear Wave Propagation in Cold Plasmas

We present a new implicit Maxwell solver, based on the ‘method of lines transpose’ for the wave equation. The fields update satisfies a modified Helmholtz equation, which is solved by a Green’s function approach. Combining an ADI procedure with an accelerated 1D method, the computational cost of multi-dimensional solvers is linear in the number of unknowns. A fully implicit exponential discretization of the linear plasma response is employed, obtaining second order accuracy and unconditional stability.

Yaman Guclu
Department of Mathematics
Michigan State University
yguclu@math.msu.edu

Matthew F. Causley

New Jersey Institute of Technology (NJIT)
causleym@math.msu.edu

Andrew J. Christlieb, Yingda Cheng
Department of Mathematics
Michigan State University
christlieb@math.msu.edu, ycheng@math.msu.edu

MS5

Discontinuous Galerkin Schemes for the (Gyro) Kinetic Simulations of Plasmas

Recently developed discontinuous Galerkin (DG) algorithms are extended for the solution of kinetic and gyrokinetic simulations of plasmas. The schemes are applicable to a wide variety of systems describable by a Hamiltonian evolution equation, coupled to field equations. Conservation properties of the algorithms, specially for energy, momentum and L_2 norm are discussed. Special basis functions for velocity space are explored in the context of representing small deviations from a Maxwellian distribution efficiently.

Ammar Hakim
Plasma Physics Laboratory
Princeton University
ammar@princeton.edu

MS5

A High-order Unstaggered Constrained Transport Method for the 3D Ideal Magnetohydrodynamic Equations based on the Method of Lines

We present finite volume methods for the 3D ideal magnetohydrodynamic (MHD) equations using a constrained transport (CT) technique, in which an evolution equation for the magnetic potential is solved during each time step. A divergence free magnetic field is obtained by computing the curl of the magnetic potential. In contrast to the 2D case, the evolution equation for the magnetic potential is not unique but instead depends on the choice of the gauge condition. By using the so-called Weyl gauge, we obtain a weak hyperbolic system for the evolution of the magnetic potential that requires an appropriate discretization.

Christiane Helzel
Department of Mathematics
Ruhr-University-Bochum
christiane.helzel@rub.de

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

Bertram Taetz
Department of Mathematics
Ruhr-University Bochum
bertram.taetz@rub.de

MS5

Two-Fluid Higher-Moment Modeling of Fast Magnetic Reconnection

Higher-moment fluid models efficiently approximate kinetic models for moderate Mach and Knudsen numbers. We seek minimal modeling requirements to resolve fast magnetic reconnection with a fluid model. For 2D problems invariant under 180-degree rotation about the out-of-plane axis, nonsingular steady-state magnetic reconnection is impossible in a conservative fluid model that lacks heat flux. To avoid the difficulties of diffusive closures, we use higher-moment hyperbolic models and compare simulations of magnetic reconnection with kinetic simulations.

Evan A. Johnson
UW-Madison
e.alec.johnson@gmail.com

MS6

Parallel Multiscale Modeling of High Explosives for Transportation Accidents

The Uintah Computational Framework provides a scalable architecture for simulating fires and detonation of large arrays of explosives in transportation accidents. The simulations utilize Kraken for these complex fluid-structure interaction problems. Uintah combines validated reaction models with robust multi-material mechanics and scaling to 250k cores to provide the predictive capability needed to improve transportation safety on the nation's highways and railways for the first time.

Jacqueline Beckvermit
Department of Chemsitry
University of Utah
beckvermit@gmail.com

Qingyu Meng
SCI Institute
Univeristy of Utah
qymeng@cs.utah.edu

Todd Harman
Department of Mechanical Engineering
University of Utah
t.harman@utah.edu

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

Charles Wight
Dept. of Chemistry
University of Utah
chuck.wight@utah.edu

MS6

Simulations of the Universe: Toward Petascale Cosmology

Numerical simulations are now the dominant tool in theoretical cosmology. As advances in computer power and software algorithms continue into the petascale regime, it becomes possible to simulate the entire visible universe at high mass and spatial resolutions. We will present a petascale-optimized version of the gravity/hydrodynamics code GADGET using an efficient and low-risk incremental strategy. Hybrid shared memory, load balancing and other improvements allow the new p-GADGET to run on petascale systems.

Tiziana DiMatteo
Department of Physics
Carnegie Mellon University
tiziana@lemo.phys.cmu.edu

MS6

Petascale Challenges in Astrophysics: Core-Collapse Supernovae

Understanding the explosive deaths of massive stars has been a major focus in stellar astrophysics for five decades. Among many other roles, these explosions are the source of many of the elements essential for life. Despite ∼50 years of concerted effort, we still do not know how massive stars explode. I will review the multi-physics, multi-scale, multi-dimensional challenges of core-collapse supernova modeling and describe the massively parallel, state-of-the-art efforts to unveil the mechanism of explosion.

Joshua C. Dolence
Department of Astrophysical Sciences
Princeton University
jdolence@astro.princeton.edu

MS6

Correlated Electronic Wave-function Calculations for Large Molecules at the Petascale Level

We describe how correlated electronic wave function calculations on large molecular systems may be carried out using a massively parallel algorithm. The presented algorithm employs two levels of parallelization and allows for efficient parallization over many thousand nodes for large molecule systems. With the presented algorithm the high-order polynomial scaling of a standard correlation wave functions calculations is reduced to linear, while error control compared to a conventional calculation is maintained.

Kasper Kristensen
Department of Chemistry
Aarhus University
kasperk@chem.au.dk

MS7

Random Light Polarization Dynamics in an Active Optical Medium

Resonant interaction of light with a randomly-prepared, lambda-configuration active optical medium is described by exact solutions of a completely-integrable, system of random partial differential equations, thus combining the opposing concepts of integrability and disorder. An optical pulse passing through such a material will switch randomly between left- and right-hand circular polarizations. Exact probability distributions of the electric-field envelope variables describing the light polarization and their switching times will be presented together with stochastic simulations.

Gregor Kovacic
Rensselaer Polytechnic Inst
Dept of Mathematical Sciences
kovacg@rpi.edu

Katherine Newhall
Courant Institute of Mathematical Science
New York University
newhall@cims.nyu.edu

Peter R. Kramer
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
kramep@rpi.edu

Ildar R. Gabitov
Department of Mathematics, University of Arizona
gabitov@math.arizona.edu

Ethan Akins
University of California, Berkeley
atkins@cims.nyu.edu

MS7

Thermally Induced Magnetization Reversals

Driving nanomagnets by spin-polarized currents offers exciting prospects in magnetoelectronics, but the actual response of the magnets to such currents remains poorly understood. The underdamped dynamics require exceedingly long simulations of stochastic trajectories to compute thermally-induced reversal times and other relevant features. Rather, we derive an averaged equation describing the diffusion of energy on a graph that could be used to analyze the behavior of a broad range of other nongradient, underdamped systems.

Katherine Newhall
Courant Institute of Mathematical Science
New York University
newhall@cims.nyu.edu

Eric Vanden-Eijnden
Courant Institute
NYU
eve@cims.nyu.edu

MS7

Limitations in Reduction of Wind Power Intermittency with Storage Technologies

Stochastic variations and unpredictability of wind energy are the major concerns of power industry and hinder the wide scale adoption of wind power. Compensation of short term variability is one of the major challenges that the industry will face in the coming years. Our study focuses on statistical analysis of fluctuations of wind power on the minute to hour time scales. Using the publicly available wind measurement data we show that the statistics of fluctuations is strongly non-Gaussian and highly correlated in this time frame. Specifically we show that traditional Gaussian can underestimate the probability of rare events by several orders of magnitude. In the second part of our work we analyze the potential impact of advanced control and storage technologies in reducing the intermittency of wind power. Using the convex optimization techniques we study the theoretical limits on the performance of storage technologies. Specifically we analyze the interplay between the statistics of electric power fluctuations and the characteristics of storage available in the system. We quantify the trade-off between the reduction in power intermittency, storage capacity, and charging rate. In the end we present a general approach to the intermittency mitigation problem that incorporates multiple objectives and system constraints.

Konstantin Turitsyn
Massachusetts Institute of Technology
turitsyn@mit.edu

MS7

Advances in Parallel Tempering

Parallel tempering (replica exchange Monte Carlo) is an important tool for exploring free-energy landscapes with many metastable minimas. For nearly-degenerate free-energy wells, the equilibration time is slow; it is diffusive in the number of replicas the system. Here we discuss benefits of MC schemes that violate detailed balance in accelerating the convergence of parallel tempering.

Marija Vucelja
Courant Institute
New York University
vucelja@cims.nyu.edu

Jon Machta
Department of Physics
University of Massachusetts Amherst
machta@physics.umass.edu

MS8

Stochastic Reduced--Order Models for Multi--Scale Simulation of Laser Weld Failure

Micro laser welds are ubiquitous in intricate engineering systems. These partial penetration welds result in sharp, crack-like notches at their root. Further, the geometry and constitutive behavior of these welds are complex and subject to significant variability and uncertainty. Modeling these welds in large systems is necessary to predict reliability; however, traditional modeling approaches cannot adequately resolve the sources of uncertainty. To this end, a surrogate model, based on stochastic reduced order models (SROMs) and detailed finite element solutions, is developed to represent the laser welds in the system-level models for the goal of achieving efficient and accurate estimates of system reliability. *Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energys National Nuclear Security Administration under contract DE-AC04-94AL85000.

John M. Emery
Sandia National Laboratories
jmemery@sandia.gov

Mircea Grigoriu
School of Civil Engineering
Cornell University
mdg12@cornell.edu

Richard Field
Sandia National Laboratories
rvfield@sandia.gov

MS8

Multifidelity Simulation of Large Scale Mass Flow

Large scale mass flows are modeled using a number of modeling assumptions to make the models computationally tractable. The first set of modeling assumptions are based on simplifications inherent in the physics (e.g. depth averaging) and numerical methodology. For ensemble based hazard analysis of such flows we create even simpler statistical surrogates (Gaussian process etc.). In this paper we will characterize the impact of each such assumption on the simulation outcome.

Abani K. Patra
SUNY at Buffalo
Dept of Mechanical Engineering
abani@eng.buffalo.edu

Bruce Pitman
Dept of Mathematics,
SUNY at Buffalo
pitman@buffalo.edu

Dinesh Kumar
Dept of Mechanical Engineering
SUNY at Buffalo
dkumar@buffalo.edu

Ramona Stefanescu
University at Buffalo
ers32@buffalo.edu

MS8

Continuum Scale Constitutive Laws Extracted from Atomistic Simulations Using Bayesian Inference and Uncertainty Quantification Methods

In atomistic-to-continuum simulations, it is crucial to derive constitutive relationships from atomistic simulations to compensate for the unresolved degrees of freedom. In this study, we extract the heat conduction law from molecular dynamics (MD) simulations, together with its associated uncertainty. This latter is due to the intrinsic noise in MD simulations. We use Bayesian inference to build the constitutive law. We present two approaches based on polynomial chaos expansions (parametric) and Gaussian processes(non-parametric), respectively. We then propagate the obtained constitutive law into a continuum scale simulation and compare with an equivalent MD simulation.

Maher Salloum
Sandia National Laboratories
Scalable and Secure Systems Research Department
mnsallo@sandia.gov

Jeremy Templeton
Sandia National Laboratories
jatempl@sandia.gov

MS8

Comparing Multiple Sources of Epistemic Uncertainty in Geophysical Simulations

Epistemic uncertainty – uncertainty due to a lack of model refinement – arises through assumptions made in physical models, numerical approximation, and imperfect statistical models. We will compare the effects of model parameter uncertainty to those introduced by numerical approximations of geophysical model output. With the objective of model-based geophysical hazard mapping in mind, we propose methodology which can account for these sources of uncertainty and quickly assess their impact on resulting hazard maps.

Elaine Spiller
Marquette University
elaine.spiller@marquette.edu

MS9

Optimal Control of Partial Differential Algebraic Equations with Application to Fuel Cell Plants

Realistic multiphysics models describing Molten Carbonate Fuel Cells (MCFC) incorporate several aspects from electrochemistry, reaction kinetics, fluid dynamics as well as heat convection and diffusion. This gives rise to a coupled system of Partial Differential and Algebraic Equations (PDAE) including 21 PDE of hyperbolic and parabolic type.
The ability of fast and save load changes is very important for stationary power plants based on MCFC. The task of controlling the gas inflows of the system to perform a fast load change constitutes on optimal control problem with PDAE and control constraints. Based on the gradient information provided by the solution of the adjoint PDAE system, we solve the problem with a quasi-Newton method of BFGS-type. We conclude with numerical experiments.

Armin Rund
Institute for Mathematics and Scientific Computing
University of Graz
armin.rund@uni-graz.at

MS9

Shape Calculus in Optics

Fabrication based perturbations, in nano-optics, can be expressed in terms of a PDE constraint shape optimization. The resulting minimization problem is solved with a steepest descent algorithm based on the method of mapping. Considering the regularity requirements of shape gradient, we study and compare the so called continuous and fully-discrete methods based on a finite element discretization. We investigate both theoretical and implementation aspects in the framework of a 2D transmission problem.

Sahar Sargheini
Department Information Technology and Electrical Engineering
ETH, Zurich
sargheini@ifh.ee.ethz.ch

Alberto Paganini
seminar for applied mathematics
ETH, Zurich
alberto.paganini@sam.math.ethz.ch

MS9

Multidisciplinary Shape and Topology Optimization

Shape optimization is a challenging issue of high practical interest. Mostly, it is treated in a single disciplinary context. But many industrial problems like aeroelastic design or instationary thermo-elastic design depend on a multidisciplinary framework. The coupling of multiple disciplines in shape and topology optimization leads to additional mathematical issues, which are to be solved in numerical approaches. Efficient implementations as well as results in practical applications will be presented.

Roland Stoffel
University of Trier
Trier, Germany
stoffel@uni-trier.de

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

MS9

Design Optimization for Wave Propagation Problems

By using adjoint based numerical deign optimization it is possible to design highly efficient components for professional audio. This talk presents an overview of the typical steps in the design optimization process. The design of loudspeaker horns and subwoofer boxes illustrates this process. Here, the Helmholtz equation models the wave propagation, shape and topology optimization algorithms design the components, and, as a final step, the components are build and analyzed in an anechoic chamber.

Eddie Wadbro
Department of Computing Science
Umeå University
eddiew@cs.umu.se

MS10

Adaptive Algorithms for Simulations on Extreme Scales

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. The focus of this talk will be to describe how these fast methods are adaptive to high performance computing.

Rick Archibald
Computational Mathematics Group
Oak Ridge National Labratory
archibaldrk@ornl.gov

MS10

Adaptive Strategies for Random Elliptic Partial Differential Equations

Solutions of elliptic boundary value problems with random operators admit efficient approximations by polynomials on the parameter domain. I will give an overview of adaptive strategies for constructing suitable spaces of polynomials and computing approximate solutions on these spaces, and discuss the merits of high-order finite elements for random partial differential equations.

Claude J. Gittelson
Purdue University
cgittels@purdue.edu

MS10

Adaptive Sequential Design for Efficient Emulation using Gaussian Processes

We provide enhancements of a sequential design approach to build Gaussian Process surrogates, where points are chosen according to the entropy-based mutual information criterion. Points are selected according to updated parameters of the covariance structure to boost efficiency. We prove that this method is close to optimal. We illustrate the benefits of our improvements on several synthetic data sets and on the investigation of parametrizations in the Whole Atmosphere Community Climate Model.

Joakim Beck, Serge Guillas
University College London
joakim.back.09@ucl.ac.uk, s.guillas@ucl.ac.uk

MS10

High Dimensional Multiphysics Metamodeling for Combustion Engine Stability

Cycle-to-cycle variations in power output of combustion engines is a major obstacle to increasing fuel efficiency, however, the causes of such variations are not fully understood. We take a multiphysics computationally expensive engine model that combines turbulence, thermodynamics and chemistry. Using high dimensional interpolation techniques, we create a cheap model approximation that is used to study the correlation between the various parameters of engine operations and the cycle-to-cycle power variations.

Miroslav Stoyanov
Florida State University
Department of Scientific Computing
mkstoyanov@gmail.com

Clayton G. Webster, Charles Finney
Oak Ridge National Laboratory
webstercg@ornl.gov, finneyc@ornl.gov

Sreekanth Pannala
Computer Science and Mathematics Division
Oak Ridge National Laboratory
pannalas@ornl.gov

Stuart Daw, Robert Wagner, Kevin Edwards, Johney Green
Oak Ridge National Laboratory
dawcs@ornl.gov, wagnerrm@ornl.gov, edwardskd@ornl.gov, greenjbjr@ornl.gov

MS11

A Fast Algorithm for Biharmonic Equation and Applications

We present an accurate fast algorithm to solve the Biharmonic equation within a unit disk of the complex plane. This accurate and fast algorithm is derived through exact analysis and properties of convolution integrals using Green's function method in complex plane. The fast algorithm has an asymptotic operation count $O(\ln N)$ per point. It has been applied to several Stokes' and Navier-Stokes flow problems.

Aditi Ghosh
Texas A&M University
amathematics@gmail.com

Prabir Daripa
Texas A&M University
Department of Mathematics
prabir.daripa@math.tamu.edu

MS11

A Parallel Fast Summation Algorithm for Volume Potentials

Traditional FMM based approaches for evaluating potentials due to continuous source distributions use large number of source points, require smoothing functions, and need frequent remeshing for time varying distributions. We instead use piecewise polynomials and adaptive octree to represent source distribution and output potential. Our method uses kernel independent FMM, enabling use of wide range of kernels (Laplace, Stokes, Helmholtz etc.). The implementation scales to the largest supercomputers available and demonstrates excellent per core performance.

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

MS11

BEM++ – a new C++/Python Library for Boundary-element Calculations

BEM++ is a new general-purpose open-source library for boundary-element calculations. All standard kernels for Laplace, Helmholtz and modified Helmholtz are available and more (Maxwell, elasticity) are being added. Accelerated assembly based on the adaptive cross approximation (ACA) algorithm is supported on CPUs via an interface to the AHMED library, and implementation of ACA on GPUs is ongoing. The library is tightly integrated with Trilinos to allow access to its iterative solver capabilities. Python wrappers allow users to rapidly develop boundary-element codes and easily visualise and postprocess results of their calculations. More information is available at www.bempp.org.

Wojciech Smigaj
Department of Mathematics
University College London
w.smigaj@ucl.ac.uk

Simon Arridge
Department of Computer Science
University College London
simon.arridge@cs.ucl.ac.uk

Timo Betcke
University College London
t.betcke@ucl.ac.uk

Joel Phillips
Department of Mathematics
University College London
joel.phillips@ucl.ac.uk

Martin Schweiger
Department of Computer Science
University College London
m.schweiger@cs.ucl.ac.uk

MS11

Fast Numerical Greens Functions with Application

to Magnetic Resonance Imaging

A common problem in a variety of electromagnetic applications is to determine fields from a variety of sources in the presence of a given scatterer. If there is an analytic representation for the Greens function given the scatterer, then the standard approach is to use that Greens function to simplify the field calculations. For general scatterer geometries, the Greens function must be computed numerically and tabulated, an intractable process in three space dimensions as the table must be six-dimensional. In this talk we describe combined SVD-plus-sampling method for constant-time Green's function evaluation, and then demonstrate its use in optimizing coil excitation in low-frequency Magnetic Resonance Imaging.

Amit Hochman
Research Lab. of Electronics
MIT
hochman@mit.edu

Jorge Fernandez Villena
Inst Engn Sistemas & Comp Invest & Desenvolviment, P-100002
jorge.fernandez@inesc-id.pt

L. Miguel Silveira
Univ Tecn Lisboa, IST, P-1000029 Lisbon, Portugal
lms@inesc-id.pt

Luca Daniel
M.I.T.
Research Lab in Electronics
luca@mit.edu

Jacob White
Dept. of Elec. Eng. and Comp. Sci and Res. Lab. Elec.
MIT
white@mit.edu

MS12
Why Julia?

Julia is a language designed from the beginning to efficiently execute the kinds of programs users write in Matlab, R, Python/SciPy, and similar environments. Performance is a major motivation, but there is also ongoing work to speed up existing languages. I will discuss what we might want other than performance, and what advantages we gain by designing a new language. Fortunately, performance does not trade off against all other desirable features.

Jeff Bezanson
Computer Science
MIT
jeff.bezanson@gmail.com

Stephan Karpinski
MIT
stefan@karpinski.org

Virah Shah
viral@mayin.org

Alan Edelman
ISC Inc.
edelman@math.mit.edu

MS12
Numba, Compiling Numpy Functions from Python with Llvm

Abstract not available at time of publication.

Travis Oliphant
Continuum Analytics
teoliphant@gmail.com

MS12
Rllvm and RLLVMCompile

R is increasingly used for both data analysis and also developing new statistical methods. There are various efforts to make this interpreted language faster, with technical and "practical" complications. Instead, we have created bindings to the LLVM libraries/APIs to allow R programmers compile their own code and parts of the R language or new similar DSLs. Rather than a "single" compiler within R, others can use the bindings to develop different compilation strategies that target run-time, memory usage and parallelism for different hardware, e.g GPUs, cores with shared memory and distributed memory clusters.

Duncan W. Temple Lang
Statistics
UC Davis
duncan@wal.ucdavis.edu

MS12
Introduction to JIT Compiling and Code Generation in Scientific Computing

Domain specific languages, code generation, and just-in-time compiling are all becoming standard part of scientific computing. Originally the target for high level languages was to produce more expressive ways of writing code over the traditional languages, i.e. Fortran and C/C++. Now these high level codes are starting to compete, and sometimes beat, low level languages using advanced compiler techniques. This introduction will overview the methods and tradeoffs used in modern high level scientific languages.

Andy R. Terrel
Texas Advanced Computing Center
University of Texas
aterrel@tacc.utexas.edu

MS13
A Narrow-band Gradient-Augmented Level Set Method for Incompressible Two-phase Flow

We have incorporated the gradient-augmented level set method (GALSM) for use in two-phase incompressible flow simulations by interpolating velocity values and introducing a re-initialization procedure. The method is conducted on a narrow band around the interface, reducing computational effort, while maintaining an optimally local advection scheme and providing sub-grid resolution. Ocean wave simulation is the primary motivation, and numerous benchmarks have been conducted, including a new comparison with wave tank data.

Curtis Lee, John Dolbow
Duke University
calee181@gmail.com, jdolbow@duke.edu

Peter J. Mucha
University of North Carolina
mucha@unc.edu

MS13

A Particle-enhanced Gradient-augmented Level Set Method

The original gradient-augmented level set methods use a projection on an interpolant space that is based only on the information provided on the grid. In this presentation, we show how we can incorporate information from Lagragian particles to improve the accuracy of the method while preserving the projection properties. We will address the questions of topological changes and show applications to multiphase Navier-Stokes.

Olivier Mercier
McGill University
oli.mercier@gmail.com

MS13

Model Equations for Jet-schemes

Model equations describe the behavior of numerical schemes in the asymptotic limit when the mesh-size vanishes, and provide information about the nature of the algorithm error in this limit. In particular, they can be used to detect the causes of long wave instabilities (if present), and devise ways to correct them. Because jet-schemes use an advect–and–project approach in function space, standard methods for obtaining model equations cannot be used. The numerical solution by a jet scheme implicitly carries (small amplitude) grid size structure (as given by the local polynomial interpolant) — modulated on a longer scale by the solution's variations. In this talk we will discuss how homogenization techniques can be used to obtain model equations for jet schemes, and apply it to some illustrative examples.

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

Benjamin Seibold
Temple University
seibold@temple.edu

Jean-Christophe Nave
McGill University
jcnave@math.mcgill.ca

MS13

An Overview of Gradient-augmented Methods and Jet Schemes

This talk provides an introduction to the philosophy and methodology of jet schemes and gradient-augmented level set methods, and an overview of how they relate to other computational approaches for advection and interface evolution problems.

Benjamin Seibold
Temple University
seibold@temple.edu

MS14

A Method for Non-Intrusive Model Reduction and Adjoint-based Optimization

We propose a method to enable existing model reduction methods (POD, DEIM, etc) and adjoint-based optimization without modification to the simulation source code. We infer a twin model to mirror the underlying model governing a PDE simulation. Our method features: Dimensionality reduction of inputs by exploiting physics invariance; Independence with the numerical scheme implemented in the original PDE simulation; A goal-oriented approach to minimize the discrepancy between the twin model and the original simulation model.

Han Chen
MIT
voila@mit.edu

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

Hector Klie
ConocoPhillips Company
Hector.M.Klie@conocophillips.com

MS14

Adaptive Proxies for Integrated Oil Field Production Optimization under Uncertainty

The treatment of geological parameter uncertainty in numerical reservoir simulation is necessary to qualify the results achieved. Similarly, the back-pressure effects and dynamics of a full transient system (from source to sink) can only be elicited by treating the interdependent boundary conditions between the component reservoir, surface and process models. However, as the resulting coupled simulation is then extremely costly to evaluate, adaptive proxy methods are required to effectively optimize the composite model under uncertainty.

Benoit Couet, Kashif Rashid, William Bailey
Schlumberger-Doll Research
couet1@slb.com, krashid@slb.com, wbailey@slb.com

MS14

Local-Global Model Reduction Techniques for Porous Media Flow Simulation

The development of algorithms for modeling and simulation of heterogeneous porous media poses challenges related to the variability of scales and requires efficient solution methodologies due to their large-scale nature. This is especially daunting in simulating unconventional reservoirs that are embedded in fractured media. In this talk, I will describe a local-global model reduction method that combines multiscale techniques with the reduced-order methods in a seamless fashion. The aim is to reduce the degrees of freedoms in the state-space by computing global reduced order models written on a coarser space.

Eduardo Gildin
Petroleum Engineering Department
Texas A&M University
eduardo.gildin@pe.tamu.edu

MS14

Controllability and Observability in Two-phase Porous Media Flow

Large-scale nonlinear flow simulation models are frequently used in oil and gas reservoir engineering to make decisions on well locations, depletion strategies, production scenarios, etc. The quality of these decisions is, among other aspects, determined by the controllability and observability properties of the reservoir model at hand. We use empirical Gramians to analyze the controllability and observability of two-phase reservoir flow. We conclude that the position of the wells and the dynamics of the front between reservoir fluids determine the controllability and observability properties of the state variables (pressures and saturations). Therefore, for fixed well positions, reduced-order models should focus on modeling the fluid front(s).

Jan Dirk Jansen
Delft University of Technology
Department of Geotechnology
j.d.jansen@tudelft.nl

Jorn van Doren
Delft University of Technology (TU Delft)
j.f.m.vandoren@tudelft.nl

Paul Van den Hof
TU Eindhoven and TU Delft
p.m.j.vandenhof@tue.nl

MS15

Fast and Reliable Trust-region Eigensolvers

Error resilient solvers are necessary to address the reduced reliability of future high-performance computing systems. Surrogate-based methods, including trust-region methods, are naturally suited for this context, because the analyses and mechanisms surrounding surrogate error are a starting point for addressing soft errors and expensive sanity checks. We present an implementation and analysis of a trust-region-based eigenvalue solver adapted to low reliability computing due to soft-errors or low precision types.

Christopher G. Baker
Oak Ridge National Laboratory
bakercg@ornl.gov

MS15

Parallel Implementations of the Trace Minimization Scheme TraceMIN for the Sparse Symmetric Eigenvalue Problem

Large scale sparse symmetric eigenvalue problems arise in many computational science and engineering applications. Often, the large size of these problems requires the development of eigensolvers that scale well on parallel computing platforms. We compare the effectiveness and robustness of our eigensolver for the symmetric generalized eigenvalue problem, the trace minimization scheme TraceMIN, against the well-known sparse eigensolvers in Sandia's Trilinos library. Our results show that TraceMIN is more robust and has higher parallel scalability.

Alicia Klinvex
Department of Computer Science
Purdue University
aklinvex@purdue.edu

Faisal Saied
Purdue University
fsaied@purdue.edu

Ahmed Sameh
Department of Computer Science
Purdue University
sameh@purdue.edu

MS15

Absolute Value Preconditioning for Symmetric Indefinite Linear Systems

We introduce a novel strategy, absolute value preconditioning, for constructing symmetric positive definite (SPD) preconditioners for linear systems with symmetric indefinite matrices. We consider a model problem with a shifted discrete negative Laplacian, and suggest a geometric multigrid (MG) preconditioner, where the inverse of the matrix absolute value appears only on the coarse grid, while operations on finer grids are based on the Laplacian. Our numerical tests demonstrate practical effectiveness of the new preconditioner. [http://arxiv.org/abs/1104.4530]

Andrew Knyazev
Mitsubishi Electric Research Laboratories
CU-Denver
knyazev@merl.com

Eugene Vecharynski
University of Colorado Denver
eugene.vecharynski@gmail.com

MS15

A MATLAB Interface for PRIMME for Solving Eigenvalue and Singular Value Problems

The software package PRIMME for solving large, sparse, Hermitian eigenvalue problems implements many state-of-the-art preconditioned eigenvalue iterative methods and embeds expert knowledge into dynamically choosing techniques. As stand-alone or through a SLEPc interface, its robustness, efficiency, and user-friendliness have made it the first choice for hundreds of groups around the world. First, we present primme_eigs(), a MATLAB interface to PRIMME's full functionality with an interface as simple as eigs(). Second, we present primme_svds(), which solves the same SVD problem as svds() but with the power of PRIMME's methods. For smallest singular values it is possible to dynamically switch between A^TA and augmented matrix techniques.

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

Lingfei Wu
Department of Computer Science
College of William & Mary
lwu@email.wm.edu

MS16

Computing Least Squares Condition Numbers

The notion of condition number provides us with a theoretical framework to measure the numerical sensitivity of a problem solution to perturbations in its data. We derive condition numbers and estimates for linear least squares and total least squares problems using various metrics to measure errors. We present numerical experiments to compare exact values and statistical estimates. We also propose performance results using new routines on top of the multicore-GPU library MAGMA.

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

MS16

Numerical Issues in Testing Linear Algebra Algorithms

We discuss a variety of numerical issues that arise in testing linear algebra algorithms. These include how to choose an error measure (backward, forward, normwise, or componentwise) and interpret the observed errors; why tiny normwise relative errors (orders of magnitude smaller than the unit roundoff) sometimes appear and how to deal with them when producing performance profiles; and the implications for testing of a possible lack of bitwise reproducibility of results on HPC systems.

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

Nicholas Dingle
School of Mathematics, University of Manchester
nicholas.dingle@manchester.ac.uk

MS16

Accuracy and Stability Issues for Randomized Algorithms

The advantage of randomized algorithms is speed and simplicity. For very large problems, they can be faster than deterministic algorithms, and they are often simple to implement. We will discuss the numerical sensitivity and stability of randomized algorithms, as well as the error due to randomization, and the effect of the coherence and leverage scores of the matrix. Algorithms under consideration include matrix multiplication, least squares solvers, and low-rank approximations.

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

MS16

Towards a Reliable Performance Evaluation of Accurate Summation Algorithms

Recent floating point summation algorithms compute the best accurate sum even for arbitrary ill-conditioned problems. So the run-time efficiency of these algorithms becomes the criteria to decide which is the best. Neither the classic flop count nor experimental timings are relevant measures of the actual performance of such core numerical algorithms. We justify these claims and we present a reliable performance evaluation of some accurate summation algorithms thanks to an automatic instruction level parallelism analysis.

Philippe Langlois, Bernard Goossens, David Parello
University of Perpignan
France
langlois@univ-perp.fr, goossens@univ-perp.fr, david.parello@univ-perp.fr

MS17

Mortar Methods for Flow in Heterogeneous Porous Media

For an elliptic problem with a heterogeneous coefficient in mixed form, nonoverlapping mortar domain decomposition is efficient in parallel if the mortar space is small. We define a new multiscale mortar space using homogenization theory. In the locally periodic case, we prove the method achieves optimal order error estimates in the discretization parameters and heterogeneity period. Numerical results show that the method works well even when the coefficient is not locally periodic.

Todd Arbogast
Dept of Math; C1200
University of Texas, Austin
arbogast@ices.utexas.edu

Hailong Xiao
ICES
The University of Texas at Austin
xiaohl@ices.utexas.edu

MS17

A Multiscale Discontinuous Galerkin Method for the Schrodinger Equation in the Simulation of Semiconductor Devices

We develop and analyze a multiscale discontinuous Galerkin method for the stationary Schrödinger equation in the simulation of nanoscale semiconductor structures. The solution of the Schrödinger equation has a small wave length and oscillate at much smaller space scale for high electron energies. We incorporate the oscillatory behavior of the solution into the multiscale basis of the discontinuous Galerkin finite element method so that the problem can be solved on coarser meshes.

Bo Dong
University of Massachusetts Dartmouth
bdong@umassd.edu

Wei Wang
Department of Mathematics and Statistics
Florida International University
weiwang1@fiu.edu

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

MS17

A Multiscale Mixed Method Based on a Nonoverlapping Domain Decomposition Procedure

We use a nonoverlapping iterative domain decomposition procedure based on the Robin interface condition to develop a multiscale mixed method for calculation of the velocity field in heterogeneous porous media. Hybridized mixed finite elements are used for the spatial discretization of the equations. We define local, multiscale mixed basis functions to represent the discrete solutions in the subdomains. In the numerical approximations, subspaces of the vector space spanned by these basis functions can be appropriately chosen, which determines the balance between numerical accuracy and efficiency. Several numerical experiments are discussed to illustrate the important features of the procedure.

Alexandre Francisco
Universidade Federal Fluminense
afrancisco@metal.eeimvr.uff.br

Victor E. Ginting
Department of Mathematics
University of Wyoming
vginting@uwyo.edu

Felipe Pereira
Center for Fundamentals of Subsurface Flow
University of Wyoming
lpereira@uwyo.edu

Joyce Rigelo
University of Wyoming
jrigelo@uwyo.edu

MS17

Relaxing the CFL Number of the Discontinuous Galerkin Methods

Discontinuous Galerkin methods applied to hyperbolic problems have a CFL number that is inversely proportional to the order of approximation. We propose a modification to the scheme which allows for a significant time step increase while preserving spatial accuracy. We further discuss how a CFL can be relaxed on non-uniform grids. These algorithms can be especially useful when discontinuities and under-resolved scales are present in the solution.

Lilia Krivodonova
University of Waterloo
lgk@math.uwaterloo.ca

MS18

A Systematic Construction of Superconvergent HDG Methods

Abstract not available at time of publication.

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

Ke Shi
University of Minnesota, MPLS
shixx075@math.umn.edu

MS18

Overview of the Discontinuous Petrov-Galerkin Methods

Abstract not available at time of publication.

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

MS18

Effect of Inexact Test Functions in the DPG Method

The discontinuous Petrov Galerkin (DPG) method uses standard trial spaces for approximating solutions of boundary value problems. However, its test functions are locally defined so as to get optimal inf-sup stability. Although it is now clear how to apply hybridization techniques to localize the problem of computing these optimal test functions to a single mesh element, the resulting problem on one element can often be infinite dimensional. We prove that replacing this infinite dimensional problem by a practical finite dimensional problem does not affect convergence rates for many examples. Thus "inexactly" computed test spaces are enough in the DPG framework. The analysis proceeds by adapting an abstract Fortin operator to the Petrov-Galerkin context. The same operator also turns out to be an essential tool in a posteriori error analysis of a natural residual-based estimator provided by the DPG formalism. Recent results in this direction will also be presented.

Jay Gopalakrishnan
Portland state university
gjay@pdx.edu

MS18

Hybridizable Discontinuous Galerkin Methods for Wave Propagation Problems at High Frequency

We present a class of hybridizable discontinuous Galerkin (HDG) methods for wave propagation problems. The methods are fully implicit and high-order accurate in both space and time, yet computationally attractive owing to their distinctive features. First, they reduce the global degrees of freedom to the numerical trace. Second, they provide optimal convergence rates for both the field variable and their gradient. Third, they allow us to improve the numerical solution by performing a simple local post-processing.

Cuong Nguyen, Jaime Peraire
Massachusetts Institute of Technology
cuongng@mit.edu, peraire@MIT.EDU

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

MS19

An Immersed Finite Element Method for Fluid-Structure Interaction Problems and Applications to Hydrocephalus

We present a fully variational formulation of an immersed method for FSI problems based on the finite element method. Our formulation does not require Dirac-δ distributions. The immersed solid can be viscoelastic of differential type or hyperelastic, but is not otherwise restricted in this constitutive class. Our discrete formulation is shown to have the natural stability estimates of the companion continuum problem. We include standard test cases and

some preliminary applications to hydrocephalus.

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

Saswati Roy
Department of Engineering Science and Mechanics
The Pennsylvania State University
sur164@psu.edu

Francesco Costanzo
Department of Engineering Science and Mechanics
Pennsylvania State University
costanzo@engr.psu.edu

MS19

An Immersed Boundary Energy-Based Method for Incompressible Viscoelasticity

Fluid-structure interaction problems, in which the dynamics of a deformable structure is coupled to the dynamics of a fluid, are prevalent in biology. For example, the heart can be modeled as an elastic boundary that interacts with the blood circulating through it. The immersed boundary method is a popular method for simulating fluid-structure problems. The traditional immersed boundary method discretizes the elastic structure using a network of springs. This makes it difficult to use material models from continuum mechanics within the framework of the immersed boundary method. In this talk, I present a new immersed boundary method that uses continuum mechanics to discretize the elastic structure, with a finite-element-like discretization. This method is first applied to a warm-up problem, in which a viscoelastic incompressible material fills a two-dimensional periodic domain. Then we apply the method to a three-dimensional fluid-structure interaction problem.

Dharshi Devendran
University of Chicago
pdevendran@gmail.com

MS19

An Approach to using Finite Element Elasticity Models with the Immersed Boundary Method

The immersed boundary (IB) method treats problems in which a structure is immersed in a fluid flow. The IB method was first introduced for the case in which the structure is composed of systems of elastic fibers, but subsequent extensions have treated increasingly general mechanics models. This talk will describe one version of the IB method that uses elasticity models approximated via standard finite element methods, and will describe applications to cardiac and cardiovascular mechanics.

Boyce Griffith
New York University School of Medicine
griffith@cims.nyu.edu

MS19

Augmenting the Immersed Boundary Method with Rbfs: Applications to Modeling of Platelets in Hemodynamic Flows

The Immersed Boundary (IB) Method is a widely-used numerical methodology for the simulation of fluid-structure interaction problems. Most IB simulations represent their structures with piecewise-linear approximations and utilize Hookean spring models to approximate structural forces. Our specific motivation is the modeling of platelets in hemodynamic flows. In this talk, we present our attempts at using Radial Basis Functions (RBFs) as the means by which to represent platelets in an Immersed Boundary simulation. We will report results for two-dimensional and three-dimensional platelet flows. Furthermore, we will report other possible extensions/additions that can be made in the Immersed Boundary context using RBFs.

Robert Kirby
University of Utah
kirby@sci.utah.edu

Varun Shankar
School of Computing
University of Utah
shankar@cs.utah.edu

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

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

MS20

$\mathcal{H}_\in$-Optimal Reduced Models for Structured Systems

We present new necessary conditions for $\mathcal{H}_\in$-optimal approximation of linear dynamical systems by reduced order systems having additional special structure such as second order systems or port-Hamiltonian systems. This gives rise to distributed interpolatory conditions similar to those arising in weighted $\mathcal{H}_\in$ model reduction which suggests, in turn, natural computational strategies for resolution.

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

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

MS20

Preserving Lagrangian Structure in Nonlinear Model Reduction

Nonlinear mechanical systems described by Lagrangian dynamics arise in structural dynamics and molecular dynamics. Standard nonlinear model-reduction methods (e.g., discrete empirical interpolation) destroy Lagrangian structure, and with it critical properties such as energy conservation and symplecticity. This talk presents an efficient nonlinear model-reduction methodology that preserves Lagrangian structure. The method approximates the system's 'Lagrangian ingredients': the Riemannian metric, potential-energy function, dissipation function, and external force. These approximations preserve salient properties while ensuring computational efficiency.

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

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

Paul Boggs
Sandia National Lab
ptboggs@sandia.gov

MS20

Passivity-preserving Algorithm for Multiport Parameterized Modeling in the Frequency Domain

Given a collection of frequency response datasets, swept over design and geometrical parameters of interest, we identify a closed-form parameterized dynamical model using constrained fitting. The parameterized models are developed by polynomial or rational multivariate approximation between non-parameterized transfer matrices. The user of the generated models will be able to instantiate reduced models guaranteed to be stable and passive for any values of design parameters. Examples will include systems with multiple ports and parameters.

Zohaib Mahmood
MIT
zohaib@mit.edu

Luca Daniel
M.I.T.
Research Lab in Electronics
luca@mit.edu

MS20

Hierarchical Structure-Preserving Phase Modelling of Oscillators

It has long been known that coupled oscillators, when synchronized, respond to external inputs as a single "macro' oscillator. Reliance on this phenomenon is common in many domains, including biology (circadian rhythms, heart pacemaking systems) and nanotechnology (eg, spin-torque based microwave sources). To our knowledge, no proof of this remarkable phenomenon has been available. We provide such a proof, along with effective computational methods for extracting quantitative characterizations of "macro' oscillators from its individual.

Jaijeet Roychowdhury
University of California at Berkeley
jr@berkeley.edu

MS21

Parallelization and Performance Issues in Adaptive Simulation of Wave Propagation

In this presentation, we will discuss questions of computational performance that occur in the adaptive simulation of wave propagation problems - in our case, parallel adaptive mesh refinement techniques for tsunami simulation and simulation of seismic wave propagation on unstructured meshes. For tsunami simulation, we will present parallelization approaches (based on the Sierpinski space-filling curve) for efficient dynamical refinement and coarsening of structured, but fully adaptive triangular grids. For example, we use dynamical scheduling of subpartitions, together with a split&join approach, to allow varying computational load of partitions and even varying number of compute cores. For SeisSol, a Discontinuous-Galerkin-based package for seismic wave propagation and dynamic rupture simulation, we will present first results on using code generation approaches to speed up the basic computational kernels. We will also discuss options to improve memory access and load balancing of the adopted ADER-DG method.

Michael Bader, Alexander Breuer
Technische Universität München
Department of Informatics
bader@in.tum.de, breuera@in.tum.de

Alexander Heinecke
TU München
heinecke@in.tum.de

Martin Schreiber
Technische Universität München
Department of Informatics
martin.schreiber@in.tum.de

Christian Pelties
LMU Muenchen
pelties@geophysik.uni-muenchen.de

MS21

Finite-Volume and Discontinuous Galerkin Algorithms for Multifluid Simulations of Plasmas

Plasmas display a rich variety of behavior, from kinetic to fluid, and span vast spatial and temporal scales. This talk will give an overview of explicit finite-volume and discontinuous Galerkin algorithms for the solution of multi fluid plasma models. The framework of weakly nonlinear hyperbolic equations will be developed and applied to parasitic decay of waves from parametric resonance arising from quadratic nonlinearities. The application of this theory to understanding losses in radio-frequency heating of plasmas will be discussed.

Ammar Hakim
Plasma Physics Laboratory
Princeton University
ammar@princeton.edu

MS21

Simulation of Poroelastic Wave Propagation in CLAWPACK for Geophysical and Medical Applications

We use the CLAWPACK (Conservation LAWs PACKage) finite volume method code to solve Biot's equations for dynamics of a porous, fluid-saturated elastic medium. These equations were developed to model fluid-saturated rock formations, but are also applicable to other porous solids, such as in vivo bone. At low frequency Biot's equations are a system of hyperbolic PDEs with a relaxation source term, which may be stiff depending on the time scales associated with wave propagation. This talk gives a brief introduction to the class of numerical methods used, including issues associated the incorporation of the stiff source term. Numerical results are shown Cartesian grids, comparing against recent discontinuous Galerkin results for geophysical problems, and on logically rectangular mapped grids capable of modeling moderately complex geometry similar to bones.

Grady I. Lemoine
Department of Applied Mathematics
University of Washington
gl@amath.washington.edu

MS21

Parallelization of Hybridizable Discontinuous Galerkin Methods for a Nonhydrostatic Free Surface Primitive Equation Ocean Model

Higher-order finite element methods are becoming increasingly common in ocean modelling. Hybridizable discontinuous Galerkin (HDG) methods are attractive because they are highly parallelizable, and the high computational cost associated with standard DG and LDG methods can be alleviated. Studies on parallel numerical schemes for solving the 2D/3D nonhydrostatic Navier–Stokes equations coupled to a free surface equation are presented, along with accuracy and performance implications.

Chris Mirabito, Mattheus Ueckermann, Patrick Haley, Pierre Lermusiaux
Massachusetts Institute of Technology
mirabito@mit.edu, mpuecker@mit.edu, phaley@mit.edu, pierrel@mit.edu

MS22

Optimal Low Rank Matrix Inverse Approximation for Image Processing

In many applications, the desired solution of an inverse problem can be well-represented using only a few vectors of a certain basis, e.g., the singular vectors. In this work, we design an optimal low-rank matrix inverse approximation by incorporating probabilistic information from training data and solving an empirical Bayes risk minimization problem. We propose an efficient update approach for computing a low-rank regularized matrix, and provide numerical results for problems from image processing.

Julianne Chung
Virginia Tech
jmchung@vt.edu

Matthias Chung
Department of Mathematics
Texas State University
mcchung@vt.edu

MS22

Windowed Regularization for Image Deblurring via Operator Approximation

Recent work by Chung, Easley and O'Leary (*SIAM J. Sci. Comput.*, 2011), has shown that by windowing the components in the spectral domain of the blurring operator and employing different regularization parameters for each window, superior image restorations can be obtained. We consider suitable surrogate representations of the blurring operator that allow us to make this new regularization technique, as well as other known methods, applicable when the SVD cannot be computed.

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

Julianne Chung
Virginia Tech
jmchung@vt.edu

Dianne P. O'Leary
University of Maryland, College Park
Department of Computer Science
oleary@cs.umd.edu

MS22

Parallel Algorithms for Iterative Image Reconstruction

Abstract not available at time of publication.

James G. Nagy
Department of Mathematics & Computer Science
Emory University
nagy@mathcs.emory.edu

MS22

Tracking Objects Using 3D Edge Detectors

Edge detection determines the boundary of objects in an image. In some applications a sequence of images records a 2D representation of a scene changing over time, giving 3D data. New 3D edge detectors, particularly ones we developed using shearlets and hybrid shearlet-Canny algorithms, identify edges of complicated objects much more reliably than standard approaches, especially under high noise conditions. We also use edge information to derive position and velocity via an optimization algorithm.

Dianne P. O'Leary
University of Maryland, College Park
Department of Computer Science
oleary@cs.umd.edu

Glenn Easley
System Planning Corporation
geasley@sysplan.com

David Schug
Department of Mathematics
University of Maryland at College Park
david.schug@gmail.com

MS23

Flexible BiCGStab and Its Use in Practice

BiCGStab is one of the de facto methods of choice for solving linear systems in many application domains. Motivated by recent development of flexible preconditioners in high performance computing, we study BiCGStab under flexible preconditioning. In this talk, we will present some analysis results of the convergence behavior of flexible BiCGStab, and show its successful use in practice. In an important application, PFLOTRAN, we demonstrate that the run time of flexible BiCGStab preconditioned by multigrid is significantly improved over the currently known fastest record.

Jie Chen
Argonne National Laboratory
jiechen@mcs.anl.gov

Lois Curfman McInnes
Argonne National Laboratory

Mathematics and Computer Science Division
curfman@mcs.anl.gov

Hong Zhang
Argonne National Lab
hzhang@mcs.anl.gov

MS23

Krylov Solvers for Singular Hermitian, Complex Symmetric, and Skew Hermitian Linear Systems

Most existing Krylov subspace algorithms for linear systems assume non-singularity of the matrices or operators. MINRES-QLP (Choi, Paige, and Saunders 2011) is a MINRES-like algorithm but designed for computing the unique pseudoinverse solution of a singular symmetric or Hermitian problem using short recurrence relations in solution updates and stopping conditions. On a nonsingular system, MINRES-QLP is more stable than MINRES (Paige and Saunders 1975). We design similarly stable algorithms for singular (or nonsingular) complex-symmetric, skew-symmetric, and skew-Hermitian linear systems or linear least-squares problems. Our goal is to provide one efficient implementation prototyped in Matlab for these different classes of linear systems. We present extensive numerical experiments to demonstrate the scalability and robustness of these algorithms.

Sou-Cheng Choi
Stanford University
sctchoi@uchicago.edu

MS23

LAMG: Fast Multilevel Linear Solver and Eigensolver of the Graph Laplacian

Graph Laplacians arise in large-scale applications including machine learning, clustering, transportation networks, and CFD. We present Lean Algebraic Multigrid (LAMG): a linear solver of $Ax = b$, where A is a graph Laplacian. It combines novel methodologies: piecewise-constant interpolation; aggregation via a novel node proximity measure; multilevel acceleration; and perturbed Galerkin coarsening. LAMG is demonstrated to scale linearly with the number of edges for real-world graphs with up to 47,000,000 edges. Finally, we present a linear-scaling LAMG-based eigensolver for computing several of the lowest eigenpairs of A.

Oren E. Livne
Department of Human Genetics
University of Chicago
livne@uchicago.edu

Achi Brandt
Weizmann Institute of Science
Microsoft, Inc.
acbrandt@microsoft.com

MS23

Flexible and Multi-Shift IDR Algorithms for Solving Large Sparse Linear Systems

The Induced Dimension Reduction algorithm, IDR(s), is one of the most efficient methods for solving large sparse nonsymmetric linear systems of equations. We present two useful extensions of IDR(s), namely a flexible variant and a multi-shift variant. The algorithms exploit the underlying Hessenberg decomposition computed by IDR(s) to generate basis vectors for the Krylov subspace. The approximate solutions are computed using a quasi-minimization approach. Numerical examples are presented to show the effectiveness of these IDR(s) variants compared to existing ones and to other Krylov subspace methods.

Martin B. van Gijzen
Delft Inst. of Appl. Mathematics
Delft University of Technology
m.b.vangijzen@tudelft.nl

Gerard Sleijpen
Department of Mathematics
Utrecht University
sleijpen@math.uu.nl

Jens-Peter M. Zemke
Inst. of Numerical Simulation
TU Hamburg-Harburg
zemke@tu-harburg.de

MS24

Solution Methods for Phase Field Methods with Arbitrary Numbers of Variables

A unique characteristic of some phase field models, for example typical phase field models of grain growth, is that they can have an arbitrary number of variables, depending on the the number of grains being modeled. Finite element approaches typically focus on problems with a large number of nodes, but only between one to five variables. In this work, we present techniques for solving systems of equations with anywhere from two to 1000 variables, focusing on determining the method that results in the shortest computation time and the least memory usage. This effort is a part of the MARMOT phase field framework.

Derek Gaston, Michael Tonks
Idaho National Laboratory
derek.gaston@inl.gov, michael.tonks@inl.gov

MS24

Optimal Control of a Semi-discrete Cahn-Hilliard / Navier-Stokes System

Optimal boundary control of a time-discrete Cahn-Hilliard Navier- Stokes system is studied. For the Cahn-Hilliard part of the system, a general class of free energy potentials is considered including the double-obstacle potential. The existence of an optimal solution to the time discrete control problem as well as an approximate version thereof is established. Moreover a C-stationarity system for the original problem is derived and numerically realized.

Michael Hintermueller
Humboldt-University of Berlin
hint@math.hu-berlin.de

MS24

Time-Stepping Methods for Phase Field Models

Phase field models are flexible approaches to represent complex materials microstructure evolution. This comes, however, at the cost of including singular terms such as infinite potential and logarithmic components for expressing certain phenomena, which pose significant difficulties for numerical approximation schemes. In this work we describe time-stepping approaches that accommodate such singularities, including ones that use variational inequalities in the subproblem. We present analytical and numerical results to demonstrate the properties of such methods.

Jungho Lee, Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
julee@mcs.anl.gov, anitescu@mcs.anl.gov

MS24

Computational Challenges Posed by Phase Field Methods

The phase field method is a novel approach to solving free-boundary problems without explicitly tracking the location of an interface. However, the diffuse interface used in the method introduces a new length scale that presents significant computational challenges when performing three-dimensional simulations over realistic length and time scales. Even more challenging is the recently developed phase field crystal method. A survey of the computational challenges posed by phase field methods will be given.

Peter Voorhees
Northwestern University
Dept. of Material Science and Engineering
p-voorhees@northwestern.edu

MS25

Soft Error Resilience for One-sided Dense Linear Algebra Algorithms

Soft errors threaten computing systems by producing data corruption that's hard to detect and correct. Current research of soft-error resilience for dense linear solvers offers limited capability on large-scale systems, and suffers from both soft error and round-off error. This work proposes a fault tolerant algorithm that can recover the solution of a dense linear system from multiple spatial and temporal soft errors. Experimental results on Kraken confirm scalability and negligible overhead in solution recovery.

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

Peng Du
The University of Tennessee
du@eecs.utk.edu

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

MS25

Quantifying the Impact of Bit Flips on Numerical Methods

The collective surface area and increasing density of components has led to an increase in the number of observed bit flips. This effort works towards rigorously quantifying the impact of bit flips on floating point arithmetic, using both analytical modeling and Monte Carlo sampling, with the eventual goal to utilize these results to provide insight into the vulnerability of leadership-class computing systems to silent faults, and ultimately provide a theoretical basis for future silent data corruption research.

James Elliott
North Carolina State University
jjellio3@ncsu.edu

MS25

Application Robustification: Fortifying (Even Discrete) Applications Against Hardware Errors

Many applications need to solve combinatorial or discrete problems. Examples include sorting and optimization problems on graphs such as maximum flow. Most discrete algorithms were not designed to tolerate hardware faults which cause incorrect arithmetic or corruption of values in memory. We will talk about how to transform or relax discrete problems into forms which can be solved using fault-tolerant numerical algorithms, including some of those discussed in this minisymposium. We will also discuss algorithmic detection and correction schemes for linear algebra problems.

Joseph Sloan, Rakesh Kumar
University of Illinois at Urbana-Champaign
jsloan@illinois.edu, rakeshk@illinois.edu

MS25

Classifying Soft Error Vulnerabilities in Extreme-Scale Scientific Applications Using Bifit

Extreme-scale scientific applications are at a significant risk of being hit by soft errors on future supercomputers. To better understand this risk, we have built an empirical fault injection tool - BIFIT. BIFIT is designed with capability to inject faults at specific targets: execution point and data structure. We apply BIFIT to three scientific applications and investigate their vulnerability to soft errors. We are able to identify relationships between vulnerabilities and classes of data structures.

Jeffrey S. Vetter
Oak Ridge National Laboratory
Georgia Institute of Technology
vetter@computer.org

Dong Li
Oak Ridge National Laboratory
lid1@ornl.gov

Weikuan Yu
Auburn University
wkyu@auburn.edu

MS26

Presentations To Be Announced

Abstract not available at time of publication.

MS27

An Energy- and Charge-conserving, Implicit, Electrostatic Particle-in-Cell Algorithm in Mapped Meshes

Recently, a fully implicit algorithm for electrostatic plasma simulation has been proposed in 1D. The approach employs a Jacobian-free Newton-Krylov solver, which is made practical by the nonlinear elimination of particle quantities in favor of field quantities. Its fully implicit character enables exact charge and energy conservation, lending the approach superior accuracy properties. In this talk, we will introduce the approach, and its extention to mapped meshes (for mesh adaption) in a multidimensional setting.

Luis Chacon
Los Alamos National Laboratory
chacon@lanl.gov

Guangye Chen
ORNL
cheng2@ornl.gov

Daniel Barnes
Coronado Consulting
coronadocon@msn.com

MS27

Gyrokinetic Edge Plasma Simulation Using Continuum Methods

Understanding the edge plasma region of a tokamak is important to maintaining the confinement of the burning plasma in the core. Because of highly non-equilibrium behavior, a kinetic plasma description is required over a significant fraction of the edge region. We will discuss algorithmic advances for the 4D gyrokinetic edge plasma code COGENT, which is based on a fourth-order finite volume formulation for mapped, multiblock grids. Results will be presented for real magnetic geometries.

Jeffrey A. Hittinger, Milo Dorr
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
hittinger1@llnl.gov, dorr1@llnl.gov

Phillip Collela
Applied Numerical Algorithms Group
Lawrence Berkeley National Laboratory
pcollela@lbl.gov

Peter Mccorquodale
Lawrence Berkeley National Laboratory
PWMcCorquodale@lbl.gov

MS27

Positivity-Preserving Hybrid Semi-Lagrangian DG Schemes for Vlasov-Poisson

The Vlasov-Poisson equations describe the evolution of a collisionless plasma. The large velocities of the system create a severe time-step restriction from which the dominant approach in the plasma physics community is the particle-in-cell (PIC) method. In this work, we present a discontinuous-Galerkin method which utilizes semi-Lagrangian methods together with classical RKDG methods bootstrapped through local time-stepping. Our method utilizes unstructured physical grids which accommodate complicated geometries. Our method conserves mass and is positivity preserving.

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

David C. Seal
Department of Mathematics
Michigan State University
seal@math.msu.edu

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

MS27

Toward Gyrokinetic Particle-in-cell Simulations of Fusion Energy Dynamics at the Extreme Scale

Abstract: The Gyrokinetic Particle-in-cell (PIC) method has been successfully applied in studies of low-frequency microturbulence in magnetic fusion plasmas. While the excellent scaling of PIC codes on modern computing platforms is well established, significant challenges remain in achieving high on-chip concurrency for the new path to exascale systems. In addressing associated issues, it is necessary to deal with the basic gather-scatter operation and the relatively low computational intensity in the PIC method. Significant advancements have been achieved in optimizing gather-scatter operations in the gyrokinetic PIC method for next-generation multi-core CPU and GPU architectures. In particular, we will present on new techniques that improve locality, reduce memory conflict, and efficiently utilize shared memory on GPUs. Performance benchmarks on two high-end computing platforms – the IBM BlueGene/Q (Mira) system at the Argonne Leadership Computing Facility (ALCF) and the Cray XK6 (Titan) with the latest GPU at Oak Ridge Leadership Computing Facility (OLCF) will be presented.

Bei Wang
Princeton Institute of Computational Science and Engineering
Princeton University
beiwang@princeton.edu

Stephane Ethier
Princeton Plasma Physics Laboratory
ethier@pppl.gov

William Tang
Princeton University
PPPL
tang@pppl.gov

MS28

Simulating Flows in 2020: Challenges for Traditional CFD Applications and Weather/ocean/climate Predictions

We present an analysis of the needs of Computational and Geophysical Fluid Dynamics applications in the time horizon of Exascale Era machines. Empasis is given to the the scientific/engineering targets and to the the algorithms expected to be used in relation to the stresses they are expected to place on hardware and software. A discussion of various architectural options is provided along with their expected advantages and disadvantages. We also consider the question of uncertainty quantification and how that transforms many problems from pure "capability" ones to mixed "capacity-capability" ones and discuss how that changes the proposed machine balance.

Constantinos Evangelinos

Computational Science Center
IBM T.J. Watson Research
cevange@us.ibm.com

MS28

The Importance of Modeling Solvers: A Case Study Using Algebraic Multigrid

As the computers used to solve CSE problems become larger and more massively parallel, it is becoming increasingly important to develop performance models of the solvers. In this talk, we present a performance model of the solve cycle of algebraic multigrid and illustrate how it has guided our efforts to improve its scalability on emerging parallel machines. We additionally use the model to make predictions for solving larger problems on future machines.

Hormozd Gahvari
University of Illinois at Urbana-Champaign
gahvari@illinois.edu

William D. Gropp
University of Illinois at Urbana-Champaign
Dept of Computer Science
wgropp@illinois.edu

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

Martin Schulz, Ulrike Meier Yang
Lawrence Livermore National Laboratory
schulzm@llnl.gov, yang11@llnl.gov

MS28

Driving to Exascale - Challenges in Systems and Applications

A new day is dawning for High Performance Computing (HPC) Systems as we see core counts continue to rise. In this talk, we will briefly describe the direction we see computer systems taking as we move to exascale, using the IBM Blue Gene/Q as an early exemplar. We describe some of the hardware and software challenges that the computational science community can address through algorithm development for exascale, overviewing this minisymposium.

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

Constantinos Evangelinos
Computational Science Center
IBM T.J. Watson Research
cevange@us.ibm.com

Vipin Sachdeva
IBM Research, Austin
vsachde@us.ibm.com

MS28

Cloud Computing Pratices for Scientific Computing Applications

Cloud computing provides on-demand access to computing utilities, an abstraction of unlimited computing resources, and support for on-demand scale up, scale down and scale out. Furthermore, dynamically federated infrastructure can support heterogeneous applications requirements. Clouds are also joining HPC systems as a viable platform for scientific exploration and discovery. Therefore, understanding application formulations and usage modes that are meaningful in such a hybrid infrastructure, and how application workflows can effectively utilize it, is critical.

Moustafa AbdelBaky, Manish Parashar
Electrical and Computer Engineering
Rutgers University
moustafa.a@rutgers.edu, parashar@rutgers.edu

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

MS29

Probabilistic Predictions of Micro-anomalies from Macro-scale Response

We will present a probabilistic approach to characterize continuum (macro-level) constitutive properties of heterogeneous materials from a limited amount of microstructural information as typically available in practice. This approach is particularly amenable to include the effects of micro-cracks, that are not discernible by naked eyes, into the continuum constitutive material properties. Such microlevel defects can be potentially dangerous for structural systems due to fatigue cyclic loading that results in *initiation* of fatigue cracks. Distinct difference in probabilistic characteristics of macro-level responses of computer simulated model is observed depending on presence or absence of micro-cracks. This finding opens up the possibility of detecting micro-cracks ($\sim$10–100 microns size) from measurements of macro-level experimental responses (say, of a structure of dimension $\sim$20 m).

Sonjoy Das, Sourish Chakravarty
University at Buffalo
sonjoy@buffalo.edu, sc267@buffalo.edu

MS29

A Stochastic Multiscale Method for the Elastodynamic Wave Equation Arising from Fiber Composites

The long-term structural degradation of composite structures in fiber-reinforced materials is largely influenced by micro mechanical events. Moreover, due to the random character of the fiber locations and diameters, material properties, and fracture parameters, most mechanical quantities must be expressed in statistical terms. We therefore consider a multiscale problem governed by the linear stochastic elastodynamic wave equation and propose a numerical method for computing statistical moments of some given quantities of interest in regions of relatively small size. The method uses the homogenized global solution to construct a good local approximation that captures the microscale features of the real solution. We present numerical examples to verify the accuracy and efficiency of the method.

Mohammad Motamed
KAUST
mohammad.motamed@kaust.edu.sa

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

MS29

Quantification of Subscale Effects in Homogenized Models through Adaptive Information Measures

We consider the problem of computing bulk behavior of particle systems that are far away from equilibrium. Continuum-level expressions of conservation are closed by sampling the particle dynamics in a predictor-corrector formulation. The continuum predictor step constrains the probable final states of the particle ensemble. We concentrate on the problem of minimizing prediction errors through variational formulations for the particle probability density function based on adaptive information measures.

Anil Shenoy
University of North Carolina
ashenoy@live.unc.edu

Sorin Mitran
University of North Carolina Chapel Hill
mitran@unc.edu

MS29

Building Surrogates of Very Expensive Computer Codes: Applications to Uncertainty Quantification

To account for the epistemic uncertainty induced by the finite number of evaluations of a computer code, we define a probability measure over the space of possible surrogates. Each sample from this measure is a candidate surrogate for the code. By quantifying the informational content of the input space, we devise active learning schemes that are able to enhance the quality of the surrogate for particular tasks. Non-stationarity (localized features, discontinuities) can be captured by employing binary tree models.

Ilias Bilionis
Center for Applied Mathematics
Cornell University, Ithaca
ib227@cornell.edu

Nicholas Zabaras
Cornell University
zabaras@cornell.edu

MS30

Quantifying Uncertainty using a-posteriori Enhanced Sparse Grid Approximations

Adaptive sparse grids are popular for approximating outputs from high-dimensional simulation codes. When these codes solve both a forward and adjoint PDE formulation, a posteriori error estimation techniques can be used to achieve significant increases in the observed rate of convergence. The gains in efficiency are achieved by using the error estimate to both guide adaptivity and enhance the final approximation.

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

Tim Wildey
Sandia National Laboratory
tmwilde@sandia.gov

MS30

Sparse Grid Data Mining for Approximating High-dimensional Functions

Sparse grid data mining techniques can be used to construct mesh-based approximations of high-dimensional functions from randomly positioned data, unlike traditional approaches based on numerical discretizations that use interpolation. In this talk we employ sparse grid data mining to construct approximations of functions from ungridded legacy data. Unlike traditional regression-based Polynomial Chaos methods such an approach can be used to approximate both smooth and discontinuous functions, the latter ones requiring suitable adaptive refinement.

Dirk Pflüger
Universität Stuttgart, SimTech-IPVS
Simulation of Large Systems
Dirk.Pflueger@ipvs.uni-stuttgart.de

MS30

Model Order Reduction for Complex Dynamical Systems in Uncertainty Quantification

We consider systems of ordinary differential equations or differential algebraic equations, which result from the modeling of technical applications. Typically, uncertainties appear in some parameters of the systems. An uncertainty quantification is based on the substitution of uncertain parameters by random variables. The corresponding problems become complex in case of a huge dimension of the dynamical systems and/or a large number of random parameters. We investigate approaches for model order reduction to reduce the complexity of the systems, i.e., efficient numerical methods shall be achieved. On the one hand, parameterized model order reduction provides a reduction of the computational effort in sampling schemes. On the other hand, stochastic Galerkin methods yield deterministic systems of an even larger dimension to be reduced. We present numerical simulations of corresponding applications.

Roland Pulch
University of Wuppertal
pulch@math.uni-wuppertal.de

E. Jan W. ter Maten
Technical University of Eindhoven
e.j.w.ter.maten@tue.nl

MS31

An Implicit Maxwell Solver based on the Method of Lines Transpose

We present a novel method for solving the wave equation implicitly, to address scale separation and complex geometries when simulating plasma phenomenon. We apply the method of lines transpose (MOL^T) to the wave equation to obtain a boundary integral solution. Additionally, we develop a fast wave solver in higher dimensions which utilizes an ADI splitting to extend the fast one-dimensional algorithm that we have developed.

Matthew F. Causley

New Jersey Institute of Technology (NJIT)
causleym@math.msu.edu

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

MS31

Second Kind Integral Equation Formulation for the Modified Biharmonic Equation and its Applications

A system of Fredholm second kind integral equations is constructed for the modified biharmonic equation in two dimensions with gradient boundary conditions. Such boundary value problem arises naturally when the incompressible Navier-Stokes equations in pure stream-function formulation are solved via a semidiscretization scheme. The advantages of such an approach are two fold: first, the velocity is automatically divergence free, and second, complicated (nonlocal) boundary conditions for the vorticity are avoided. Our scheme can be coupled with standard fast algorithms such as FFT, fast multipole methods, or fast direct solvers to achieve optimal complexity, bringing accurate large-scale long-time fluid simulations within practical reach.

Mary-Catherine Kropinski
Simon Fraser University
mkropins@math.sfu.ca

Shidong Jiang
Department of Mathematical Sciences
New Jersey Institute of Technology
shidong.jiang@njit.edu

Bryan D. Quaife
Institute for Computational Engineering and Sciences
University of Texas at Austin
quaife@ices.utexas.edu

MS31

High Volume Fraction Simulations of Two-Dimensional Vesicle Suspensions

A boundary integral equation method for simulating inextensible vesicles in 2D viscous fluid was developed by Veerapaneni et al. I will discuss extensions that allow us to consider suspensions with a high concentration of vesicles. If time permits, results for simulating tracers, and computing pressures and tensors will be presented.

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

MS31

Simulations of Fiber and Particle Suspensions Accelerated by a Spectrally Accurate FFT-based Ewald Method

Numerical methods for 3D simulations of fiber and particle suspensions as based on boundary integral discretizations will be discussed. The simulations are accelerated by a spectrally accurate FFT based Ewald method, as applied to different fundamental solutions and periodic boundary conditions. This new spectral Ewald method compares favorably to the established so-called SPME method, especially regarding memory. This reduced memory use is a key feature in allowing for the simulations that are presented.

Anna-Karin Tornberg
KTH
Stockholm, Sweden
akto@kth.se

MS32

Using High Level Languages in Petascale Applications with PyClaw

Development of scientific software involves tradeoffs between ease of use, generality, and performance. We emphasize the importance of Pythonic interfaces to low-level languages and libraries. We then describe how an open source scientific Python stack allows us to implement PyClaw, a general hyperbolic PDE solver with simply written application scripts that scale to large supercomputers.

Aron Ahmadia
King Abdullah University of Science and Technology
aron.ahmadia@kaust.edu.sa

MS32

Using R with Open MPI and CUDA

At SuperComputing 2011, the University of Houston - Downtown received a Little Fe Cluster for teaching and research purposes. The software is the Bootable Cluster CD, which contains Open MPI and CUDA. We will discuss the history of the cluster and the software. We will then consider the uses of R with the cluster and present some empirical results to demonstrate speed up.

Erin M. Hodgess
University of Houston - Downtown
hodgesse@uhd.edu

MS32

LOO.PY: Transformation-Based Programming for Loops

Creating peak-performance compute codes on CPUs and GPUs often requires accomodating machine granularities such as vector widths, core counts, and on-chip memory sizes. Transforming even mathematically simple algorithms to suit these restrictions is onerous and error-prone. Loo.py is a tool that automates these transformations, allowing practitioners to experiment quickly while helping to ensure correct code. Loo.py is built for run-time code generation infrastructure using PyOpenCL, which interfaces Python with many types of compute devices.

Andreas Kloeckner
Courant Institute of Mathematical Sciences
New York University
kloeckner@cims.nyu.edu

Tim Warburton
Department of Computational and Applied Math
Rice University
timwar@caam.rice.edu

MS32

Bringing Exploratory Analytics to Big Data on Leadership Class HPC Platforms

The need for high-level languages in big data analysis comes from two challenges. First, we need to easily prototype complex analyses with a syntax close to high-level mathematical expressions and have a diverse toolbox of known analytics. Second, portability and optimization of codes on large computing platforms is increasingly difficult, calling for higher-level approaches. We begin to solve both of these challenges with pbdR (r-pbd.org). This project elevates R, a high-level language for data with arguably the most diverse analytics toolbox, to leadership class HPC platforms. We do this by a tight coupling with scalable and portable HPC libraries and by building additional infrastructure to allow high-level management of big data on complex architectures.

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

Drew Schmidt
University of Tennessee Knoxville
schmidt@math.utk.edu

Wei-Chen Chen
Computer Science and Mathematics Division
Oak Ridge National Laboratory
wccsnow@gmail.com

Pragneshkumar Patel
University of Tennessee Knoxville
pragnesh@utk.edu

MS33

Optimal Runge-Kutta-type Smoothers for Time Dependent Problems

We consider implicit time integration schemes for time dependent nonlinear PDEs. The appearing nonlinear systems are typically solved using either preconditioned JFNK or the FAS variant of multigrid, where the latter are typically designed for steady problems. This leads to suboptimal schemes. With regards to parallel computations, low storage preconditioners respectively multigrid smoothers are important. In this talk, we discuss possibilities of finding optimal low storage Runge-Kutta type smoothers, using time dependent model problems.

Philipp Birken
Department 10, Mathematics and Natural Sciences
University of Kassel, Germany
birken@mathematik.uni-kassel.de

Antony Jameson
Professor, Department of Aeronautics & Astronautics
Stanford University
jameson@baboon.stanford.edu

MS33

Efficient Exponential Integrators for Large Stiff Systems of ODEs

We discuss construction, analysis and implementation of exponential propagation iterative (EPI) methods for solving large stiff nonlinear systems of ODEs. These schemes allow integration with large time steps compared to explicit methods and computational savings per time step compared to implicit and other exponential schemes. We present high-order adaptive Runge-Kutta-type (EPIRK) integrators and compare their performance with other state-of-the-art methods on serial and parallel machines.

Mayya Tokman
University of California, Merced
School of Natural Sciences
mtokman@ucmerced.edu

John Loffeld
University of California
Merced
jloffeld@ucmerced.edu

MS33

Discontinuous Galerkin Methods and Implicit Wave Propagation

We consider discontinuous Galerkin discretizations in space for first-order hyperbolic linear systems with heterogeneous coefficients and full upwind flux. In time we compare standard explicit schemes with sufficiently small time steps, stable implicit Runge-Kutta methods, and exponential integrators, where the linear systems are solved by a parallel multigrid method with block smoothing. Numerical examples for acoustic, elastic, and electro-magnetic waves are presented. We show that the implicit methods are more efficient since the same accuracy in can be achieved with larger time steps.

Christian Wieners
Institute for Applied and Numerical Mathematics
Karlsruhe Institute of Technology, Germany
christian.wieners@kit.edu

MS33

Efficient Unsteady Flow Simulation using GMRES-E with Rosenbrock Time Integration Schemes

In this contribution the computational efficiency of Rosenbrock time integration schemes is compared to multi-stage DIRK and ESDIRK schemes. The computational benefits of Rosenbrock should come from the linear implicitness and constant system matrix for all stages within a time step. Solving one time step therefore reduces to solving the same linear system with a different right-hand-side per stage. In solving the linear systems with GMRES, both the retained effectiveness of the preconditioner and the reuse of Krylov vectors through enrichment improve computational efficiency.

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

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

Hester Bijl
Faculty Aerospace Engineering
Delft University of Technology, NL
h.bijl@tudelft.nl

MS34

Global/Local Surrogate Based Reservoir Management Optimization

In the waterflooding optimal management problem the objective is to maximize the NPV using as controls the rates of injector and producer wells. The duration of each cycle may also be included as design variables. As the numerical simulation has high computational cost we use kriging based surrogate models. The trust region based Sequential Approximate Optimization technique is used for local optimization. As the problem may be multimodal we use a hybrid Genetic Algorithm strategy.

Bernardo Horowitz
Universidade Federal de Pernambuco - UFPE
Civil Engineering Department
horowitz@ufpe.br

Silvana Afonso, Leonardo Oliveira
Universidade Federal de Pernambuco
Civil Engineering Department
smb@ufpe.br, leonardo.cdo@gmail.com

MS34

Efficient EnKF Conditioning with Reduced-Order Modeling

In this work we propose to use reduced-order modeling techniques to reduce the computational burden of the ensemble Kalman filter for large-scale problems. The reduced-order model proposed here can perform dimension reduction in both state and parameter spaces. The high-dimensional state space is reduced by the proper orthogonal decomposition, while the high-dimensional parameter space is reduced by the discrete cosine transform. The efficiency and robustness of the reduced-order model are demonstrated by an uncertainty quantification example of subsurface transport.

Binghuai Lin
MIT
Department of Civil and Environmental Engineering
blin@mit.edu

Dennis McLaughlin
Civil Engineering MIT
dennism@mit.edu

MS34

Using Geological Uncertainty to Simplify History-Match Optimization

Geological realism in history-matching is fundamental to accurate reservoir model predictions. We frame this as a Bayesian optimization problem, fitting the observables while conforming to prior uncertainties. Principal component analysis greatly reduces the number of optimization variables by describing only the intrinsic degrees of freedom allowed by the prior. Using the SPE Brugge Challenge model as a test case, we demonstrate improvement over previously published data, with higher NPV and lower predictive error.

Michael Prange, William Bailey
Schlumberger-Doll Research
prange@slb.com, wbailey@slb.com

Thomas Dombrowsky
Longhorn Technology Limited
tom@longhorntech.co.uk

MS34

Efficient Surrogate Surface Global Optimization for Estimating Carbon Sequestration Plumes with Sparse Observations

Estimation of sequestered CO2 and pressure plumes is important but difficult because monitoring data is very sparse and inverse optimization problem has multiple local optima. Each objective function evaluation requires expensive forward simulation of 3-D, highly nonlinear, multiphase, multi-constituent set of PDEs. We get good current estimates and forecasts of plumes with a surrogate response surface global optimization algorithm Stochastic RBF with small number of original model simulations. Variants of the algorithm are presented.

Christine A. Shoemaker, Antoine Espinet
Cornell University
cas12@cornell.edu, aje46@cornell.edu

Christine Doughty
Lawrence Berkeley National Laboratory
DOE
doughty@lbl.gov

MS35

A Stochastic Dimensional Reduction Multiscale Finite Element Method with Applications for Subsurface Flows in Random Porous Media

Stochastic modeling has become a widely accepted approach to quantify uncertainty of flows in random porous media. To efficiently treat the high-dimensionality of the stochastic space and the inherent heterogeneity of porous media, we propose a new truncated high dimensional model representation technique and combine it with a mixed multiscale finite element method. To capture the non-local spatial features of the media and the effects of important random variables, we hierarchically incorporate some global information individually from each of random parameters. Numerical experiments are carried out for subsurface flows in random porous media.

Lijian Jiang
IMA, University of Minnesota
ljjiang16@gmail.com

David Moulton
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
moulton@lanl.gov

MS35

Geometric Integration and Analysis of General Multiscale Systems

In order to accelerate computations, improve long time accuracy of numerical simulations, and sample statistics distribution by dynamics, we develop multiscale geometric integrators. These integrators employ coarse steps that do not resolve the fast timescale in the system; nevertheless, they capture the effective contribution of the fast dynamics. Distinct from existing approaches, an identification of underlying slow variables is not required, and intrinsic geometric structures (e.g., symplecticity, conservation laws,

and invariant distribution) can be preserved by the numerical simulation.

Molei Tao
Courant Institute
mtao@cims.nyu.edu

Houman Owhadi
Applied Mathematics
Caltech
owhadi@caltech.edu

Jerrold Marsden
Caltech
(deceased)

MS35

Multiscale Discontinuous Galerkin Method for Elliptic Equations with Rapidly Oscillatory Coefficients

In this talk, a special discontinuous Galerkin method is proposed for a class of second order elliptic problems with rough coefficients, whose local oscillating directions vary smoothly within the domain. The key ingredient of the method lies in choosing special approximation space to capture the multiscale solutions without having to resolve the finest scale therein. Theoretical proof and numerical examples would be presented for the second order method in two dimensional case.

Yifan Zhang
Brown University
yifan_zhang@brown.edu

Wei Wang
Department of Mathematics and Statistics
Florida International University
weiwang1@fiu.edu

Johnny Guzman
Brown University
johnny_guzman@brown.edu

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

MS35

Numerical Homogenization: From Higher Order Poincare Inequalities to Optimally Localized Basis

We introduce a new method for the numerical homogenization of divergence form elliptic equations with rough (L^∞) coefficients. Our method does not rely on concepts of ergodicity or scale-separation but on the property that the solution operator is compact from L^2 to H^1. The approximation space is generated as an interpolation space (over a coarse mesh of resolution H) minimizing the L^2 norm of source terms; its (pre-)computation involves solving $\mathcal{O}(\mathcal{H}^{-\lceil})$ bi-harmonic PDEs on localized sub-domains of size $\mathcal{O}(\mathcal{H}(\ln \mathcal{H})^{\in})$; its accuracy ($\mathcal{O}(\mathcal{H})$ in H^1 norm) is established via the introduction of a new classs of higher-order Poincaré inequalities. The method naturally generalizes to time dependent problems.

Lei Zhang
University of Oxford
lzhang2012@sjtu.edu.cn

MS36

Parallel Constrained Minimization Methods

Abstract not available at time of publication.

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

MS36

Sparse Quadrature Approach to Bayesian Inverse Problems

In this talk, we present a novel, deterministic approach to inverse problems for identification of parameters in differential equations from noisy measurements. Based on the parametric deterministic formulation of Bayesian inverse problems with input parameter from infinite-dimensional, separable Banach spaces, we develop a practical computational algorithm for the efficient approximation of the infinite-dimensional integrals with respect to the posterior measure.

Claudia Schillings
ETH Zurich
Seminar for Applied Mathematics
claudia.schillings@math.ethz.ch

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

MS36

Multigrid Optimization on GPGPU

GPGPU is a cost-effective means for the usage of a parallel computer with several hundred cores. In this talk, the computational potential of this architecture for optimization problems is investigated. In particular, a specific implementation of multi grid methods for large optimal control problems on several stream processors are discussed, where a domain decomposition approach is employed for the distribution on multiple GPU.

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

Christian Wagner
Department of Mathematics
University of Trier
christian.wagner@uni-trier.de

MS36

GPU Accelerated Discontinuous Galerkin Methods for Simulation and Optimization

Recent studies have shown advantages of the application of streaming processors for highly arithmetic expensive problems arising from partial differential equations. This talk introduces a high-order discontinuous Galerkin implemen-

tation on GPUs for Euler equations with focus on its parallel performance. In particular, GPU acceleration for the corresponding discrete adjoint is addressed and its application in optimization is discussed. The importance of accurate boundary resolutions is demonstrated and a mesh curvature approach dealing with complex domains is presented.

Martin Siebenborn
University of Trier
siebenborn@uni-trier.de

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

MS37

A *hp*-nonconforming Discontinuous Galerkin Spectral Element Method: Analysis and Application to Large Scale Seismic Inversions

We analyze the consistency, stability, and convergence of an *hp* discontinuous Galerkin spectral element method. Our analytical results are developed for both conforming and non-conforming approximations on hexahedral meshes. A mortar-based non-conforming approximation is developed to treat both h and p non-conforming meshes simultaneously. We demonstrate the scalability and accuracy of the proposed method on large-scale global seismic wave inversion.

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

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

MS37

Output-based Space-time Adaptation for DG Simulations with Moving Meshes

We address the question of how to adapt a mesh for problems with deformable domains, like a wing in flapping flight, to obtain accurate outputs. We present a space-time adaptation strategy that uses output error estimates attained from a set of discrete, unsteady adjoint solutions. We derive an additional adjoint for the geometric conservation law, and show the accuracy and efficiency of our strategy for Navier-Stokes problems solved within a DG framework.

Steve Kast, Krzysztof Fidkowski
University of Michigan
kastsm@umich.edu, kfid@umich.edu

MS37

Discontinuous Petrov-Galerkin Finite Element Method for the Analysis of Plate and Shell Structures

Stress recovery and numerical locking present common difficulties encountered in the finite element analysis of plate and shell structures. One way to address these issues is to base the numerical discretization on a mixed variational principle where the stresses are declared as independent unknowns next to the displacements. The discontinuous Petrov-Galerkin (DPG) finite element methodology provides a systematic approach to devise such a variational formulation and ensure its stability also in the discrete (finite dimensional) setting. In this talk, we discuss recent developments in the analysis of plate and shell deformations utilizing the DPG framework.

Antti H. Niemi
Aalto University, Finland
antti.h.niemi@aalto.fi

MS37

A High-Order Implicit-Explicit Discontinuous Galerkin Scheme for Fluid-Structure Interaction

We present a discontinuous Galerkin scheme for fully coupled fluid-structure interaction problems. A high-order DG-ALE formulation is used for the fluid, and standard CG schemes for beams and membranes are used for the solids. Using implicit-explicit Runge-Kutta time integrators with explicit predictors for the coupling variables, we can integrate each problem domain independently using efficient parallel implicit solvers while retaining the high formal order of accuracy.

Per-Olof Persson
University of California Berkeley
Dept. of Mathematics
persson@berkeley.edu

Bradley Froehle
University of California, Berkeley
Dept. of Mathematics
bfroehle@berkeley.edu

MS38

**Comparative
Review of Fractional-step Approaches to the Immersed Boundary Method**

In the various immersed boundary methods based on a fractional-step approach, there are either unjustified boundary conditions or overlooked conservation properties. Often, this is caused by mixing continuum and discrete equations in the analysis. With a careful, fully discrete derivation, we construct a unifying perspective of IB methods that resolves these issues. We have developed an open-source GPU code, cuIBM, that provides a common framework for IB methods integrating these improvements.

Lorena A. Barba
Department of Mechanical Engineering
Boston University
labarba@bu.edu

Anush Krishnan
Boston University
anush@bu.edu

MS38

Increased Accuracy of Immersed Boundary Methods Using Fourier Approximations of Delta Functions

In immersed boundary methods, the fluid and structure communicate through smoothed approximate delta functions with small spatial support. We take a different approach and construct highly accurate approximations to the delta function directly in Fourier space. This method leads to high-order accuracy away from the boundary and significantly smaller errors near the boundary. We present accuracy tests and simulation results from an application in cell biology in which large errors in the traditional IB method produce unphysical results.

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

David Hartenstine
Department of Mathematics
Western Washington University
david.hartenstine@wwu.edu

Wanda Strychalski
Department of Mathematics
University of California, Davis
wanda@math.ucdavis.edu

MS38

Immersed Boundary Method for the Incompressible Navier-Stokes Equations Based on the Lattice Green's Function Method

A parallel, three-dimensional immersed boundary method is developed to solve external, viscous incompressible flows on an infinite domain. The equations are formally discretized on an infinite staggered Cartesian grid. The Lattice Green's Function method is used to reduce the problem to a finite region. The fully discrete equations are obtained by combining an integrating factor technique and a half-explicit Runge-Kutta scheme, and solved using a nested projection technique. Results for test problems are presented.

Sebastian Liska
Mechanical and Civil Engineering
California Institute of Technology
sliska@caltech.edu

Tim Colonius
Division of Engineering and Applied Science
California Institute of Technology
colonius@caltech.edu

MS38

Dynamics of an Elastic Rod with Curvature and Twist: Stokes Formulation

We develop a Lagrangian numerical algorithm for an elastic rod immersed in a viscous, incompressible fluid at zero Reynolds number. The elasticity of the rod is described by a version of the Kirchhoff rod model. The coupling to the fluid is accomplished by the use of the method of regularized Stokeslets for the force, and regularized rotlets and dipoles for the torque. This method will be compared to the generalized IB method and asymptotic results.

Sarah D. Olson
Worcester Polytechnic Institute
sdolson@wpi.edu

Sookkyung Lim
University of Cincinnati
sookkyung.lim@uc.edu

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

MS39

Linear Algebra Libraries with DAG Runtimes on GPUs

Nowadays many clusters integrate GPUs accelerators in their architectures that provide a huge amount of computational units rarely fully exploited. We present in this talk how tile algorithms and DAG schedulers as PaRSEC or StarPU can allow the programmer to integrate GPUs in their algorithms. We will present dense linear algebra algorithms as Cholesky or LU factorizations that exploit distributed architectures equipped with GPUs.

George Bosilca
University of Tennessee - Knoxville
bosilca@eecs.utk.edu

Aurelien Bouteiller
University of Tennessee
bouteill@eecs.utk.edu

Mathieu Faverge
U. of Tennessee - Knoxville
mfaverge@eecs.utk.edu

Thomas Herault
Univ. of Tennessee - Knoxville
herault@eecs.utk.edu

MS39

Parallel LU Factorizations on GPUs in AORSA

We describe the performance of a parallel dense matrix solver in AORSA, the All ORders Spectral Algorithm fusion application for modeling the response of plasma to radio frequency waves in a tokamak device, which takes advantage of Graphics Processing Unit (GPU) acceleration and is compatible with ScaLAPACK LU factorization routines. The left-looking out-of-core algorithm factors matrices that are larger than the available GPU memory and achieves 170 GFlops (double precision) per GPU-equipped node on the CrayXK6.

Judith Hill, Ed D'Azevedo, David Green
Oak Ridge National Laboratory
hilljc@ornl.gov, dazevedoef@ornl.gov, greendl1@ornl.gov

MS39

A Performance Study of Solving a Large Dense Matrix for Radiation Heat Transfer

This talk presents a performance study of solving a large dense matrix arising from the thermal radiation problem. The radiosity exchanged among a number of radiating surfaces depends on their view factors. The radiosity matrix is a strictly diagonally dominant matrix. But in a limited number of applications the radiosity matrix is also SPD. We will present the LU and LLT procedures and results used in solving the matrix on a GPU cluster.

Kwai L. Wong

Joint Institute for Computational Science
University of Tennessee/ORNL
wong@jics.utk.edu

Edurado D'Azevedo
ORNL
e6d@ornl.gov

Harvy Hu
Chinese University of Hong Kong
haz8816415@gmail.com

Shiquan Su
University of Tennessee
ssu2@utk.edu

MS39

Mult-GPU Tridiagonalzation on Shared-and-distributed-memory Systems

Tridiagonalization of a symmetric dense matrix is an important computational kernel in many symmetric eigenvalue problems. In this talk, we describe our extension of one-stage tridiagonalization to use multiple GPUs on shared- and distributed-memory computers. We present experimental results to demonstrate that the tridiagonal reduction time of LAPACK or ScaLAPACK can be reduced significantly using GPUs.

Tingxing Dong
UTK
tdong@eecs.utk.edu

Jack J. Dongarra
University of Tennessee
dongarra@eecs.utk.edu

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

Ichitaro Yamazaki
UTK
ic.yamazaki@gmail.com

MS40

A Structured Quasi-Arnoldi Procedure for Model Order Reduction of Second-order Systems

Most Krylov subspace-based structure-preserving model order reduction methods for the second-order systems consist of two stages. The first stage is to generate basis vectors of the underlying Krylov subspaces. The second stage is to perform explicit subspace projections via matrix-matrix multiplications to generate the reduced-order models. For very large scale systems, the second stage could be prohibitively expensive due to the costs of data storage and communication. In this talk, we discuss a Structured Quasi-Arnoldi (SQA) procedure to avoid the second stage. The SQA procedure is based on a Krylov-type decomposition that allows us to derive the structure-preserving reduced-order model directly from the computed Krylov-type decomposition. Numerical examples will be presented to illustrate accuracy and efficiency of the SQA procedure.

Yung-Ta Li
Fu Jen Catholic University, Taiwan
yungtali@gmail.com

Wen-wei Li
National Chiao-Tung University, Taiwan
wwlin@math.nctu.edu.tw

Zhaojun Bai
Departments of Computer Science and Mathematics
University of California, Davis
bai@cs.ucdavis.edu

MS40

A New Approach to Model Order Reduction of the Navier-Stokes Equations

A new approach to model order reduction of the Navier-Stokes equations is proposed. Unlike traditional approaches, this method does not rely on empirical turbulence modeling or modification of the Navier-Stokes equations. It provides spatial basis functions different from the usual proper orthogonal decomposition basis function in that, in addition to optimally representing the solution, the new basis functions also provide stable reduced-order models. The proposed approach is illustrated with two test cases: two-dimensional flow inside a square lid-driven cavity and a two-dimensional mixing layer.

Maciej Balajewicz, Earl Dowell
Duke University
maciej.balajewicz@stanford.edu, earl.dowell@duke.edu

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

MS40

Reduced Order Models Preserving Spectral Continuity for Wave Propagation in Unbounded Domains

We address the main issues obstructing the application of the Krylov subspace based model reduction to the time-domain solution of the exterior wave problems. To avoid spurious reflections and instability, the ROM should preserve in some sense delicate spectral properties of the original problem, i. e., absolutely continuous spectral measure supported on real negative semiaxis. We introduce such a ROM, give its rigorous justification, discuss approaches to preconditioning and present large scale seismic examples.

Vladimir L. Druskin
Schlumberger-Doll Research
druskin1@slb.com

Leonid Knizhnerman
Schlumberger, consultant
lknizhnerman@gmail.com

Olga Podgornova
Schlumberger Moscow Research
opodgornova@slb.com

Rob Remis
Circuits and Systems Group

Delft University of Technology
r.f.remis@tudelft.nl

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

MS40

A Hyper-Reduction Method for Nonlinear Dynamic Finite Element Models

A model order reduction method for nonlinear systems arising from the discretization of second-order hyperbolic problems using a finite element method in space and a finite difference method in time is presented. Its main components are a Galerkin projection onto a reduced-order basis constructed using snapshots and the proper orthogonal decomposition method, a hyper-reduction scheme that preserves important properties of both the physical system and its finite element representation, and a fast computer implementation.

Charbel Farhat
Stanford University
CFarhat@stanford.edu

MS41

Automating Adjoints for Earth Science Applications

Adjoint models are crucial ingredients in solving inverse problems in the geosciences. However, deriving adjoint models is typically very difficult. I present a new technique for rapidly deriving adjoints of finite element models via symbolic analysis and code generation (based on the FEniCS project). I give examples of its application in ocean circulation, mantle convection and ice modelling.

Patrick Farrell
Department of Earth Science and Engineering
Imperial College London
patrick.farrell@imperial.ac.uk

Simon W. Funke
Imperial College London
Department of Earth Science and Engineering
s.funke09@imperial.ac.uk

David Ham
Imperial College
david.ham@imperial.ac.uk

Marie E. Rognes
Simula Research Laboratory
meg@simula.no

MS41

Second-order Adjoints in Seismic Tomography

We present an introductory and general review of second-order adjoints for the efficient computation of Hessian-vector products. Using the specific example of large seismic tomographic inverse problems, we demonstrate how sencond-order adjoints can be used to accelerate convergence towards an optimal Earth model, and to analyse resolution and trade-offs in the final solution.

Andreas Fichtner
Utrecht University
fichtner@geo.uu.nl

Naiara Korta
Barcelona Center for Subsurface Imaging
nkorta86@gmail.com

Jeannot Trampert
Utrecht University
j.a.trampert@uu.nl

MS41

Estimating the Rheology of the Earth's Mantle: An Application of the Adjoint Method in Geodynamics

Determining rheologic parameters of the Earth with the help of numerical simulations of mantle flow is one of the essential tasks in geodynamics. But where standard forward simulations suffer from the problem of an unknown initial condition and instantaneous models are non-unique with respect to rheology (Schaber et al. (2009)), the adjoint method in geodynamics can be used to create a time-series of the temperature distribution of the Earth's mantle over the last 40 Ma that is both sensitive to rheology and consistent with present-day observations.

Andre Horbach
Department of Earth and Environmental Sciences, LMU Munich
horbach@geophysik.uni-muenchen.de

Hans-Peter Bunge
LMU Munich, Germany
bunge@geophysik.uni-muenchen.de

MS41

Calibration of Stratigraphical Models

Calibration of the model is a type of inverse problem that estimates coefficients on the basis of observations. Depositional models were chosen because of their significance in the basin research and the available observations: seismic studies provide the thickness of the the deposited layers, while well-log data give the information on the type of sediments and the history of deposition.

Lyudmyla Vynnytska, Stuart Clark
Simula Research Laboratory
lyudav@simula.no, stuart@simula.no

MS42

An Algorithm for Shape Detection in Computed Tomography

Computed tomography (CT) is essential to modern medicine. Usually practitioners' interest in CT is to identify and quantify structures found in CT images. Traditionally, this is done by post-processing of reconstructed images (i.e. segmenting the images). In this talk, we propose to address this problem directly; we extract the geometric structures directly from measured data, bypassing the reconstruction phase. We formulate this task as a shape optimization problem and devise an efficient numerical algorithm to perform the optimization.

Gunay Dogan
Theiss Research, NIST
gunay.dogan@nist.gov

MS42

Image Segmentation Methods Based on Centroidal Voronoi Tessellation and Its Variants

In this talk we review some recent progresses on image segmentation methods based on centroidal Voronoi tessellation and its variants. The classic CVT model, the edge-weighted CVT (EWCVT) model, the local variation and edge-weighted CVT (LVEWCVT) model, and their implementation algorithms will be discussed in detail. We will also illustrate and compare these interesting segmentation methods through extensive numerical examples.

Jie Wang
Arizona State University
jie.wang.2@asu.edu

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Xiaoqiang Wang
Florida State University
wwang3@fsu.edu

MS42

Image Registration for the Future: Fast, Scalable, and Highly Memory-Efficient Algorithms

Image registration is a fundamental problem in many imaging applications. The ever growing image sizes and new modalities pose a significant challenge for efficient algorithms. In the first part of the talk, we present techniques for fully matrixfree implementations of multimodal parametric registration algorithms. Besides of virtually reducing memory consumption to zero, an excellent scalability on multi-core architectures is achieved. The second part of deals with selected computational aspects of nonlinear approaches yielding diffeomorphic transformations.

Jan Rühaak
Institute of Mathematics and Image Computing
University of Lübeck
jan.ruehaak@mevis.fraunhofer.de

MS42

A Mesh Warping Algorithm for Brain Biomechanics Boundary Evolution Tracking

Hydrocephalus is a neurological disease which causes ventricular dilation due to abnormalities in the cerebrospinal fluid circulation. Similarly, a brain tumor involves the abnormal growth of brain cells. To aid neurosurgeons in treatment planning, we propose an automated geometric computational approach for tracking evolution of the applicable brain boundaries via the level set method and a mesh warping technique. Our method uses brain geometries taken from medical images and may incorporate other relevant clinical information.

Suzanne M. Shontz
The Pennsylvania State University
sshontz@math.msstate.edu

Corina Drapaca
Department of Engineering Science and Mechanics
Pennsylvania State University
csd12@psu.edu

MS43

Pulling Fibers

Abstract not available at time of publication.

Linda Cummings
Department of Mathematical Sciences
New Jersey Institute of Technology
linda.j.cummings@njit.edu

MS43

An Application of Matrix Theory to the Evolution of Coupled Modes

In order to overcome loss in optical fibers, experimentalists are interested in employing parametric amplifiers using four-wave mixing. A variety of other signal-processing functions also use parametric amplifiers based on four-wave mixing. Upon linearizing the nonlinear Schrödinger equation typically used to model such amplifiers, one obtains a system of ODEs for the complex amplitude. The solution of this system can be expressed as the product of transfer matrices with the initial condition and its conjugate. Physical insight into the fiber-optic system can be obtained by examining the theoretical properties of these matrices. This presentation explores these properties.

Joseph D. Fehribach
Worcester Polytechnic Institute
bach@wpi.edu

David A. Edwards
University of Delaware
Newark, DE
edwards@math.udel.edu

Richard O. Moore
New Jersey Institute of Technology
rmoore@njit.edu

Colin J. McKinstrie
Bell Laboratories, Alcatel-Lucent
colin.mckinstrie@alcatel-lucent.com

MS43

Modeling Glass Temperature in a Tempering Furnace

Equations governing the temperature distribution in a glass plate as it progresses through a tempering furnace are formulated. Heat transfer due to natural convection, forced convection from an array of impinging gas jets, conduction due to contact with supporting rollers, and radiative heating from heating elements is described. Two and three dimensional time-dependent numerical simulations allow us to explore how to control the glass temperature profile by adjusting heating conditions.

Harrison Potter
Duke University
Department of Mathematics
hdp2@math.duke.edu

MS43

A Homogenization Analysis of the Compressible Flow Between a Slider and a Moving Rough Sur-

face

The compressible flow between a slider and a moving rough surface is examined asymptotically and numerical in the limit of very small gap height. The amplitude and wavelength of the roughness are assumed to be of the order of the gap height. A two-scale homogenization analysis is employed to determine a nonlinear elliptic partial differential equation governing the leading-order pressure in the gap on the scale of the slider. The equation involves coefficient functions which are determined numerically by averaging Stokes flows on the scale of the roughness. Comments and a brief analysis is given on the reduction of the governing equation for pressure in the limit of long wavelength of the surface roughness.

Burt S. Tilley
Mathematical Sciences Department
Worcester Polytechnic Institute
tilley@wpi.edu

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

Colin Please
School of Mathematics
University of Southampton
cpp@maths.soton.ac.uk

Ferdinand Hendriks
HGST, a Western Digital Corporation
ferdi.hendriks@hgst.com

MS44

Performance Modeling of the Eigen-K Dense Eigensolver on Massively Parallel Machines

We aim for developing an efficient dense eigensolver on massively parallel machines such as the K computer in Japan. In our consideration, performance modeling of the solver is helpful to performance tuning. In this talk, we introduce an approach to modeling the performance of the Eigen-K solver developed by Imamura, where the parallel execution time of the solver is estimated from the single node performance and the internode communication performance.

Takeshi Fukaya
Kobe University
Japan
fukaya@people.kobe-u.ac.jp

Toshiyuki Imamura
The University of Electro-Communications
imamura@im.uec.ac.jp

Yusaku Yamamoto
Kobe University
yamamoto@cs.kobe-u.ac.jp

MS44

A Hybrid CPU-GPU Generalized Eigensolver for Electronic Structure Calculations

The adoption of hybrid GPU-CPU nodes in traditional supercomputing platforms opens acceleration opportunities for electronic structure calculations in materials science and chemistry applications, where medium-sized generalized eigenvalue problems must be solved many times. We developed a novel, architecture aware algorithm that is state-of-the-art in HPC, significantly outperforming existing libraries. We describe the algorithm and analyze its performance impact on applications of interest when different fractions of eigenvectors are needed by the host electronic structure code.

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@utk.edu

Raffaele Solcá
ETH
Eidgen{̈o}ssische Technische Hochschule Z{̈u}rich
rasolca@student.ethz.ch

Stan Tomov
Computer Science Department
University of Tennessee
tomov@eecs.utk.edu

Jack Dongarra
University of Tennessee
dongarra@eecs.utk.edu

Thomas Schulthess
ETH
Eidgenössische Technische Hochschule Zürich
schultho@ethz.ch

MS44

ELPA: A Highly Scalable Eigensolver for Petaflop Applications

The symmetric eigenproblem is of great interest in many scientific disciplines (e.g., quantum chemistry, network analysis) and represents a severe bottleneck for several numerical simulations. ELPA is a new direct eigensolver, which has been consequently designed towards massively parallel systems. Compared to state-of-the-art libraries for distributed memory systems, it leads to significant improvements in both, efficiency and scalability. In this talk we present the ELPA library and address the most recent developments and results.

Thomas Auckenthaler
TU München Institut für Informatik
Germany
auckenth@in.tum.de

Alexander Heinecke
TU München
heinecke@in.tum.de

Hermann Lederer
Rechenzentrum Garching
lederer@rzg.mpg.de

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

MS44

Eigensolvers in Combustion Simulations

Thermoacoustic instabilities are an important concern in the design of gas turbine combustion chambers. Their study leads to the numerical solution of large sparse unsymmetric eignevalue problems. In this talk we discuss various numerical techniques to exploit a priori information on the sought eigenmodes. We illustrate the features of these different techniques on small toy examples as well as on large scale industrial calculations.

Luc Giraud
INRIA
Toulouse, France
luc.giraud@inria.fr

Pablo Salas
Inria - Bordeaux Sud Ouest
jont Inria-CERFACS lab on HPC
salas@cerfacs.fr

Vasseur Xavier
CERFACS
vasseur@cerfacs.fr

MS45

Guaranteed Stability and Passivity of Reduced Order Models

Some efficient model reduction methods fail to guarantee important properties of the reduced system, such as stability and passivity. In this contribution we will show how methods from control theory can be combined with model reduction methods in order to guarantee these properties. We will concentrate on the application of such methods to interpolatory reduction approaches without deterioration of the interpolation properties. We will apply the results to challenging examples from S-parameter design in electronic devices.

Athanasios C. Antoulas
Dept. of Elec. and Comp. Eng.
Rice University
aca@rice.edu

MS45

Fast Solver for Large Scale Inverse Problems using Data-driven Reduced Order Models

A novel data-driven model reduction technique is developed for solving large-scale inverse problems. The proposed technique exploits the fact that the solution of the inverse problem often concentrates on a low dimensional manifold. Unlike typical MCMC approaches for solving the inverse problem, our approach avoids repeated evaluation of expensive forward models by coupling the inference algorithm with the reduced-order model. This maintains the accuracy of the inference and also results in a lower-dimensional reduced model than obtained with the typical POD approach.

Tiangang Cui
MIT
tcui@mit.edu

Youssef M. Marzouk, Karen E. Willcox
Massachusetts Institute of Technology
ymarz@mit.edu, kwillcox@MIT.EDU

MS45

A Reproducing-Kernel Framework for $\mathcal{H}_\in$ Model Order Reduction

The Iterative Rational Krylov (IRKA) algorithm for model order reduction has recently attracted attention because of its effectiveness in real world applications. The key idea is to construct a reduced order model that satisfies a set of necessary optimality conditions formulated as interpolatory conditions: the reduced r-th order transfer function interpolates the full n-th order function and its derivative(s) in the reflected (about the imaginary axis) images of the reduced order poles. This is formulated as a fixed point problem, and the interpolation nodes are generated as $\sigma^{(k+1)} = \phi(\sigma^{(k)})$. Here $\phi(\cdot)$ computes the eigenvalues of the Petrov-Galerking projection of the state matrix to rational Krylov subspaces computed at $\sigma^{(k)} = (\sigma_1^{(k)}, \ldots, \sigma_r^{(k)})$. The most expensive part of the IRKA is computing the transfer function at a sequence of dynamically generated points in the right half-plane $\mathbf{C}_+$. Also, this setting is not suitable for a data driven framework, since we are given only the transfer function values along the imaginary axis. On the other hand, if we interpret the necessary optimality conditions as orthogonality (in $\mathcal{H}_\in$) of the residual to the tangent space at the reduced order model to the manifold of the models of given reduced order, we can deploy rich theory of the geometry of linear systems, as well as reproducing kernel space property of $\mathcal{H}_\in$. All action takes place on $\mathbf{iR}$, and, with a suitable kernel, function evaluation can be achieved by a linear functional, implemented using numerical quadrature. This new framework is well suited for data driven applications. It also offers efficient implementation and deeper insight in the behavior and properties of the numerical algorithm.

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

MS45

Model Reduction for Indoor-Air Behavior in Energy-Efficient Building Design

We present a two-step approach for generating reduced order models for the indoor-air environment in control design for energy efficient buildings. Using a data-driven model reduction approach, we first construct an intermediate reduced-model directly from input and output measurements. We then apply optimal interpolatory model reduction techniques to this intermediate model to obtain the final reduced model. Numerical results illustrate that the reduced model accurately represents the input-output behavior of the full-order system.

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

Eugene Cliff
Virginia Tech
Interdisciplinary Center for Applied Mathematics
ecliff@vt.edu

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

MS46

Issues in the Management and Analysis of Large Data Sets Arising in Complex Problems

Reliable fast and accurate estimation of useful quantities of interest from analysis of large computational and experimental data sets and from modeling and simulation of complex systems involving multi-physics and multi-scale pose many significant challenges that require sophisticated mathematical and statistical tools. In this talk, we will discuss some of these challenges and their solutions against the backdrop of some complex problems arising in heterogeneous environment such as turbulence, porous media flows and so on. This talk will be based on an ongoing work fine details of which will have to wait until the talk.

Prabir Daripa
Texas A&M University
Department of Mathematics
prabir.daripa@math.tamu.edu

MS46

Numerical Investigation of the Effective Properties of Tight Fractured Porous Shale Rock

We focus on the effect of micro-pores and micro-fractures and their interactions on the overall effective mechanical, thermal and hydrological flow and transport properties of tight fractured porous shale rock under several conditions of uncertainties. Uncertainty associated with the physical and chemical aggregates of tight shale rock and their parametrization is investigated using brute-force Monte-Carlo method.

Souheil M. Ezzedine
LLNL
ezzedine1@llnl.gov

MS46

Covariance Models Based on Local Interaction (Spartan) Functionals

We derive two-dimensional covariance functions corresponding to Gibbs random fields based on quadratic energy functionals with local spatial interactions. We then develop the Karhunen-Loéve representation and explicit solutions for simple boundary geometries. We illustrate the parametric dependence of the realizations using two-dimensional Karhunen-Loéve expansions. We develop inverse covariance functions in $d \in Z_+$-dimensions using a frequency cutoff and approximating the respective spectral integral. Finally, we discuss potential applications in spatial interpolation and simulation problems.

Dionissios T. Hristopulos
Technical University of Crete
dionisi@mred.tuc.gr

MS46

The Effect of Material Heterogeneity in Computing Local Deformation Effects

Model predictions of large deformation geologic structures are based on constitutive models that are typically calibrated using experimental data that assume material homogeneity at a micro- to meso-scale. This work shows computational representations of the kinematics of granular materials using Discrete and Finite Elements as an attempt to identify and characterize the main sources of material heterogeneity as stochastic processes, and the impact these have on the simulation of the material's mechanistic performance.

Zenon Medina-Cetina
Zachry Department of Civil Engineering
Texas A&M University
zenon@tamu.edu

Ahran Song
Stochastic Geomechanics Laboratory
Texas A&M University, TX, USA
arsong@tamu.edu

Patrick R. Noble
Texas A&M University
pnoble10@neo.tamu.edu

Tam Duong
Stochastic Geomechanics Laboratory
Texas A&M University, TX, USA
meokeen@neo.tamu.edu

MS47

Robust Integral Solver for 3D Acoustic Scattering from Doubly-Periodic Media

We construct a new class of high-order accurate integral equation solvers for the scattering of time-harmonic scalar waves from a doubly-periodic array of smooth obstacles, or a surface with a doubly-periodic array of bumps. This combines recent QBX surface quadratures with a simple quasi-periodizing scheme based on matching on unit cell walls. No quasi-periodic Green's function nor lattice sums are needed; the solver is thus robust for all scattering parameters including Wood's anomalies.

Alex Barnett
Department of Mathematics
Dartmouth College
ahb@math.dartmouth.edu

Leslie Greengard, Zydrunas Gimbutas
Courant Institute
New York University
greengard@courant.nyu.edu,
zydrunas.gimbutas@gmail.com

MS47

A Fast Direct Solver for Quasi-periodic Scattering Problems

In this talk, we present an integral equation based direct solver for the two dimensional scattering of time-harmonic plane waves from an infinite periodic array of obstacles in a homogeneous background medium. By coupling a recently developed periodizing scheme, potential theory and "fast"

direct linear algebra, the direct solver has $O(N)$ complexity and is robust for all incident angles. For design problems where multiple incident angles are needed, the solver is extremely efficient.

Adrianna Gillman
University of Colorado at Boulder
Department of Applied Mathematics
adrianna.gillman@dartmouth.edu

Alex Barnett
Department of Mathematics
Dartmouth College
ahb@math.dartmouth.edu

MS47

Efficient Representations for the Fundamental Solutions of Stokes Flow

In this talk, we present a simple and efficient method for the evaluation of the fundamental solutions for Stokes flow in a half space. We show that both the direct and image contributions can be decomposed using harmonic functions with a physically meaningful interpretation of the corresponding components. This decomposition is easy to incorporate with existing FMM libraries. This is joint work with Leslie Greengard.

Zydrunas Gimbutas
Courant Institute
New York University
zydrunas.gimbutas@gmail.com

MS47

A Direct Solver for Variable Coefficient Elliptic PDEs

The talk describes a highly accurate technique for solving elliptic PDEs with variable coefficients and smooth solutions. The domain is tessellated into squares, and the differential operator is discretized via high order (p=10 or 20) spectral differentiation on each square. A hierarchical direct solver is used to solve the resulting discrete system. The method is very efficient; e.g., a Helmholtz problem on a domain of size 200x200 wavelengths is solved to ten digits of accuracy in ten minutes on a standard laptop (using 6M degrees of freedom).

Gunnar Martinsson
Univ. of Colorado at Boulder
per-gunnar.martinsson@colorado.edu

MS48

A Semi-implicit Gradient-augmented Level Set Method

Here a semi-implicit formulation of the gradient-augmented level set method is presented. Stability is enhanced by the addition of a high-order smoothing term added to the gradient-augmented level set equations. The new approach is a hybrid Lagrangian-Eulerian method which allows for the investigation of flows based on the curvature of the interface and the intrinsic Laplacian of the curvature. In this talk the method will be outlined and sample results will be presented. The influence of the smoothing term on the stability and accuracy of the method will also be shown.

Ebrahim M. Kolahdouz
SUNY Buffalo
mkolahdo@buffalo.edu

David Salac
University at Buffalo - SUNY
davidsal@buffalo.edu

MS48

Efficient Synchronous Update of Multiple Level Set Functions using Jet Schemes

In this talk we will present a technique to evolve a reference map. This map can then be used to reference more than one level set function, and thus provide for an efficient way to advect multiple level set functions. Time permitting, we will present various applications for this approach.

Jean-Christophe Nave, Olivier Mercier
McGill University
jcnave@math.mcgill.ca, oli.mercier@gmail.com

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

Benjamin Seibold
Temple University
seibold@temple.edu

MS48

An Augmented Fast Marching Method for Level Set Reinitialization

Here a new gradient-augmented fast marching method for reinitialization of level set functions is presented. This method will calculate the signed distance function and up to the second-order derivatives of the signed distance function for arbitrary interfaces. Sample results in both two- and three-dimensions will be show that the resulting level set and curvature field are smooth even for coarse grids.

David Salac
University at Buffalo - SUNY
davidsal@buffalo.edu

MS48

Jet Schemes for Hamilton-Jacobi Equations using an Evolve-and-project Framework

Jet schemes are based on tracking characteristics and using suitable Hermite interpolations to achieve high order. For Hamilton-Jacobi equations, the characteristic equations are in general nonlinear, and explicit schemes for solving characteristic equations can yield incorrect results. We therefore propose an implicit update rule that is based on solving a constrained polynomial optimization problem in each grid cell, and then reconstructing the solution from Hermite interpolations and evolving it in time.

Dong Zhou
Temple University
dong.zhou@temple.edu

MS49

Fluctuating Hydrodynamics for Dynamic Simulations of Coarse-Grained Implicit-Solvent Models of

Lipid Bilayer Membranes

Many coarse-grained models have been developed for efficient equilibrium studies of lipid bilayer membranes by treating implicitly interactions mediated by the solvent. However, for problems involving dynamics, such CG studies require accounting for the momentum transfer through the missing solvent degrees of freedom. We introduce a new thermostat for such CG models and computational methods based on fluctuating hydrodynamics. We show our approach yields results significantly different than conventional Langevin dynamics. We then present a number of simulation results both for self-assembled planar-bilayer sheets and vesicles.

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

MS49

Modeling of Thermal Fluctuations in Multicomponent Reacting Systems

We present a multicomponent version of the fluctuating Navier-Stokes equations that includes detailed transport and chemical reactions. The resulting system includes stochastic flux terms for momentum, energy and species transport and a Langevin-type model for fluctuations in the chemical reactions. We discuss isses in the numerical solution of the resulting systems and illustrate the impact of fluctuations on numerical solutions Turing patterns.

John B. Bell
CCSE
Lawrence Berkeley Laboratory
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

Kaushik Balakrishnan
CCSE
Lawrence Berkeley Laboratory
kaushikb@lbl.gov

MS49

Low Mach Number Fluctuating Hydrodynamics of Diffusively Mixing Fluids

We formulate low Mach number fluctuating hydrodynamic equations appropriate for modeling diffusive mixing in mixtures of fluids of unequal density. These equations eliminate the fast isoentropic fluctuations in pressure associated with the propagation of sound waves by replacing the equation of state with a local thermodynamic constraint. We demonstrate that the low Mach number model preserves the spatio-temporal spectrum of the slower diffusive fluctuations in the linearized setting. We develop a strictly conservative finite-volume spatial discretization of the low Mach number fluctuating equations in both two and three dimensions. We construct several explicit Runge-Kutta temporal integrators that strictly maintain the equation of state constraint. The resulting spatio-temporal discretization is second-order accurate deterministically and maintains fluctuation-dissipation balance in the stochastic setting. We apply our algorithms to model the development of giant concentration fluctuations in the presence of concentration gradients, and investigate the validity of common simplifications neglecting the spatial non-homogeneity of density and transport properties. As a validation of the low Mach number fluctuating equations and our algorithm, we perform simulations of diffusive mixing of two fluids of different densities in two dimensions and compare the results of low Mach number continuum simulations to hard-disk molecular dynamics simulations. Excellent agreement is observed between the particle and continuum simulations of giant fluctuations during time-dependent diffusive mixing.

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
aleks.donev@nyu.edu

Andy Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

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

Alejandro Garcia
San Jose State University
algarcia@algarcia.org

Thomas Fai
Courant Institute
tfai@cims.nyu.edu

MS49

Implicit and Explicit Solvent Models for the Simulation of a Single Polymer Chain in Solution: Lattice Boltzmann vs Brownian Dynamics

We present a comparative study of two computer simulation methods to obtain static and dynamic properties of dilute polymer solutions. The first approach is a recently established hybrid algorithm based upon dissipative coupling between Molecular Dynamics and lattice Boltzmann (LB), while the second is standard Brownian Dynamics (BD) with fluctuating hydrodynamic interactions. Applying these methods to the same physical system (a single polymer chain in a good solvent in thermal equilibrium) allows us to draw a detailed and quantitative comparison in terms of both accuracy and efficiency. It is found that the static conformations of the LB model are distorted when the box length L is too small compared to the chain size. Furthermore, some dynamic properties of the LB model are subject to an L^{-1} finite size effect, while the BD model directly reproduces the asymptotic $L \to \infty$ behavior. Apart from these finite size effects, it is also found that in order to obtain the correct dynamic properties for the LB simulations, it is crucial to properly thermalize all the kinetic modes. Only in this case, the results are in excellent agreement with each other, as expected. Moreover, Brownian Dynamics is found to be much more efficient than lattice Boltzmann as long as the degree of polymerization is not excessively large.

Burkhard Duenweg

Max Planck Institute Mainz
duenweg@mpip-mainz.mpg.de

MS50

Computational Methods for Parametric Sensitivities in the Chemical Kinetic Context

I will discuss methods for the computation of parametric sensitivities for stochastically modeled biochemical systems. In particular, I will discuss a new finite difference method that arises from a non-trivial coupling of a nominal and perturbed process. The method is easy to implement and produces an estimator with a much lower variance than the previous state of the art for a wide array of systems.

David F. Anderson
Department of Mathematics
University of Wisconsin Madison
anderson@math.wisc.edu

MS50

Infinite Swapping Schemes for Accelerated Monte Carlo Approximation

Abstract not available at time of publication.

Paul Dupuis
Division of Applied Mathematics
Brown University
dupuis@dam.brown.edu

MS50

Parameterisation and Multilevel Approximations of Coarse-grained Dynamics in KMC Simulations

We present an information-theoretic approach to parameterisation of coarse-grained dynamics defined as continuous time Markov chains. Rates of the coarse-grained process are parameterised and optimal parameters are selected by minimisation of the relative entropy on the path space. This approach extends techniques also known as inverse Monte Carlo to models with non-equilibrium stationary states, for example, systems driven by external parameters or reaction-diffusion systems in catalysis. This a joint work with P. Plechac.

Markos A. Katsoulakis
UMass, Amherst
Dept of Mathematics
markos@math.umass.edu

MS50

Off-lattice KMC Simulation of Heteroepitaxial Growth

In this talk I will review previous work on weakly off-lattice KMC simulation of quantum dots, including recent progress on a two-scale domain decomposition approach. This method is used to study various phenomena that take place during heteroepitaxial growth. For example, it is demonstrated that faceted quantum dots occur via the layer-by-layer nucleation of pre-pyramids on top of a critical layer with faceting occurring by anisotropic surface diffusion. It is also shown that the dot growth is enhanced by the depletion of the critical layer which leaves behind a wetting layer. Capping simulations provide insight into the mechanisms behind dot erosion and ring formation. I will then discuss efforts to extend these methods to fully off-lattice simulations. This is joint work with Peter Smereka and Henry Boateng.

Tim Schulze
University of Tennessee
Department of Mathematics
schulze@math.utk.edu

MS51

An Efficient and Scalable Lanczos-based Eigensolver for Multi-core Systems

We describe an efficient scalable symmetric iterative eigensolver for finding few lowest eigenpairs on distributed multi-core platforms. We show over 80% computational efficiency by major reductions in communication overhead for the SpMV and basis orthogonalization tasks. In particular, we present strategies for hiding communication on multi-core platforms. We demonstrate the effectiveness of these techniques by reporting the performance improvements in large-scale eigenvalue problems arising in nuclear structure calculations.

Hasan Metin Aktulga, Chao Yang
Lawrence Berkeley National Lab
hmaktulga@lbl.gov, CYang@lbl.gov

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

Pieter Maris, James Vary
Iowa State University
pmaris@iastate.edu, jvary@iastate.edu

MS51

Density Functional Electronic Band Structure Calculations with a Complex Moment Based Eigensolver

First-principles electronic band structure calculations based on the density functional theory is one of the best choices for understanding and predicting phenomena in material sciences. In electronic band structure calculations of large systems, one needs to solve large sparse interior eigenproblems. We present an efficient approach for solving these eigenproblems using a complex moment based eigensolver. Numerical experiments on the K computer are also presented.

Yasunori Futamura
University of Tsukuba
futamura@mma.cs.tsukuba.ac.jp

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

Shinnosuke Furuya, Jun-Ichi Iwata
The University of Tokyo
furuya@comas.t.u-tokyo.ac.jp, iwata@ap.t.u-tokyo.ac.jp

MS51

Computing a Large Number of Eigenpairs on Mulit-/many-core Systems

We examine two techniques for compute a relatively large

number of eigenpairs of a sparse symmetric matrix. One is based on multiple shift-invert Lanczos. The other is based on a projection method that makes use of a contour integral representation of a projection operator. We compare the pros and cons of both approaches and discuss a number of practical issues of implementing these methods on multi/many core systems

Chao Yang, Hasan Metin Aktulga
Lawrence Berkeley National Lab
CYang@lbl.gov, hmaktulga@lbl.gov

Christopher Haine
University of Versaille
christopher.haine0@gmail.com

Lin Lin
Lawrence Berkeley Lab
linline@lbl.gov

MS51

Computing Eigenspace by a Penalty Approach

Abstract not available at time of publication.

Yin Zhang
Rice University
Dept of Comp and Applied Math
yzhang@rice.edu

Xin Liu
Academy of Mathematics and Systems Science
Chinese Academy of Sciences
liuxin@lsec.cc.ac.cn

Zaiwen Wen
Department of Mathematics
Shanghai Jiaotong University
zw2109@sjtu.edu.cn

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

MS52

Limited Data-Driven Uncertainty Quantification

In this talk, we will present work on limited data-driven stochastic collocation approach to include the effect of uncertain design parameters during complex multi-physics simulation of microsystems. The proposed framework comprises of two key steps: firstly, probabilistic characterization of the input uncertain parameters based on available experimental information, and secondly, propagation of these uncertainties through the predictive model to relevant quantities of interest. The uncertain input parameters are modeled as independent random variables, for which the distributions are estimated based on available experimental observations, using a nonparametric diffusion mixing based estimator. The diffusion based estimator derives from the analogy between the kernel density estimation (KDE) procedure and the heat dissipation equation, and constructs density estimates that are smooth and asymptotically consistent. The diffusion model allows for the incorporation of the prior density and leads to an improved density estimate, in comparison to the standard KDE approach, as demonstrated through several numerical examples. Following the characterization step, the uncertainties are propagated to the output dependent variables using the stochastic collocation approach, based on sparse grid interpolation. The stochastic multiphysics framework is used to study uncertainties in microsystems.

Narayana R. Aluru
Department of Mechanical Science and Engineering
Beckman Institute, University of Illinois
aluru@illinois.edu

MS52

A Decomposition Approach to Uncertainty Analysis of Multidisciplinary Systems

Uncertainty analysis for complex systems can become cumbersome and computationally intractable. This talk describes an approach for decomposing and distributing the uncertainty analysis task amongst the various components comprising a complex system. The approach draws on multidisciplinary analysis and optimization methods, density estimation, and sequential Monte Carlo methods. The distributed multidisciplinary uncertainty analysis approach is provably convergent and is compared to a traditional all-at-once Monte Carlo uncertainty analysis approach for several examples.

Sergio Amaral
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
samaral@mit.edu

Doug Allaire, Karen E. Willcox
Massachusetts Institute of Technology
dallaire@mit.edu, kwillcox@MIT.EDU

MS52

A Finite Element Method for Density Estimation with Gaussian Priors

A variational problem characterizing the density estimator defined by the maximum a posteriori method with Gaussian process priors is derived. It is shown that this problem is well posed and can be solved with Newton's method. Numerically, the solution is approximated by a Galerkin/finite element method with piecewise multilinear functions on uniform grids. Error bounds for this method are given and numerical experiments are performed for one-, two-, and three-dimensional examples.

Markus Hegland
Australian National Unversity
Markus.Hegland@anu.edu.au

MS52

Density Estimation for Large Datasets with Sparse Grids

Kernel density estimation can become computationally expensive for large data sets. Furthermore, its performance highly depends on the choice of the kernel bandwidth. Our sparse-grid-based method overcomes these drawbacks to some extent. We give details on how to estimate density functions on sparse grids and show numerical experiments with large data sets to demonstrate that our method is competitive with respect to accuracy and superior to conventional approaches with respect to computational complexity.

Benjamin Peherstorfer

SCCS, Department of Informatics
Technische Universität München
pehersto@in.tum.de

MS53

Analysis of Ship Structural Hydroelasticity by using a Fully-coupled Higher-order BEM and FEM

This study considers the problem of ship hydroelasticity, which is an important technical issue in the design of ultra-large marine vessels. To analyze the ship structural hydroelasticity, two initial boundary value problems should be solved simultaneously: hydrodynamic problem for ship motion and dynamic structural problem. In the present study, a partitioned method is applied for the coupled fluid-structure interaction problem. The fluid domain surrounding a flexible body is solved using a Rankine panel method based on higher-order B-spline basis function, and the structural domain is handled with a three-dimensional finite element method. The two distinct methods are fully coupled in the time domain by using an implicit iterative scheme. The detailed numerical method with stable time-marching is described, and the validation of the developed method is introduced. The numerical results include a real ship application.

Yonghwan Kim
Seoul National University
Naval Architecture and Ocean Engineering
yhwankim@snu.ac.kr

MS53

Flexible Ring Flapping in a Uniform Flow

An improved version of the immersed boundary method for simulating an initially circular or elliptic flexible ring pinned at one point in a uniform flow has been developed. A penalty method derived from fluid compressibility was used to ensure the conservation of the internal volume of the flexible ring. A new bistability phenomenon was observed: for certain aspect ratios, two periodically flapping states coexist with different amplitudes in a particular Reynolds number range.

Boyoung Kim, Hyung Jin Sung
Department of Mechanical Engineering
KAIST
kbyshaker@kaist.ac.kr, hjsung@kaist.ac.kr

MS53

Numerical Modeling of the Interaction between Moving Solid Structures and Two-phase Fluid Flows: Application in Ocean Wave Energy Converters

We are developing a computational tool for interaction analysis of two-phase fluid flows with a moving solid object. Two-step projection method with GPU acceleration solves the flow equations. The volume-of-fluid method tracks the fluid interfaces while the fast-fictitious-domain method simulates the interactions between a moving solid object and two-phase fluid flows. A geometrical reconstruction is employed to handle liquid-gas-solid interfaces. We will present results of canonical tests and preliminary results on ocean wave energy converters.

Amirmahdi Ghasemi, Ashish Pathak
University of Massachusetts Dartmouth
aghasemi@umassd.edu, apathak@umassd.edu

Mehdi Raessi
Department of Mechanical Engineering
University of Massachusetts Dartmouth
mraessi@umassd.edu

MS53

Generalized Fictitious Methods for Fluid-structure Interactions

We present two methods for fluid-structure interaction (FSI) problems: the fictitious pressure method and the fictitious mass method. For simplified problems we obtain expressions for the convergence rate index, which demonstrates a similarity of fictitious methods to the FSI approach with Robin boundary conditions. In numerical tests, we verify the selection of optimal values for the fictitious parameters, and develop an empirical analysis for complex geometries and apply it to 3D patient-specific flexible brain arteries.

Yue Yu
Applied Mathematics
Brown
yueyucarol@gmail.com

Hyoungsu Baek
Massachusetts Institute of Technology
Mathematics Department
hbaek@math.mit.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS54

Parallel Computing for Long-Time Simulations of Calcium Waves in a Heart Cell

The flow of calcium ions in a heart cell is modeled by a system of time-dependent reaction-diffusion equations. The large number of calcium release sites requiring high-resolution meshes and the need for large final times require parallel computing for effective simulations. Using Krylov subspace methods offers the opportunity for efficient parallel computing, with choices in parallel algorithms driven by the special considerations of the available CPU, GPU, or hybrid CPU-GPU architecture.

Matthias K. Gobbert
University of Maryland, Baltimore County
Department of Mathematics and Statistics
gobbert@umbc.edu

Xuan Huang
Department of Mathematics and Statistics
University of Maryland, Baltimore County
hu6@umbc.edu

Stefan Kopecz
Institute for Mathematics
University of Kassel
kopecz@mathematik.uni-kassel.de

Philipp Birken
Department 10, Mathematics and Natural Sciences
University of Kassel, Germany
birken@mathematik.uni-kassel.de

Andreas Meister
University of Kassel
Institute of Mathematics
meister@mathematik.uni-kassel.de

MS54

High End of Mesh-based PDE Simulations

Abstract not available at time of publication.

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

MS54

A Memory Efficient Finite Volume Method for Advection-Diffusion-Reaction Systems

Advection-Diffusion-Reaction Systems occur in a wide variety of applications. The use of matrix-free parallel methods allows for simulations on high resolution meshes, since no memory is needed to store system matrices and the paralellism distributes the workload among several processes. We present a method of this type based on a finite volume discretization and demonstrate its performance simulating calcium flow in human heart cells.

Jonas Schäfer
Institute of Mathematics
University of Kassel
jonas.schaefer@t-online.de

Xuan Huang
Department of Mathematics and Statistics
University of Maryland, Baltimore County
hu6@umbc.edu

Stefan Kopecz
Institute for Mathematics
University of Kassel
kopecz@mathematik.uni-kassel.de

Philipp Birken
Department 10, Mathematics and Natural Sciences
University of Kassel, Germany
birken@mathematik.uni-kassel.de

Matthias K. Gobbert
University of Maryland, Baltimore County
Department of Mathematics and Statistics
gobbert@umbc.edu

Andreas Meister
University of Kassel
Institute of Mathematics
meister@mathematik.uni-kassel.de

MS54

Asynchronous Multilevel Algorithms

Iterative algorithms for solving elliptic PDE are traditionally formulated in an imperative style that prescribes a fixed order of operations. In a parallel environment, this implies a rigid sequence of communication and synchronization steps. In the talk we will present alternative asynchronous strategies based on a greedy and randomized version of the Gauss-Southwell algorithm and the theory of stable multilevel splittings.

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ruede@informatik.uni-erlangen.de

MS55

Reduced-space inexact-Newton-Krylov for High-dimensional Optimization

Abstract not available at time of publication.

Jason E. Hicken
postdoctoral fellow, Stanford University
hickej2@rpi.edu

MS55

A Matrix-Free Augmented Lagrangian Approach to Structural Optimization

In structural optimization, it is common practice to aggregate the failure constraints to reduce the cost of computing the constraint Jacobian. We present an alternative approach in which the constraint Jacobian is never formed; only matrix-vector products are needed. We illustrate our approach using an augmented Lagrangian optimization algorithm and several example problems. We also show how adopting the matrix-free approach can lead to lower-mass structural designs compared to the aggregation approach.

Andrew Lambe
University of Toronto
lambe@utias.utoronto.ca

Joaquim Martins
University of Michigan
jrram@umich.edu

Sylvain Arreckx
Ecole Polytechnique de Montreal
sylvain.arreckx@gmail.com

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

MS55

Adjoint-Based Equivalent Area Methods for Supersonic Low-Boom Design

Abstract not available at time of publication.

Francisco Palacios
Stanford University
fpalacios@stanford.edu

MS55

Practical Experience with a Multi-Objective Model-Management Framework Optimization Method

Solving multi-objective optimization problems that involve computationally expensive functions is a normal part of

many engineering applications. Runtime issues are magnified in a multidisciplinary design optimization setting. Normal-Boundary Intersection (NBI) is a multi-objective optimization method. Running NBI directly requires prohibitive computational cost. Running on surrogate approximations to the simulation fails to produce sufficiently accurate solutions. We combine the use of surrogates and simulations in a way similar to the general surrogate management framework. We conclude with a representative aircraft design.

Joseph P. Simonis
The Boeing Company
joseph.p.simonis@boeing.com

Evin Cramer
Boeing
Math & Computing
evin.j.cramer@boeing.com

Joerg M. Gablonsky
Mathematics & Engineering Analysis
Phantom Works - The Boeing Company
Joerg.M.Gablonsky@boeing.com

Laura Lurati, Paul Sellers
The Boeing Company
laura.b.lurati@boeing.com, paul.s.sellers@boeing.com

MS56

Four Modules for Teaching CUDA in a Computational Science Context

Four modules developed for the UPEP project deal with using NVIDIA's CUDA programming environment to perform scientific computations on graphics cards. The first is a straightforward introduction to CUDA in the context of matrix multiplication. Two modules introduce dynamic programming and show how such problems can be solved within CUDA, and the last shows how to integrate CUDA with OpenGL in order to do high performance scientific visualization. We will give an overview and demonstration.

Robert Hochberg
University of Dallas
Department of Mathematics
hochberg@udallas.edu

MS56

Supporting Petascale Education: The Blue Waters Undergraduate Petascale Education Program

The Blue Waters project, in collaboration with the National Computational Science Institute and national HPC programs, has launched a coordinated effort to prepare current and future generations of students for the growing complexity of computing paradigms. To support this effort, the Blue Waters Undergraduate Petascale Education Program (BW-UPEP) seeks to promote understanding and interest in petascale computing and its applications among undergraduate students and faculty by supporting the development of high quality undergraduate curricular materials.

Jennifer K. Houchins
Shodor Education Foundation Inc.
jhouchins@shodor.org

MS56

Classroom Explorations of N-body Gravitational Simulations using GalaxSeeHPC

GalaxSeeHPC is the most recent release of the GalaxSee N-body solver designed for classroom study of gravitational systems, adding periodic boundary conditions, increased control over input, and additional force calculation methods to the previous GalaxSee-MPI code. GalaxSeeHPC and its corresponding two Blue Waters Petascale Education Project modules allow students to study both the techniques required to scale N-body solutions to millions of particles as well as the new science enabled by larger N.

David A. Joiner
Kean University
Assistant Professor
djoiner@kean.edu

MS56

Biofilms: Linked for Life (Understanding Biofilms through Modeling and Simulation)

Most microbial organisms do not exist as individuals, but within communities of interconnected members. In this presentation, we discuss development of a simulation of such biofilm structural growth appropriate for modeling, simulation, or high performance computing courses. Consideration of cellular automaton simulations, boundary conditions, and diffusion can empower students to develop similar simulations for other applications. Moreover, extensions of the basic model can illustrate and motivate the need for high performance computing in computational science.

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

MS57

Bernstein-Bezier Techniques in High Order Finite Element Analysis

Algorithms are presented that enable the element matrices for the standard finite element space, consisting of continuous piecewise polynomials of degree n on simplicial elements in R^d, to be computed in optimal complexity $O(n^2 d)$. The algorithms (i) take account of numerical quadrature; (ii) are applicable to non-linear problems; and, (iii) do not rely on pre-computed arrays containing values of one-dimensional basis functions at quadrature points (although these can be used if desired). The elements are based on Bernstein-Bezier polynomials and are the first to achieve optimal complexity for the standard finite element spaces on simplicial elements.

Mark Ainsworth
University of Strathclyde
Department of Mathematics
mark.ainsworth@live.com

MS57

High-order Methods for Fractional Differential Equations

Modeling of non-classic phenomena in science, finance, biology and engineering increasing utilizes fractional differential equations, highlighting the need for robust, accurate and efficient computational models for such operators. However, despite fractional calculus being almost as old as classic calculus, the development of computational techniques is less advanced. In this talk we introduce discontinuous Galerkin and spectral penalty methods for fractional PDEs, consider accuracy and stability, and illustrate performance on test problems.

Jan S. Hesthaven
Brown University
Division of Applied Mathematic
Jan.Hesthaven@Brown.edu

Weihua Deng
Department of Mathematics
Lanzhou University, China
dengwh@lzu.edu.cn

Qinwu Xu
Division of Applied Mathematics
Brown University, Providence, RI
qinwu_xu@brown.edu

MS57

High-order Virtual Element Methods

The Virtual Element spaces are just like the usual Finite Element spaces with the addition of suitable non polynomial functions. This is not a new idea; the novelty here is to choose the degrees of freedom in such a way that the elementary stiffness matrix can be computed without actually computing these non polynomial functions. In doing that we can easily deal with complicated element geometries and higher order methods.

Alessandro Russo
Milan University
alessandro.russo@unimib.it

MS58

Mint: A User-fiendly C-to-CUDA Code Translator

Aiming at automated source-to-source code translation from C to CUDA, we have developed the Mint framework. Users only need to annotate serial C code with a few compiler directives, specifying host-device data transfers plus the parallelization depth and granularity of loop nests. Mint then generates CUDA code as output, while carrying out on-chip memory optimizations that will greatly benefit 3D stencil computations. Several real-world applications have been ported to GPU using Mint.

Xing Cai
Simula Research Laboratory
1325 Lysaker, Norway
xingca@simula.no

Didem Unat
Lawrence Berkeley National Laboratory
dunat@lbl.gov

Scott Baden
University of California, San Diego
baden@eng.ucsd.edu

MS58

PyOP2 - An Abstraction for Performance-portable Simulation Software

OP2 is an abstraction which expresses mesh-based simulation code in terms of numerical kernels applied in parallel over a mesh. This enables performance portability and frees the developer from the increasingly complex details of parallel programming. Here we present PyOP2, a just-in-time compiled OP2 which delays kernel compilation and parallel strategy formation until runtime. PyOP2 presents a high-level programmable interface. Alternatively PyOP2 code for finite elements may be generated automatically by the FEniCS system.

Carlo Bertolli
Imperial College London
c.bertolli@imperial.ac.uk

Mike Giles
University of Oxford
mike.giles@maths.ox.ac.uk

David Ham
Imperial College
david.ham@imperial.ac.uk

Paul Kelly, Nicolas Loriant, Graham Markall
Imperial College London
p.kelly@imperial.ac.uk, n.loriant@imperial.ac.uk, graham.markall08@imperial.ac.uk

Lawrence Mitchell
University of Edinburgh
lawrence.mitchell@ed.ac.uk

Giham Mudalige
University of Oxford
gihan.mudalige@oerc.ox.ac.uk

Florian Rathgeber
Imperial College London
f.rathgeber10@imperial.ac.uk

MS58

Achieving High Performance and Portability in Stencil Computations

Physis is an application framework for stencil computations that is designed to achieve both performance and productivity in large-scale parallel computing systems, in particular heterogeneous GPU-accelerated supercomputers. The framework consists of an implicitly parallel domain-specific language for stencil computations and its translators for various target parallel architectures. This talk presents the current status of the framework and its performance studies using several application benchmarks.

Naoya Maruyama
RIKEN Advanced Institute for Computational Science
nmaruyama@riken.jp

Satoshi Matsuoka
Tokyo Insitute of Technology
matsu@is.titech.ac.jp

MS58

Automating the Communication-computation Overlap with Bamboo

MPI has been a popular standard for implementing HPC applications. However, codes running at scale would require significant optimizations, one of which is the code restructuring to mask communication. Such optimizations require aggressive effort and complicate the software maintenance. We present Bamboo, a translator that automates the optimization by transforming MPI into a data-driven form. Experiments on up to 98304 processors demonstrate that Bamboo significantly speeds up its input, meeting/exceeding the performance of hand-optimized variants.

Tan Nguyen, Scott Baden
University of California, San Diego
nnguyenthanh@eng.ucsd.edu, baden@eng.ucsd.edu

MS59

Assessing the Predictive Capabilities of Miniapps

Under what conditions does an application proxy represent a key performance characteristic in a full application? In this talk we define a methodology for comparing applications and miniapps, providing a framework for reasoning about important performance-impacting issues. This methodology will be illustrated using miniapps from the Mantevo project, configured to represent the runtime characteristics of some key mission application codes.

Richard Barrett, Paul Crozier, Douglas Doerfler
Sandia National Laboratories
rfbarre@sandia.gov, pscrozi@sandia.gov,
dwdoerf@sandia.gov

Simon D. Hammond
Scalable Computer Architectures
Sandia National Laboratories
sdhammo@sandia.gov

Michael A. Heroux, Paul Lin, Heidi K. Thornquist,
Timothy G. Trucano, Courtenay T. Vaughan
Sandia National Laboratories
maherou@sandia.gov, ptlin@sandia.gov,
hkthorn@sandia.gov, tgtruca@sandia.gov,
ctvaugh@sandia.gov

MS59

Programming Model Exploration and Efficiency Modeling using Mini and Proxy Applications

Application proxies and kernels of applications that are small and easily run can serve as a useful test bed for a variety of new languages. As languages are developed to combat the increasingly complex architectural changes, they become a crucial part of the development effort. In this talk we discuss some recent results of various programming models applied to kernel applications.

Alice Koniges
Lawrence Berekely Laboratory
aekoniges@lbl.gov

MS59

Programming Models using Workflows from Proxies

Abstract not available at time of publication.

James Sexton
IBM Research
sextonjc@us.ibm.com

MS59

Examples of Codesign Using Application Proxies

Architectures are undergoing a sea change with a transition from cheap-memory-expensive-flops to cheap-flops-expensive memory/power. This requires us to rethink our applications to get the most out of a differently balanced machine than we have been used to for the past 15 years. Yet, rewriting entire applications is not practical. Proxy applications have emerged as one possible solution to this thorny problem. In this talk we explore the meaning and role of proxy applications in co-design and their effectiveness as a vehicle for reducing many possible paths forward to a manageable few.

Sriram Swaminarayan
Los Alamos National Laboratory
sriram@lanl.gov

MS60

Accuracy of Some Finite-difference and Finite-element Methods for Wave Propagation at a Fluid-solid Interface

The problem of modeling wave propagation in media with solid and fluid regions has many applications in geophysics, engineering and medicine. We have performed a grid-dispersion analysis and numerical computations to compare the performance of the following numerical methods in handling fluid-solid interfaces: the spectral-element method, and several versions of the the finite-difference method. We conclude that a first-order velocity stress formulation can be used in dealing with fluid-solid layers without using staggered grids necessarily.

Jonas D. De Basabe
The University of Texas at Austin
Institute for Computational Engineering and Sciences
jonas@cicese.mx

Mrinal Sen
Institute of Geophysics
University of Texas
mrinal@ig.utexas.edu

MS60

The Discontinuous Enrichment Method for Wave Propagation

Wave propagation problems in the medium frequency regime are computationally challenging. One avenue of research pursues higher-order discretization methods that can deliver both accuracy and computational efficiency at smaller mesh resolutions. The Discontinuous Enrichment Method (DEM) is one example which distinguishes itself from competing approaches in the additional information it incorporates in the approximation method. It has shown a significant promise for acoustic and structural acoustic problems and therefore is reviewed here, together with new applications to multi-scale problems.

Charbel Farhat, Radek Tezaur
Stanford University
CFarhat@stanford.edu, rtezaur@stanford.edu

MS60

Dispersion Reducing Techniques for FDTD Schemes

Abstract not available at time of publication.

Bezalel Finkelstein
Sami Shamoon College of Engineering, Israel
fibk100@bezeqint.net

MS60

Analysis of High Order FDTD Methods for Maxwell's Equations in Dispersive Media

We consider models for electromagnetic wave propagation in linear dispersive media which include ordinary differential equations for the electric polarization coupled to Maxwell's equations. We discretize these models using high order finite difference methods and study the properties of the corresponding discrete models. In this talk we will present the stability, dispersion and convergence analysis for a class of finite difference methods that are second order accurate in time and have arbitrary (even) order accuracy in space in one and two dimensions.

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

MS60

Dispersion Reduction for Acoustic Wave Equation using m-adaptation

We present a novel discretization strategy, dubbed m-adaptation, for acoustic wave equation in time domain in 2D and 3D. This strategy is based on the optimization within a family of second-order accurate Mimetic Finite Difference (MFD) discretizations. The optimized scheme has a computational complexity of second-order scheme and is shown to be fourth-order accurate in dispersion on rectangular and cubic meshes. On square meshes the anisotropy is shown to be sixth-order accurate.

Vitaliy Gyrya
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
vitaliy_gyrya@lanl.gov

Konstantin Lipnikov
Los Alamos National Laboratory
lipnikov@lanl.gov

MS61

Upwind Methods for Second-order Wave Equations

In this talk we discuss a newly developed class of high-order accurate upwind schemes for the wave equation in second-order form. By working directly on the equations in second-order form, we avoid issues of compatibility that can arise when converting from second- to first-order formulations. The schemes are based on embedding the solution to a local Riemann-type problem that uses d'Alembert's exact solution. High-order accuracy is obtained using a single-step space-time procedure. The result is a method that is highly efficient in both memory and speed, and has very attractive accuracy and robustness characteristics. Stability and accuracy are discussed using normal-mode theory, and the efficacy of the approach is demonstrated for a set of challenging test problems.

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banks20@llnl.gov

MS61

A Performance Study of a Massively Parallel Water Wave Model for Engineering Applications

We present results of a unique study of several computational techniques suitable for improving performance of a fully nonlinear and dispersive water wave model intended for use in coastal engineering applications and performance critical real-time ship simulator software. The model is attractive for description of a broad range of wave phenomena, flow kinematics and wave-structure interactions. A flexible-order finite difference method is used for accurate discretization and efficient and scalable mapping to modern many-core hardware.

Allan P. Engsig-Karup
Technical University of Denmark
Department of Informatics and Mathematical Modeling
apek@imm.dtu.dk

Stefan L. Glimberg
DTU Informatics
Technical University of Denmark
slgl@imm.dtu.dk

MS61

DG Schemes for Quadrature-based Moment-closure Models of Plasma

The dynamics of collisionless plasma can be simulated using kinetic or fluid models. Kinetic models are valid over most of the spatial and temporal scales that are of physical relevance in many application problems; however, they are computationally expensive due to the high-dimensionality of phase space. Fluid models have a more limited range of validity, but are generally computationally more tractable than kinetic models. One critical aspect of fluid models is the question of what assumptions to make in order to close the fluid model. In this work we study so-called *quadrature-based moment closure models* for collisionless plasma. We develop high-order discontinuous Galerkin (DG) schemes for these models, and in particular, we consider several important issues in the discretization of these models, including hyperbolicity, positivity, and moment-realizability. The resulting schemes are tested on several standard collisionless plasma test cases.

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

Yongtao Cheng
University of Wisconsin - Madison
Department of Mathematics
cheng@math.wisc.edu

MS61

ManyClaw: Slicing and Dicing Riemann Solvers for Next Generation Highly Parallel Architectures

Next generation computer architectures, e.g. Intel MIC and GPUs, include an order of magnitude increase in intranode parallelism over traditional CPUs. We present ManyClaw, a project that explores the exploitation of this intranode parallelism in hyperbolic PDE solvers via the Clawpack software package. Our goal is to separate the low level parallelism and the physical equations thus providing users the capability to leverage intranode parallelism without explicitly writing code to take advantage of newer architectures.

Andy R. Terrel
Texas Advanced Computing Center
University of Texas
aterrel@tacc.utexas.edu

Kyle T. Mandli
University of Texas at Austin
ICES
kyle@ices.utexas.edu

MS62

GPU-Accelerated Implementations of B-Spline Signal Processing Operations for FFD-Based Image Registration

B-spline signal processing operations are widely used in the analysis of two and three-dimensional images. In this talk, we investigate GPU-accelerated implementations of basic B-spline signal processing operations in CUDA, including direct transformations, indirect transformations, computation of partial derivatives, and interpolation. We then illustrate how these operations can be used to design a free-form deformation based image registration algorithm that enables dense control point spacing.

Nathan D. Cahill
School of Mathematical Sciences
Rochester Institute of Technology
nathan.cahill@rit.edu

Alex Karantza, Sonia Lopez Alarcon
Department of Computer Engineering
Rochester Institute of Technology
alex.karantza@gmail.com, slaeec@rit.edu

MS62

Model Mis-specification - in Search of the Missing Link

Frequently, since our ability to simulate data is limited, only a discrete, incomplete observation operator can be specified. It will benefit a lot if we could supplement such reduced model somehow. In this study, we developed a nuclear norm stochastic optimization technique to supplement the model with rank restriction given a training set of models and data.

Ning Hao
Tufts University
Department of Mathematics
ning.hao@tufts.edu

Lior Horesh
IBM Research
lhoreshus.ibm.com

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

MS62

Title Not Available at Time of Publication

Abstract not available at time of publication.

Harald Koestler
University of Erlangen-Nuremberg
harald.koestler@informatik.uni-erlangen.de

MS62

Ultra-low-dose Method for Lung Cancer Screening

In the United States, lung cancer is the leading cause of cancer death. The screening CT and the follow-up CTs expose the patient to ionizing radiation, which carries with it an increased risk of malignancy. To maximize the benefit of CT lung cancer screening, we need to minimize the radiation dose as low as reasonably achievable. The compressive sensing theory has shown its potential in CT reconstruction for dose reduction. In contrast to the popular total variation (TV) based functions, a redundant dictionary is more specific for a particular application and more effective in terms of a sparse representation. In this work, we develop a dictionary learning based ultra-low dose method for lung cancer screening. Very supportive results are obtained from preclinical sheep lung study.

Hengyong Yu
Department of Radiology
Wake Forest University
hyu@wakehealth.edu

MS63

Pipelining the Fast Multipole Method over a Runtime System

Fast multipole methods (FMM) usually require a careful tuning of the algorithm for both the targeted physics and the hardware. In this talk, we propose a new approach that achieves high performance across architectures. We express the FMM algorithm as a task flow and employ a runtime system to process the tasks on the different processing units. We carefully design the task flow, the mathematical operators, their CPU and GPU implementations and their schedule.

Emmanuel Agullo, Berenger Bramas, Olivier Coulaud
INRIA
emmanuel.agullo@inria.fr, berenger.bramas@inria.fr, olivier.coulaud@inria.fr

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

Matthias Messner
INRIA
matthias.messner@inria.fr

Toru Takahashi
Nagoya University

ttaka@nuem.nagoya-u.ac.jp

MS63

Achieving High Performance with Multiple-GPU Non-symmetric Eigenvalue Solver

The first step in the non-symmetric eigenvalue problem is the reduction to upper Hessenberg form. Dependencies caused by applying an orthogonal matrix on both sides make this reduction difficult to parallelize. We present a multiple-GPU implementation of the Hessenberg reduction that uses the GPUs' superior memory bandwidth to achieve high performance on memory-bound operations. Our algorithm includes optimizations to overlap work on the CPU and GPU. We also explore accelerating the computation of eigenvectors.

Mark Gates
Computer Science Department
University of Tennessee
mgates3@eecs.utk.edu

MS63

An Applications Perspective on Multi-core, Massive Multi-threading, and Hybrid Systems

Multi/many-core and hybrid architectures that support massive multi-threading have raised considerable uncertainty as to what programming models might be appropriate. We will discuss applications that have been developed within the Swiss High-Performance and High-Productivity Computing platform (www.hp2c.ch) and are exploiting emerging computer architecture quite successfully. We will see what architectural aspects are important for the various algorithmic motifs that appear in applications, and what new programming models seem to find broad acceptance among scientific programmers.

Thomas Schulthess
Swiss National Supercomputing Center
schulthess@cscs.ch

MS63

Adapting to the Heterogeneous HPC Phenomena in the Industry

The lions share of Shells global HPC capacity is consumed by geophysical seismic imaging and reservoir engineering fluid flow simulations for oil and gas reservoirs. Legacy algorithms and software must be replaced with fundamentally different ones that scale to 1000s of possibly heterogeneous- cores. Geophysical Reverse Time Migration is an example. In this talk, we present how we're adapting our algorithms to tackle this phenomena.

Amik St-Cyr
Computation and Modeling,
Shell International E&P, Inc.
amik.st-cyr@shell.com

Chaohui Chen
Computation and Modeling
Shell International E&P, Inc.
chaohui.chen@shell.com

Detlef Hohl
Shell Global Solutions (US)
detlef.hohl@shell.com

MS64

Front Tracking and Spring Fabric Model for Parachute FSI

We use the front tracking method on a spring system to model the dynamic evolution of parachute canopy. The canopy surface of a parachute is represented by a triangulated surface mesh with preset equilibrium side length. This mechanical structure is coupled with the Navier-Stokes solver through the Impulse Method. The numerical solutions are compared with the experimental data and there are good agreements in the terminal descent velocity and breathing frequency of the parachute.

Xiaolin Li
Department of Applied Math and Stat
SUNY at Stony Brook
linli@ams.sunysb.edu

Joungdong Kim, Yan Li
SUNY at Stony Brook
Stony Brook, NY 11794
selom114@gmail.com, yli@ams.sunysb.edu

MS64

Implicit Schemes for Fluid-Fluid and Fluid-Structure Interaction Problems in an Eulerian Framework

The simulation of compressible multi-fluid flows interacting with flexible structures is important in many areas, especially in problems involving implosions and explosions. The numerical methods utilized for these problems have historically used explicit time integration. In this talk we describe an implicit time integration scheme for these simulations within an Eulerian framework (FIVER). The presented scheme shows speedups of up to 40x on many problems, while preserving the same level of accuracy.

Alex Main
Stanford University
alexmain@stanford.edu

Kevin G. Wang
Department of Aeronautics and Astronautics
Stanford University, Stanford, CA
icmewang@stanford.edu

Charbel Farhat
Stanford University
CFarhat@stanford.edu

MS64

Hybridization Techniques in Coupled Multiphysics Flow Problems

Coupled problems, involving multiple scales and/or physics, arise in many applications; e.g., in fluid-structure interaction or the coupling of porous media with free flow. An important question in such coupled problems is the choice of flexible, stable, physically meaningful and well-approximating interface couplings. We show that hybridizable methods offer some interesting properties for addressing these points and illustrate the proposed coupling concepts by examples.

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

MS64

Computational Methods for Multi-Material Fluid-Structure Interaction with Dynamic Fracture

This talk addresses the interaction of multi-material compressible fluid flows and elastic-plastic structures subject to large deformation and fracture. We present a high-fidelity, fluid-structure coupled computational framework based on an embedded boundary method for CFD and an extended finite element method (XFEM) for CSD. We illustrate this framework, highlight its features, and assess its performance with the three-dimensional simulation of pipeline explosions and underwater implosions for which test data is available.

Kevin G. Wang
Department of Aeronautics and Astronautics
Stanford University, Stanford, CA
icmewang@stanford.edu

Patrick Lea
Northwestern University
patrick.lea@u.northwestern.edu

Charbel Farhat
Stanford University
CFarhat@stanford.edu

MS65

An Adaptive Patchy Method for the Numerical Solution of the Hamilton-Jacobi-Bellman Equation

We present an adaptive numerical method to compute solutions to the Hamilton-Jacobi-Bellman PDE arising from an infinite-horizon optimal control problem. The method is based on a patchy technique and builds local polynomial solutions using a continuation algorithm. A level set method is used that adaptively changes the step-size of the cost levels depending on the magnitude of the error incurred by the computed solution. Numerical experiments are presented that illustrate the accuracy of the method.

Cesar O. Aguilar, Arthur J. Krener
Naval Postgraduate School
Department of Applied Mathematics
coaguila@nps.edu, ajkrener@nps.edu

MS65

An Iterative POD (I-POD) based Approach to Solving the Fokker-Planck Equation with Application to Nonlinear Filtering

The Fokker-Planck-Kolmogorov Equation (FPK) is a partial differential equation that governs the evolution of uncertainty through a stochastically perturbed nonlinear dynamical system. The FPK is a linear parabolic PDE, however, it suffers from the curse of dimensionality owing to its high dimensional state space. We propose an iterative proper orthogonal decomposition (I-POD) based approach to obtaining reduced order models (ROM) of the FPK equation such that it is suitable for online implementation. The I-POD is a data based technique that can be used to extract the dominant eigenmodes of large scale linear systems such as those arising from the FPK equation, which can then be used to construct a suitable ROM. These ROMs are then utilized to solve the nonlinear filtering/ data assimilation problem in a computationally tractable, real-time fashion.

Suman Chakravorty
Texas A&M University
College Station, TX
schakrav@tamu.edu

MS65

Computational Aspects of Optimal Control for Nonlocal Problems

We consider a control problem for the nonlocal Poisson equation. We study existence and uniqueness of the optimal control and we show the convergence of the nonlocal solution to the classical (local) one as the horizon approaches zero. Also, we introduce a finite element (FE) discretization, we derive a priori error estimates and we discuss advantages and disadvantages in using different FE spaces. Theoretical results are confirmed by numerical computations for one-dimensional test cases.

Marta D'Elia
Florida State University
mdelia@fsu.edu

Max Gunzburger
Florida State University
School for Computational Sciences
mgunzburger@fsu.edu

MS65

Robust Control via Optimal Uncertainty Quantification

This talk will cover recent results in the optimal quantification of uncertainties with incomplete information on probability measures and response functions and in presence of sample data (and also, possibly, with incompletely specified priors). We will show how these results can be applied to robust control. Various parts of this talk are joint work with Clint Scovel (Caltech), Tim Sullivan (Caltech, Warwick), Mike McKerns (Caltech) and Michael Ortiz (Caltech).

Houman Owhadi
Applied Mathematics
Caltech
owhadi@caltech.edu

MS66

Model Velocity Estimation based on Seismogram Registration

Seismograms from seismic surveys contain subsurface velocity and structure information. We show that optimal mapping between a measured seismogram and its corresponding predicted seismogram can be obtained and effectively used for full waveform inversion. Highly non-convex image registration problems due to the oscillatory wavelets are solved in a multiscale manner. Such mapping enables us to modify the measured seismogram close to the predicted one so that a low frequency velocity update is com-

puted.

Hyoungsu Baek
Massachusetts Institute of Technology
Mathematics Department
hbaek@math.mit.edu

Laurent Demanet
Mathematics, MIT
demanet@gmail.com

MS66

Time-domain Seismic Imaging and Inversion

Seismic reflection imaging can be formulated in the depth domain (original Cartesian coordinates of the subsurface) or in the time domain (coordinates defined by the so-called image rays). The time-domain formulation admits efficient imaging and inversion algorithms thanks to the existence of effective approximations, which provide prior information for seismic velocity estimation. I will describe recent advances in developing time-domain algorithms and in solving the problem of time-to-depth conversion.

Sergey Fomel
University of Texas at Austin
sergey.fomel@beg.utexas.edu

MS66

Algorithms for Seismic Imaging with Multiply Scattered Waves

The single scattering assumption is ubiquitous in seismic imaging; almost all images are formed under the Born approximation. There is ample evidence, however, that a significant amount of energy is multiply scattered within the Earth. When this is the case, applying algorithms that assume single scattering results in images with artifacts. We will discuss extensions that allow for the inclusion of multiply-scattered waves in imaging algorithms and discuss their performance.

Alison Malcolm
MIT
amalcolm@mit.edu

Alan Richardson
Department of Earth, Atmospheric and Planetary Sciences
Massachusetts Institute of Technology
alan_r@mit.edu

MS66

Frozen Gaussian Approximation for High Frequency Wave Propagation

Abstract not available at time of publication.

Xu Yang
Courant Institute of Mathematical Sciences
New York University
xuyang@cims.nyu.edu

MS67

Scientific Computing Projects for Undergraduates

Technology today is changing so dramatically and so quickly that there is certainly a need to educate students in computer literacy and applications. Computer experiments are as valuable to the mathematics student as lab experiments are to a chemistry student. I have integrated computer projects in all levels of courses to help develop the students' mathematical skills that are necessary to function in this growing technological world.

James Baglama
University of Rhode Island
jbaglama@math.uri.edu

MS67

An Innovative Scientific Computation Course for Undergraduates

Over the past five years, we have developed a problem-driven Python-based course in scientific computation for our undergraduate math majors, which replaces a conventional Java (or C-C^{++}) programming course requirement. The course is taught in a computer-equipped classroom where students can obtain in-class help with their programming and mathematical errors. During my talk, we will discuss our rationale for this curriculum change and provide examples of classroom activities and student's projects.

Adam O. Hausknecht
University of Massachusetts Dartmouth
North Dartmouth, MA
ahausknecht@umassd.edu

MS67

Running an Undergraduate Summer Research Program in Parallel Computing: a Challenging but Rewarding Experience

Parallel computing, the future of computing, is rarely taught at undergraduate level due to the complexity of parallel programming and the cost and maintenance expenses of massively parallel systems. EXERCISE- Explore Emerging Computing in Science and Engineering is a new Research Experiences for Undergraduates (REU) site hosted at Salisbury University funded by NSF. This presentation will focus on how to overcome the challenges of running such an undergraduate research program in a primarily undergraduate institution.

Enyue Lu
Salisbury University
Salisbury, MD
ealu@salisbury.edu

MS67

The Impact of Undergraduate Research in Scientific Computing on Undergrads at the University of Massachusetts

In this talk, we will describe undergraduate mathematical research that we have supervised at the University of Massachusetts Amherst. This includes honor's theses as well as summer research. We will give descriptions of the projects, which were primarily in the area of modeling tumor angiogenesis. We also discuss the impact of these projects on the direction and the future success of the students.

Nathaniel Whitaker
Department of Mathematics and Statistics
University of Massachusetts, Amherst
whitaker@math.umass.edu

MS68

Data-driven Optimal $\mathcal{H}_\in$ Model Reduction

We present a combined $\mathcal{H}_\in$ trust-region descent and Iterative Rational Krylov Algorithm (IRKA) approach for optimal $\mathcal{H}_\in$ model reduction of multiple-input/multiple-output (MIMO) linear dynamical systems. The proposed approach is formulated in terms of transfer function evaluations only, without requiring access to any particular realization. We consider this approach in the context of data-driven model reduction applied to extreme MIMO systems (large number of inputs or outputs).

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

Tim Campbell
Naval Research Laboratory
Stennis Space Center
Tim.Campbell@nrlssc.navy.mil

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

MS68

Model Order Reduction fo Un(Steady) Aerodynamic Applications

In this talk the fast simulation of air flowing past an airfoil is addressed using POD-based model order reduction. In particular, reduced order models are used to predict (un)steady aerodynamic flows. Numerical results for a high-lift wing-flap configuration are presented. Such configurations are used in the take-off and landing phase of flight.

Heike Fassbender
TU Braunschweig
h.fassbender@tu-braunschweig.de

MS68

Index-preserving MOR for Nonlinear DAE Systems

We introduce a model order reduction procedure for differential-algebraic equations, which is based on explicitly splitting the DAE into the intrinsic differential equation contained in the original system and on the remaining algebraic constraints. The index 1 case will be discussed extensively, as well as extensions to higher index. We implement numerically this procedure and show numerical evidence of its validity, as well as advances over existing MOR techniques.

Wil Schilders
Technical University of Eindhoven
w.h.a.schilders@TUE.nl

Nicodemus Banagaaya
TU Eindhoven
n.banagaaya@tue.nl

Giuseppe Ali
Univ. of Calabria
giuseppe.ali@unical.it

Caren Tischendorf
Humboldt University of Berlin
caren.tischendorf@math.hu-berlin.de

MS68

Data-Driven Model Order Reduction via Convex Optimization: Improved Bounds for Fitting Stable Nonlinear Models

In this talk we present a convex optimization framework for optimizing stable nonlinear state-space models to match recordings from experiment or simulation. The approach consists of a parameterization of stable nonlinear models and a convex upper bound on the long-term mismatch between simulated response and recordings. The resulting optimization problem is implementable as a semidefinite program. We demonstrate the approach on the task of identifying reduced order models of examples drawn from the circuits domain.

Mark Tobenkin, Yan Li, Alexandre Megretski
MIT
mmt@mit.edu, liyan@mit.edu, ameg@mit.edu

MS69

Estimation of Uncertain Parameters Through Parallel Inversion

We compare an optimization-based method for identifying stochastic parameters with a sampling-based approach. Unlike Bayesian methods, the sampling-based approach we consider is based on solving a sequence of deterministic inverse problems combined with ideas from sparse grid collocation. Knowledge of the parameters at collocation points allows us to approximate statistical quantities of interest via numerical quadrature. We present numerical results for an elliptic PDE where the conductivity parameter is stochastic.

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

Hans-Werner Van Wyk
Department of Scientific Computing
The Florida State University
hvanwyk@fsu.edu

MS69

Karhunen-Loeve Expansion for Multiple Correlated Stochastic Processes

In this work, we propose two numerical techniques to model and simulate multiple correlated random processes. The two techniques find the appropriate expansion for each correlated random process by generalizing the Karhunen-Loève expansion to multiple processes, while attaining the entire correlated stochastic structure. The primary difference between the two methods lies in whether the random variables used in the expansion are independent or correlated. Some explicit formulae and analytical results are presented for exponentially correlated random processes. In addition, these two methodologies are compared to each other in terms of convergence and computational efficiency and to other methods including mixtures of probabilistic PCA. We also simulate the tumor cell model induced by colored noise and limit cycle oscillator with our techniques.

Heyrim Cho
Brown University
Providence, RI
heyrim_cho@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

Daniele Venturi
Brown University
daniele_venturi@brown.edu

MS69

Predictive Simulations for Problems with Solution Nonuniqueness

Large eddy simulations of turbulent mixing show examples of nonunique but apparently converged simulations. Theory allows a zero parameter selection of a unique solution, within the context of front tracking and dynamic subgrid models. Experimental data confirms this result. Extrapolation beyond the experimental range to infinte Reynolds numbers is shown to be mild. Results are interpretated in the context of inertial confinement fusion simulations.

James Glimm
State University of New York, Stony Brook
glimm@ams.sunysb.edu

MS69

Stochastic (w*) Convergence for Turbulent Combustion

We test two fundamental ideas for numerical simulation of turbulent combustion: (1) Finite rate chemistry for Large Eddy Simulations; (2) Stochastic (w*) convergence based on probability distribution functions and mathematical ideas associated with Young measures. Convergence is measured in terms of L_1 norms for CDFs. Our verification, validation and uncertainty quantification test platform is the combustion within the engine of a scram jet, which is a Mach 7 experimental aircraft under study at Stanford University PSAAP Center.

Tulin Kaman
SUNY Stony Brook
tkaman@ams.sunysb.edu

James Glimm
State University of New York, Stony Brook
glimm@ams.sunysb.edu

MS70

Hermite Methods for Hyperbolic Systems: Applications and Extensions

This talk discuss applications and extensions of the Hermite methods described in the previous talk. In particular, propagation of discontinuities, hybridization with discontinuous Galerkin methods will be considered. Applications of the Hermite methods to compressible flow and electromagnetics will also be presented.

Daniel Appelo
Department of Mathematics and Statistics
University of New Mexico
appelo@math.unm.edu

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

MS70

Hermite Methods for Hyperbolic Systems: Basic Theory

Hermite methods are spectral element methods defined on staggered computational cells of cuboids whose degrees-of-freedom are tensor-product Taylor polynomials defined at each vertex. When they are applied to hyperbolic systems, the time step is restricted only by the CFL constraint, independent of the polynomial degree, enabling memory-efficient implementations. In this talk we outline the theory of Hermite discretizations and compare their resolving power to that of competitive schemes.

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

Chang-Young Jang
Southern Methodist University
cjang@smu.edu

Daniel Appelo
University of New Mexico
appelo@unm.edu

Ronald Chen
University of Arizona
xqroy@gmail.com

MS70

Accurate Solution of Diffusive Problems in Immersed Domains

In this talk we discuss the Correction Function Method (CFM), a general framework used to devise highly accurate numerical schemes to discretize diffusion dominated phenomena in immersed domains. The combination of the CFM with Gradient Augmented Level Set Methods results in a powerful tool that can be applied to a variety of situations in which both immersed domains and high accuracy are required. Here we present results for the Poisson and the incompressible Navier-Stokes equations.

Alexandre N. Marques
MIT
noll@alum.mit.edu

Jean-Christophe Nave
McGill University
jcnave@math.mcgill.ca

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

MS71

Efficient Simulation of Multiscale Kinetic Transport

We discuss a new class of approaches for simulating multiscale kinetic problems, with particular emphasis on applications related to small-scale transport. These approaches are based on an algebraic decomposition of the distribution function into an equilibrium part, that is described deterministically (analytically or numerically), and the remainder, which is described using a particle simulation method. We show that it is possible to derive evolution equations for the two parts from the governing kinetic equation, leading to a decomposition that is dynamically and automatically adaptive, and a multiscale method that seamlessly bridges the two descriptions without introducing any approximation. Our discussion pays particular attention to stochastic particle simulation methods that are typically used to simulate kinetic phenomena; in this context, algebraic decomposition can be thought of as control-variate variance-reduction formulation, with the nearby equilibrium serving as the control. Such formulations can provide substantial computational benefits in a broad spectrum of applications because a number of transport processes and phenomena of practical interest correspond to perturbations from nearby equilibrium distributions. In many cases, the computational cost reduction is sufficiently large to enable otherwise intractable simulations. The proposed methods will be illustrated with applications to a variety of problems of engineering interest, such as microscale/nanoscale gas flows and microscale/nanoscale solid-state heat transfer as mediated by phonon transport.

Nicolas Hadjiconstantinou
Massachussets Institute of Technology
ngh@mit.edu

MS71

Modeling and Simulation of Suspensions with a Large Number of Interacting Micro-swimmers

Microorganisms play an important role in nature. Understanding aspects of their locomotion and collective behavior is essential to understanding many biological and physical phenomena, as well as how to best use them in technological applications. Designing mathematical and computational models to help scientists in these endeavors is paramount. We present a new mathematical model and simulation method to compute the collective dynamics of a large colony of micro-swimmers that interact with each-other and the fluid they are suspended in and can affect by their locomotion. This fast computational method uses the immersed boundary framework and enables us to efficiently simulate thousands of interacting motile particles. We illustrate the method by showing examples of collective dynamics in large suspensions of "pusher" and "puller" micro-swimmers. The model satisfactorily captures macroscopic structures of observed in experiments of bacterial baths. Lastly, applications of the method will be discussed, e.g. in for bacteria suspensions or synthetic chemically-powered particles.

Enkeleida Lushi
Imperial College London
and Courant NYU
lushi@cims.nyu.edu

Charles S. Peskin
Courant Institute of Mathematical Sciences, New York Univers
peskin@cims.nyu.edu

MS71

Lattice-Boltzmann-Langevin Simulations of Binary Mixtures and Wetting Instabilities in Thin Fluid Films

I will describe a hybrid numerical method for the solution of the Model H fluctuating hydrodynamic equations for binary mixtures. The momentum conservation equations with Landau-Lifshitz stresses are solved using the fluctuating lattice Boltzmann equation while the order parameter conservation equation with Langevin fluxes is solved using stochastic method of lines ensuring that fluctuation-dissipation theorem is satisfied. The method is benchmarked by comparing static and dynamic correlations and excellent agreement is found between analytical and numerical results. Thermally induced capillary fluctuations of the interface are captured accurately, indicating that the model can be used to study nonlinear fluctuations. I will also discuss the performance of such an method to understand the instabilities of confined binary mixtures and the role of dynamic wetting.

Ignacio Pagonabarraga Mora
Departament de Fsica Fonamental
Universitat de Barcelona, Barcelona
ipagonabarraga@ub.edu

MS71

Simulation of Osmotic Swelling by the Stochastic Immersed Boundary Method

We present a numerical model that employs the stochastic immersed boundary method to study the osmotic swelling of a microscopic vesicle. The scale is so small that individual solute molecules are tracked explicitly. This is an important biological problem because water movement inside cells is generally driven by osmotic forces. The time-dependent Stokes equations are discretised using the stochastic immersed boundary method, and we allow the fluid to slip through the vesicle wall thus making the model lipid bilayer permeable to water. The elastic energy of the membrane has been modeled by a sum of three terms: a term proportional to area that generates surface tension, a surface neo-Hookean energy that resists shear, and a Helfich bending energy proportional to the integral of the square of the sum of the principal curvatures. In additon to the membrane energy, we also employ an energy of interaction between the membrane and the explicitly tracked solute particles to keep the solute within the vesicle. We find that the vesicle swells or shrinks (depending on its initial size) and eventually fluctuates about an equilibrium size that depends on the number of solute particles contained within the vesicle.

Charles S. Peskin
Courant Institute of Mathematical Sciences
New York University
peskin@cims.nyu.edu

Chen-Hung Wu
Courant Institute of Mathematical Sciences
chenhung@cims.nyu.edu

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

MS72

On the Evaluation of the Singular Integrals of Scattering Theory

I will describe an efficient method for the numerical evaluation of the singular integrals which arise in the discretization of certain integral operators of scattering theory given on surfaces. Standard techniques for evaluating these integrals, while adequate in simple cases, become prohibitively expensive when applied to integral operators given on complicated surfaces. The scheme I will describe, by contrast, is largely insensitive to the geometry of the underlying surface.

James Bremer
UC Davis
bremer@math.ucdavis.edu

MS72

Fast Volume Integral Equation Solver for Layered Media

Wave scattering in layered media is studied via the Green's function. The Green's function is developed by the scattering matrix technique and Sommerfeld-type integral. Then, the application of the integral operator for the Helmholtz equation in layered media is accelerated with the fast multipole method (FMM) and local-expansion treecode for the Bessel function. The same methods are applied to the vector potential Green's function for Maxwell's equations in layered media.

Min Hyung Cho
Department of Mathematics
Dartmouth College
min.h.cho@dartmouth.edu

Wei Cai
University of North Carolina, Charlotte
wcai@uncc.edu

MS72

Quadrature by Expansion: A New Method for the Evaluation of Layer Potentials

We present a systematic, high-order approach to the computation of layer potentials that works for any singularity (including hypersingular kernels), based only on the assumption that the field induced by the integral operator is smooth when restricted to either the interior or the exterior of the bounding curve or surface. The scheme, denoted QBX (quadrature by expansion), is easy to implement and compatible with the fast multipole method.

Andreas Kloeckner
Courant Institute of Mathematical Sciences
New York University
kloeckner@cims.nyu.edu

Alexander Barnett
Department of Mathematics
Dartmouth College
ahb@gauss.dartmouth.edu

Leslie Greengard, Michael O'Neil
Courant Institute of Mathematical Sciences
New York University
greengar@cims.nyu.edu, oneil@cims.nyu.edu

MS72

Fast Algorithms for Layer Heat Potentials

We will describe a new fast algorithm for evaluating layer heat potentials. A new recurrence relation is derived which reduces the computational complexity of the heat potentials from quadratic (of direct evaluation) to linear in the number of time-steps. The key advantages over other fast methods are its adaptivity and insensitivity to the time-step size. When combined with high-order product integration rules that overcome geometrically-induced stiffness, this algorithm can be used for solving the diffusion equation accurately and efficiently on moving geometries with Dirichlet or Neumann data. This is joint work with Leslie Greengard and Shidong Jiang.

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

MS73

A First Passage Time Algorithm for Reaction-Diffusion Processes on a 2D Lattice

Abstract not available at time of publication.

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

MS73

Parallelization and Error Analysis in Lattice Kinetic Monte Carlo

n this talk we explain an operator splitting approach to parallel implementation of kinetic Monte Carlo simulations of spatially distributed particle systems on a lattice. We discuss the error analysis of the algorithm that gives approximations of the conventional, serial kinetic Monte Carlo (KMC). A key aspect of our analysis relies on emphasizing a goal-oriented approach for suitably defined macroscopic observables (e.g., density, energy, correlations, surface roughness), rather than focusing on strong topology estimates for individual trajectories. One of the key implications of our error analysis is that it allows us to address systematically the processor communication of different parallelization strategies for KMC by comparing their (partial) asynchrony, which in turn is measured by their respective fractional time step for a prescribed error tolerance. This is a joint work with G. Arampatzis and M. A. Katsoulakis.

Petr Plechac
University of Delaware
Department of Mathematical Sciences
plechac@math.udel.edu

MS73

Efficient Algorithms and Parallel Issues for Kinetic Monte Carlo Modeling in Materials Science

Making kinetic Monte Carlo (KMC) applications run scalably in parallel involves trade-offs in performance versus accuracy, particularly with respect to relaxing the requirement that the underlying KMC algorithm be "exact" versus acceptably approximate. I'll discuss serial and parallel algorithms we've developed that come down on both sides of this issue and illustrate how we've implemented them in our parallel KMC simulator SPPARKS (http://spparks.sandia.gov). One novel feaature of SPPARKS is that it allows users to add new KMC models by writing functions that calculate energy changes and enumerate possible "events". The code achieves parallelism by spatially decomposing the simulation domain, which assumes the energetics of an event depends only on nearby information. This is a good assumption for many materials science problems, but I'll also highlight examples for surface growth and nuclear fuel sintering where this is not the case.

Steven J. Plimpton
Sandia National Laboratory
sjplimp@sandia.gov

MS73

Simulation of Strained Epitaxial Growth with Kinetic Monte Carlo

We present weakly-off-lattice and off-lattice kinetic Monte Carlo models for strained epitaxial growth. Both formulations are based on the observation that near equilibrium, the chemical potential is the thermodynamic driving force for film evolution. For the weakly off-lattice system, the energy is taken from a bonding counting, ball and spring model whereas in the off-lattice case an intermolecular potential is used. This is joint work with Henry Boateng and Tim Schulze.

Peter Smereka
Department of Mathematics
University of Michigan
psmereka@umich.edu

MS74

Stable Product of Matrices and Application in Quantum Monte Carlo Simulation on GPU Accelerated Multicore Systems

The product singular value decomposition is one way of circumventing numerical instability in calculatioing the product of many matrices. However, it is prohibitively expensive in terms of floating point operations, storage requriements and data communication. This is particularly true for large-scale many-body Monte Carlo simulations in computational materials science. In this talk, we present a different approach and work with a graded decomposition of the product. Furthermore, we use a pre-pivoting scheme for the grade revealing to reduce the data communication cost and efficiently exploit highly optimized primitive matrix kernels on hybrid CPU and GPU systems.

Zhaojun Bai
Departments of Computer Science and Mathematics
University of California, Davis
bai@cs.ucdavis.edu

Andrés Tomás
Department of Computer Science
University of California, Davis
andres@cs.ucdavis.edu

Richard Scalettar
Department of Physics
University of California, Davis
scaletter@physics.ucdavis.edu

MS74

Numerical Behavior of Two-step Splitting Iteration Methods

In this contribution we study numerical behavior of several stationary or two-step splitting iterative methods for solving large sparse systems of linear equations. We show that inexact solution of inner systems associated with the splitting matrix may considerably influence the accuracy of computed approximate solutions computed in finite precision arithmetic. We analyze several mathematically equivalent implementations and find the corresponding componentwise or normwise forward or backward stable implementations. The theory is then illustrated on the class of efficient two-step iteration methods such as Hermitian and skew-Hermitian splitting methods. This is a joint work with Zhong-Zhi Bai.

Miro Rozloznik
Academy of Sciences of the
Czech Republic, Prague
miro@cs.cas.cz

MS74

Using High-precision Arithmetic in the Design of a Stopping Criterion for Lanczos

The classical Lanczos method is the most memory-efficient way to compute all the eigenvalues of a large sparse symmetric matrix. Its convergence can be slow, but it does converge to all the eigenvalues and it is possible to determine which eigenvalues have converged. However, unless all the eigenvalues are distinct, it is not possible to determine when all of them have been found. To address this issue, we show that multiple eigenvalues can be dispersed by adding to A a random matrix with a small norm. By using high-precision arithmetic, we can perturb the eigenvalues by an amount that does not affect the accuracy of double-precision computed eigenvalues.

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

Alexander Alperovich, Alex Druinsky
Tel-Aviv University
sashaal@tau.ac.il, alexdrui@tau.ac.il

MS74

Computational Noise, Derivatives, and Optimization

Efficient simulation of complex phenomena often results in computational noise. Noise destroys underlying smoothness that otherwise could benefit optimization algorithms. We present a non-intrusive method for estimating computational noise and show how this noise can be used to derive finite-difference estimates with provable approximation guarantees. Building upon these results, we show how step sizes for model minimization and improvement can be selected in derivative-free optimization.

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

Jorge J. Moré
Argonne National Laboratory
Div of Math & Computer Science

more@mcs.anl.gov

MS75

Towards a Fine-Grained Parallel Implementation of the Nonsymmetric QR Algorithm

We present the first step towards a fine-grained parallel implementation of the non-symmetric QR algorithm targeting multi-core processors and shared memory. Our primary goal is high performance on the full range from small to large problems. To reach this goal, we overlap the critical path with delayed updates and parallelize both the bulge chasing and the aggressive early deflation kernels. In addition, different scheduling techniques and matrix storage formats are evaluated.

Lars Karlsson
Umeå University
Computing Science and HPC2N
larsk@cs.umu.se

MS75

Parallel Multishift QR and QZ Algorithms with Advanced Deflation Strategies - Recent Progress

Key techniques used in our novel distributed memory parallel QR and QZ algorithms include multi-window bulge chain chasing and distributed aggressive early deflation (AED), which enable level-3 chasing and delayed update operations as well as improved eigenvalue convergence. Recent progress includes a multi-level recursive approach for performing AED in a parallel environment leading to communication avoiding algorithms via data redistribution. Application and test benchmarks confirm the superb performance of our parallel library software.

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

MS75

Designing Fast Eigenvalue Solvers on Manycore Systems

Abstract not available at time of publication.

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

MS75

The Parallel Nonsymmetric QR Algorithm with Aggressive Early Deflation

We present the new parallel nonsymmetric QR algorithm in ScaLAPACK v2.0. The multiwindow bulge chain chasing approach ensures that most computations in the bulge chasing stage are performed in level 3 BLAS. We also develop multilevel aggressive early deflation algorithms which decrease the total amount of communications. These techniques make the new approach significantly outperform the pipelined QR algorithm in ScaLAPACK v1.8.0. Both performance models and numerical experiments demonstrate the efficiency of our new approach.

Meiyue Shao
MATHICSE
EPF Lausanne
meiyue.shao@epfl.ch

MS76

Lessons Learned from Managing the Open Source Library Deal.II

We will review our experience with managing deal.II, an open source project which today has 600,000 lines of code and hundreds of users around the world. This will include the technical side, such as ensuring quality in the long term. It will also include social aspects of managing a community, with the challenges of attracting and retaining volunteer contributors, maintaining quality in contributions of newcomers, and adequate documentation and training.

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

Timo Heister
Texas A&M University
Department of Mathematics
heister@math.tamu.edu

Guido Kanschat
Department of Mathematics
Texas A&M University
kanschat@tamu.edu

MS76

The Development and Adoption of the TriBITS Lifecycle Model in CSE Projects

We describe a proposal for a well-defined software lifecycle process based on modern Lean/Agile software engineering principles for research-driven CSE software. What we propose is appropriate for many CSE software projects that are initially heavily focused on research but also are expected to eventually produce usable high-quality capabilities. We describe the motivations for this work as well as the efforts to get it adopted in CSE projects including Trilinos and CASL.

R. A. Bartlett
Oak Ridge National Laboratory
bartlettra@ornl.gov

Michael A. Heroux, James Willenbring
Sandia National Laboratories
maherou@sandia.gov, jmwille@sandia.gov

MS76

IPython: a Tool for the Lifecycle of Computational Ideas

IPython is an open source environment for interactive and parallel computing that supports all stages in the lifecycle of a scientific idea: individual exploration, collaborative development, large-scale production using parallel resources, publication and education. The web-based IPython Notebook supports multiuser collaboration and allows scientists to share their work in an open document format that is a true "executable paper": notebooks can be version controlled, exported to HTML or PDF for publication, and used for teaching.

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

MS76

libMesh: Lessons in Distributed Collaborative Design and Development

Science is naturally suited to open source software development. Accurate, replicable publication requires complete description of all algorithms in numerical experiments. Source code publication allows true review and reuse of algorithms therein, enabling greater collaboration and cooperation between researchers. We review advantages and challenges of open source scientific applications. Efficiency versus usability, compatibility versus innovation, and centralized versus distributed development are discussed in the context of the MOOSE [**?**] simulation framework and libMesh [**?**] finite element library.

Roy Stogner
University of Texas at Austin
roystgnr@ices.utexas.edu

MS77

An Enhanced Derivative-free Approach to Energy Applications

It has been well documented that derivative-free algorithms are immensely useful for the optimization of black-box functions. In this talk, I will explain how their effectiveness can be augmented by the inclusion of sensitivity calculations. I will describe an algorithm that monitors local sensitivities throughout the optimization process and uses this information as a stopping criteria. I will discuss the applicability of this approach to energy relevant applications. Specifically, I will discuss planning and operations of the electrical grid and radioactive waste disposal options for nuclear power plants.

Genetha Gray
Sandia National Laboratories
gagray@sandia.gov

MS77

Results of Design Studies Using Derivative-free Optimization for Multi-Layered Filters

Filtration applications appear in a variety of physical settings; among them are industrial filtration for polymer processing, protein separation in pharmaceutical drug purification, and oil and air filtration in the automotive industry. Effective filters remove large amounts of debris, but cost considerations warrant filters that have long lifetimes. Thus, one must balance the need for effective filters against the costs of replacement; filters that trap everything would have short life spans. Alternatively, one could make a filter last forever by trapping nothing. Filter design can be evaluated using computational simulators and optimization tools that balance these competing objectives. In this talk,we summarize the optimization results obtained using a variety of different algorithms. We discuss the use of different objective functions, provide analysis of the designs returned by the algorithms, and give directions for future work.

Lea Jenkins
Department of Mathematical Sciences
Clemson University
lea@clemson.edu

MS77

Sparse Interpolatory Models for Molecular Dynamics

We describe a method for using interpolatory models to accurately and efficiently simulate molecular excitation and relaxation. We use sparse interpolation for efficiency and local error estimation and control for robustness and accuracy. The objective of the project is to design an efficient algorithm for simulation of light-induced molecular transformations. The simulation seeks to follow the relaxation path of a molecule after excitation by light. The simulator is a predictive tool to see if light excitation and subsequent return to the unexcited or ground state will produce a different configuration than the initial one. The goals of the simulation are not only to identify the end point, but to report the entire path in an high-dimensional configuration space so that one can look for nearby paths to interesting configurations and examine the energy landscape near the path to see if low energy barriers make jumping to a different path possible

Carl T. Kelley
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu

David Mokrauer
BAE Systems
david.mokrauer@baesystems.com

James Nance
North Carolina State Univ
jdnance2@ncsu.edu

MS77

Parameter Estimation for Modeling Threat Detection in the Brain using Derivative-free Methods

Several neural network architectures have been constructed modeling the human brain's attentive response to stimuli. These neural network architectures are optimized to fit data from physical experiments. Using optimization tools, parameters are fit to these data. Parameters can vary from connection strength in the architecture to the existence of a connection between nodes. We examine both which model architecture best fits experimental electrophysiological data and which optimization tool achieves the best fit for each model.

Benjamin Ritz
Department of Mathematics
Clarkson University
ritzbd@clarkson.edu

MS78

Output-Based Adaptation for Hybridized Discontinuous Galerkin Methods

We present output-based adaptive solutions of the steady compressible Euler and Navier-Stokes equations using a hybridized discontinuous Galerkin (DG) discretization. Error estimates are obtained using a discrete adjoint that takes into account both element-interior and interface approximations. Adaptation consists of changing the approximation orders independently on elements and interfaces. Both hybridized and lower-cost embedded DG methods are considered, and results are compared to standard DG adaptive runs.

Krzysztof Fidkowski, Johann Dahm, Peter Klein
University of Michigan
kfid@umich.edu, jdahm@umich.edu, kelinpn@umich.edu

MS78

Recent Developments in the Flux Reconstruction Method and Extensions to Large Eddy Simulation

Theoretical studies and numerical experiments suggest that unstructured high-order methods can provide solutions to otherwise intractable fluid flow problems. However, existing high-order schemes are less robust and more complex to implement than their low-order counterparts. Such issues have limited the adoption of high-order techniques in both academia and industry. In this talk our efforts to address a range of issues currently pacing the adoption of unstructured high-order schemes will be discussed.

Antony Jameson
Professor, Department of Aeronautics & Astronautics
Stanford University
jameson@baboon.stanford.edu

MS78

High-Order Flux Reconstruction Schemes: Theory and Implementation

The Flux Reconstruction (FR) approach to high-order methods is simple to implement and allows various high-order schemes, such as nodal discontinuous Galerkin methods, and Spectral Difference methods, to be cast within a single unifying framework. In this talk, new theoretical aspects of FR schemes will be discussed, and efficient implementation strategies for novel hardware platforms will be presented.

Peter E. Vincent
Department of Aeronautics and Astronautics
Stanford University
p.vincent@imperial.ac.uk

Freddie Witherden, Antony Farrington
Department of Aeronautics
Imperial College London
freddie.witherden08@imperial.ac.uk,
antony.farrington07@imperial.ac.uk

MS78

Finite Spectral Element Method for Incompressible Flows

Finite spectral method is a category of pointwise or cellwise local spectral schemes based on Fourier integral. We combine the finite spectral basis function with the finite element method. We can not only use element discretization in the computation domain, but also increase the exactness in each element. We first introduce the finite spectral interpolation functions and compare them with the Lagrange interpolation functions. Then, we combine the finite spectral interpolations functions with finite element method and give the 2D incompressible Navier-Stokes equations in terms of the stream function and the vorticity. Finally, we solve the benchmark lid-driving cavity problem and flow around a cylinder with unstructured mesh.

Jian-Ping Wang
Peking University
wangjp@pku.edu.cn

MS79

Exploring Co-Design in Chapel using LULESH

Chapel is an emerging parallel programming language whose design and development are led by Cray Inc. LULESH is an unstructured Lagrangian explicit shock hydrodynamics code developed by Lawrence Livermore National Laboratory. In this talk, we describe a collaboration between Cray and LLNL to explore the expression of LULESH within Chapel. This codesign effort has resulted in an improved version of LULESH while also providing valuable feedback from users on Chapel's feature set and implementation priorities.

Bradford L. Chamberlain
Cray Inc.
bradc@cray.com

MS79

Programming Model Support Necessary for Adapting High Performance Code to Differing Platforms

Portable performance of million line multi-physics codes requires programming techniques and compiler support that enable optimization without substantial code changes. In this talk, we describe tuning techniques used to improve the performance of ALE hydrodynamics mini-apps. The LULESH and LUAU mini-aps are presented as a case study for performance portability, architectural impact on code performance and vendor co-design. We present a style and abstractions for writing multi-physics codes that can be auto-tuned for multiple platforms.

Ian Karlin, Jim McGraw, Jeff Keasler, Bert Still
Lawrence Livermore National Laboratory
karlin1@llnl.gov, mcgraw1@llnl.gov, keasler1@llnl.gov, still1@llnl.gov

MS79

Leveraging the Cloud for Materials Proxy Applications

One oft-discussed approach to exascale development is MPI+X. This two-part model envisions using X to accelerate code on local, heterogeneous node architectures and MPI to construct the overall application and communicate between the nodes running X. Today, X includes languages such as CUDA, OpenCL, Cilk+, TBB, HMPP, OpenMP, etc.-more will follow in the future. We describe the addition of 'Cloud' technologies that enable the rapid development of dynamic multi-scale materials science proxy applications. Our discussion will focus on NoSQL databases (e.g. Riak, MongoDB) and non-traditional programming languages (e.g. Erlang).

Christopher Mitchell
Los Alamos National Laboratory

mitchell@lanl.gov

MS79

High-Level Abstractions for Portable Performance using LULESH

Exascale machines will be a significant departure from traditional architectures, for which multiphysics simulations codes have been developed, and will require an accompanying change in the programming model. Obtaining portable performance necessitates introducing high-level abstractions above the architecture-specific details. We report on explorations into these paradigms using LULESH, a simple shock hydrodynamics proxy application representative of the data structures and numerics used in larger applications of interest.

Charles H. Still
Lawrence Livermore National Laboratory
still1@llnl.gov

Zach Devito
Stanford University
zdevito@stanford.edu

MS80

Hybrid Subspace Methods for Dimensionality Reduction in Nonlinear Multi-Physics Models

Recent developments on hybrid subspace methods have introduced a new general approach to performing dimensionality reduction for nonlinear multi-physics models. The basic idea is to construct linear transformation to fewer degrees of freedom replacing the original I/O streams of the coupled physics models. We demonstrate the application to a typical nonlinear nuclear engineering model with many input parameters and quantities of interest, currently intractable with the state-of-the-art methods due to the computational cost required.

Hany S. Abdel-Khalik
North Carolina State Univ.
abdelkhalik@ncsu.edu

MS80

Decomposition Methods for Multidisciplinary Uncertainty Analysis

The focus of this talk is a decomposition approach for multidisciplinary uncertainty analysis of systems governed by partial differential equations. The main idea of our approach is to perform uncertainty analysis independently on local components in an "offline" phase, and then to assemble global uncertainty quantification with pre-computed local information in an "online" phase. In this talk, we show how this is achieved through a combination of domain decomposition methods, reduced basis methods, and sampling importance re-sampling.

Qifeng Liao
MIT
qliao@mit.edu

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

Tiangang Cui
MIT
tcui@mit.edu

MS80

Stochastic Dimension Reduction of Multi Physics Systems through Measure Transformation

Uncertainty quantification of multiphysics systems represents numerous mathematical and computational challenges. Indeed, uncertainties that arise in each physics in a fully coupled system must be captured throughout the whole system, the so-called curse of dimensionality. We present techniques for mitigating the curse of dimensionality in network-coupled multiphysics systems by using the structure of the network to transform uncertainty representations as they pass between components. Examples from the simulation of nuclear power plants will be discussed.

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

Paul Constantine
Stanford University
paul.g.constantine@gmail.com

John Red-Horse
Validation and Uncertainty Quantification Processes
Sandia National Laboratories
jrredho@sandia.gov

Tim Wildey
Sandia National Laboratory
tmwilde@sandia.gov

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

Maarten Arnst
Universite de Liege
maarten.arnst@ulg.ac.be

MS80

Gradient-based Model Reduction for High-dimensional Uncertainty Quantification

This talk focuses on a Reduced Order Modeling approach for Uncertainty Quantification, where we use gradient information to partition the uncertainty domain into "active' and "passive' subspaces, where the "passive' subspace is characterized by near constant value of the quantity of interest. We project the model onto the low dimensional "active' subspace and solve the resulting problem using conventional techniques. We derive rigorous error bounds for the projection algorithm and show convergence in L^1 norm.

Miro Stoyanov, Clayton G. Webster
Oak Ridge National Laboratory
stoyanovmk@ornl.gov, webstercg@ornl.gov

MS80

A Generalized Adjoint Framework for Sensitivity

and Global Error Estimation in Burnup Calculations

We develop an abstract framework for computing the adjoint of a general set of differential-algebraic equations, and we apply the framework to the transport/depletion equations which are used to model advanced nuclear reactor designs. The framework efficiently generates both parametric sensitivity information, which we use for calibration of the (up to thousands of) random or uncertain input parameters, as well as estimates for numerical discretization errors.

Hayes Stripling
Texas A&M University
h.stripling@tamu.edu

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

Marvin Adams
Texas A&M University
mladams@ne.tamu.edu

MS81

A Hybrid Adaptive Mesh Framework for Wave Propagation Algorithms on a Forest of Locally Refined Cartesian Meshes

We describe current efforts to develop ForestClaw, a hybrid AMR finite volume code based on wave propagation algorithms in which non-overlapping fixed-size Cartesian grids are stored as leaves in a forest of quad- or oct-trees. The tree-based code p4est manages the multi-block connectivity and is highly scalable in realistic applications. In joint work with researchers at the Cascade Volcanic Observatory (L. Mastin and H. Schwaiger, CVO, Vancouver, WA), we will present results from our efforts to use ForestClaw for modeling the transport of volcanic ash in the atmosphere.

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

MS81

Title Not Available at Time of Publication

Abstract not available at time of publication.

David George
U.S. Geological Survey
Cascades Volcano Observatory
dave.jorge@gmail.com

MS81

Optimal Strong-Stability-Preserving Runge-Kutta Methods for Discontinuous Galerkin Spatial Discretizations of Hyperbolic Problems

In this talk, we present the construction of explicit strong-stability-preserving Runge–Kutta methods with optimal stability regions for discontinuous Galerkin spatial discretizations applied to hyperbolic problems.

Ethan Kubatko
Department of Civil and Environmental Engineering
The Ohio State University
kubatko.3@osu.edu

Benjamin Yeager
The Ohio State University
yeager.51@osu.edu

MS81

High Order Accurate RKDG Methods for the Shallow Water Equations on Unstructured Triangular Meshes

Shallow-water equations with a non-flat bottom topography have been widely used to model flows in rivers and coastal areas. These equations have steady-state solutions in which the flux gradients are non-zero but exactly balanced by the source term. Therefore extra care must be paid to approximate the source term numerically. Another important difficulty arising in these simulations is the appearance of dry areas, and standard numerical methods may fail in the presence of these areas. In this presentation, we will talk about recently developed high-order discontinuous Galerkin methods on unstructured triangular meshes, which can preserve the steady-state exactly, and at the same time are positivity preserving without loss of mass conservation. Some numerical tests are performed to verify the positivity, well-balanced property, high-order accuracy, and good resolution for smooth and discontinuous solutions.

Yulong Xing
Department of Mathematics
Univeristy of Tennessee / Oak Ridge National Lab
xingy@math.utk.edu

MS82

A New Three-Field Stabilized Finite Element Method for Fluid-Structure Interactions

We present some advancement towards a monolithic solution procedure for the numerical solution of fluid-structure interactions. A new three-field stabilized finite element formulation is presented for modelling the interactions between the fluid (laminar or turbulent) and the structure (rigid or elastic). We combine this method with anisotropic mesh adaptation to ensure an accurate capturing of the discontinuities at the interface. The accuracy of the formulation is demonstrated by means of 2D and 3D numerical examples.

Elie Hachem, Thierry Coupez
MINES ParisTech
elie.hachem@mines-paristech.fr,
thierry.coupez@mines-paristech.fr

Ramon Codina
Universtitat Politecnica de Catalunya (UPC)
ramon.codina@upc.edu

MS82

Three Dimensional Optimal Transportation Meshfree (OTM) Simulations of Human Arterial Blood

Flow

We present a monolithic Lagrangian solution for the fluid-structure interaction problems involving large deformation structure and free-surface flows. In our approach, both the fluid and structure are modeled by the Optimal Transportation Meshfree (OTM) method. The performance of the proposed method is demonstrated by a three-dimensional simulation of patient-specific modeling of arterial blood flow. Especially an anisotropic hyperelastic constitutive model with fiber reinforcement is developed to model the three layers of the human arterial wall.

Bo Li, Stefanie Heyden
California Institute of Technology
libo@caltech.edu, heyden@caltech.edu

Anna Pandolfi
Politecnico di Milano
anna.pandolfi@polimi.it

Michael Ortiz
California Institute of Technology
ortiz@aero.caltech.edu

MS82

Deforming Composite Grids for Fluid-structure Interaction

In this talk, we discuss recent work on overlapping grid discretizations of compressible fluid-structure interaction problems. Both elastic and rigid structures are considered. Deforming composite grids (DCGs) are used to address geometric complexity in a highly efficient and flexible manner. The FSI coupling is partitioned so that the fluid and solid solvers remain independent. Added-mass instabilities are addressed through the use of a newly developed interface projection technique.

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

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banks20@llnl.gov

William D. Henshaw
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
henshaw@llnl.gov

MS82

A Semi-local Solver for h-p Discretization of the Structure Equations

We develop a semi-local method for structure solvers, which decouples the three directions of displacement, enabling the use of an efficient low energy preconditioner for the conjugate gradient solver. Based on spectral element with Jacobi modal basis, we demonstrate high parallel efficiency for structure simulations on a 3D patient-specific flexible brain arteries test problem. The new solver combined with spectral element for the fluid discretization, improves greatly the computational efficiency of the FSI solver.

Yue Yu
Brown University
yue_yu_1@brown.edu

Marco Bittencourt
Departamento de mecanico
Universidade de Campinas
mlb@fem.unicamp.br

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS83

Computational Issues for Boundary Control Problems with Actuator Dynamics

The problem of boundary control in systems governed by partial differential equations often leads to abstract control systems with unbounded input operators and very weak state spaces. Moreover, in most practical settings the input at the boundary $v(t)$ is typically the output of a dynamic "actuator' so that $v(t) = \mathbf{Hx_a(t)}$ where the actuator state $x_a(t)$ is defined by a finite dimensional system. Although the inclusion of actuator dynamics is a more realistic representation of the complex system, certain abstract formulations of the composite can bring additional complexity to the control problem and does not provide a practical approach to the development of computational methods. However, if one begins with the fundamental PDE system and treats the actuator dynamics as boundary dynamical systems, then it is possible to simplify both the theoretical and computational challenges. We illustrate this approach with a simple boundary control problem defined by a nonlinear parabolic PDE.

John A. Burns
Virginia Tech
Interdisciplinary Center for Applied Mathematics
jaburns@vt.edu

MS83

A Combined Controls and Computational Fluids Approach for Estimation of a Moving Gaseous Source

A combined controls and computational fluids dynamics approach is applied to the estimation of gas concentration associated with an emitting moving source. The state estimator uses a filter gain parameterized by the sensor position whose motion dynamics is incorporated into the spatial process. The estimator is based on a 3D adaptive, multi-grid, multi-step finite-volume method with upwind and flux limiting. The grid is adapted with local refinement and coarsening during the process-state estimation.

Michael A. Demetriou
Worcester Polytechnic Institute
mdemetri@wpi.edu

Nikolaos Gatsonis
Worcester Polytechnic Institute
gatsonis@wpi.edu

MS83

Challenges in Computational Nonlinear Control Theory: a Perspective from the Air Force Office

of Scientific Research

Modeling, design and control of high performance complex Air Force Systems have put increasing demand on computational resources and theoretical and algorithmic development to meet the scientific challenges in dealing with the underlying high-dimensional, nonlinear, and stochastic problems. While there has been serious research on development of algorithms specifically designed for control or optimization of these problems (e.g. flow control, vibration control, design of structures), we have yet to meet the needs of scalability, efficiency and accuracy required to meet the real-world application requirements. In this talk, a brief discussion of Air Force needs in computational control theory will be presented and the scientific challenges will be addressed.

Fariba Fahroo
Air Force Office of Scientific Research
fariba.fahroo@afosr.af.mil

MS83
Simultaneous Actuator Placement and Controller Design

Many control systems of interest, for example, active noise control and control of structural vibrations, are modelled by partial differential equations. Because of the distribution of the system in space the location of actuators, and sensors, in these systems is a variable in the design of a control system. The performance of the controlled system is strongly dependent on these locations. Thus, actuator and sensor locations are important variables in controller design and should be considered as part of controller synthesis. Conditions for well-posedness of the optimal actuator location problem and also for convergence of optimal locations chosen using approximations in the cases where the cost is linear-quadratic (or H-2) and also for the H-infinity situation have been obtained. Algorithms have been developed to solve these problems, but challenges for large-scale problems remain. The best actuator locations do not always agree with physical intuition, even for simple examples. This supports the need for further research into this area.

Kirsten Morris
Dept. of Applied Mathematics
University of Waterloo
kmorris@uwaterloo.ca

MS84
Analysis and Numerical Solutions of Quasi-steady State Poroelasticity Problems

Abstract not available at time of publication.

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

MS84
Adaptive Sparse Reconstruction for Prior Selection and Robust Geophysical Inversion

Abstract not available at time of publication.

Behnam Jafarpour
Chemical and Electrical Engineering
University of Southern California
behnam.jafarpour@usc.edu

MS84
Sparse Regularized Seismic Inverse Problem

In this talk, we frist address the sparsity-promoting sesmic data reconstruction from L1-norm to nuclear-norm minimization. Then we apply low-rank matrix completion (MC) with a designed texture-patch pre-transformation to three-dimensional seismic data reconstruction. An efficient L_1-norm minimizing algorithm, named approximate message passing (AMP), is extended to use for our nuclear-norm minimization problem. Numerical experiments show promising performance of the proposed method in comparison to traditional singular value thresholding (SVT) of MC and recent tensor completion.

Jianwei Ma
Department of Mathematics
Harbin Institute of Technology
jma@hit.edu.cn

MS85
Thoughts on Preparation: How to Lower the Barrier for using Computational Tools and Learning to Program?

A successful undergraduate research experience has the student making a contribution to the research group and becoming better prepared for a graduate program in CSE. To accomplish this, the student must quickly develop her computational-thinking and coding skills. Ideally, coding ability should be developed integrally, with computing present across the curriculum. In practice, this is rare. A research group hosting undergraduates can offer a "boot camp" of coding tasks of incremental nature to develop skills.

Lorena A. Barba
Department of Mechanical Engineering
Boston University
labarba@bu.edu

MS85
A Five Year Experiment on Developing an Undergraduate Research Computing Program

In 2006 we received a CSUMS grant from the NSF. According to the solicitation, CSUMS was to enhance the education and training of math undergraduates and to better prepare them for fields that require integrated strengths in computation and mathematics. This talk will describe what we did, what we got right, what we modified, and what we botched. It will also address the question as to whether computing helps reinforce, or complements, or interferes with the more mathematical aspects of the projects.

Mark Holmes
Rensselaer Polytechnic Institute
holmes@rpi.edu

MS85
Writing and Publishing a Scientific Paper with Undergraduate Students

I had opportunities to write journal papers with undergraduate students. Some papers were co-authored with one undergraduate student and some with a group of students. In this talk, I will share my experience in writing a journal paper with undergraduate students, e.g. how I started my research with undergraduate students and what difficulties and joy I found while working with them.

Jae-Hun Jung
SUNY at Buffalo
jaehun@buffalo.edu

MS85

The PRISM Interdisciplinary Program at Northeastern

The NSF-funded PRISM program at Northeastern University is run jointly by faculty from the Mathematics, Physics, Biology, Engineering and Education Departments. Students participate in faculty-led exploration and discovery courses, involving a mixture of theory, hands-on activities, and data analysis, as well as virtual environments in the Action Lab. Matlab is used throughout the program and is taught at multiple levels.

Christopher King
Northeastern University
king@neu.edu

MS86

A Model Reduction Approach for Partitioned Treatment of Uncertainty in Coupled Systems

We present a stochastic model reduction framework for partitioned treatment of the uncertainty space in domain coupling problems. In particular, we expand the solution of each domain in a stochastic basis that is constructed adaptively and by calling the domain solvers separately. The novelty of the proposed approach is that the propagation of uncertainty is achieved through a sequence of approximations with respect to the dimensionality of each individual domain and not the combined dimensionality.

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
doostan@colorado.edu

Mohammad Hadigol
Aerospace Engineering Sciences
University of Colorado, Boulder
mohammad.hadigol@colorado.edu

Hermann Matthies, Rainer Niekamp
Technische Universität Braunschweig
wire@tu-braunschweig.de, rniekamp@tu-bs.de

MS86

Stochastic Polynomial Chaos Basis Selection in the Bayesian Framework

Generalised polynomial chaos (gPC) expansions allow the representation of the stochastic solution of a system whose input parameters are random variables. It is useful to select a smaller set of gPC basis to reduce the computational cost of making measurements or avoid over-fitting that leads to inaccurate solutions especially when the number of the available gPC basis is much bigger that the number of the model evaluations. In some cases, the solution is sparse at the stochastic level and can be expanded with a gPC of only a few terms. E.g., under weak conditions, in elliptic stochastic partial differentiable equations (SPDE) with high-dimensional random coefficients, the solutions admit sparse representations with respect to gPC basis, while the deterministic solver required is expensive. We propose a fully Bayesian stochastic search strategy for the selection of the important gPC basis and the evaluation of the associated coefficients. The proposed method combines Bayesian model selection and regularised regression methods. Therefore, it accomplishes both shrinkage and basis selection while it takes into account the model uncertainty. For the evaluation of the required posterior quantities we propose an MCMC sampler. Some of the main advantages: (1) it provides interval estimates, (2) it quantifies the importance of each gPC base, (3) it allows the computation of Bayesian model average estimates or Bayes-optimal predictors. We show that the proposed method is able to detect the significant gPC basis and recover potential sparse patterns in gPC expansion on a toy example and an 1D elliptic SDE with high-dimensional random diffusion coefficients.

Guang Lin
Pacific Northwest National Laboratory
guang.lin@pnnl.gov

Georgios Karagiannis
pacific Northwest National Laboratory
georgios.karagiannis@pnnl.gov

MS86

Uncertainty Propagation in Finite Element Simulation of Particle Driven Flow

In this talk we address the initial conditions uncertainty propagation in the variational multiscale finite element simulation of particle driven flow. We use sparse grid stochastic collocation methods managed by a scientific workflow engine on a HPC environment. Particular interest is devoted to a quantity of interest related to the spatial pattern of sediment deposition.

Fernando A. Rochinha
The Federal University of Rio de Janeiro
Brazil
faro@mecanica.coppe.ufrj.br

Gabriel Guerra
Federal University of Rio de Janeiro
gguerra@ufrj.br

Alvaro Coutinho, Jonas Dias
The Federal University of Rio de Janeiro
Brazil
alvaro@coc.ufrj.br, faro@mecanica.coppe.ufrj.br

Marta Mattoso, Eduardo Ogasawara
Fedral University of Rio de Janeiro
marta@cos.ufrj.br, eduardo.ogasawara@gmail.com¿

Erb Lins
Federal University of Para
erb@ufpa.br

MS86

Noise Propagation in the Multiscale Simulation of

Coarse Fokker-Planck Equations

We present a numerical procedure to compute the solution of a FokkerPlanck equation for which the drift and diffusion parameters are unknown, but can be estimated using appropriately chosen realizations of a fine-scale, individual-based model. If the latter model is stochastic, the estimation procedure introduces noise on the coarse level. We investigate stability conditions for this procedure and present an analysis of the propagation of the estimation error in the numerical solution of the FokkerPlanck equation.

Yves Frederix
Dept. of Computer Science, K.U. Leuven
Yves.Frederix@cs.kuleuven.be

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

Dirk Roose
KU Leuven
Dept. of Computer Science
Dirk.Roose@cs.kuleuven.be

MS87

Interoperability of PETSc and Chombo

Abstract not available at time of publication.

Mark Adams
Columbia University
ma2325@columbia.edu

MS87

BoxLib: Overview and Applications

BoxLib is a publicly available software framework for building parallel block-structured AMR applications. It supports grid-based and particle-mesh operations on adaptive hierarchical meshes. BoxLib-based codes use both MPI and OpenMP, have demonstrated excellent scaling behavior on today's largest machines, and are in active use in a number of research areas. The BoxLib distribution contains an extensive User's Guide as well as straightforward tutorials which demonstrate how to build parallel adaptive application codes using BoxLib.

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

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

Michael Lijewski
Lawrence Berkeley National Laboratory
mjlijewski@lbl.gov

MS87

Region-Based AMR: A New AMR Paradigm in BoxLib

Region-based AMR (RAMR) is a new paradigm in BoxLib in which different regions at the same level of spatial refinement may have different temporal refinement. When combined with optimal subcycling, which selects a timestep for each region based on maximizing the efficiency of the overall algorithm, the use of RAMR can result in significant computational savings over traditional AMR approaches. Examples will be given from the BoxLib-based Nyx code for cosmology.

Ethan Van Andel, Ann S. Almgren
Lawrence Berkeley National Laboratory
evanan@math.berkeley.edu, ASAlmgren@lbl.gov

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

Michael Lijewski
Lawrence Berkeley National Laboratory
mjlijewski@lbl.gov

MS88

Building Envelope Parameter Estimation through Heat Transfer Modeling

In order to reduce energy consumption in buildings, one needs to understand and be capable of modeling the underlying heat transfer mechanisms, characteristics of building structures, operations and occupant energy consumption behaviors. Often, some of the crucial building envelope parameters are unknown. In order to infer thermal parameters associated with buildings envelope, sensor data is utilized within an inversion procedure. The forward problem involves a system of differential equations capturing the governing heat transfer model.

Raya Horesh
IBM T.J. Watson Research Center
rhoresh@us.ibm.com

Lianjun An, Young M. Lee, Young T. Chae, Rui Zhang
IBM T.J. Watson Reseach Center
alianjun@us.ibm.com, ymlee@us.ibm.com, ychae@us.ibm.com, zhangrui@us.ibm.com

MS88

Simulating Heat Transfer and Environmental Conditions in Buildings Equipped with Sensor Networks

We consider the problem of modeling air flow, heat transfer, and humidity in buildings, specifically in environments where natural convection is dominant. Buildings are nowadays often equipped with a network of sensors and a data management system which gathers the sensor data in real time. Such sensor measurements serve as input data for the boundary value problems for systems of partial differential equations comprising the physical models. We will discuss modeling techniques that facilitate performing simulations of the physical phenomena at hand and which are suitable for coupling with measurements gathered via sensor networks.

Vanessa Lopez-Marrero
IBM T. J. Watson Research Center
lopezva@us.ibm.com

MS88

A Reduced-order Energy Performance Modeling

Approach for Buildings

The reduced-order building and HVAC system models with parameter estimation methods are generic and will eliminate the labor intensive, time consuming task of developing and calibrating specific detailed physics-based models. In this talk, we will present a circuit-equivalent 3R2C thermal network of building, which is a widely used and simple reduced order building model. The underlying physics of the 3R2C model will be introduced firstly, followed by the validation of the 3R2C model with ASHRAE Standard 140 and a case study of a reduced-order 3R2C model of a real building.

Zheng ONeill
United Technologies Research Center
oneillz@utrc.utc.com

MS89

Three-dimensional Acoustic Scattering from Obstacles in a Half-space with Impedance Boundary Conditions

A classical problem in acoustic and electromagnetic scattering concerns the evaluation of the Greens function for the Helmholtz equation subject to impedance boundary conditions on a half-space. We will discuss a hybrid representation of this Greens function which combines images and a rapidly converging Sommerfeld-like integral. The representation is valid at arbitrary source and target locations, and is amenable to evaluation using fast-multipole methods. This is joint work with Leslie Greengard.

Michael O'Neil
Courant Institute
New York University
oneil@courant.nyu.edu

MS89

Quadrature Methods for the Sommerfeld Integral and their Applications

Abstract not available at time of publication.

Josef Sifuentes
New York University
Courant Institute
sifuentes@cims.nyu.edu

MS89

Fast Fourier Transforms of Piecewise Polynomials

We will construct a fast transform, based on low-rank approximation, which evaluates Fourier coefficients of piecewise-polynomial generalized functions. These generalized functions are supported on d-dimensional simplices such as points, lines, triangles, or tetrahedra, in D-dimensional space. The transform employs a stable new dimensional recurrence and a tree-based butterfly scheme.

John A. Strain
Department of Mathematics
University of California at Berkeley
strain@math.berkeley.edu

MS89

High-order Solvers for Scattering Problems in Domains with Geometric Singularities

Abstract not available at time of publication.

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

MS90

Uncertainty in Turbulent Flows with Empirical Sub-cooled Boiling Models

A sensitivity study is performed by varying tunable parameters in four different implementations of Eulerian multiphase boiling flow models. The modeling differences between the codes allows for a more complete study of multiphase boiling flow, but requires coupling the codes to meaningfully interpret the results. Parameter ranges are compiled from the literature. A few models with relatively little theoretical and empirical development have a significant impact on the quantities of interest.

Isaac Asher, Krzysztof Fidkowski
University of Michigan
isaaca@umich.edu, kfid@umich.edu

MS90

Determination of Nitridation Reaction Parameters Using Bayesian Inference

In this work, we present a computational study of a flow tube experiment in order to infer reaction parameters for graphite nitridation. We construct a two-dimensional representation of the experimental setup and model the flow using a reacting low-Mach number approximation to the Navier-Stokes equations. We employ a Bayesian approach in order to produce probability distributions that can used to quantify uncertainty in models where surface nitridation is important.

Paul T. Bauman
ICES
The University of Texas at Austin
pbauman@ices.utexas.edu

MS90

Uncertainty Modeling with Stochastic PDEs for Turbulent Channel Flow

Validation of and UQ for extrapolative predictions made by RANS turbulence models are necessary to properly inform decisions based on such predictions. Here, we explore the use of stochastic PDEs for this purpose. In particular, multiple stochastic PDEs describing the modeling errors observed in the Reynolds stress are coupled with multiple deterministic turbulence models to make uncertain predictions of channel flow. These predictions are compared with DNS data to assess their credibility.

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

MS90

Calibration of Stochastic non-Boltzmann Kinetic Models using EAST Shock Tube Data

We use Bayesian inference to calibrate the physical and stochastic model parameters using shock tube radiation measurements from NASA. We propose a novel formulation of the stochastic model based on the physical intuition that the discrepancy between predictions and experimental data is due to assumption of equilibrium population of the energy levels. This formulation will enable us to propagate model uncertainty to both the observable for the parameter calibration and the quantity of interest.

Marco Panesi
University of Illinois at Urbana-Champaign
marco.panesi@gmail.com

MS91

H-FaIMS: A Hierarchical Fast Inverse Medium Solver

We consider the inverse medium problem for the wave equation with broadband and multi-point illumination. We use a nonlinear least-squares formulation. If M is the number of illuminations, a Hessian matrix-vector multiplication requires 2M wave solves. We have developed H-FaIMS, a scheme based on hierarchical decompositions that, asymptotically, enables Hessian approximations that scale independently of the number of the sources. We present results for the 3D Helmholtz problem.

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

MS91

Interferometric Waveform Inversion via Lifting and Semidefinite Relaxation

In seismic and SAR imaging, fitting cross-correlations of wavefields rather than the wavefields themselves can result in much improved robustness vis-a-vis model uncertainties. This approach however raises two challenges: (i) new spurious local minima may complicate the inversion, and (ii) one must find a good subset of cross-correlations to make the problem well-posed. I the talk I will explain how to address these two problems with lifting, semidefinite relaxation, and expander graphs.

Laurent Demanet, Vincent Jugnon
Mathematics, MIT
demanet@gmail.com, vjugnon@math.mit.edu

MS91

Bayesian Uncertainty Quantification in FWI

We adopt the Bayesian inference formulation for seismic inversion: given observational data and their uncertainty and a prior probability distribution describing uncertainty in the parameter field, find the posterior probability distribution over the parameter field. We exploiting the relation between the covariance matrix for the Gaussian approximation of the posterior probability density function (pdf) and the inverse of the Hessian to study this posterior pdf. A low-rank representation of the misfit Hessian is used to make our approach computationally feasible.

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

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

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

Georg Stadler
University of Texas at Austin
georgst@ices.utexas.edu

Lucas Wilcox
HyPerComp
me@lucaswilcox.com

MS91

Optimal Design of Simultaneous Source Encoding

Many parameter estimation problems involve the collection of an excessively large number of observations N, but it has been observed that similar results can often be obtained by considering a far smaller number K of multiple linear superpositions of experiments. To find the optimal weights of these superpositions, we formulate the problem as an optimal experimental design problem and show that the weights can be determined by using techniques from this field.

Kees van den Doel
University of British Columbia
kvdoel@cs.ubc.ca

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

Lior Horesh
Business Analytics and Mathematical Sciences
IBM TJ Watson Research Center
lhoresh@us.ibm.com

Kai Rothauge
University of British Columbia
rothauge@math.ubc.ca

MS92

GPU Computing with QUDA

The exponential growth of floating point power in GPUs and high memory bandwidth, has given rise to an attractive platform upon which to deploy HPC applications. However, deploying such computations on GPUs can be non-trivial because pre-existing applications cannot be recompiled and run while maintaining high performance. We review the QUDA library which is a domain-specific library designed to accelerate legacy lattice quantum chromodynamics applications through providing a library of the common performance-critical algorithms.

Michael Clark
Harvard-Smithsonian Center for Astrophysics
mclark@nvidia.com

MS92

VexCL: Vector Expression Template Library for OpenCL

VexCL is modern C++ library created for ease of OpenCL developement. VexCL strives to reduce amount of boiler-plate code needed to develop OpenCL applications. The library provides convenient and intuitive notation for vector arithmetic, reduction, and sparse matrix-vector multiplication. Multi-device and even multi-platform computations are supported. This talk is a brief introduction to VexCL interface.

Denis Demidov
Lobachevsky Institute of Mathematics and Mechanics
Kazan Federal University
ddemidov@ksu.ru

MS92

Developing Numerical Algorithms on Heterogeneous Architectures with High Productivity in Mind

Porting existing or developing new scientific applications on today's heterogeneous architectures can be a very challenging and time-consuming process. The necessity of abstracting the underlying hardware from the numerical developers becomes a crucial approach to effectively use the available processing units. This separation of concerns further allows the mathematician and the computer scientist to respectively concentrate on what they are good at. This talk will describe how some of the linear algebra community are tackling complex GPU-based systems when it comes to implementing high performance numerical libraries.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Jack Dongarra
University of Tennessee
dongarra@eecs.utk.edu

Hatem Ltaief
KAUST Supercomputing Laboratory
Thuwal, KSA
hatem.Ltaief@kaust.edu.sa

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

MS92

ViennaCL: GPU-accelerated Linear Algebra at the Convenience of the C++ Boost Libraries

In order to provide simple access to the vast computing resources in graphics processing units (GPUs) for general purpose scientific computing, the open source linear algebra library ViennaCL is presented. ViennaCL is written in C++ and used like existing CPU-based linear algebra libraries, thus it can be integrated into existing code easily. Moreover, the generic implementations of algorithms such as iterative solvers allow for code reuse beyond device and library boundaries.

Karl Rupp
Institute for Analysis and Scientific Computing, TU Wien
Institute for Microelectronics, TU Wien
rupp@mcs.anl.gov

MS93

A Multi-Scale Model for Capillary Driven Contact-Line Dynamics

We present a multi-scale method to simulate the flow of two immiscible incompressible fluids in contact with solids. The macro model is a level set method. The contact line is tracked explicitly and moves according to a slip velocity that depends on the wall contact angle of the interface with the solid. The relation between wall contact angle and slip velocity is determined in a micro model based on the phase field method. The phase field method seeks for an equilibrium slip velocity in a box around the contact point, prescribed a static contact angle at the solid and the wall contact angle in the far field. The dimensions of the box are chosen such that physical diffusion processes around the contact point are fully represented. We present numerical results for capillary-driven flows which demonstrate the convergence of results in the macro model and compare the behavior with other approaches in contact line dynamics.

Gunilla Kreiss, Martin Kronbichler
Division of Scientific Computing
Uppsala University
gunilla.kreiss@it.uu.se, kronbichler.martin@gmail.com

MS93

A Model for Simulating the Wrinkling and Buckling Dynamics of a Multicomponent Vesicle

Abstract not available at time of publication.

John Lowengrub
Department of Mathematics
University of California at Irvine
lowengrb@math.uci.edu

MS93

High-resolution Solver for the Poisson-Nernst-Planck Equations and its Applications

In this talk we present a high-resolution finite-volume method for solving the Poisson-Nernst-Planck equations on adaptive grids and for complicated geometries. We will highlight the importance of local charge conservation, at the coarse-fine grid interface, on the overall accuracy of the solver. Next, we utilize the solver to study the charging dynamics of super-capacitors at high voltages where nonlinear effects can lead to new charging mechanism previously unknown. Finally we will discuss possible future directions.

Mohammad Mirzadeh
Dept. Mechanical Engineering
UCSB
m.mirzadeh@engineering.ucsb.edu

MS93

Mechanical Simulation of Mammalian Acini

Acini are small groups of cells that form hollow compartments and serve multiple biological functions in different organs. They are a common source of various types of cancer, and recent work suggests that mechanical interactions of the cells with their local environment may play a role in cancer development. This talk will describe two mechanical simulation studies of acini that make use of new computational techniques for large-strain nonlinear elasticity and multiple deforming boundaries.

Chris Rycroft
Department of Mathematics
Lawrence Berkeley National Laboratory
chr@math.berkeley.edu

MS94

Avoiding Communication in Parallel Bidiagonalization of Band Matrices

Successive band reduction is a technique for reducing a symmetric band matrix to tridiagonal form for the symmetric eigenproblem. We have shown that a careful reformulation of the technique can asymptotically reduce the communication costs (i.e., data movement) on a sequential machine, compared to standard algorithms. In this talk, we will present the application of this approach to the reduction of a non-symmetric band matrix to bidiagonal form (for the SVD) in the distributed-memory parallel setting.

Grey Ballard, Nicholas Knight
UC Berkeley
ballard@cs.berkeley.edu, knight@cs.berkeley.edu

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

MS94

Improved Accuracy for MR3-based Eigensolvers

A number of algorithms exist for the dense Hermitian eigenproblem. In many cases, MRRR is the fastest one, although it does not deliver the same accuracy as Divide&Conquer or the QR algorithm. We demonstrate how the use of mixed precisions in MRRR-based eigensolvers leads to an improved orthogonality, even surpassing the accuracy of DC and QR. Our approach comes with limited performance penalty, and increases both robustness and scalability.

Paolo Bientinesi
AICES, RWTH Aachen
pauldj@aices.rwth-aachen.de

MS94

Spectral Divide-and-conquer Algorithms for Generalized Eigenvalue Problems

Spectral divide-and-conquer (SDC) algorithms solve eigenvalue problems by computing an invariant subspace for a subset of the spectrum to decouple the problem into two smaller subproblems. Recently Nakatsukasa and Higham introduced a stable and efficient SDC algorithm for the symmetric eigenvalue problem and the SVD. We propose an SDC algorithm for generalized eigenvalue problems that minimizes communication (its main computational kernels are QR factorization and matrix-multiplication) and has operation count within a small constant factor of that for the standard QZ algorithm.

Yuji Nakatsukasa
Department of Mathematics
University of Manchester
yuji.nakatsukasa@manchester.ac.uk

MS94

Restructuring the Symmetric QR Algorithm for Performance

Abstract not available at time of publication.

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

MS95

Model Reduction for Parameter Estimation in Computational Hemodynamics

We present a variational Data Assimilation procedure to the estimation of the Young modulus of an artery. The application of this approach to real problems is computationally intensive. We present a Proper Orthogonal Decomposition-based strategy for the reduction of the computational costs of the inverse problem associated with the parameter estimation. We address the role of surrogate modelling (such as 1D Euler equations) and surrogate solutions in this context.

Luca Bertagna
Dept. Math & CS
Emory University
lbertag@emory.edu

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

MS95

Energy-stable Galerkin Reduced Order Models for Prediction and Control of Fluid Systems

The focus of this talk is the construction of POD/Galerkin ROMs for real-time prediction and control of fluid systems. An energy stability analysis reveals that the inner product employed in the Galerkin projection step of the model reduction dictates the ROMs stability. For linearized compressible flow, a symmetry transform leads to a stable formulation for the inner product. Stability of the proposed ROM is demonstrated on several model problems. Extensions involving flow control are described.

Irina Kalashnikova, Srinivasan Arunajatesan, Bart G. Van Bloemen Waanders
Sandia National Laboratories
ikalash@sandia.gov, sarunaj@sandia.gov, bartv@sandia.gov

MS95

Computation of Periodic Steady States with the

Harmonic Balance Reduced Basis Method

In many applications, the flow solution reaches a periodic steady state. To compute the PSS, the harmonic balance method expands the solution and the operator of the problem as Fourier series. These expansions are truncated and the problem reduces to solving a set of coupled nonlinear equations. The harmonic balance method is coupled with the reduced basis method for spatial reduction. Error estimates and stability issues for the complex-valued reduced basis systems are considered.

Toni Lassila
Mathematics Institute of Computational Science and Eng
EPFL
toni.lassila@epfl.ch

Gianluigi Rozza
SISSA, International School for Advanced Studies
Trieste, Italy
gianluigi.rozza@sissa.it

MS95

Reduced Basis Methods for Coupled Transport-reaction Problems

We consider a parameter-dependent multiphysics problem which consists of horizontal transport and coupled bivariate convection-reaction. Such problems occur e.g. in catalysts. The coupling is nonlinear through Robin-type boundary conditions. Parameters may include inflow as well as chemical properties of the reaction. The truth discretization leads to a nonlinear generalized saddle-point problem. A Reduced Basis Method is developed and a-posteriori error estimates are developed. In addition, we consider other features that are often present in real-world CFD-problems such as parameter functions, long-time horizons, time-periodicity or stochastic influences. The talk is based upon joint work with Masayuki Yano, Anthony T. Patera (MIT), Dominik Lechler, Antonia Meyerhofer, Kristina Steih, Bernhard Wieland and Oliver Zeeb (all Ulm).

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

MS96

High Performance Computational Models of Coastal and Hydraulic Processes in an Interactive Python Environment

The Proteus toolkit is a software package for research on models for coastal and hydraulic processes and improvements in numerics. The models considered include multiphase flow, shallow water flow, turbulent free surface flow, and various flow-driven processes. We will discuss the objectives of Proteus and recent evolution of the toolkit's design as well as present examples of how it has been used used to construct computational models for the US Army Corps of Engineers.

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

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

MS96

Developing Open Source Software: Lessons Learned from Clawpack

The Clawpack (Conservation Laws Package) open source software project has grown substantially since 1994 and has recently branched into several related projects, with developers scattered at many institutions. I will give a brief overview of the current development process, with a focus on some aspects that may be of most interest to other researchers who are contemplating sharing their own code in this manner, such as choice of licenses, version control systems, public repositories and hosting, and the use of virtualization to facilitate use of the software.

Randall J. LeVeque
Applied Mathematics
University of Washington (Seattle)
rjl@uw.edu; rjl@amath.washington.edu; rjl@washingt

MS96

Feel++, a Library and Language in C++ for Galerkin Methods, from Rapid Prototyping to Large Scale Multi-physics Applications

Abstract not available at time of publication.

Christophe Prud'homme
University of Strasbourg
France
christophe.prudhomme@math.unistra.fr

MS96

The waLBerla/PE Parallel Multiphysics Framework

WaLBerla is a large scale software framework for multi-physics-applications based on kinetic methods. It is scalable to beyond 100 000 cores. This talk will focus on how to deal with conflicting goals like high node performance and scalability on the one side, and clean software structure, maintainability, and flexibility on the other side.

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ruede@informatik.uni-erlangen.de

Christian Feichtinger
Chiar for System Simulation
University of Erlangen-Nuremberg, Germany
christian.feichtinger@informatik.uni-erlangen.de

Harald Koestler
University of Erlangen-Nuremberg
harald.koestler@informatik.uni-erlangen.de

Tobias Preclik
University Erlangen-Nuremberg
tobias.preclik@informatik.uni-erlangen.de

Florian Schornbaum
University of Erlangen-Nuremberg
florian.schornbaum@informatik.uni-erlangen.de

MS97

Optimal Decision Making in Network Security Under Uncertainty

With the advent of modern cyber warfare techniques, the field of cyber security has been forced to adapt to a changing climate of malicious activity. Complicated multi-stage attacks that constantly change have become the new norm and have created a perpetual state of uncertainty for network administrators, who struggle to adapt their defenses to these new and ever-changing threats. In this work, we create a model that combines the methods of a Partially Observable Markov Decision process (POMDP) and Hidden Markov Models (HMM) in an attempt to optimize the action and strategies of a Network Administrator.

Michael J. Fowler
Department of Mathematics
Clarkson University
fowlermj@clarkson.edu

MS97

Optimization to Understand Trade-offs in Agricultural Practices

Seawater intrusion along coastal California threatens to destroy freshwater resources, which is especially detrimental to the berry industry in that region. Surface water analysis and crop simulation can further aid decision makers in planting cycles. We seek to reduce the aquifer draw by analyzing alternative farming techniques while simultaneously meeting demands and maximizing profits. This is accomplished using a farm model and optimization strategies to analyze approaches that meet a sustainable water yield constraint.

Kathleen Fowler
Clarkson University
Department of Mathematics
kfowler@clarkson.edu

MS97

Exploiting Expert Knowledge for Enhanced Simulation-Based Optimization

Safeguarding water supplies from contaminated sites is aided by linking models and global optimizers. This research augmented selected derivative-free optimizers to leverage non-traditional information like site-specific knowledge and practitioner rules-of-thumb. A benchmarking application was used in which extraction wells must intercept pollutants at a contaminated site. A rules engine adjusted candidate wells to place them within plume boundaries with a bias toward areas of rapid pollutant migration. Incorporating this expert knowledge significantly improved optimizer performance.

L. Shawn Matott, Camden Reslink
University at Buffalo
lsmatott@buffalo.edu, camdenre@buffalo.edu

MS97

Revealing the Difficulties in Managing Coastal Aquifer Supply Problems

This research examines the difficulties encountered when attempting to use optimization techniques to manage coastal ground water aquifers that are utilized as a source of fresh water. Coastal groundwater aquifers general consist of fresh water underlain by salt water. When fresh water is extracted from the aquifer, this mechanism draws the salt water into the fresh water zone. If the salt water is drawn into the well, then the well is contaminated and can no longer be utilized as a source of fresh water. The movement of the fresh water- salt water interface in response to ground water extraction can be predicted using mathematical models that are solved using computationally intensive algorithms. In this research the ground water flow model MODFLOW coupled with the principles of the Ghyben-Herzberg approach are used to model the fresh water- salt water interface. An optimization problem is developed that provides a fixed supply of fresh water while minimizing the threat of salt water contaminating the extraction wells in the model. Derivative free methods of optimization are used to solve this optimization problem, first using genetic algorithms, followed by a pattern search for refining the solutions determined. The results from this optimization exercise reveal multiple management solutions within the boundaries of acceptance for maintaining the utility of a coastal aquifer. Statistical tools are used to examine the solutions to the optimization problems, revealing symmetries and balancing that occurs in the solutions to coastal management systems.

Karen L. Ricciardi
University of Massachusetts Boston
Karen.Ricciardi@umb.edu

MS98

A Method-of-lines Approach to Computing on General Surfaces

The Closest Point Method is a set of mathematical principles and associated numerical techniques for solving partial differential equations (PDEs) posed on curved surfaces or other general domains. The method works by embedding the surface in a higher-dimensional space and solving the PDE in that space using simple finite difference and interpolation schemes. This presentation outlines some current work on formulating the algorithm as a method of lines and on solving elliptic problems.

Colin B. Macdonald
Oxford University
macdonald@maths.ox.ac.uk

Ingrid Von Glehn, Yujia Chen, Tom März
University of Oxford
ingrid.vonglehn@maths.ox.ac.uk,
yujia.chen@maths.ox.ac.uk, maerz@maths.ox.ac.uk

MS98

Strong Stability Preserving Methods for Time Evolution of Hyperbolic PDEs

In this talk we present the development of and recent research in strong stability preserving (SSP) methods for time evolution of hyperbolic PDEs. We will discuss SSP Runge–Kutta and multistep methods as well as multi-step multistage methods, and present the barriers and bounds that these methods have.

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

MS98

An ODE and PDE Test Suite for Time-stepping Methods

We have created a MATLAB test suite for ODEs and PDEs which allow users to test different time-integration methods in a simple way. In this talk we present the details of the test suite and its current capabilities.

Daniel L. Higgs
University of Massachusetts, Dartmouth
DanielHiggs@gmail.com

MS98

Matrix-free Integrators for Large Discretized PDEs

We present a new class of Rosenbrock integrators based on a Krylov space solution of the linear systems. We develop a framework for the derivation of order conditions for the new Rosenbrock-K methods. This new class of methods require only a small number of basis vectors, determined by the order of accuracy, but independent of the ODE under consideration. Numerical results show favorable properties of Rosenbrock-K methods when compared to current Rosenbrock and Rosenbrock-W methods.

Adrian Sandu, Paul Tranquilli
Virginia Polytechnic Institute and
State University
sandu@cs.vt.edu, ptranq@vt.edu

MS99

Very High Order Residual Distribution Schemes for Laminar and Turbulent Compressible Flow

We consider the numerical approximation of compressible viscous fluid problems by residual distribution schemes that can be considered as a non linear version of the stabilized finite element method. They need conformal unstructured (hybrid) meshes. We first show how to approximate scalar steady viscous, and show that optimal order can be reached. Then we extend this to fluid problems. We show how to deal with unsteady problems. Meshing issues will be also considered: optimal order on complex geometries require the use of curved meshes. In particular we foccus on the boundary layers geometrical approximation

Remi Abgrall
INRIA Bordeaux Sud-Ouest
Universite de Bordeaux
remi.abgrall@inria.fr

Dante de Santis, Mario Ricchiuto
INRIA
dante.de_santis@inria.fr, mario.ricchiuto@inria.fr

Cecile Dobrzynski
IMB, IPB, Université de Bordeaux, France
Bacchus, INRIA Bordeaux Sud-Ouest, France
cecile.dobrzynski@math.u-bordeaux1.fr

MS99

Robust Untangling of Curvilinear Meshes

We present a technique that allows to untangle high order/curvilinear meshes. The technique makes use of unconstrained optimization where element Jacobians are constrained to lie in a prescribed range through moving log-barriers. The untangling procedure starts from a possibly invalid curvilinear mesh and moves mesh vertices with the objective of producing elements that all have bounded Jacobians. Provable bounds on Jacobians are computed adaptively for any kind of elements, both for surface, volume, hybrid or boundary layer meshes.

Christophe Geuzaine
University of Liège
cgeuzaine@ulg.ac.be

MS99

H-to-P Efficiently: A Progress Report

The spectral/hp element method can be considered as bridging the gap between the traditionally low order finite element method on one side and spectral methods on the other side. Consequently, a major challenge which arises in implementing the spectral/hp element methods is to design algorithms that perform efficiently for both low- and high-order spectral/hp discretisations, as well as discretisations in the intermediate regime. In this talk, we explain how the judicious use of different implementation strategies – combined with an understanding of the architecture on which one plans to run – can be employed to achieve high efficiency across a wide range of polynomial orders. We examine both static and time-dependent problems, and examine differences between using continuous Galerkin versus discontinuous Galerkin discretizations.

Robert Kirby
University of Utah
kirby@sci.utah.edu

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

MS99

Simulation of An Oscillating-Wing Power Generator Using a High-Order CFD Method

We developed a massively parallel 3D code of compressible Navier-Stokes equations on moving and deforming domains. The code is based on an efficient high-order correction procedure via reconstruction (CPR). We will report our recent study of unsteady turbulent flow past an oscillating-wing power generator using this code. We apply a technique of active flow control to gain improved efficiency of this power generator.

Chunlei Liang
George Washington University
chliang@gwu.edu

MS100

Exploring Code Performance Issues on Many-core Devices using the Multifluid PPM Code as a Representative CFD Application

The PPM (piecewise-Parabolic Method) code has been used in a detailed study of code performance issues on many-core devices, represented by Intels new MIC coprocessor, with over 50 cores. The different implementation strategies that were tried will be described along with the implications of the observed performance for recommended programming techniques for many-core devices. In particular a key technique for increasing the computational intensity of the algorithm will be described.

Paul R. Woodward
Laboratory for Computational Science and Engineering
University of Minnesota
paul@lcse.umn.edu

Jagan Jayaraj, Pei-Hung Lin, Michael Knox
University of Minnesota
jaganj@cs.umn.edu, phlin@cs.umn.edu,
mikeknox@lcse.umn.edu

Simon D. Hammond
Scalable Computer Architectures
Sandia National Laboratories
sdhammo@sandia.gov

MS100

Accelerating Mini-FE, a Finite Element Proxy Application, on GPUs

The Mantevo performance project is a collection of self-contained proxy applications that illustrate the main performance characteristics of important algorithms. miniFE is intended to be and approximation to an unstructured implicit finite element or finite volume application. This talk will focus on GPU algorithms for assembling a finite element matrix. Results on both NVIDIAs Fermi and Kepler GPUs will be presented.

Justin Luitjens
NVIDIA Corporation
juitjens@nvidia.com

MS100

Developing a Multi-Architecture Implicit Particle-in-Cell Proxy

PlasmaApp3D is a 3D fully implicit Particle-in-Cell proxy app developed to explore co-design issues associated with various levels of hardware and algorithmic abstraction on heterogeneous architectures with multiple levels of parallelism. The guiding philosophy behind this proxy app is that the physics should only be implemented once, and should be independent of the underlying hardware specific operations and data-management. This notion of isolating the physics leads to a proxy app that allows for rapid testing of various data-management and architecture specific optimizations. The current goal of the PlasmaApp3D developement is a proxy app optimized for shared multi-core and GPU system, although this may be later extended to include other architectures. Performance results as well as design and development methods will be presented.

Joshua Payne, Dana Knoll, Allen McPherson, William Taitano, Luis Chacon
Los Alamos National Laboratory
payne@lanl.gov, nol@lanl.gov, mcpherson@lanl.gov, taitano@lanl.gov, chacon@lanl.gov

MS100

UMMA: Unstructured Mesh Micro Apps

We begin with a discussion of our target application, Chicoma. We then cover the conception and development of Unstructured Mesh Micro-Applications (UMMA) on a variety of architectures and finish by relating our experiences introducing ideas from UMMA back into our large application.

Dean Risinger
Los Alamos National Laboratory
ldr@lanl.gov

MS101

Hybrid FOSLS/FOSLL*

Abstract not available at time of publication.

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

MS101

Least-Squares Finite Element Methods for Coupled Generalized Newtonian Stokes-Darcy Flow

The coupled problem for a generalized Newtonian Stokes flow in one domain and a generalized Newtonian Darcy flow in a porous medium is considered in this talk. The flows are treated as a stress-velocity formulation for the Stokes problem and a volumetric flux-hydraulic potential formulation for the Darcy problem. The coupling is done by using the well known Beavers-Joseph-Saffman interface condition. A least-squares finite element method is used for the numerical approximation of the solution.

Steffen Münzenmaier
Gottfried Wilhelm Leibniz University of Hannover
muenzenm@ifam.uni-hannover.de

MS101

Goal-Oriented Least-Squares Finite Element Methods

We present an approach to augment least-squares finite element formulations with user-specified goals. The method inherits the global approximation properties of the standard formulation with increased resolution of the goal. Several theoretical properties such as optimality and enhanced convergence under general assumptions are discussed and we present an adaptive approach that results in efficient, locally refined approximations that hone in on the quantity-of-interest with a range of numerical examples to support the approach.

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

Jehanzeb Chaudhry
Department of Mathematics
Colorado State University
jehanzeb@colostate.edu

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Kuo Liu, Tom Manteuffel
Department of Applied Mathematics
University of Colorado at Boulder
liukuo99@gmail.com, tmanteuf@colorado.edu

Lei Tang
University of Colorado at Boulder
lei.tang@colorado.edu

MS101

Least Squares Finite Element Methods for Non-Newtonian Fluids with Application to Blood Flow

Due primarily to the presence of red blood cells, whole blood may exhibit significant non-Newtonian behavior, and consequently, accurate numerical models for blood flow should reflect the appropriate physics. This talk examines three categories of nonlinear constitutive models applicable to blood: viscoelastic fluids of Oldroyd type, shear-thinning fluids, and yield-stress behavior. In particular, we focus on the least-squares finite element approach as the basis for an accurate and flexible discretization technique for these challenging problems.

Chad Westphal
Wabash College
westphac@wabash.edu

MS102

Discretization and Solvers for the Stokes Equations of Ice Sheet Dynamics at Continental Scale

Ice sheets exhibit incompressible creeping flow with shear-thinning rheology. On a continental scale, the flow is characterized by localized regions of fast flow that are separated from vast slow regions by thin transition zones. We use a parallel, adaptive mesh, higher-order finite element discretization and an inexact Newton method for the solution of the nonlinear Stokes equations modeling ice sheet dynamics. Preconditioned Krylov methods for the scalable solution of the discretized system, and effects of a highly anisotropic discretization are discussed.

Tobin Isaac, Georg Stadler, Omar Ghattas
University of Texas at Austin
tisaac@ices.utexas.edu, georgst@ices.utexas.edu, omar@ices.utexas.edu

MS102

Resolving Grounding Line Dynamics Using the BISICLES Adaptive Mesh Refinement Model

Ice sheet dynamics span a wide range of scales. Extremely fine spatial resolution is required to resolve the dynamics of features like grounding lines, while such fine resolution is unnecessary over large quiescent regions. BISICLES is a scalable adaptive mesh refinement (AMR) ice sheet model built on the Chombo framework, with a dynamical core based on the vertically-integrated model of Schoof and Hindmarsh (2010). We will present results demonstrating the effectiveness of our approach.

Daniel Martin
Lawrence Berkeley National Laboratory
dfmartin@lbl.gov

Stephen Cornford
University of Bristol
s.l.cornford@bristol.ac.uk

William Lipscomb
Los Alamos National Laboratory
lipscomb@lanl.gov

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

Stephen Price
Los Alamos National Laboratory
sprice@lanl.gov

MS102

A Finite Element Implementation of Higher-order Ice-sheet Models: Mathematical and Numerical Challenges

Several models, characterized by different complexity and accuracy, have been proposed for describing ice-sheet dynamics. We introduce a parallel finite element implementation for some of these models which constitutes part of the land ice dynamical core for MPAS library. We present large-scale simulations of the Greenland ice-sheet, and explore methods for the solution of the resulting linear and nonlinear systems. We also address the estimation of model parameters such as the friction coefficient at the ice-bedrock interface.

Max Gunzburger
Florida State University
School for Computational Sciences
mgunzburger@fsu.edu

Matt Hoffman
Los Alamos National Lab
mhoffman@lanl.gov

Wei Leng
Laboratory of Scientific and Engineering Computing
Chinese Academy of Sciences, Beijing, China
lengweee@gmail.com

Mauro Perego
Florida State University
mperego@fsu.edu

Stephen Price
Los Alamos National Laboratory
sprice@lanl.gov

MS102

Analysis of Convergence and Performance Variability for Continental Ice Sheet Modeling at Scale

A prototype python-based verification and validation package to assess the significant development work occurring in continental-scale ice sheet models is presented. A key aspect is a performance V&V capability to quantify algorithms that are efficient but also sensitive and variable, performance variability of leadership class hardware, and the impact of developments on both speed and robustness. The assessment of alternative dynamical core performance in terms of simulation value (speed versus accuracy) is discussed.

Patrick H. Worley, Katherine J. Evans
Oak Ridge National Laboratory
worleyph@ornl.gov, evanskj@ornl.gov

Adrianna Boghozian
University of Tennessee
aboghozi@utk.edu

MS103

Fourier Continuation Methods for Therapeutic Ultrasound

HIFU is an ultrasound therapy which focuses destructive acoustic energy on a target (e.g. a cancerous tumor), leaving the surrounding tissue unharmed. For HIFU simulation, wherein nonlinear waves propagate many times their fundamental wavelength, the use of high-order methods is vital for accurate and efficient simulation. We present a class of high-order solvers, based on the Fourier Continuation method, that are capable of accurate and efficient HIFU simulation in large, complex domains.

Nathan Albin
Kansas State University
albin@math.ksu.edu

MS103

Jet Noise DNS using Arbitrary-order Hermite Methods

We discuss the application of Hermite methods to direct numerical simulations of compressible turbulence with special attention to jet noise.

Daniel Appelo
Department of Mathematics and Statistics
University of New Mexico
appelo@math.unm.edu

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

Tim Colonius
Division of Engineering and Applied Science
California Institute of Technology
colonius@caltech.edu

Chang Young Jang
Southern Methodist University
cjang@mail.smu.edu

MS103

Stable High Order Finite Difference Methods for Wave Propagation Problems

During the last decade, stable high order finite difference methods as well as finite volume methods applied to initial-boundary-value-problems have been developed. The stability is due to the use of so-called summation-by-parts operators (SBP), penalty techniques for implementing boundary and interface conditions, and the energy method for proving stability. In this talk we discuss some aspects of this technique including the relation to the initial-boundary-value-problem. By reusing the main ideas behind the recent development, new coupling procedures for multi-physics applications have been developed. We will present the theory by analyzing simple examples and apply to complex multi-physics problems such as fluid flow problems, elastic and electromagnetic wave propagation, fluid-structure interaction and conjugate heat transfer.

Jan Nordstrom
professor in scientific computing
Department of Mathematics, Linkoping University
jan.nordstrom@liu.se

MS103

Stability of Interacting Solitary Water Waves, Standing Waves, and Breathers

We develop a high-performance shooting algorithm to compute new families of time-periodic and quasi-periodic solutions of the free-surface Euler equations involving breathers, traveling-standing waves, and collisions of counter-propagating and unidirectional solitary waves of various types. The wave amplitudes are too large to be well-approximated by weakly nonlinear theory, yet we often observe behavior that resembles elastic collisions of solitons in integrable model equations. A Floquet analysis shows that many of the new solutions are stable to harmonic perturbations.

Jon Wilkening
UC Berkeley Mathematics
wilken@math.berkeley.edu

MS104

2D FTLE in 3D Flows: The Accuracy of using Two-dimensional Data for Lagrangian Analysis in Three-dimensional Fluid Flows

In experimental, three-dimensional vortex-dominated flows, common particle image velocimetry (PIV) data is often collected in only the plane of interest due to equipment constraints. For flows with significant out of plane velocities or velocity gradients, this can create large discrepancies in Lagrangian analyses that require accurate particle trajectories. A Finite Time Lyapunov Exponent (FTLE) analysis is one such example, and has been shown to be very powerful at examining vortex dynamics and interactions in a variety of aperiodic flows. In this work, FTLE analysis of a turbulent channel simulation was conducted using both full three-dimensional velocity data and modified planar data extracted from the same computational domain. When the out-of-plane velocity component is neglected the difference in FTLE fields is non-trivial. A quantitative comparison and computation of error is presented for several planes across the width of the channel to determine the efficacy of using 2D analyses on the inherently 3D flows.

Melissa Green
Syracuse University
greenma@syr.edu

MS104

Dimensionality Reduction in Neuro-sensory Systems

Abstract not available at time of publication.

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

MS104

Model Reduction for Large-scale Systems using Balanced POD and Koopman Modes

Abstract not available at time of publication.

Clancy W. Rowley
Princeton University
cwrowley@princeton.edu

MS104

Recent Advances in Discrete Empirical Interpolation for Nonlinear Model Reduction

Abstract not available at time of publication.

Danny C. Sorensen
Rice University
sorensen@rice.edu

MS105

A Geometric Load-Balancing Algorithm for Multicore Parallel Computers

Geometric partitioning is fast and effective for load-balancing dynamic applications, particularly those requiring geometric locality of data (particle methods, crash simulations). We present, to our knowledge, the first parallel implementation of a multidimensional-jagged geometric partitioner. Its MPI+OpenMP implementation makes it compatible with hybrid MPI+threads applications on multicore-based parallel architectures. By keeping data in place throughout the algorithm, we minimize data movement compared to recursive bisection methods. We demonstrate the algorithm's scalability and quality relative to recursive bisection.

Mehmet Deveci
THe Ohio State University
mdeveci@bmi.osu.edu

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
umit@bmi.osu.edu

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

MS105

Parallel Anisotropic Mesh Adaptation with Specific Consideration of Flow Features

Effective PDE-based simulations in a number of fields require meshes with strong anisotropy and controlled meshes that isolate specific solution features. For example, fluid simulations require appropriate anisotropic mesh layouts near boundaries (e.g., boundary layers) and at interior portions (e.g., flows with shocks or free shear layers). In this talk, we will present recent advances made in local mesh modification procedures for adaptively creating such meshes. Parallelization of these procedures will also be discussed.

Onkar Sahni, Ryan Molecke, Mark Shephard
Rensselaer Polytechnic Institute
sahni@scorec.rpi.edu, molecr@rpi.edu, shephard@rpi.edu

Kedar Chitale
Univeristy of Colorado Boulder
kendar.chitale@colorado.edu

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

Saurabh Tendulkar
Simmetrix Inc.
saurabh@simmetrix.com

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

MS105

Predictive Load Balancing Using Mesh Adjacencies for Mesh Adaptation

Parallel mesh adaptation on unstructured meshes requires that the mesh be distributed across a large number of processors such that the adapted mesh fits within memory. The goal of ParMA it to dynamically partition unstructured meshes directly using the existing mesh adjacency information to account for multiple criteria. Results will demonstrate the ability of ParMA to rebalance large meshes (greater than 100,000,000 regions) on large core count machines (greater than 32,000) accounting for multiple criteria.

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

Onkar Sahni, Mark Shephard
Rensselaer Polytechnic Institute
sahni@rpi.edu, shephard@rpi.edu

MS105

Parallel Mesh Generation and Adaptation on CAD Geometry

Effective parallel generation and adaptivity of meshes of CAD geometry add complexity due to the need to ensure the mesh conforms to the CAD geometry. We will discuss procedures to support unstructured meshes, with or without boundary layers, on massively computers including the ability to distribute the geometric model during mesh adaptation. This leads to substantial memory usage reductions on high processor counts by not requiring a complete copy of the model on each process.

Saurabh Tendulkar
Simmetrix Inc.
saurabh@simmetrix.com

Ottmar Klaas
Simmetrix, Inc.
oklaas@simmetrix.com

Rocco Nastasia
Simmetrix Inc.
nastar@simmetrix.com

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

MS106

Towards Real-Time Power Grid Dynamics Simula-

tion using PETSc

The need for real-time power grid dynamics simulation has been a primary focus of the power system community in recent years. In this talk we present our experiences, and present preliminary results, to simulate power grid dynamics in real-time using the high performance computing library PETSc. With the range of scalable linear, nonlinear, and time-stepping solvers, PETSc has the potential to be the mathematical and computing framework needed for an online dynamics simulator.

Shrirang Abhyankar
Argonne National Laboratory
abhyshr@mcs.anl.gov

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

Alexander Flueck
Illinois Institute of Technology
flueck@iit.edu

MS106

Exploitation of Dynamic Information in Power System Phasor Measurements

The power grid is becoming more dynamic with high penetration of intermittent renewable sources and responsive loads. It is essential to establish a dynamic operation paradigm in contrast to todays operation built on static modeling. Recent developments in phasor technology provide such an opportunity with high-speed time-synchronized measurement data. This talk will discuss the mathematical and computational challenges and recent advancements in extracting and utilizing dynamic information in phasor measurements.

Zhenyu Huang
Pacific Northwest National Laboratory
zhenyu.huang@pnnl.gov

MS106

Next Generation Modeling and Simulation of Building Energy and Control Systems

Building energy and control systems can be represented by systems of coupled stiff differential equations, algebraic equations and discrete equations. These equations couple models from multiple physical domains. We will present how these systems can be modeled using the equation-based, object-oriented Modelica language. We discuss numerical challenges and opportunities to use such models for analysis that is outside the capabilities of conventional building simulation programs.

Michael Wetter
LBNL
mwetter@lbl.gov

MS106

Scalable Dynamic Optimization

We present an approach for nonlinear programming (NLP) based on the direct minimization of an exact diffferentiable penalty function using trust-region Newton techniques. The approach provides all the features required for scalability: it can effifficiently detect and exploit directions of negative curvature, it is superlinearly convergent, and it enables the scalable computation of the Newton step through iterative linear algebra. Moreover, it presents features that are desirable for parametric optimization problems that have to be solved in a latency-limited environment, as is the case for model predictive control and mixed-integer nonlinear programming.

Victor Zavala
Argonne National Laboratory
vzavala@mcs.anl.gov

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

MS107

High Resolution Simulation of Two-phase Flows on Quadtree Grids

Abstract not available at time of publication.

Arthur Guittet
UCSB
arthur.guittet@gmail.com

MS107

High-Order Interface Tracking Methods for Compressible and Incompressible Two-Phase Flow

We present new high-order interface tracking methods for compressible and incompressible two-phase flow. Both systems represent domains with a signed distance function, and PDEs use embedded boundary finite volume discretizations. For hyperbolic problems the level set is advected using a velocity from Riemann problems. For elliptic problems, the interface velocity comes from coupled implicit solves for viscosity and pressure. We will demonstrate the methods using two-phase compressible Euler equations, and two-phase Navier-Stokes with surface tension.

Mehdi Vahab
University of California Davis
Department of Applied Science
mvahab@ucdavis.edu

Greg Miller
University of California
Department of Applied Science
grgmiller@ucdavis.edu

MS108

A High-order Discontinuous Galerkin Method with Lagrange Multipliers for Advection-diffusion Problems

A high-order Discontinuous Galerkin method with Lagrange Multipliers (DGLM) is presented for the solution of advection-diffusion problems on unstructured meshes. Unlike other hybrid discretization methods for transport problems, it operates directly on the second-order form the advection-diffusion equation. Like the Discontinuous Enrichment Method (DEM), it chooses the basis functions among the free-space solutions of the homogeneous form of the governing partial differential equation, and relies on Lagrange multipliers for enforcing a weak continuity of the approximated solution across the element interface boundaries. However unlike DEM, the proposed hybrid discontinuous Galerkin method approximates the Lagrange multipliers in a subspace of polynomials, instead of a subspace of traces on the element boundaries of the normal derivatives of a subset of the basis functions. For homogeneous problems, the design of arbitrarily high-order elements based on this DGLM method is supported by a detailed mathematical analysis. For non-homogeneous ones, the approximated solution is locally decomposed into its homogeneous and particular parts. The homogeneous part is captured by the DGLM elements designed for homogenous problems. The particular part is obtained analytically after the source term is projected onto an appropriate polynomial subspace. This decoupling between the two parts of the solution is another differentiator between DGLM and DEM with attractive computational advantages. An a posteriori error estimator for the proposed method is also derived and exploited to develop an automatic mesh refinement algorithm. All theoretical results are confirmed by numerical simulations that furthermore highlight the potential of the proposed high-order hybrid DG method for transport problems

Sebastien Brogniez, Charbel Farhat
Stanford University
brogniez@stanford.edu, CFarhat@stanford.edu

MS108

Entropy Stability and High-order Approximation of the Compressible Euler Equations

This talk will discuss questions regarding parabolic regularization of the Euler equations and entropy stability. In particular a sub-class of parabolic regularizations is identified that yields a minimum entropy principle. The consequences of this property will be illustrated on a high-order finite elements method.

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

Bojan Popov
Department of Mathematics
Texas A&M University
popov@math.tamu.edu

Murtazo Nazarov
TAMU
murtazo@math.tamu.edu

MS108

Discontinuous Galerkin Spectral Element Approximation for Wave Scattering from Moving Objects

Accurate computation of wave scattering from moving, perfectly reflecting objects, or embedded objects with material properties that differ from the surrounding medium, requires methods that accurately represent the boundary location and motion, propagate the scattered waves with low dissipation and dispersion errors, and don't introduce errors or artifacts from the movement of the mesh. We describe development of a discontinuous Galerkin spectral element approximation to satisfy these requirements. Examples are taken from aeroacoustics and electromagnetics.

David A. Kopriva
Department of Mathematics
The Florida State University
kopriva@math.fsu.edu

MS108

Discontinuous Galerkin Methods for Vlasov Maxwell Equations

The Vlasov-Maxwell system is one of the important models to study collisionless magnetized plasmas. It couples the Vlasov equation satisfied by the distribution function of the particle(s) and the Maxwell system, and has wide applications such as in space and laboratory plasmas, and fusion. The challenges in simulation come from high-dimensionality, multiple scales, and conservation. We will present our recent progress in developing and analyzing discontinuous Galerkin methods for Vlasov-Maxwell equations.

Fengyan Li
Rensselaer Polytechnic Institute
lif@rpi.edu

MS109

Stochastic Collocation Techniques for Uncertainty Quantification in Reactor Criticality Problems

This talk focuses on stochastic collocation method applied to a radiation transport problem. We consider a one dimensional reactor model and we study the effects of material cross-section uncertainty on the reactor criticallity, specifically uncertainty in the properties of the nuclear fuel, control rod and potential crud deposit. We use the DENOVO radiation transport simulator and we demonstrate convergence of the stochastic collocation technique even for very large ranges of the uncertainty.

Nick Dexter
University of Tennessee
ndexter@utk.edu

MS109

Local Sensitivity Derivative Enhanced Monte Carlo

We present a Local Sensitivity Derivative Enhanced Monte Carlo (LSDEMC) method that utilizes derivative information in Voronoi cells around sample points to construct surrogate response functions, then uses MC integration to correct the bias in the surrogate-based evaluations. A new estimator for the mean is developed using this strategy, which has lower variance than a simple average. We illustrate the use of LSDEMC in uncertainty propagation for an SPDE system and present some theoretical results.

Vikram Garg
Massachusetts Institute of Technology
vikram.v.garg@gmail.com

Roy Stogner
University of Texas at Austin
roystgnr@ices.utexas.edu

MS109

Adaptive ANOVA-based Probabilistic Collocation Kalman Filter for Inverse Uncertainty Quantifica-

tion

Probabilistic Collocation Kalman Filter (PCKF) is an approach solving inverse uncertainty quantification problems. PCKF resembles the Ensemble Kalman Filter (EnKF) except that it represents and propagates model uncertainty with the Polynomial Chaos Expansion (PCE) instead of an ensemble of model realizations. The accuracy and efficiency of PCKF depends on the selection of PCE bases. We present an adaptive algorithm to build the PCE expression for PCKF. The adaptivity is based on two techniques. First, we adopt the idea of adaptive functional ANOVA (Analysis of Variance) decomposition, which approximates a high dimensional function with the summation of a set of low dimensional functions. Thus instead of expanding the original model into PCE, we implement the PCE expansion on much lower dimensions. Second, for those inverse problems where the observations are assimilated sequentially, we propose the method of Dynamic Parameterization which controls the dimensionality of the problem (the number of random variables used to represent parameter uncertainty) in different Kalman Filter loops. We give two illustrative examples to demonstrate our algorithm.

Guang Lin
Pacific Northwest National Laboratory
guang.lin@pnnl.gov

Weixuan Li
University of Southern California
weixuan23@gmail.com

Dongxiao Zhang
USC
donzhang@usc.edu

MS109

Intrusive Analysis and Uncertainty Quantification for Nuclear Engineering Models

We investigate uncertainty analysis of nuclear engineering simulation models. Given that sampling at high-resolution, over large parameter space, is computationally prohibitive, we argue for the use of compact representations of uncertainty and for extraction of additional information from each model run. We accelerate the standard response surface learning techniques by using gradient information, and also calibrated lower-resolution data; goal-oriented model reduction is an important tool. Our work has been tested on codes developed within NEAMS/SHARP.

Oleg Roderick
Argonne National Laboratory
roderick@mcs.anl.gov

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

MS110

Graph 500 Performance on a Distributed-Memory Cluster

The Graph 500 benchmark ranks high-performance computers based on speed of memory retrieval. We report on our experience implementing this benchmark on the distributed-memory cluster tara in the UMBC High Performance Computing Facility. Our best run to date using 64 nodes achieved a GTEPS rate that would have put tara at rank 58 on the June 2012 Graph 500 list. We have submitted an official benchmark run for publication in the Graph 500 list. Advisor: Matthias Gobbert, UMBC

Jordan Angel
East Tennessee State University
zjba15@goldmail.etsu.edu

MS110

Comparative Transcriptome Analysis to Identify Genes Regulating Elastogenesis

Arterial smooth muscle cells produce significantly more elastic fibers when overexpressing the V3 variant of the proteoglycan versican. To identify possible causes of this increase, microarray and qPCR analyses were used to measure associated changes in gene expression. We refined the set of V3 influenced genes by analyzing their expression in two other models of elastogenesis. Twenty-nine genes were consistently upregulated or downregulated in all three models, suggesting these genes importance in elastogenesis.

Sharon Guffy
Medical University of South Carolina
guffysl@email.wofford.edu

MS110

ACTIVE: A Bayesian Approach to Asset Covariance Estimation

An improved stock return covariance estimation procedure is developed using an implicit Bayesian method known as shrinkage. We extend the work of Ledoit (2003) in covariance matrix shrinkage by incorporating the VIX, a market implied volatility index, into the estimation process and name our strategy ACTIVE. ACTIVE's ability to forecast market volatility makes it more dynamic by allowing it to respond to changing financial regimes, making it a viable alternative to industry standard techniques. Advisors: Marcel Blais, Worcester Polytechnic University and Carlos Morales, Wellington

Daniel Helkey
Emmanuel College
helkeyd@emmanuel.edu

Tejpal Ahluwalia
New Jersey Institute of Technology
tba@siam.org

Robert Hark
Montana Tech
tba@siam.org

Nicholas Marshall
Clarkson University
tba@siam.org

MS111

Multigrid For Divergence-Conforming Discretizations of the Stokes Equations

Recent years have seen renewed interest in the numerical solution of the Stokes Equations. Divergence-conforming discretization methods are of interest because they allow for the strong solution of the incompressibility condition. Monolithic multigrid methods offer the possibility of solving the resulting systems with better overall efficiency than block-factorization methods. We explore a BDM-based discretization, BDM-P_0, and present preliminary numerical experiments using overlapping Vanka smoothers within GMG, as well as analysis of the method.

James H. Adler
Tufts University
james.adler@tufts.edu

Thomas Benson, Scott Maclachlan
Department of Mathematics
Tufts University
thomas.benson@tufts.edu, scott.maclachlan@tufts.edu

MS111

Modeling and Numerical Simulations of Immiscible Multi-Phase Fluids

We introduce a multi-phase fluid model derived via an energetic variational approach. We present simulations of several applications of the model, including the dynamics of a buoyant air bubble or a falling solid in a two-phase Stokes flow. Possible slip effects at the interfaces of the fluids are also considered. Numerical simulations are presented, including results of pressure Schur complement approaches for solving the fluid system and multilevel solvers for solving the pressure system in each iteration of these schemes.

James Brannick
Department of Mathematics
Pennsylvania State University
brannick@psu.edu

MS111

Performance of Efficient AMG Based Preconditioners for Implicit Resistive MHD

The resistive magnetohydrodynamics model describes the dynamics of charged fluids in the presence of electromagnetic fields. This model is strongly coupled, highly nonlinear, and characterized by physical mechanisms that span a wide-range of interacting time scales. Implicit time stepping is an attractive choice for simulating a wide range of physical phenomena. Such methods require an effective preconditioner. We explore the performance of several candidate AMG based preconditioners including fully-coupled AMG projection methods and approximate block factorization methods.

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Luis Chacon
Los Alamos National Laboratory
chacon@lanl.gov

MS111

An Investigation of Multigrid Smoothers for Fully Implicit 2D Incompressible Resistive MHD

We consider a number of different smoothers for use within a multigrid procedure applied to a resistive MHD application. The smoothers include a generalization of a Braess-Sarazin procedure, a generalization of a Vanka smoother, and the use of operator split (or approximation block factorizations) techniques to derive a new smoother. A number of numerical results are presented illustrating the merits of the different smoothers in terms of scalability, setup, and cost per iteration.

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

Tom Benson
Tufts University
thomas.benson@tufts.edu

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

MS112

State and Parameter Sensitivities and Estimation in Coupled Ocean/ice Shelf/ice stream Evolution

A challenge in cryospheric research is the quantification of what controls coupled ice shelf/ice sheet evolution in response to sub-ice shelf cavity circulation changes and melting. Here we present a modeling infrastructure for simulating the evolution of ice sheet, ice shelf, and sub-ice shelf circulation, for inferring high-dimensional fields of state and parameter sensitivities through adjoint model components, and for gradient-based optimal estimation using available observations. Examples provided include idealized setups as well as realistic simulations for Pine Island Ice Shelf, West Antarctica.

Patrick Heimbach
Massachusetts Institute of Technology
heimbach@mit.edu

Daniel Goldberg
MIT
dgoldber@mit.edu

Martin Losch
Alfred Wegener Institute
martin.losch@awi.de

MS112

Hierarchical Error Estimates for Adaptive Shallow Ice Sheet and Ice Shelf Models

Mesh adaptation methods are necessary to face the sharp changes of the ice flow regime in the narrow transition zones between ice sheets and ice shelves. However, existing

refinement criteria are often empirical. To optimize automatically the number and the position of mesh nodes, we implement a new hierarchical error estimate for the shallow shelf approximation equation. The efficiency of the resulting re-meshing procedure is shown on the MISMIP exercices.

Guillaume Jouvet
Free University of Berlin
Institut für Mathematik
guillaume.jouvet@fu-berlin.de

Carsten Gräser
Freie Universität Berlin
Institut für Mathematik
graeser@math.fu-berlin.de

MS112

A Trust Region Stochastic Newton MCMC Method with Application to Inverse Problems Governed by Ice Sheet Flows

We address the problem of quantifying uncertainty in the solution of inverse problems governed by Stokes models of ice sheet flows within the framework of Bayesian inference. The posterior probability density is explored using a stochastic Newton MCMC sampling method that employs local Gaussian approximations based on gradients and Hessians (of the log posterior) as proposal densities. The method is applied to quantify uncertainties in the inference of basal boundary conditions for ice sheet models.

Noemi Petra
Institute for Computational Engineering and Sciences (ICES)
The University of Texas at Austin
noemi@ices.utexas.edu

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

Georg Stadler, Omar Ghattas
University of Texas at Austin
georgst@ices.utexas.edu, omar@ices.utexas.edu

MS112

Initialization Strategy for Short Term Projections using the Ice Sheet System Model

Accurate projections of ice sheet contribution to sea level rise require numerical modeling. However, large-scale modeling of ice sheets remains scientifically and technically challenging. The Ice Sheet System Model (ISSM) is a three-dimensional thermo-mechanical model that relies on the finite element method. It includes higher-order and full-Stokes ice flow approximations, adaptive mesh refinement, and data assimilation among other capabilities. We present here the initialization strategy adopted to improve short-term simulations of ice sheet dynamics.

Helene Seroussi
Jet Propulsion Laboratory
helene.seroussi@jpl.nasa.gov

Mathieu Morlighem
University of California - Irvine
mathieu.morlighem@jpl.nasa.goc

Eric Larour
Jet Propulsion Laboratory
eric.larour@jpl.nasa.gov

Eric Rignot
University of California - Irvine
Jet Propulsion Laboratory
eric.rignot@jpl.nasa.gov

Ala Khazendar
Jet Propulsion Laboratory
ala.khazendar@jpl.nasa.gov

MS113

Local Time-stepping and Spectral Element Methods for Wave Propagation

We consider the numerical solution of the wave equation on locally refined meshes. To overcome the severe stability restriction of explicit time-stepping methods, due to possibly just a few small elements in the mesh, we consider leapfrog based local time-stepping schemes, combined with a high-order spectral element spatial discretization. Numerical experiments illustrate the efficiency of the proposed method.

Loredana Gaudio
University of Basel
loredana.gaudio@unibas.ch

Marcus Grote
Universität Basel
marcus.grote@unibas.ch

MS113

High Efficiency Algorithms for Incompressible Flows and Moving Geometry

This talk will discuss some new high-order accurate algorithms for simulating incompressible flows using overlapping grids for complex, possibly moving geometry. The approach is based on an approximate-factored compact scheme for the momentum equations together with a fourth-order accurate multigrid solver for the pressure equation. The scheme will be described and results will be presented for some three-dimensional (parallel) computations of flows with moving rigid-bodies.

William D. Henshaw
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
henshaw@llnl.gov

MS113

Debye Potentials for the Time Dependent Maxwell Equations

The explicit solution to the scattering problem of time dependent Maxwell equations on a sphere is derived. The derivation of the explicit solution is based on a generalization of Debye potentials for the time harmonic case and reduces the problem to two scalar wave problems - one with Dirichlet condition and the other with Robin condition. A spectrally accurate discretization scheme is discussed and numerical examples are presented. The solution can be served as a reference for checking the accuracy of other numerical methods. More importantly, it will provide some insight towards better integral equation formulations for

the wave equation and time-dependent Maxwell equations in a general domain.

Leslie Greengard
Courant Institute
New York University
greengard@courant.nyu.edu

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

Shidong Jiang
Department of Mathematical Sciences
New Jersey Institute of Technology
shidong.jiang@njit.edu

MS113

Symmetry-preserving High-order Schemes

In this talk we will show how to systematically construct numerical schemes that preserve the geometrical structure of the underlying PDE. We will then show how to extend these ideas to generate high-order versions of these schemes.

Jean-Christophe Nave
McGill University
jcnave@math.mcgill.ca

Alexander Bihlo
CRM
McGill
alexander.bihlo@univie.ac.at

MS114

Combining Domain Decomposition with Multigrid Optimization for Hierarchical Problems Arising from Nanoporous Material Design

Multigrid optimization algorithms have been shown to improve the convergence on nanoporous material design problems, owing to their multiscale nature. However, for more realistic applications, the large problems corresponding to the finest resolution levels represent a performance bottleneck. Well-known domain decomposition techniques can accelerate the solution of the PDE constraints, i.e. the inner loop of the optimization algorithm. In the present work, these ideas are used to introduce parallelism directly in the multigrid framework.

Julien Cortial
Sandia National Laboratories (Livermore)
Quantitative Modeling and Analysis
jcortia@sandia.gov

Paul Boggs
Sandia National Lab
ptboggs@sandia.gov

Stephen G. Nash
George Mason University
Systems Engineering & Operations Research Dept.
snash@gmu.edu

Kevin Long
Texas Tech University
kevin.long@ttu.edu

David Gay
AMPL Optimization, Inc.
dmg@acm.org

R. Michael Lewis
College of William & Mary
rmlewi@wm.edu

MS114

Network Heuristics for Initial Guesses to Nanoporous Flow Optimization Problems

Designing channel structures to optimize suitably constrained flow in nanoporous materials leads to a large-scale, hierarchical, PDE-constrained optimization problem that invites solution by multigrid optimization. The problem has many local solutions, and starting guesses strongly affect the solutions found by a descent algorithm. We sketch the problem and describe a greedy heuristic, possibly improved by a Metropolis-Hastings algorithm, for choosing starting guesses. We show some test results on regular and Delaunay triangulations.

David Gay
AMPL Optimization, Inc.
dmg@acm.org

Paul T. Boggs
Sandia National Labs
ptboggs@sandia.gov

Robert H. Nilson
Sandia National Labs (retired)
rhnilson@yahoo.com

MS114

Preliminary Design of Nanoporous Materials Via Semidefinite Programming

We discuss the preliminary design of nanoporous materials to obtain initial designs for detailed optimization. The goal is a material that has specified storage capacity (say, for electric charge) while allowing sufficiently rapid discharge. Assuming transport is diffusive, standard energy estimates allow us to pose the problem as a semidefinite program. We discuss the formulation of the problem, its relation to the fastest mixing Markov chain problem, and present numerical results.

Robert M. Lewis
College of William and Mary
rmlewi@wm.edu

David Phillips
United States Naval Academy
dphillip@usna.edu

MS114

Software for Automating PDE-constrained Optimization

Sundance is a package in the Trilinos suite designed to provide high-level components for the development of high-performance PDE simulators with built-in capabilities for

PDE-constrained optimization. We review the implications of PDE-constrained optimization on simulator design requirements, then survey the architecture of the Sundance problem specification components. These components allow immediate extension of a forward simulator for use in an optimization context.

Kevin Long
Texas Tech University
kevin.long@ttu.edu

Paul Boggs
Sandia National Lab
ptboggs@sandia.gov

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS115
"Wide or Tall" and "Sparse Matrix Dense Matrix" Multiplications

This talk explores sparse matrix dense matrix (SMDM) multiplications, useful in block Krylov or block Lanczos methods. SMDM computations are AU, VA, with A sparse, U with a few columns or V with a few rows. Another interesting operation is a combined AU, VA, AUV operation. In a block Lanczos or Krylov algorithm, matrix matrix multiplications with the "tall" U and "wide" V are also needed, as demonstrated here in computation of singular values of sparse matrices. All of these operations can run significantly faster than BLAS-1 (vector vector) or BLAS-2 (matrix vector) operations, but are not always well implemented in vendor BLAS. Tuning on multi-core and NUMA architectures is discussed.

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

Masha Sosonkina
Ames Laboratory/DOE
Iowa State University
masha@scl.ameslab.gov

MS115
BLAS Specification Revisited

Certain Blas were not included in the 2002 updated Blas because their significance was not appreciated. Examples include the simultaneous application of different Givens transformations to independent rows and columns and parallel matrix- matrix multiplication with different matrices. For the indefinite symmetric update operation, the 2002 proposal suggested the rarely implemented $A = A+XJX^T$ where J was a symmetric tridiagonal matrix. A simpler update that would cover this case would be updating the triangular portion of $A = A + XY^T$, as proposed by John Lewis.

Linda Kaufman
Livingston
lkaufmang@comcast.net

MS115
Communication-Avoiding Oblique QR Factorizations

We study parallel QR factorizations where Q has orthogonal columns with respect to an oblique inner product. We present stable, communication-avoiding algorithms that cast most of their computations to Level-3 BLAS routines. We also provide performance comparisons with communication-intensive algorithms. For sparse inner product, communication-avoiding algorithms show an increase in performance while performing more FLOPS than communication-intensive algorithms.

Bradley Lowery
University of Colorado, Denver
bradley.lowery@ucdenver.edu

Julien Langou
University of Colorado Denver
julien.langou@ucdenver.edu

MS115
Search Strategies for Empirical Autotuning in Linear Algebra

We present a compiler that translates a linear algebra specification into a parallel and cache-efficient implementation. The compiler achieves high performance through autotuning, trying many code forms to find the fastest code for the architecture. The ability to specify linear algebra operations that are larger than those in the BLAS gives us greater opportunities for optimization. This talk focuses on the challenges of empirical search in autotuning and compares BTO's results with optimized BLAS.

Thomas Nelson
University of Colorado at Boulder
Argonne National Laboratory
thomas.nelson@colorado.edu

Geoffrey Belter
Dept. of Electrical, Computer, and Energy Engineering
University of Colorado at Boulder
geoffrey.belter@colorado.edu

Elizabeth Jessup
University of Colorado at Boulder
Department of Computer Science
jessup@cs.colorado.edu

Boyana Norris
Argonne National Laboratory
norris@mcs.anl.gov

Jeremy Siek
Department of Electrical and Computer Engineering
University of Colorado at Boulder
jeremy.siek@colorado.edu

MS116
Automating DEIM using Automatic Differentiation

The Discrete Empirical Interpolation method(DEIM) uses snapshots to interpolate a nonlinear vector valued function by knowing only a few selected components of the function. For DEIM to be practical, the user must be able to evaluate these selected components without having to evaluate all of the components of the function. We propose using the operator overloading approach of automatic differentiation to generate these subfunctions. We demonstrate our approach on an example involving Miscible Flow.

Russell Carden
Computational and Applied Mathematics
Rice University
russell.l.carden@rice.edu

Danny C. Sorensen
Rice University
sorensen@rice.edu

MS116

Reduced Basis Methods for Nonlinear Diffusion Problems

We present reduced basis approximations and associated *a posteriori* error estimation procedures for quadratically nonlinear diffusion equations. We develop an efficient computational procedure for the evaluation of the approximation and bound. The method is thus ideally suited for many-query or real-time applications. Numerical results are presented to confirm the rigor, sharpness and fast convergence of our approach.

Martin Grepl, Mohammad Rasty
RWTH Aachen University
grepl@igpm.rwth-aachen.de, rasty@igpm.rwth-aachen.de

MS116

Combining a Data-driven Loewner Approach with $\mathcal{H}_\in$ Optimal Interpolation

We present an idea to combine $\mathcal{H}_\in$ optimal model order reduction with a data-driven approach based on Loewner matrices. This method can be extended to parametric model order reduction by using radial basis function interpolation for specially defined functions over the parameter domain. We give an estimate on the $\mathcal{H}_\in$ error of this method.

Sara Grundel
MPI Magdeburg
grundel@mpi-magdeburg.mpg.de

Nils Hornung
Fraunhofer SCAI Bonn
nils.hornung@scai.fraunhofer.de

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

MS116

Model Reduction using Snapshot based Realizations

Abstract not available at time of publication.

Dick Luchtenburg
Princeton University
dluchten@princeton.edu

MS117

A Multiresolution Adjoint Sensitivity Analysis of Time-lapse Seismic Data

We present a method to speed up the computation of adjoint sensitivities during the estimation of reservoir parameters from time-lapse seismic data. The seismic data, in form of change in saturation over time, is first transformed into wavelets and a small subset of the wavelets is retained for parameter estimation. The method reveals how to minimize computational time by modifying the standard adjoint equation to compute directly the adjoint sensitivities of only the wavelets retained.

Abeeb Awotunde
King Fahd University of Petroleum & Minerals
awotunde@stanfordalumni.org

MS117

Quantifying the Effect of Observations in 4D-Var Data Assimilation

Data assimilation combines information from an imperfect numerical model, noisy observations and error statistics to produce a best estimate of the state of a physical system. The contribution of individual observations in reducing the error can provide useful insight for pruning redundant measurements and designing sensor networks. The following presentation describes a novel systematic approach to accomplish this in the context of 4D-Var data assimilation, with emphasis put on the computational challenges.

Alexandru Cioaca
Virginia Tech
alexgc@vt.edu

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

MS117

Conjugate Unscented Transform based Approach for Uncertainty Characterization and Data Assimilation

This talk will focus on recent development of mathematical and algorithmic fundamentals for uncertainty characterization, forecasting, and data assimilation for nonlinear systems. Recently developed Conjugate Unscented Transformation (CUT) methodology will be presented to compute multi-dimensional expectation integrals efficiently. CUT methodology offers the systematic development of non-product cubature rules for accurately evaluating multi dimensional expectation integrals with respect to a symmetric pdf. By accurately characterizing the uncertainty associated with both process and measurement models, this work offers systematic design of low-complexity data assimilation algorithms with significant improvement in nominal performance and computational effort. The applicability and feasibility of these new ideas will be demonstrated on benchmark problems and some real world problems such as tracking resident space objects and forecasting toxic plume.

Nagavenkat Adurthi
MAE Deaprtment
University at Buffalo
n.adurthi@gmail.com

Puneet Singla

Mechanical & Aerospace Engineering
University at Buffalo
p.singla@gmail.com

Tarunraj Singh
Dept of Mechanical Engineering
SUNY at Buffalo
tsingh@buffalo.edu

Abani K. Patra
SUNY at Buffalo
Dept of Mechanical Engineering
abani@eng.buffalo.edu

MS117

Variational Data Assimilation of Chaotic Dynamical Systems over Climate Timescales

In this talk, we first demonstrate difficulties encountered by conventional data assimilation methods, when the data is from observation of a chaotic dynamical system over climate timescales. We then introduce a new method, the Least Squares Sensitivity Analysis method, that overcomes the difficulties of climate data assimilation. The Least Squares Sensitivity method formulates an adjoint problem that aims at computing derivatives of statistics of a chaotic dynamical system, which is used in our variational data assimilation. Our new method is demonstrated on the Lorenz 63 system.

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

MS118

Statistical Perspectives

Abstract not available at time of publication.

George Casella
Department of Statistics
University of Florida
casella@stat.ufl.edu

MS118

Spectral/hp Element and Discontinuous Galerkin Methods for Response-Excitation PDF 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 driven by colored random noise. This talk presents an efficient numerical method for this evolution equation of REPDF by considering the response and excitation spaces separately. For the response space, a non-conforming adaptive discontinuous Galerkin method is used to resolve both local and discontinuous dynamics while a probabilistic collocation method is used for the excitation space. We propose two fundamentally different adaptive schemes for the response space using either the local variance combined with the boundary flux difference or using particle trajectories. The effectiveness of the proposed new algorithm is demonstrated in two prototype applications dealing with randomly forced nonlinear oscillators. The resulted PDF is compared against the one obtained from kernel density estimation based on Monte Carlo simulation and the solution of the effective Fokker-Planck equation. The framework we develop here is general and can be readily extended to first order stochastic PDEs subject to random boundary conditions, random initial conditions, or random forcing terms.

Heyrim Cho
Brown University
Providence, RI
heyrim_cho@brown.edu

Daniele Venturi
Brown University
daniele_venturi@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS118

UQ and High-Performance Computing

Abstract not available at time of publication.

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

MS118

Hierarchical Preconditioners for the Stochastic Galerkin Finite Element Methods

Use of the stochastic Galerkin finite element methods leads to large systems of linear equations. These systems are typically solved iteratively. We propose a family of preconditioners that take an advantage of (the recursion in) the hierarchy of the global system matrices. Neither the global matrix, nor the preconditioner need to be formed explicitly. The ingredients include only the number of stiffness matrices from the truncated Karhunen-Loève expansion and a preconditioner for the mean-value problem. The performance is illustrated by numerical experiments.

Bedrich Sousedik, Roger G. Ghanem
University of Southern California
sousedik@usc.edu, ghanem@usc.edu

ERIC T. Phipps
Sandia National Laboratories
etphipp@sandia.gov

MS119

Error Estimate for Nonlinear Model Reduction using Discrete Empirical Interpolations

This work provides an error analysis of parametric nonlinear model reduction using Discrete Empirical Interpolation Method together with standard projection-based model reduction methods, such as Proper Orthogonal Decomposition. The analysis will be given in the setting of ODE systems arising from spatial discretizations of parabolic PDEs and in the setting of discretized steady-state problems. The conditions under which the stability of the original system is preserved and the reduction error is uniformly bounded will be discussed. Some numerical examples will be used to illustrate the applicability of this error analysis.

Saifon Chaturantabut

Virginia Tech
saifonc@math.vt.edu

Danny C. Sorensen
Rice University
sorensen@rice.edu

MS119

Error Estimation for Nonlinear Reduced Basis Methods based on Empirical Operator Interpolation

Many applications based on numerical simulations of PDEs depend on expensive computations and therefore use low-dimensional surrogate models. Here, it is important to control the error of the surrogate model. If the surrogate is obtained with the reduced basis method, efficient operator evaluations are required for nonlinear problems. These can be obtained through the empirical operator interpolation method. We present error estimators for both the interpolated operators and reduced basis models for parametrized nonlinear evolution equations.

Martin Drohmann
University of Münster
Institute for Computational and Applied Mathematics
martin.drohmann@uni-muenster.de

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

Mario Ohlberger
Universität Münster
Institut für Numerische und Angewandte Mathematik
mario.ohlberger@uni-muenster.de

MS119

Certified Reduced Basis Approximation for Component-Based Problems

We discuss the SCRBE Method [Huynh et.al., A Static Condensation Reduced Basis Element Method, M2AN, accepted 2012] for large component-based (e.g. 3D structural) problems. We combine RB model reduction for the high-fidelity component-local bubble spaces with a "port reduction" procedure [Eftang et.al., Adaptive Port Reduction in Static Condensation, Proceedings of MATHMOD 2012] to reduce the Schur system size.

The additional port reduction procedure is necessary in particular in 3D when the number of port degrees of freedom is large. The errors introduced by both the RB approximations and port reduction procedure can be rigorously bounded with respect to an underlying FE approximation.

Jens Eftang
MIT
Department of Mechanical Engineering
eftang@mit.edu

MS119

Error Estimates for Some Galerkin Reduced-order Models of the Semi-discrete Wave Equation

Reduced-order models for first-order dynamical systems have been extensively analyzed. Even though the semi-discrete wave equation can be formulated as a first-order dynamical system, model-order reduction may destroy the original structure of the wave equation. This talk will present an error analysis for Galerkin-based reduced systems preserving the structure of the wave equation. Error bounds are derived in the continuous setting and when the Newmark scheme is used. Numerical experiments illustrate the theoretical results.

David Amsallem
Stanford University
amsallem@stanford.edu

Ulrich Hetmaniuk
University of Washington
Department of Applied Mathematics
hetmaniu@uw.edu

MS120

The Effective Combination of Mesh Adaptation and Non-linear Thermo-mechanical Solution Components for the Modeling of Weld Failures

Computational studies of the transient failure of laser welds require maintaining a high quality mesh for use in a fully coupled thermo-mechanical, finite deformation, simulation of the necking and subsequent unloading of the specimen. Large deformations encountered during necking necessitate methods to locally adapt the mesh. The method used to effectively combine FASTMaths MeshAdapt and Sandia's Albany thermo-mechanical solution framework as well as its application in the parallel simulation of weld failures will be presented.

Glen Hansen
Sandia National Laboratories
gahanse@sandia.gov

Ryan Molecke
Rensselaer Polytechnic Institute
molecr@rpi.edu

James Foulk III
Sandia National Laboratories
jwfoulk@sandia.gov

Mark Shephard
Rensselaer Polytechnic Institute
shephard@rpi.edu

MS120

Parallel Interface Preservation in Mesquite for Bubble-Shock Interaction Problems

We present a parallel material interface preservation scheme developed in the context of the Mesquite mesh quality improvement toolkit. This scheme allows the movement and tracking of interior and exterior boundaries and significantly improves the robustness of an ALE shock/bubble interaction test problem over standard techniques.

Brian Miller
Lawrence Livermore National Laboratory
miller125@llnl.gov

MS120

Advances of an Unstructured Mesh Infrastructure

to Support Massively Parallel Adaptive Simulations on High Core Count Machines

The flexible distributed Mesh Database (FMDB) is FASTMath iMesh/iMeshP compliant unstructured mesh infrastructure to represent and manipulate mesh data for parallel adaptive simulations. FMDB has been extended to fully support multiple parts per process, faster migration, and full interaction with ZOLTAN partitioning. Recent FMDB developments include making it architecture-aware and structural changes to the partition model to support a partition taking advantage of node level shared memory and threading, while maintaining message passing between nodes.

Seegyoung Seol
Rensselaer Polytechnic Institute
Scientific Computation Research Center
seol@scorec.rpi.edu

Daniel Ibanez
RPI
dan.a.ibanez@gmail.com

Mark Shephard
Rensselaer Polytechnic Institute
shephard@rpi.edu

MS120

Interoperable Solution Transfer Tool for Coupled Multi-physics Simulations

I will describe a MOAB-based tool for solution transfer that simplifies connection of physics codes and supports parallel coupled multi-physics simulation; examples will include problems in Nuclear Energy (coupling Nek5000 and UNIC) and reservoir simulation (coupled reservoir & seismic codes)

Timothy J. Tautges
Argonne National Labory
tautges@mcs.anl.gov

MS121

Accurate Solutions to Scattering Problems in Nonsmooth Planar Domains

Recursively Compressed Inverse Preconditioning (RCIP) is a method for obtaining accurate solutions to integral equations stemming from elliptic PDEs in piecewise smooth domains. Tractable features include: corners, cusps, multiple junctions, small edge separation, mixed boundary conditions, barely integrable solutions, and composed integral operators. In a postprocessor, the solution to the PDE is recovered in the entire domain using high-order analytic product rules close to the boundary. A range of numerical examples is presented.

Johan Helsing
Centre for Mathematical Sciences
Lund University
helsing@maths.lth.se

MS121

Fast Direct Solvers for the Lippmann-Schwinger Equation

Recent years have seen the development of a number of fast direct solvers for integral equations that are extremely effective for systems involving many right-hand sides. Here, we adapt a recursive skeletonization-based solver to the Lippmann-Schwinger equation for high-frequency wave scattering in 2D. For a domain discretized with N elements, the algorithm incurs an $\mathcal{O}(\mathcal{N}^{\ni/\in})$ precomputation cost, after which each incident wave can be analyzed in only $\mathcal{O}(\mathcal{N}\log\mathcal{N})$ operations.

Kenneth L. Ho
Courant Institute of Mathematical Sciences
New York University
ho@courant.nyu.edu

MS121

Robust Charge-current Formulations for Electromagnetic Scattering

In the electromagnetics literature, significant attention has been paid to the integral equation formulation of scattering problems. Certain problems as low frequency breakdown, high-density mesh breakdown, resonances, and catastrophic cancellation arise when dealing with the most common integral equations as MFIE, EFIE, CFIE... In this paper we present some new formulations that avoid most of the problems above mentioned. The formulations proposed are based on second kind integral equations and current and charge representation for the sources.

Felipe Vico
Departamento de Comunicaciones
Universidad Politecnica de Valencia
felipe.vico@gmail.com

Leslie Greengard
Courant Institute of Mathematical Sciences
New York University
greengar@cims.nyu.edu

Zydrunas Gimbutas
Courant Institute
New York University
zydrunas.gimbutas@gmail.com

Miguel Ferrando Bataller
Departamento de Comunicaciones
Universidad Politecnica de Valencia
mferrand@dcom.upv.es

MS121

High Frequency FMMs and Improved WKB Approximations

WKB approximation is ubiquitous in several areas of science and engineering. It approximates solutions of second order ODEs by a simple formula; a major nuisance is that it is only asymptotic, as opposed to an exact expression.It turns out that such exact expressions can be constructed via the classical Kummer's equation. Among other things, it leads to approximations of special functions in the Fresnel regime, with obvious applications in fast multipole methods, etc.

Bogdan G. Vioreanu
Applied Mathematics
Yale University
bgv@umich.edu

Vladimir Rokhlin
Yale University
rokhlin@cs.yale.edu

MS122

Full Waveform Inversion for Diffraction Focusing

Seismic diffractions are an important part of seismic reflection data. Diffractions represent the direct response of small objects that are often the subject of geophysical exploration: faults, cavities, channels, etc. By separating diffractions from specular reflections, it is possible to formulate the seismic inverse problem as the problem of focusing diffraction energy in the imaging domain. We demonstrate the connection between this formulation and the conventional full waveform inversion and highlight differences and similarities.

Sergey Fomel, Siwei Li
University of Texas at Austin
sergey.fomel@beg.utexas.edu, siwei.li@utexas.edu

MS122

Title Not Available at Time of Publication

Abstract not available at time of publication.

Loukas F. Kallivokas
The University of Texas at Austin
loukas@mail.utexas.edu

MS122

On Parameterization of Full Waveform Inversion in Exploration Geophysics

Over the last five years, acoustic full waveform inversion has been applied to large surface seismic data sets to image the first kilometers of the earth crust. The problem is however severely ill-posed, non-linear and expensive. Certain parameterizations, that may be case dependent, should help to reduce the number of iterations and some artifacts. During this talk, we will discuss some possible parameterizations or pre-conditioning and their challenges and show some examples.

Rene-Edouard Plessix
Shell Global Solutions International
reneedouard.plessix@shell.com

MS122

Relaxation of Nonlinear Full Waveform Inversion via Extended Modeling

Least squares data fitting has become known as "full waveform inversion" in computational seismology, and is the subject of very active research in industry and academia over the last couple of years. The objective function of full waveform inversion tends to be highly multimodal, especially for data rich in reflected wave energy. Reformulation via extended modeling and soft physicality constraints has proven able to convexify this difficult optimization problem, at least to some extent: for example, the Marmousi benchmark problem has recently yielded to extended inversion. This presentation will review the theoretical basis for extended inversion, and outline some of the remaining computational challenges.

William Symes
Rice University
symes@caam.rice.edu

MS123

Competitive Product Differentiation in a Three Level Market

Semiconductor chips are not sold directly to the end customer but to Original Equipment Manufacturers (OEM) who assemble a computer and sell to the end market. The OEM acts as a filter for the customer demand. Selecting the optimal number and the optimal spacing between products in price and performance is known as the vertical differentiation problem. The interaction between this optimization problem and the game theoretic competition in this three level market is analyzed.

Dieter Armbruster
Arizona State University
Department of Mathematics
armbruster@asu.edu

Tulin Inkaya
Arizona State University
tinkaya@asu.edu

Karl Kempf
Intel Corporation
karl.g.kempf@intel.com

Hongmin Li
Arizona State University
hongmin.li@asu.edu

Morgan Dempsey
Intel Corporation
morgan.dempsey@intel.com

MS123

Interactions between Product Development and Manufacturing Planning within INTEL

Abstract not available at time of publication.

Karl Kempf
Intel Corporation
karl.g.kempf@intel.com

MS123

Process and Production Optimization for Industrial Applications

ABB Corporate Research tackles real world problems arising in industry applying methods from mathematical optimization and optimal control. In this talk, we will present two research projects: Energy-aware enterprise-wide production scheduling in the metals industry and a solution for plant-wide asset-performance management. The underlying problems can be solved by hybrid methods combining mathematical optimization with production intelligence and using asset performance information (diagnosis) to generate operations and maintenance actions (therapy) in steel and paper mills.

Sleman Saliba
ABB AG Corporate Research Center
sleman.saliba@de.abb.com

Iiro Harjunkoski
ABB AG Corporate Research Center
iiro.harjunkoski@de.abb.com

MS123

Production Planning with Nonlinear Clearing Functions: A Review of Recent Results

Nonlinear clearing functions that represent the relationship between the expected throughput of a planning resource in a planning period and the workload of the resource have shown promising results in production planning. We review basic concepts and computational results, focusing on a study of a large semiconductor wafer fabrication facility.

Reha Uzsoy
Purdue
uzsoy@ecn.purdue.edu

Lars Moench
Fernuniversitat Hagen
lars.moench@fernuni-hagen.de

N. Baris Kacar
SAS Institute
baris.kacar@sas.com

MS124

Large Eddy Simulation of Pathological and Medical Device Hemodynamics using a Novel Multiblock Immersed Boundary Method Approach

Large eddy simulations of several flows involving pathological or medical device hemodynamics will be presented. A high-order finite-difference method is used to integrate the incompressible Navier-Stokes equations on a structured Cartesian grid. Complex geometries are handled using a novel multiblock immersed boundary method. Results are compared to experimental measurements to assess accuracy and gain insight into these flowfields.

K Anupindi, Y Delorme, D Shetty, Steven H. Frankel
School of Mechanical Engineering
Purdue University
kanupind@purdue.edu, ydelorme@purdue.edu, dshetty@purdue.edu, frankel@purdue.edu

MS124

Multiscale Modeling and Optimization in Single Ventricle Heart Disease

Single ventricle patients typically undergo three palliative surgical procedures before the age of three. We will present a framework for multiscale blood flow simulations and surgery optimization. We will discuss challenges of modeling the coupled system of blood flow and circulatory dynamics that arises in these patients, and how simulations may impact their clinical care. We present an efficient and modular method to couple our custom finite element solver to a closed loop circulatory model.

Alison Marsden
Department of Mechanical and Aerospace Engineering
University of California, San Diego
amarsden@ucsd.edu

MS124

Modeling of Blood Flow in Normal and Diseased Left-ventricles

The current research focuses on the modeling and analysis of flow in normal and diseased left-ventricles; disease conditions include myocardial infarction, diastolic dysfunction, and hypertrophic obstructive cardiomyopathy. Anatomical models of the left ventricles are derived from high resolution CT as well as echo images, and the simulations employ an immersed boundary flow solver. The effect of flow patterns on the performance of the left ventricle is explored in detail.

Vijay Vedula, Jung Hee Seo
Mechanical Engineering
Johns Hopkins University
vvsn.vijay@gmail.com, junghseo@gmail.com

Albert Lardo, Theodore Abraham
Medical School
Johns Hopkins University
al@jhmi.edu, tabraha3@jhmi.edu

Rajat Mittal
Department of Mechanical Engineering
The Johns Hopkins University
mittal@jhu.edu

MS124

Numerical Study of Pulse Wave Propagation Patterns in Stiffened Arteries

A 2-D axis-symmetric model of aortic pulse-wave-traveling is simulated using our modified Immersed Finite Element Method (mIFEM), a numerical method that accurately simulates the dynamics of fully-coupled fluid-structure interactions. This model is to validate pulse wave imaging measurement, a noninvasive and quantitative way of measuring arterial stiffness. In our model, pulse wave velocity (PWV) is measured by calculating the spatiotemporal variation of the pulse wave-induced displacement of the arterial wall. Boundary conditions and material properties are to follow the experiments. The results are in good agreements with PWV measured from aortic phantoms.

Lucy Zhang, Jubiao Yang
Department of Mechanical, Aerospace and Nuclear Engineering
Rensselaer Polytechnic Institute
zhanglucy@rpi.edu, yangj9@rpi.edu

Danial Shahmirzadi
Department of Biomedical Engineering
Columbia University
ds3031@columbia.edu

Elisa Konofagou
Columbia University
ek2191@columbia.edu

MS125

Data Analysis of Cascading Power System Failures

We present analyses of large-scale simulations of extreme events in power grids, in particular (a) cascading failures and (b) random fluctuations of renewable power injections. In the case of cascading failures we use models that capture some of the short-term grid dynamics (including voltage and frequency deviations) in approximate form, as well as generator ramp-up behavior and frequency response (broadly interpreted). Several interesting phenomena can be observed, in particular dominant modes of cascading, and observable features in the early stages of the cascade. In the case of fluctuating renewable injections we consider various conditional distributions of wind power (conditional on short-term forecasts) and analyze resulting PDFs of line overloads.

Daniel Bienstock
Columbia University IEOR and APAM Departments
IEOR Department
dano@columbia.edu

MS125

Scalable Stochastic Optimization for Power Grid Systems

Stochastic optimization of complex energy systems results in extremely large optimization problems that can be solved only by means of high-performance computers. We present the latest algorithmic developments and implementations for the scalable solution of continuous and integer stochastic programming problems with recourse. We also discuss the numerical results obtained on Argonne's "Intrepid" BG/P platform.

Cosmin G. Petra, Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
petra@mcs.anl.gov, anitescu@mcs.anl.gov

MS125

Economic Dispatch with Renewable Power Generation

A two-stage non-linear stochastic formulation for the economic dispatch problem under renewable-generation uncertainty is investigated. Certain generation decisions are made only in the first stage and fixed for the second stage, where the actual renewable generation is realized. The uncertainty in renewable output is captured by a finite number of scenarios. Any resulting supply-demand mis-match must then be alleviated using high marginal-cost power sources. We present two decomposition algorithms to solve this problem to optimality.

Dzung Phan
IBM T.J. Watson Research Center
phandu@us.ibm.com

MS125

Scalable, Parallel Stochastic Unit Commitment For Reliability Operations

Modern grid reliability systems include a security-constrained unit commitment (SCUC) optimization algorithm operating in tandem with a security-constrained economic dispatch (SCED) optimization algorithm. Independent system operators (ISOs) use these models to commit and dispatch generation resources. Increasing penetration of intermittent renewable generation necessitates a shift to stochastic SCUC/ SCED algorithms. We will describe the stochastic programming SCUC/SCED formulations and algorithms we are currently developing under an effort recently funded by the ARPA-e GENI program.

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

MS126

Efficient Symmetric Positive Definite Second-order Accurate Monolithic Solver for Fluid/solid Interactions

We introduce a second-order accurate method to simulate strongly coupled (monolithic) fluid/rigid-body interactions. We take a fractional step approach, where the intermediate state variables of the fluid and of the solid are solved independently, before their interactions are enforced via a projection step. The projection step produces a symmetric positive definite linear system that can be efficiently solved using the preconditioned conjugate gradient method. Numerical results indicate that the method is second-order accurate in the maximum norm and demonstrate that its solutions agree quantitatively with experimental results.

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

Chohong Min
Ewha Womans University, Korea
chohong@ewha.ac.kr

MS126

An Adaptive Multi-material Moment-of-fluid Method for Computing Multi-phase Flows

We combine the multimaterial Moment-of-Fluid (MOF) work of Ahn and Shashkov with the work of Kwatra et al for removing the acoustic time step restriction in order to solve multimaterial flows in which each material might be compressible or incompressible. The mass weights found in the algorithm of Kwatra et al are computed directly from the multimaterial MOF reconstructed interface. Simulations for multimaterial flows are presented with applications to combustion (atomization and spray) and microfluidics.

Mark Sussman
Department of Mathematics
Florida State University
sussman@math.fsu.edu

Matt Jemison, Michael Roper
Florida State University
mjemison@math.fsu.edu, roper@chem.fsu.edu

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

Marco Arienti
Sandia National Labs, Livermore Calif.
marient@sandia.gov

Xiaoyi Li
UTRC

lixy2@utrc.utc.com

MS126

A Sharp Level-set Approach for the Dendritic Growth

We developped a sharp level-set approach for the dendritic growth in two spatial dimensions, where the inherent multi-scale nature of the physical problem is handle by the use of adaptive quadtree grids. The obtained numerical results show an excellent match with experimental observations and theoritical predictions. In particular the dendrites morphology (primary and secondary spacing) and the transition between the different regime (palanar-cellular-dendritic) were accurately captured.

Maxime Theillard
Dept. Mechanical Engineering
UCSB
maxime_theillard@umail.ucsb.edu

MS126

Computations of Mass Transfer in Bubbly Flows using a Hybrid Finite Volume Method and an Embedded Analytical Description

DNS of mass transfer from hundreds of bubbles in turbulent flows introduces resolution requirements that are currently difficult to match. Here we introduce an embedded analytical description for mass transfer in the thin boundary layer at bubble surfaces and couple it with a finite volume method for the rest of the flow. Comparisons with fully resolved simulations of simple systems and experiments shows good agreement, for relatively modest increase in the computational effort.

Gretar Tryggvason
University of Notre Dame
Department of Aerospace and Mechanical Engineering
gretar.tryggvason.1@nd.edu

Bahman Aboulhasanzadeh
Department of Aerospace and Mechanical Engineering
University of Notre Dame
baboulha@nd.edu

Jiacai Lu
Mechanical Engineering Department
Worcester Polytechnic Institute
jlu@wpi.edu

MS127

A High-order Immersed Finite Element Method for PDEs with Discontinuous Coefficients

We present an immersed finite element for interface problem with discontinuous coefficients. Here we allow the interface to cut elements and construct immersed finite element shape functions that satisfy jump conditions across the interface. The immersed finite element spaces are used with a penalized weak formulation. We discuss the existence and approximation properties of these finite element spaces and solve several interface problems to show the performance of the proposed method.

Slimane Adjerid
Department of Mathematics
Virginia Tech
adjerids@math.vt.edu

Mohamed Benromdhane
KAUST
mohamed.benromdhane@kaust.edu.sa

Tao Lin
Virginia Tech
tlin@vt.edu

MS127

Low-complexity Application of Finite Element Operators on Simplices via Bernstein Polynomials

Evaluation and differentiation of Bernstein-form polynomials on the simplex may be accomplished with low complexity, and so may computation of integral moments of data against the Bernstein basis. I will demonstrate these features via a recursive decomposition of an appropriate generalized Vandermonde-type matrix. H(div) and H(curl) are similarly handled by first converting the exterior calculus bases of Arnold, Falk, and Winther to the Bernstein basis efficiently and then using the H1 techniques. I will sketch these techniques, and also comment on shared memory parallelism of the resulting algorithms.

Robert C. Kirby
Texas Tech University
Robert_Kirby@baylor.edu

MS127

Error Analysis of the Discrete Wave Equation Based on a Full-Modal Decomposition: A Unifying Perspective for Higher-order Methods

Abstract not available at time of publication.

Pedro Moy
King Abdullah University of Science and Technology
pedro.moy@kaust.edu.sa

Victor M. Calo
Applied Mathematics and Computational Science
King Abdullah University of Science and Technology
victor.calo@kaust.edu.sa

MS127

Discontinuous Galerkin Method for Hyperbolic Equations Involving δ-functions

In this talk, we apply discontinuous Galerkin (DG) methods to solve hyperbolic equations involving δ-functions. In general, the numerical solutions are highly oscillatory near the singularities, which we refer to as the pollution region. We first analyze the size of the pollution region and the rate of convergence outside for some model equations, then apply the method to pressureless Euler equations to show the good performance of the DG method.

Yang Yang
Brown university
yang_yang@brown.edu

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

MS128

A Practically Painless Path to Petascale Parallelism: PETSc + Python

I describe the design of numerical software that is operated with the convenience of MATLAB yet achieves efficiency near that of hand-coded Fortran and scales to the largest supercomputers. Python is used for most of the code while automatically-wrapped Fortran kernels are used for computationally intensive routines. For parallelism we use PETSc via petsc4py. The software described here is PyClaw, a structured grid solver for systems of hyperbolic PDEs based on Clawpack.

David I. Ketcheson
Mathematical and Computer Sciences & Engineering
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

MS128

Constructing and Configuring Nested and Hierarchical Solvers in PETSc

We look at code architecture for constructing and also configuring deeply nested solver hierarchies. We are most concerned with

- how much code a user has to write
- how frequently code has to change
- how an implementation chooses among options

and the impact tradeoffs here have on the interface design. Our examples will be drawn from multiphysics PDE simulation, and exhibit block solvers, algebraic and geometric multigrid, and composition of both linear and nonlinear solvers. However, this is not an overview of packages capabilities, but rather a discussion about implementation strategy and user experience.

Matthew G. Knepley
University of Chicago
knepley@ci.uchicago.edu

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory
jedbrown@mcs.anl.gov

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

MS128

The FEniCS Project: Organization, Practices, Maintenance and Distribution

Abstract not available at time of publication.

Anders Logg
Simula Research Laboratory
logg@simula.no

MS128

Joining Forces: Combining FEniCS and DUNE using a High-level Form Language and Code Generation

Abstract not available at time of publication.

Steffen Müthing
Institute of Parallel and Distributed Systems
University of Stuttgart
steffen.muething@ipvs.uni-stuttgart.de

MS129

Building An Applied and Computational Math Degree Program from the Ground Up

Abstract not available at time of publication.

Jeffrey Humpherys
Brigham Young University
jeffh@math.byu.edu

MS129

Interdisciplinary Undergraduate Research in Computational Sciences as a Model for Institutional Change

Abstract not available at time of publication.

Padmanabhan Seshaiyer
George Mason University
pseshaiy@gmu.edu

MS129

Enabling Excel-lent Experiences: Computational Science Internships

Professors can help undergraduates excel through obtaining meaningful CSE research internships. Such internships can enhance students' professional and personal lives and can add new dimensions to a school's program. With students' inexperience, active involvement by professors is usually a crucial. Using numerous "success stories," this talk explores how advisors can help students obtain and benefit from computational research experiences.

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

MS129

Modeling Across the Curriculum

The August, 2012 SIAM-NSF Workshop on Modeling across the Curriculum will be described, and an introduction to the report presented. The meeting proposal predated but responded to the PCAST "Engage to Excel" report's call for one million new STEM graduates by 2020. Three themes were addressed: coordinated STEM curriculum content in K-12; undergraduate curricula in modeling and CSE as the heart of STEM; and readiness for college STEM education. The last theme is a basis for assessment of programs and how they address the "math gap."

Peter R. Turner
Clarkson University
School of Arts & Sciences
pturner@clarkson.edu

MS130

High-Order Multi-Stage Lax-Wendroff Time Integrators with Applications for Cahn-Hilliard

In this work, we present a novel class of high-order numerical integrators. These integrators utilize a combination of Lax-Wendroff time-steps and Hermite interpolation. Given their low-storage requirements and high-order accuracy, this class of methods proves to be promising for high performance computing. We explore applications of explicit and implicit formulations on ordinary and partial differential equations. The implicit scheme is a low-storage, 4th order, A-stable method, and the 4th order explicit formulation serves as an accurate error predictor for adaptive time stepping on stiff PDEs. Results for the Cahn-Hilliard equations are investigated. The method is extends to 6th order and higher.

Andrew Christlieb
Michigan State University
christlieb@math.msu.edu

David C. Seal
Department of Mathematics
Michigan State University
seal@math.msu.edu

Jaylan S. Jones
Michigan State University
jaylanjones@gmail.com

MS130

High Order Partitioned Time Stepping Methods for Stiff Problems

We discuss recent developments for multistage partitioned time stepping methods for solving ODEs and PDEs. We analyze stability and consistency properties of different operator splitting strategies driven by stiff and nonstiff terms that require integrators with different properties for stability and efficiency reasons. We focus on multistage additive implicit-explicit (IMEX) and Rosenbrock W schemes.

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS130

Conditions for Positivity and SSP Property of Linearly Implicit Methods and Applications to CFD

Positivity preservation and the diminishing of some convex functionals along the solution in time are features of many classes of initial (or initial-boundary) value problems, which ideally are maintained by their numerical discretization. These requirements can be fulfilled only under some conditions on the parameters of the discretization procedure. In this talk we investigate such conditions for linearly implicit methods, in particular when applied to some problems in CFD.

Zoltan Horvath
Széchenyi István University
horvathz@sze.hu

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory
jedbrown@mcs.anl.gov

Matthew G. Knepley
University of Chicago
knepley@ci.uchicago.edu

Tihamér A. Kocsis, Adrián Németh
Széchenyi István University
katihi@sze.hu, nemetha@sze.hu

MS130

A Comparison of High Order Explicit Runge-Kutta, Extrapolation and Integral Deferred Correction Methods

Two popular approaches for solving ODEs with very high accuracy are extrapolation and deferred correction methods. When the base method on which they are built is a one-step method (i.e., a Runge-Kutta method), both methods result in a Runge-Kutta method with very many stages. This connection between the two methods is used to study and investigate stability and accuracy properties of the two ODE solving schemes along with high order embedded Runge-Kutta pairs. These methods are compared in practice by implementing them in an adaptive error control framework and a small set of test problems, including a three-body chaotic problem, are used to derive conclusions.

Umair bin Waheed
King Abdullah University of Science and Technology
umairbin.waheed@kaust.edu.sa

David I. Ketcheson
Mathematical and Computer Sciences & Engineering
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

MS131

Development of Exact-to-Precision Generalized Perturbation Theory in Nuclear Engineering Calculations

Perturbation theory provides the most computationally efficient approach for calculating responses variations resulting from parameters variations without the need to re-execute the model. In this talk, we introduce EpGPT (Exact-to-Precision Generalized Perturbation Theory) intended to address the challenges of existing theory, including ability to calculate higher order variations and bound the variational errors, handle models with many responses.

Hany S. Abdel-Khalik
North Carolina State Univ.
abdelkhalik@ncsu.edu

MS131

Utilizing Adjoints to Improve Propagation of Uncertainties through Surrogate Response Surfaces

We consider the use of surrogate response surfaces, sometimes called emulators, to cheaply propagate probability distributions between an input parameter space and observable quantities of interest. We develop a framework for both forward and inverse propagation and utilize adjoints to improve the pointwise accuracy of the emulators and subsequently of the computed cumulative distribution functions. An error analysis is provided for concrete examples using polynomial chaos expansions to define the surrogates.

Troy Butler
Colorado State University
butler@stat.colostate.edu

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

Tim Wildey
Sandia National Laboratory
tmwilde@sandia.gov

MS131

On an Adjoint Consistent Formulation for a Coupled Problem

Microfluidic devices are usually simulated using coupled models due to the multiscale and multiphysics nature of microflows. In the Helmholtz slip electroosmotic model, coupling between physical models is enforced on the boundary of the domain through a slip boundary condition. However, analysis of this coupled formulation shows that the resulting forward problem may be ill-posed. We propose a new formulation of the coupling term to ensure that both the forward and adjoint problems are well-posed.

Serge Prudhomme
ICES
The University of Texas at Austin
serge@ices.utexas.edu

Vikram Garg
Massachusetts Institute of Technology
vikram.v.garg@gmail.com

Kris van der Zee
Multiscale Engineering Fluid Dynamics
Eindhoven University of Technology
k.g.v.d.zee@tue.nl

MS132

Fast Direct Solvers for H^2 Matrices

Abstract not available at time of publication.

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

Amirhossein Aminfar
Stanford University
aminfar@stanford.edu

MS132

Fast Multipole Method as a Preconditioner

Recent efforts to view the FMM as an elliptic PDE solver have opened the possibility to use it as a preconditioner for even a broader range of applications. Use of FMM as a preconditioner requires optimization for low accuracy, which is a challenge on many-core co-processors and GPUs. This talk addresses various optimization techniques that combine treecodes and FMMs to achieve a highly parallel FMM preconditioner.

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

MS132

A Treecode-accelerated Boundary Integral Poisson-Boltzmann Solver

We present a boundary integral method for electrostatics of solvated biomolecules described by the linear Poisson-Boltzmann equation. The method employs GMRES iteration and a Cartesian treecode which reduces the cost from $O(N^2)$ to $O(N \log N)$, where N is the number of faces in the triangulated molecular surface. We present results for the Kirkwood sphere and a protein.

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

Weihua Geng
University of Alabama
wgeng@as.ua.edu

MS132

Directional FMM for Maxwell's Equations

For boundary element methods in electromagnetics, there are a variety of fast algorithms used to accelerate the matrix-vector product with the impedance matrix. In this talk, we propose a different approach based on the directional multilevel algorithm by Engquist and Ying. Using the directional low-rank property of the Green's function in the high-frequency regime, it is shown that the algorithm achieves O(N log N) complexity for time-harmonic Maxwell problems.

Paul H. Tsuji
University of Texas at Austin
Institute for Computational Engineering & Sciences
ptsuji@gmail.com

Lexing Ying
University of Texas
Department of Mathematics
lexing@math.utexas.edu

MS133

The DSP as a High-performance, General-purpose Processor: First Experiences with FLAME

Digital Signal Processors (DSPs) are specialized architectures with low-power consumption, real-time processing support and sophisticated development tools. In this talk, we report our first experiences with the TMS320C6678 multi-core DSP from Texas Instruments. We introduce the first complete BLAS implementation for this processor. Building upon the BLAS and OpenMP, we use the flexibility of the FLAME methodology and the `libflame` library to port higher-level linear algebra functionality to the DSP parallel architectures.

Murtaza Ali
Texas Instruments
mali@ti.com

Francisco Igual
Univ. Complutense de Madrid
figual@fdi.ucm.es

MS133

Optimization of Blas Kernels: Close-to-the Metal Tricks for Multicore and Manycore Architectures

Novel computer architectures often require non-traditional optimization techniques in order to allow even routines as simple as the traditional BLAS kernels to perform as they should. This talk will focus on techniques for optimizing these kernels as well as higher-level techniques to employ them to greater advantage on large and/or new architectures.

John A. Gunnels
T.J. Watson Research Center
IBM
gunnels@us.ibm.com

MS133

MKL, BLAS, and New Architectures

The Intel(R) Math Kernel Library (MKL) is well known for its high performance for HPC applications across many scientific domains. We show some of the code generation and design behind some key techniques and models for our linear algebra routines and the performance, time, and accuracy benefits behind our tools and techniques. We focus on some of the many-core results from our newest Intel(R) Xeon Phi processors.

Greg Henry
Intel Corporation
greg.henry@intel.com

MS133

BLIS: A New Framework and Interface for the BLAS

We present a modern alternative to the Basic Linear Algebra Subprograms (BLAS) that addresses many shortcomings in the original BLAS. This framework, which we call BLIS, allows an expert to (1) rapidly instantiate BLAS-like libraries for new architectures, and (2) design entirely new operations by leveraging existing components of the framework. Preliminary performance on select architectures is highly competitive with high-performance open source and commercial products.

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

MS134

A new Monte Carlo Method for Velocity-dependent Neutrino and Photon Transport

Abstract not available at time of publication.

Ernazar Abdikamalov
California Institute of Technology
abdik@tapir.caltech.edu

MS134

Discontinuous Galerkin Methods for Vlasov-Maxwell Systems

Abstract not available at time of publication.

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

MS134

Asymptotic Preserving DG Methods for Kinetic Equations

A family of high order asymptotic preserving schemes are proposed for some kinetic equations. The methods are based on the micro-macro decomposition of the problem, and they combine high order discontinuous Galerkin spatial discretizations and second order IMEX temporal discretizations. Both theoretical and numerical studies are carried out to demonstrate the performance of the methods.

Juhi Jang
University of California, Riverside, USA
juhijang@math.ucr.edu

Fengyan Li
Rensselaer Polytechnic Institute
lif@rpi.edu

Jingmei Qiu
University of Houston
jingqiu@math.uh.edu

MS134

Asymptotic Preserving Schemes for Kinetic Equations in the High Field Regime

I will introduce numerical methods for two kinetic equations—the Vlasov -Poisson-Fokker-Planck (VPFP) system and the semiconductor Boltzmann equation, with emphasis on the connections to the high field limit. The two stiff terms under this scaling–the collision term and the force term, make the classical kinetic solver expensive to implement. For the VPFP system, the idea is to combine the two stiff terms and treat them implicitly; while for the semiconductor Boltzmann equation, we use the penalization idea to overcome another remarkable difficulty that no explicit form of the local equilibrium is available.

Li Wang
University of California, Los Angeles
liwang@math.ucla.edu

MS135

Simulating Optogenetic Control of the Heart

Optogenetics is the process of expressing light-sensitive proteins in tissue then eliciting a precise electrical response with illumination. Initial experiments in optical heart stimulation have yielded promising results, but clinical translation remains a distant goal. Developing the ability to accurately simulate optogenetics in biophysically-detailed cardiac models is an essential step towards this goal. We present a comprehensive but flexible framework for simulating cardiac optogenetics, modeling optical stimulation effects from molecule scale to organ level.

Patrick M. Boyle
Institute for Computational Medicine
Johns Hopkins University
pmjboyle@jhu.edu

John C. Williams, Emilia Entcheva
Stony Brook University
jocwilli@ic.sunysb.edu, emilia.entcheva@sunysb.edu

Natalia A. Trayanova
Institute for Computational Medicine
Johns Hopkins University
ntrayanova@jhu.edu

MS135

Reconstruction of Catheter Electrograms and 12-lead ECG in Heart-failure Patients using a Bidomain Reaction-diffusion Model of the Heart and Torso

To predict and improve response to cardiac resynchronization therapy in heart-failure patients, we are building patient-specific models based on MRI data of 12 cases and simulating cardiac activation, local electrograms, and the surface electrocardiogram. Simulations used 2048 cores of a Cray XE6. Activation order as well as electrogram and electrocardiogram morphology are compared to their measured counterparts. Here we report on our modeling techniques, challenges in data acquisition, and results in the first 4 cases.

Mark Potse, Dorian Krause, Rolf Krause
Institute of Computational Science
University of Lugano
mark@potse.nl, dorian.krause@usi.ch, rolf.krause@usi.ch

Wilco Kroon
University of Lugano
Lugano, Switzerland
wilco.kroon@usi.ch

Enrico Caiani
Politecnico di Milano
Milano, Italy
enrico.caiani@biomed.polimi.it

Frits Prinzen
Maastricht University
Maastricht, The Netherlands
frits.prinzen@maastrichtuniversity.nl

Angelo Auricchio
Fondazione Cardiocentro Ticino
angelo.auricchio@cardiocentro.org

MS135

Highly Scalable Cardiac Modeling Codes for Petascale Computing

Electrophysiology simulations of high-resolution human heart ventricles with realistic anatomy and biophysically detailed cellular models are run on Sequoia at LLNL. Maximum sustained performance is 11.8 PFlop/s (58.8% of peak) with exceptional strong scaling and time to solution. For a heart at 0.10 mm resolution, the throughput is 18.24 heartbeats per minute, over 1200 times faster than the previous state of the art. At 0.13 mm resolution, our throughput is just 12% below real-time simulation.

John J. Rice
Computational Biology Center
IBM T.J. Watson Research Center
johnrice@us.ibm.com

Erik Draeger, James Glosli
Lawrence Livermore National Laboratory
draeger1@llnl.gov, glosli1@llnl.gov

Slava Gurev, Changhoan Kim, John Magerlein
IBM Research
vgurev@us.ibm.com, kimchang@us.ibm.com,
mager@us.ibm.com

Art Mirin
Lawrence Livermore National Laboratory
mirin1@llnl.gov

Matthias Reumann
IBM Research
reumann@au1.ibm.com

David Richards
Lawrence Livermore National Laboratory
richards12@llnl.gov

MS135

Simulation of Cardiac Electrophysiology: Efficient Time Integration

Cardiac electrophysiology can be modeled by the bidomain equations, a multi-scale reaction-diffusion system of non-linear ODEs describing the ionic currents at the cellular scale coupled with a set of PDEs describing the propagation of the electrical activity at the tissue scale. To produce clinically useful data, billions of variables must be evolved. In this presentation, I discuss the most efficient time-integration methods for the bidomain equations and reveal the secrets of their success.

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

MS136

Implicit Domain Decomposition Methods for Atmospheric Flows on the Cubed-sphere

We introduce a fully implicit second-order finite volume method for solving some transport problems on the cubed-sphere. To solve the linear system of equations at each time step, we investigate some single-level and multilevel overlapping Schwarz preconditioners. Even though, Schwarz preconditioners are initially designed for elliptic problems, we show by numerical experiments that they scale quite well for this class of purely hyperbolic problems on a machine with a large number of processors.

Haijian Yang
Hunan University,Changsha, Hunan 410082, China
haijian.yang@colorado.edu

Chao Yang
Dept. of Computer Science
University of Colorado, Boulder
chao.yang@coloroado.edu

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

MS136

A New Multi-tracer-efficient Semi-Lagrangian Transport Scheme for the Community Atmosphere Model (CAM)

Recently, a conservative semi-Lagrangian multi-tracer transport scheme (CSLAM) was integrated in the High-Order Method Modeling Environment (HOMME). A new dynamical core for the Community Atmosphere Model is based on HOMME-Spectral Elements (CAM-SE) and the cubed-sphere grid. To improve computational efficiency we replace the SE transport scheme for structured grids by CSLAM but still use the velocity provided by the SE discretization. We compare this hybrid approach on the baroclinic wave test example and discuss the effects of several coupling methods.

Christoph Erath
NCAR
University of Colorado, Boulder
erath@ucar.edu

Mark A. Taylor
Sandia National Laboratories, Albuquerque, NM
mataylo@sandia.gov

MS136

Application of the Cubed-Sphere Grid to Tilted Black Hole Accretion Disks

Traditionally, general relativistic magnetohydrodynamic simulations of black hole accretion disks have used spherical-polar meshes. Such meshes are adequate for cases when the angular momentum axis of the disk is aligned with the spin axis of the black hole. However, for misaligned (i.e. tilted) disks, the cubed-sphere grid offers a number of advantages. In this talk I present some of our work using the cubed sphere grid for simulations of tilted black hole accretion disks.

P. Chris Fragile, Will DuPre, Julia Wilson
College of Charleston
fragilep@cofc.edu, whdupre@gmail.com,
julia.wilson31@gmail.com

Chris Lindner
University of Texas
lindner@astro.as.utexas.edu

MS136

Advances in Numerical Modeling of the Atmosphere on the Cubed-sphere Grid

The desire to improve our understanding of the Earth system has driven numerical models to very fine grid resolutions. Simulations at these resolutions require massively parallel systems in order to be completed over a reasonable time scale. In order to accomplish parallelism on this scale, modern numerical methods have relied on quasi-uniform grids, in particular the cubed-sphere grid. This talk will both review advances in modeling the atmosphere on the cubed-sphere grid and identify new avenues for future research on this subject.

Paul Ullrich
UC Davis
paullrich@ucdavis.edu

MS137

Computational Chemistry: Chemical Accuracy and Errors at Different Scales

Computational chemistry can be used to reliably predict the properties of compounds with density functional theory and correlated molecular orbital theory. The use of computational methods to design syntheses for new materials, for example, for catalysis, solar energy capture, and nuclear fuels, is in its infancy. We will describe the complex issues that need to be addressed for the design of new materials syntheses and initial progress on understanding basic steps in such reaction mechanisms.

David Dixon
University of Alabama
dadixon@as.ua.edu

MS137

Chemical Kinetic Model Development using Bayesian Variable Selection

We present a novel approach for tractable Bayesian inference of chemical kinetic models from noisy and indirect system-level data. Formulating the problem as variable selection and making use of point-mass mixture priors, our approach allows an exhaustive comparison of all models comprised of subsets of elementary reactions. Adaptive Markov chain Monte Carlo methods are used to efficiently explore the posterior distribution. We also present a newly-developed proposal distribution that improves MCMC chain mixing in variable selection problems containing strong posterior correlations.

Nikhil Galagali, Youssef M. Marzouk
Massachusetts Institute of Technology
nikhilg1@mit.edu, ymarz@mit.edu

MS137

Split Step Adam Moulton Method for Stiff Stochastic Differnetial Equations

We present a split-step method for solving stochastic differential equations (SDEs). The method is based on a second order split Adams-Moulton Formula for stiff ordinary differential equations and modified for use on stiff SDEs which are stiff in both the deterministic and stochastic components. Its order of strong convergence is established and stability region is displayed. Numerical results show the effectiveness of the method in the path wise approximation.

Abdul M. Khaliq
Middle Tennessee State University
Department of Mathematical Sciences
akhaliq@mtsu.edu

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

MS137
Accurate Filtering of The Navier-Stokes Equation

Filtering in the data assimilation context is often considered to be the reproduction of a deterministic point estimate of the state of a system from noisily observed data and knowledge of the underlying system, resulting in a continuous feedback control problem. This is in contrast to the probabilistic interpretation of the state as a random quantity, whose distribution reflecting our knowledge at a particular time is referred to as the filtering distribution, i.e. the distribution of state given the filtration generated by the random observations made until the given time. We focus on the former case in which the model is know, i.e. the perfect model scenario. In the perfect model scenario two ideas drive accurate filtering: (i) observe enough low frequency information, and (ii) model variance inflation: trust the observations. In this talk I will illustrate this for the simplest filter, referred to as 3DVAR (perhaps more accurately called 2DVAR, if one then adapts 4DVAR to (2+1)DVAR to avoid the necessary confusion) applied to the 2D Navier-Stokes equations, in the low and high frequency observation limits. The latter will yield an SPDE for the estimator.

Kody Law
Mathematics Institute
University of Warwick
kodylaw@gmail.com

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

Dirk Bloemker
University of Augsburg, DE
dirk.bloemker@math.uni-augsburg.de

Kostas Zygalakis
University of Southampton, UK
k.zygalakis@soton.ac.uk

MS138
An Application of Sparse Sensing to Partial Differential Equations

Abstract not available at time of publication.

Ido Bright
University of Washington
ibright@uw.edu

MS138
Reduced Order Models of Unsteady Aerodynamic Flow for Control

Abstract not available at time of publication.

Steven L. Brunton
Princeton University
sbrunton@princeton.edu

MS138
TBD - Equation Free Modeling

Abstract not available at time of publication.

I. G. Kevrekidis
Princeton University
yannis@princeton.edu

MS138
Hybrid Reduced Order Integration Using the Proper Orthogonal Decomposition and Dynamic Mode Decomposition

Abstract not available at time of publication.

Matthew O. Williams
Program in Applied and Computational Mathematics
Princeton University
mow2@Princeton.edu

MS139
Probabilistic Graphical Models: Applications to UQ

We develop a non-parametric probabilistic graphical model framework that can provide a clear and efficient way to represent correlated random variables in uncertainty quantification problems governed by stochastic PDEs. Given a set of training data, the basic objectives include structure design, graph learning and inference problem. Using a probabilistic graphical model approach to UQ allows 1) capturing the correlations between random variables; 2) inference of any unobserved variables with appropriate error bars; 3) wide application prospects due to the non-parametric nature of the model.

Peng Chen, Nicholas Zabaras
Cornell University
pc423@cornell.edu, zabaras@cornell.edu

MS139
The Ensemble Kalman Filter for Inverse Problems

We present a novel non-standard perspective of the ensemble Kalman filter (EnKF) methodology that converts it into a generic tool for solving inverse problems. We provide numerical results to illustrate the efficacy of the proposed EnKF for solving a wide range of applications. In particular, we discuss (i) inversion of hydraulic head data to determine transmissivity in a groundwater model; (ii) inversion Eulerian velocity measurements to determine the initial condition in an incompressible fluid.

Marco Iglesias
Civil and Environmental Engineering
MIT
m.a.iglesias-hernandez@warwick.ac.uk

Kody Law

Mathematics Institute
University of Warwick
kodylaw@gmail.com

Andrew Stuart
University of Warwick
andrewstuart@gmail.com

MS139

Tensor-based Uncertainty Quantification of Experimental Data

A methodology is presented for inferring a functional representation of a stochastic quantity from a collection of its realizations. This collection comes from experiments, where no specific sampling strategy such as sparse grids or low-discrepancy sequences can be assumed. We derive a new tensor-train approach based on orthogonal polynomials in order to accurately approximate the random quantity with few degrees of freedom, allowing for a stable solution even with scarce data. Provided examples support the notion of data-driven multilinear algebra as a potentially effective tool for high-dimensional uncertainty quantification.

Tarek Moselhy
MIT
tmoselhy@mit.edu

Lionel Mathelin
LIMSI - CNRS
mathelin@limsi.fr

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

MS139

Strategies for Quantification of Epistemic Uncertainty

Epistemic uncertainty refers to the uncertainty due to our lack of knowledge. In this talk, we are concerned with parametric type of uncertainty when its complete probabilistic information is not available. We describe a few numerical approaches to construct reliable system response, without using probability distribution function, and their ranges of applicabilities. We also present a method to compute both the upper bound and lower bound of the systems responses using variational inequality of relative entropy. The results bounds can serve as estimates of the "best case scenario" and "worst case scenario". We also discuss how to efficiently compute the bounds as a mere post-processing step of the traditional UQ computation.

Xiaoxiao Chen, Jing Li, Xin Qi, Dongbin Xiu
Purdue University
chen294@math.purdue.edu, li176@purdue.edu, qixin123@gmail.com, dxiu@purdue.edu

MS140

Recent Advances in Scalable Sparse Factorization Methods

Because of their robustness, factorization-based solvers and preconditioners play a significant role in developing scalable solvers for large-scale, ill-conditioned and highly-indefinite algebraic equations. We present our recent development on sparse factorization algorithms for multi/many-core architectures and exploiting low-rank structures. These include new scheduling algorithm for DAG execution to reduce processes idle time, incorporate light-weight OpenMP threads in MPI processes to reduce memory footprint, and superfast factorization algorithms based on hierarchically semi-separable (HSS) matrices for structured sparse linear systems arising from discretized PDEs.

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

MS140

Multithreaded Sparse Kernels for Solution of Sparse Linear Systems

Solving sparse linear systems is one of the most expensive parts of computational science applications. It is important to exploit the shared memory parallelism available in modern multicore architectures to develop scalable sparse linear solvers. We describe a task-based model and implement the sparse kernels using the model to improve multithreaded performance. We focus on sparse kernels that are important for both iterative and direct solvers and present performance results on various multicore architectures.

Sivasankaran Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

MS140

Accelerated Fixed Point Methods and Other Advances in SUNDIALS

The SUNDIALS software suite supplies time integration and nonlinear solvers used in implicit solution approaches to ODEs and PDEs. New capabilities going into this parallel, C language suite of codes include accelerated fixed point nonlinear solvers and IMEX Runge-Kutta time integration methods. We will overview the fixed point solvers, their interfaces in SUNDIALS, and their application in subsurface and materials sciences. Lastly, we will present information on the new time integrators soon to be added.

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

MS140

Recent Advances on Reducing Communication in AMG

Algebraic Multigrid (AMG) solvers are an essential component of many large-scale scientific simulation codes. Their continued numerical scalability and efficient implementation is critical for preparing these codes for emerging computer architectures. Previous investigations have shown that the increasing communication complexity on coarser grids combined with the effects of increasing numbers of cores lead to severe performance bottlenecks for AMG on various multicore architectures. We present recent progress on several efforts to reduce communication in AMG.

Ulrike Meier Yang
Lawrence Livermore National Laboratory
yang11@llnl.gov

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

Jacob Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

MS141

Applications and Challenges in Large-scale Graph Analysis

Emerging real-world graph problems include detecting community structure in large social networks, improving the resilience of the electric power grid, and detecting and preventing disease in human populations. We discuss the opportunities and challenges in massive data-intensive computing for applications in social network analysis, genomics, and security. The explosion of real-world graph data poses substantial challenges for software, hardware, algorithms, and application experts.

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

David A. Bader
Georgia Institute of Technology
bader@cc.gatech.edu

MS141

Large Scale Graph Analytics and Randomized Algorithms for Applications in Cybersecurity

Abstract not available at time of publication.

John R. Johnson
Pacific Northwest National Laboratory
john.johnson@pnnl.gov

Emilie Hogan
Pacific Nortwest National Lab
emilie.hogan@pnnl.gov

Mahantesh Halappanavar
Pacific Northwest National Laboratory
mahantesh.halappanavar@pnnl.gov

MS141

Combinatorial and Numerical Algorithms for Network Analysis

We report on our recent efforts in developing combinatorial and numerical algorithms for network analysis. Algorithms under consideration include community detection and seed set expansion, algebraic distances and lean algebraic multigrid. The development is driven by the rationale of combining ease-of-use and high performance.

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

Christian Staudt
Karlsruhe Institute of Technology
christian.staudt@kit.edu

MS141

Anomaly Detection in Very Large Graphs: Modeling and Computational Considerations

Graph theory provides an intuitive mathematical foundation for dealing with relational data, but there are numerous computational challenges in the detection of interesting behavior within small subsets of vertices, especially as the graphs grow larger and the behavior becomes more subtle. This presentation discusses computational considerations of a residuals-based subgraph detection framework, including the implications on inference with recent statistical models. We also present scaling properties, demonstrating analysis of a billion-vertex graph using commodity hardware.

Benjamin Miller, Nicholas Arcolano, Edward Rutledge, Matthew Schmidt, Nadya Bliss
MIT Lincoln Laboratory
bamiller@ll.mit, arcolano@ll.mit.edu, rutledge@ll.mit.edu, matthew.schmidt@ll.mit.edu, nt@ll.mit.edu

MS142

Least-squares Finite Element Methods for Incompressible Flows with Improved Mass Conservation

We present a new locally conservative LSFEM for the velocity-vorticity-pressure Stokes and Navier-Stokes equations, which uses a piecewise divergence-free basis for the velocity and standard C^0 elements for the vorticity and the pressure. Computational studies demonstrate that the new formulation achieves optimal convergence rates and yields high conservation of mass. We also propose a simple diagonal preconditioner for the dV-VP formulation, which significantly reduces the condition number of the LSFEM problem.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Luke Olson
UIUC
lukeo@illinois.edu

James Lai
Microsoft
james.h.lai@outlook.com

MS142

Asymptotically Exact a posteriori Error Estimators for LS Methods

A new asymptotically exact a posteriori error estimator is developed for first-order div least-squares(LS) finite element methods. Let $(u_h\,,\,\sigma_h)$ be equal-order LS approximate solution for $(u\,,\,\sigma=-A\nabla u)$. Then, $\mathcal{E}=\|A^{-1/2}(\sigma_h+A\nabla u_h)\|_0$ is asymptotically exact a posteriori error estimator for $\|A^{1/2}\nabla(u-u_h)\|_0$ or $\|A^{-1/2}(\sigma-\sigma_h)\|$ depending on the order of approximate spaces for σ and u. For $\mathcal{E}$ to be asymptotically for $\|A^{1/2}\nabla(u-u_h)\|_0$, we require higher order approximation property for σ, and vice versa. When both $A\nabla u$ and σ are approximated in the same order of accuracy, the estimator becomes an equivalent error estimator for both errors. Confirming numerical results are presented.

Jaeun Ku
Oklahoma State University
jku@math.okstate.edu

Zhiqiang Cai
Purdue University
Department of Mathematics
zcai@math.purdue.edu

Varis Carey
University of Texas at Austin
variscarry@googlemail.com

Eun-Jae Park
Yonsei University
ejpark@math.yousei.ac.kr

MS142
FOSLL* For Nonlinear Partial Differential Equations

For sufficiently regular, elliptic-like problems, a least-squares approach may be continuous/coercive in the H^1 or $H(div)$ norm, yielding convergence in H^1 or $H(div)$. Often, however, coarse approximations may have large L^2-error relative to the H^1 or $H(div)$ seminorm. In contrast, the first-order system LL^* (FOSLL*) approach minimizes error in a dual norm induced by the differential operator, yielding better control of the L^2-error. We extend this general framework to nonlinear problems and establish convergence theory.

Chad Westphal
Wabash College
westphac@wabash.edu

MS142
Superconvergence: Unclaimed Territories

Abstract not available at time of publication.

Zhimin Zhang
Wayne State University
Department of Mathematics
ag7761@wayne.edu

MS143
Power Flows and Stability of a Modern Distribution Feeder

In this talk I briefly review the work at LANL on ODE and PDE modeling of a distribution feeder. In particular, I plan to discuss effects of many inductive motors and of many distributed generators on the feeder stability, bifurcation diagram and potentially dangerous hysteretic response to a voltage fault.

Michael Chertkov
Los Alamos National Laboratory
chertkov@lanl.gov

MS143
Graph-based Approaches for N-x Contingency Analysis of Electric Power Grids

The goal of $N-x$ contingency selection is to pick a subset of critical cases to assess their potential for causing a severe crippling of a power grid. Since the number grows exponentially, it can be overwhelming even for a moderately-sized system. We propose a novel method for $N-x$ selection using group-betweenness centrality and show that the amount of computation can be decoupled from the problem size, thus making it feasible for large systems.

Mahantesh Halappanavar
Pacific Northwest National Laboratory
mahantesh.halappanavar@pnnl.gov

Yousu Chen, Zenyu Huang, Mark Rice
Pacific Northwest National Lab
yousu.chen@pnnl.gov, zenyu.huang@pnnl.gov, mark.rice@pnnl.gov

MS143
Random Chemistry and Dual Graphs: Two Ways to Understand Cascading Failures in Power Grids

Power systems are vulnerable to low probability, high impact failures. Because the probability distribution of blackout sizes follows a power law, new statistical methods are needed to provide good information about the risk of large blackouts. The Random Chemistry approach strategically tests a model to quickly generate large unbiased sets of cascading outage data. The dual graph method transforms simulation data into a graph that provides insight into how cascades propagate. Both methods transform models and data into meaningful information.

Paul Hines
University of Vermont
paul.hines@uvm.edu

MS143
The Role of Grid Topology in Secure Power System Optimization Problems

The electric grid is considered critical infrastructure and data concerning transmission lines and their connected topology are restricted. We consider the masking of power system optimization models for shared computing, and for distributing equivalent power system models among researchers. We examine properties of masking transformations that allow equivalent optimization models to be public, the analysis of which can be only be related back to the original model through knowledge of the transformation.

Bernard Lesieutre, Alex Borden, Daniel Molzahn, Parmesh Ramanathan
University of Wisconsin at Madison
lesieutre@engr.wisc.edu, borden.alex@gmail.com, molzahn@wisc.edu, parmesh@ece.wisc.edu

MS144
Closure Modeling for the Proper Orthogonal Decomposition of Turbulent 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, especially

for three-dimensional nonlinear problems. 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. Its original promise of computationally efficient, yet accurate approximation of coherent structures in high Reynolds number turbulent flows, however, 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. In this talk, we put forth several closure models for the POD-ROMs of structurally dominated turbulent flows. These models, which are considered state-of-the-art in large eddy simulation, are carefully derived and numerically investigated. We also discuss several approaches for an efficient and accurate numerical discretization of general nonlinear POD closure models. We numerically illustrate these developments in several computational settings, including a three-dimensional turbulent flow past a cylinder at Reynolds number Re = 1000. A rigorous numerical analysis of the new computational framework will also be presented.

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

Zhu Wang
Institute for Mathematics and its Applications
University of Minnesota
wangzhu@ima.umn.edu

MS144

Trust Region POD 4D VAR Data Assimilation of a Parabolized NavierStokes Equations Model

A reduced order model based on Proper Orthogonal Decomposition (POD) 4D VAR data assimilation for parabolized NavierStokes (PNS) equations is derived. Various approaches of POD implementation of reduced order inverse problem are compared including an ad-hoc POD adaptivity along with a trust region POD adaptivity. Numerical results show that trust region POD 4D VAR provides the best results for all error metrics for reduced order inverse problem of PNS equations.

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

Juan Du
IAP, Academia Sinica, Beijing, China
dujuan10@mail.iap.ac.cn

MS144

Model Reduction and Multi-fidelity Data in Uncertainty Quantification of Flow Models

We study the use of lower-fidelity training data for uncertainty quantification of advanced fluid mechanics models. We use POD-based dimensionality reduction to obtain approximations of the full model. This imperfect data, calibrated using stochastic-processes learning, can be used to recover the full model's response to uncertainty. Our approach can be generalized to using multiple sources of data. We performed experiments on simplified fluid flow models, as well as on a professional fluid mechanics solver nek5000.

Oleg Roderick
Argonne National Laboratory
roderick@mcs.anl.gov

MS144

Towards a Blackbox Approach for Model Reduction via EIM

We will introduce two EIM (empirical interpolation method) -based methods for nonlinear model reduction, both inspired by discrete empirical interpolation method (DEIM). These new approaches provide comparable performance with DEIM at reduced cost and are promising for blackbox model reduction. Different points selection algorithms will be compared with DEIM and numerical examples will be shown to demonstrate the performance of the new algorithms. Finally, stability issues will be discussed.

Xueyu Zhu
Division of Applied Mathematics
Brown University
xueyu_zhu@brown.edu

Ramakanth Munipalli
Hypercomp. Inc
mrk@hypercomp.net

Hesthaven Jan
Division of Applied Mathematics
Brown University
jan_hesthaven@brown.edu

MS145

Approximation and Use of Set Valued Solutions to Stochastic Inverse Problems

We discuss the use of set-valued solutions for approximate solution of stochastic inverse problems for parameter determination. The focus is on using observations on multiple quantities of interest.

Donald Estep, Troy Butler
Colorado State University
estep@math.colostate.edu or don.estep@gmail.com, butler@stat.colostate.edu

MS145

Adaptive Stochastic Collocation for PDE Optimization under Uncertainty

Many optimization problems in engineering and science are governed by partial differential equations (PDEs) with uncertain parameters. Although such problems can be formulated as optimization problems in Banach spaces and derivative based optimization methods can in principle be applied, the numerical solution of these problems is more challenging than the solution of deterministic PDE constrained optimization problems. The difficulty is that the PDE solution is a random field and the numerical solution of the PDE requires a discretization of the PDE in space/time as well as in the random variables. As a consequence, these optimization problems are substantially larger than the already large deterministic PDE constrained optimization problems.
In this talk we discuss the numerical solution of such op-

timization problems using stochastic collocation methods. We explore the structure of this method in gradient and Hessian computations. A trust-region framework is used to adapt the collocation points based on the progress of the algorithms and structure of the problem. Convergence results are presented. Numerical results demonstrate significant savings of our adaptive approach.

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

Drew P. Kouri
Mathematics and Computer Science Division
Argonne National Laboratory
dpkouri@mcs.anl.gov

MS145

A Framework for Sequential Experimental Design for Inverse Problems

Abstract not available at time of publication.

Luis Tenorio
Colorado School of Mines
ltenorio@mines.edu

MS145

A Generalized Stochastic Collocation Approach to Constrained Optimization for Random Data Identification Problems

Characterizing stochastic model inputs to physical and engineering problems relies on approximations in high-dimensional spaces, particularly in the case when the experimental data or targets are affected by large amounts of uncertainty. To approximate these high-dimensional problems we integrate a generalized adaptive sparse grid stochastic collocation method with a SPDE-constrained least squares parameter identification approach. Rigorously derived error estimates will be used to show the efficiency of the methods at predicting the behavior of the stochastic parameters.

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

Max Gunzburger
Florida State University
gunzburger@fsu.edu

Catalin S. Trenchea
Department of Mathematics
University of Pittsburgh
trenchea@pitt.edu

MS146

Portability of a Large-scale Heart Model through Hybrid Parallelization

Portability of CS&E software between diverse parallel architectures is a challenging task and important for its flexible and efficient usage. In this talk, we discuss our experiences with a large-scale heart model (monodomain and bidomain equations). The in-house code Propag was ported from its original design for shared memory systems to contemporary massively parallel multicore architectures by means of a hybrid OpenMP–MPI parallelization. We also report on the relevance of the newly introduced parallel setup phase.

Thomas Dickopf, Dorian Krause, Mark Potse, Rolf Krause
Institute of Computational Science
University of Lugano
thomas.dickopf@usi.ch, dorian.krause@usi.ch, mark@potse.nl, rolf.krause@usi.ch

MS146

ONELAB: Open Numerical Engineering LABoratory

We present the ONELAB software, a lightweight open source toolkit to interface finite element solvers used in a variety of engineering disciplines. Based on the design of the open-source CAD modeler, mesh generator and post-processor Gmsh, ONELAB is targeted for use in education and small- and medium-size businesses, which do not require the power (and cannot justify the cost) of commercial finite element software suites.

Christophe Geuzaine
University of Liège
cgeuzaine@ulg.ac.be

MS146

Commercial Open-source: An Approach to Computational Fluid Dynamics

Business and open-source are no contradiction: while this basic truth is now widely adopted, many software users still struggle to embrace the world of commercial open-source and to separate between irrational fears and healthy precautions. We present as an example the open-source project Palabos (www.palabos.org) for computational fluid dynamics, which follows the logic of a commercial business plan, yet is open-source and supports University-level research through academic collaboration. While its win-win benefits are obvious, in practice this approach leads to numerous misunderstandings and fears which have a considerable impact on the project management. On the side of commercial users, the most frequent misunderstandings relate to licensing and usage term, and to considerations about software quality and support guarantees. On the side of the academic community, fears about a future interruption of the free-software guarantee, or about otherwise unethical project management appear to be preponderant. We show that in order to overcome these difficulties, a policy of extensive and open communication with the user base is crucial, related to the project management as well as the technical aspects of the software. In particular, we show that both the programming and graphical user interfaces need to be thought in such a way as to send signals with which users from both academic and industrial communities feel comfortable.

Jonas Latt
Universite de Geneve
jonas.latt@unige.ch

Bastien Chopard
University of Geneva, Switzerland
bastien.chopard@unige.ch

MS146

Building Effective Parallel Unstructured Adaptive Simulation by In-memory Integration of Existing Software Components

Unstructured adaptive meshing driven by error estimation procedures support reliable mesh-based PDE simulations. Typically, interfacing of unstructured adaptive procedures with existing analysis components is through files. File I/O is currently the critical bottleneck in parallel adaptive analysis, even when advanced parallel I/O methods are used. An approach to eliminate this bottleneck using in-memory integration of existing analysis procedures with adaptive mesh procedures is presented. Results demonstrate strong scaling to $>$ 32,000 processors of the in-memory integration.

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

Onkar Sahni
Rensselaer Polytechnic Institute
sahni@scorec.rpi.edu

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

MS147

Implications of the Choice of SDC Nodes in the Multilevel PFASST Algorithm

The Parallel Full Approximation Scheme in Space and Time (PFASST) algorithm iteratively improves the solution on each time slice by applying Spectral Deferred Correction (SDC) sweeps to a hierarchy of discretizations at different spatial and temporal resolutions. The number and type of SDC nodes used on each level of refinement impacts several aspects of the scheme including stability, accuracy, convergence, and efficiency. We analyse various SDC nodes and their impact on several representative PDEs.

Matthew Emmett
University of North Carolina at Chapel Hill
Department of Mathematics
mwemmett@lbl.gov

Michael Minion
ICME Stanford
mlminion@gmail.com

MS147

On the Convergence of Parallel Deferred Correction Methods

We discuss novel convergence results for parallel deferred correction methods for solving initial-value problems $u'(t) = f(t, u(t))$, $u(0) = u_0$, with an emphasis on the Parareal-SDC variant proposed by Minion & Williams (2008) and Minion (2011), and a new variant based on barycentric rational interpolation. This is joint work with Martin J. Gander.

Stefan Guettel
University of Oxford
stefan.guettel@manchester.ac.uk

MS147

An Optimized RIDC-DD Space-time Method for Time Dependent Partial Differential Equations

Recently, the Revisionist Integral Deferred Correction (RIDC) approach has been shown to be a relatively easy way to add small scale parallelism (in time) to the solution of time dependent PDEs. In this talk I will show how large scale spatial parallelism can be added to RIDC using classical and optimized Schwarz methods. This results in truly parallel space-time methods for PDEs suitable for hybrid OpenMP/MPI implementation. Initial scaling studies will demonstrate the viability of the approach.

Ronald Haynes
Memorial University of Newfoundland
St. John's, NL, Canada
rhaynes74@gmail.com

Benjamin Ong
Michigan State University
ongbw@msu.edu

MS147

Parallelization in Time: How Far into the Future Can We See?

This talk will provide an overview of past and present developments in parallelization of numerical methods for ordinary and partial differential equations in the temporal direction. I will highlight some recent efforts to extend the efficiency and applicability of parallel in time methods and give some insight into the practical limitations of current approaches.

Michael Minion
ICME Stanford
mlminion@gmail.com

MS148

Efficient Methods for Wave Propagation in Layered Absorbing Media

In this talk we discuss Galerkin boundary integral formulations and their fast solution via Hierarchical matrices for the problem of wave propagation in nested, layered, absorbing media. This is the forward problem in diffuse optical tomography, where the different layers model for example a head geometry. The software implementation of the results presented in this talk is done via BEM++, a new library for the solution of boundary integral equations via Galerkin boundary elements.

Simon Arridge
Department of Computer Science
University College London
simon.arridge@cs.ucl.ac.uk

Timo Betcke
University College London
t.betcke@ucl.ac.uk

Martin Schweiger
Department of Computer Science
University College London
m.schweiger@cs.ucl.ac.uk

Wojciech Smigaj
Department of Mathematics, University College London

Faculty of Physics, Adam Mickiewicz University in Poznan
w.smigaj@ucl.ac.uk

MS148

An ADI Timestepping Preconditioner for the Helmholtz Equation

Passing to the time domain is a great way to avoid the preconditioning issues in the high-frequency Helmholtz equation, though the complexity count is typically unfavorable due the CFL condition. In the talk I will explain how this timestepping restriction can be circumvented by a specific ADI scheme, in the scope of narrowband solutions. As a result, the Helmholtz equation can be solved in low complexity in rather general variable media.

Laurent Demanet
Mathematics, MIT
demanet@gmail.com

MS148

Sweeping Preconditioners for the Helmholtz Equation

Many standard techniques for preconditioning do not work well for high frequency wave propagation. We will describe and analyze new class of preconditioners based on sweeping processes and apply them to GMRES iterative solutions of frequency domain equations. Hierarchical matrix techniques for compression and moving perfectly matched layers play important roles in the algorithm. The number of operations scales essentially linearly in the number of unknowns.

Bjorn Engquist
Department of Mathematics and ICES, UT Austin
engquist@ices.utexas.edu

MS148

Paraxial Estimates and a Direct Structured Solver for the Helmholtz Equation in the High-frequency Regime

Abstract not available at time of publication.

Maarten de Hoop
Center for Computational & Applied Mathematics
Purdue University
mdehoop@math.purdue.edu

MS149

Computational Engineering and Science Software for Nanoscale Explorations

Interactive computational engineering and science software has been developed for students to explore phenomena important at the nanoscale. The software has been designed for integration into physics, chemistry, materials science, and engineering courses. As an example, the software allows students to simulate the propagation of light through a photonic band gap material, and to design its internal nanoscale dimensions to tune the amount of light reflected and transmitted at different frequencies.

Richard Braatz
Massachusetts Institute of Technology
Department of Chemical Engineering
braatz@mit.edu

MS149

Complexity Science and Computational Modeling

Complexity Science creates an opportunity to bring to the undergraduate curriculum topics that are useful and important, accessible and engaging to students, and naturally connected to material from Mathematics, Physics, and Computer Science. At Olin College we have developed an unusual new class, Computational Modeling, with an accompanying textbook, *Think Complexity.* Students in this class read seminal papers in Complexity Science, write Python programs to run experiments, and write case studies reporting their results.

Allen Downey
Olin College of Engineering
Green Tea Press
downey@allendowney.com

MS149

Computational Laboratory Activities for Medicinal Chemistry

Lab activities for a medicinal chemistry course were developed to teach undergraduate biology and chemistry students the important role that computers play in the drug discovery process. During lab time, students learned how to use a docking program such as AutoDock Vina to conduct a virtual screen of compounds and a protein visualization program such as PyMOL. They also gained an appreciation for how much a supercomputer can speed up the virtual screening process.

Jimmy Franco
Merrimack College
francoj@merrimack.edu

MS149

Computational Quantum Mechanics in the Undergraduate Curriculum

Quantum Mechanics offers a fertile field for computation. The subject is mathematically and conceptually difficult and because Planck's constant is small undergraduates have poor quantum intuition. Computation can be invaluable in helping students build intuition and opens up new classes of problems, many of current experimental interest. I will describe a number of modules that I have developed for teaching computational quantum mechanics and how these materials are being used in the classroom.

Richard Gass
University of Cincinnati
richard.gass@uc.edu

MS150

Investigations of High Order Discontinuous Galerkin Methods for Implicit Large Eddy Simulations

We will present DG implicit LES simulations of canonical turbulent flows and investigate the potential of stabilized very high-order approximations (polynomial degree ≥ 7) in underresolved scenarios. We will focus on stabilization by polynomial de-aliasing. Thus, the only remaining parameter in such a DG formulation is the choice of the numerical flux functions. We will investigate the influence of different flux functions on the quality of iLES results for the Taylor-Green vortex.

Andrea D. Beck
University of Stuttgart
beck@iag.uni-stuttgart.de

Gregor Gassner
Institute for Aerodynamics and Gasdynamics
Universitaet Stuttgart
gassner@iag.uni-stuttgart.de

MS150

Discretely Energy Stable Discontinuous Galerkin Spectral Element Methods

In this talk, we present a skew symmetric discontinuous Galerkin collocation spectral element (DGSEM) formulation for the non-linear Burgers equation. We prove that this formulation remains discretely conservative. Furthermore, we prove that in combination with the commonly used Roe flux energy stability cannot be guaranteed, whereas in combination with the local Lax-Friedrichs flux energy stability can be proved. In combination, this yields a novel discretely energy stable and discretely conservative DGSEM formulation.

Gregor Gassner
Institute for Aerodynamics and Gasdynamics
Universitaet Stuttgart
gassner@iag.uni-stuttgart.de

MS150

A Solver for the Incompressible Navier-Stokes Equations using DG, QBX, and Integral Equations

We present a new method for the solution of the incompressible Navier-Stokes equation. Our approach combines the stiffly stable splitting scheme and a discontinuous Galerkin spatial discretization of the advection terms with a new integral-equation-based solver for the pressure and viscous stages. The latter is based on a volume generalization of Quadrature by Expansion (QBX), originally a method for the computation of Nyström discretizations of layer potentials. We demonstrate the scheme's performance in general geometries.

Andreas Kloeckner
Courant Institute of Mathematical Sciences
New York University
kloeckner@cims.nyu.edu

Tim Warburton
Department of Computational and Applied Math
Rice University
timwar@caam.rice.edu

MS150

Shock Capturing Using Artificial Viscosity and Multiscale Methods

We present some advances in the application of high order Discontinuous Galerkin methods to problems involving shock waves. First, an implicit scheme based on the divergence of the velocity as a driver for artificial viscosity will be presented. Following, a multiscale approach in which elements affected by shocks are discretized using a low order sub-mesh will also be discussed. Both approaches will be compared on the same cases to asses their performance and robustness.

Ngoc Cuong Nguyen, David Moro, Jaime Peraire
Massachusetts Institute of Technology
cuongng@mit.edu, dmoro@mit.edu, peraire@MIT.EDU

MS151

Error Control for Output Quantities of Interest in Parameterized Partial Differential Equations

In this work, we propose and study some adaptive approaches to control errors in numerical approximations of differential equations with uncertain coefficients. Adaptivity is based here on a posteriori error estimation and the approaches rely on the ability to decompose the a posteriori error estimate into contributions from the physical discretization and the approximation in parameter space. Errors are evaluated in terms of quantities of interest using well established adjoint based methodologies.

Corey M. Bryant
ICES
The University of Texas at Austin
cbryant@ices.utexas.edu

Tim Wildey
Sandia National Laboratory
tmwilde@sandia.gov

Serge Prudhomme
ICES
The University of Texas at Austin
serge@ices.utexas.edu

MS151

A Posteriori Analysis of Implicit-Explicit (imex) Methods

Implicit-Explicit (IMEX) schemes are an important and widely used class of time integration methods for parabolic or hyperbolic partial differential equations. For example, explicit scheme may be used for the convection (or reaction) term, and an implicit one for the stiff diffusion term. Such schemes may preserve monotonicity properties inherent in the equation. In this talk we cast a class of IMEX schemes in a variational format and perform a posteriori analysis to compute error in a quantity of interest for such schemes.

Jehanzeb Chaudhry
Department of Mathematics
Colorado State University
jehanzeb@colostate.edu

Don Estep, Simon Tavener
Colorado State University
estep@stat.colostate.edu, tavener@math.colostate.edu

Victor E. Ginting
Department of Mathematics
University of Wyoming
vginting@uwyo.edu

MS151

On the Application of Adjoint Methods in Subsur-

face Flow Simulations

This presentation explores an application of the adjoint method for subsurface characterization with a benchmark problem in water infiltration through soil. The idea is to employ available sparse measurements of pressure and/or water content to infer the subsurface permeability that honors the measurement. The process that is designed to achieve this goal is to devise an iterative procedure in which for every iteration a possible permeability field is proposed, then a simulated response is obtained by solving the associated boundary value problems, which is to be compared with the measurements. The aim is to minimize the difference between the simulated response and the measurements. This difference (coined as a residual) is then used as a data in the adjoint equation associated with the original boundary value problem, which in turn gives an information on the perturbation values, serving as a correction to the proposed permeability. The iterative procedure is run this way until a convergence is reached. Some numerical examples are presented to illustrate the framework.

Lawrence Bush, Victor E. Ginting
Department of Mathematics
University of Wyoming
lbush4@uwyo.edu, vginting@uwyo.edu

MS151

Adjoint Methods for Adjoint Inconsistent Formulations

Often, computing a numerically stable approximation to a PDE requires the formulation to be modified. Unfortunately, the adjoint of the modified formulation does not always correspond to a modification of the formal adjoint. We demonstrate that using the adjoint of such a formulation to compute a posteriori error estimates can lead to unstable results. We show that an alternative approach exists for DG and stabilized Galerkin methods that provides more stable and accurate error estimates.

Tim Wildey
Sandia National Laboratory
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

MS152

Large-scale Stochastic Linear Inversion using Hierarchical Matrices

Large scale inverse problems, which frequently arise in earth-science applications, involve estimating unknowns from sparse data. The goal is to evaluate the best estimate, quantify the uncertainty in the solution, and obtain conditional realizations, which are computationally intractable using conventional methods. In this talk, I will discuss the hierarchical matrix approach optimized for a realistic large scale stochastic inverse problem arising from a cross-well seismic tomography application.

Sivaram Ambikasaran
Institute for Computational and Mathematical Engineering
Stanford University
sivaambi@stanford.edu

Judith Yue Li, Peter K Kitanidis
Stanford University
yuel@stanford.edu, peterk@stanford.edu

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

MS152

Large-scale Biomolecular Electrostatics with Massively Parallel FMM

A challenge in understanding biomolecular interactions is that molecules are always in a solution (water molecules and dissolved ions). The potential is described by a Poission-Boltzmann equation which is directly solved via BEM and accelerated with FMM. We present the two dielectric formulation where the electrostatic field is calculated using continuum dielectric media: the solvent and the molecule. Using FMM, we enable large-scale calculations with millions of unknowns and advance towards solving biologically challenging problems.

Aparna Chandramowlishwaran
Georgia Institute of Technology
aparna@cc.gatech.edu

MS152

Optimizing the Black-box FMM for Smooth and Oscillatory Kernels

A black-box FMM for smooth kernels was introduced in [Fong and Darve, 2009] and has been extended to oscillatory kernels in [Messner, Schanz and Darve, 2012]. Its major advantages are the easy implementation and adaptation for new kernels. However, it requires significantly more floating point operations compared to other FMMs. In my talk, I will present ways to tackle this drawback: (1) The exploitation of symmetries allow us to reduce the pre-computation time by a factor greater than 1000. (2) Blocking schemes increase the applicability of optimized Level 3 BLAS routines and lead to a great performance of the actual matrix-vector product.

Matthias Messner
Stanford University
burgerbua@gmail.com

Berenger Bramas, Olivier Coulaud
INRIA
berenger.bramas@inria.fr, olivier.coulaud@inria.fr

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

MS152

Adaptive Parallel Scheduling of the Fast Multipole

Method

We introduce an adaptive scheduling method for parallel execution of the fast multipole method (FMM) in a dynamic computing environment. The scheduling is adaptive to variation in algorithm, change in data as well as specifics and dynamics in the computing system. We present the graph-theoretic analysis underlying the scheduling method and experimental results with parallel FMM.

Bo Zhang
Duke University
zhangb@cs.duke.edu

Nikos Pitsianis
Department of Electrical & Computer Engineering
Aristotle University
nikos.pitsianis@eng.auth.gr

Jingfang Huang
Mathematics
UNC
huang@email.unc.edu

Xiaobai Sun
Department of Computer Science
Duke University
xiaobai@cs.duke.edu

MS153

Early Experience of Adaptation of ppOpen-AT: An Auto-tuning Description Language

We will present an auto-tuning (AT) methodology with ppOpen-AT:an AT description language. One of targets of ppOpen-AT is selection between CPU and GPU. In addition to this selection, code optimization for CPU, such as loop fusion and loop separation, is also main function of the AT. The targets come from real numerical simulation codes, such as earthquake simulation by FDM and electromagnetism simulation by FVM. Performance evaluation is performed by several current CPU architectures, like the SPARC64 IXfx of Fujitsu FX10.

Takahiro Katagiri
The University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

Satoshi Ito
Information Technology Center
The University of Tokyo
sito@cc.u-tokyo.ac.jp

Satoshi Ohshima
The University of Tokyo
Information Technology Center
ohshima@cc.u-tokyo.ac.jp

MS153

Parameter Auto-tuning for a Contour Integral based Eigensolver using Stochastic Estimation of Eigenvalue Count

We consider a parallel eigensolver based on contour integral. The method requires several parameters that are related to the number of eigenvalues around the contour path. To improve the performance of the eigensolver, we present a method that selects suitable parameters in the eigensolver by using a stochastic estimation of eigenvalue count in a domain on the complex plane. The efficiency of the presented method is demonstrated by some numerical experiments.

Yasuyuki Maeda, Yasunori Futamura
University of Tsukuba
maeda@mma.cs.tsukuba.ac.jp,
futamura@mma.cs.tsukuba.ac.jp

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

MS153

Evaluation of Genetic Algorithm on Initial Vector Settings for GMRES

GMRES(m) solver is very popular for sparse matrix computations in science and engineerings, but, performs very poorly with some badly selected values of initial vectors and Krylov dimensions. This paper evaluates our newly proposed GA-based Automatic Tuning for GMRES(m) in size of populations of initial vectors. Experiments on 20 largest unsymmetric matrices in Florida collection show that the GA of initial vectors with ten populations is quite effective in convergence as well as predictive parallel performance.

Ken Naono
R&D department
Hitachi Asia Malaysia, Malaysia
knaono@has.hitachi.com.my

Nordin Zakaria
HPC service center
Universiti Teknologi PETRONAS, Malaysia
nordinzakaria@petronas.com.my

Takao Sakurai
Central Research Laboratory
Hitachi Ltd.
takao.sakurai.ju@hitachi.com

AnindyaJyoti Pal
HPC service center, Universiti Teknologi PETRONAS
anindyajp@gmail.com

Nobutoshi Sagawa
R&D center, Hitachi Asia Ltd.
nsagawa@has.hitachi.com.sg

MS153

Energy-Aware Matrix Auto-tuning using Genetic Algorithm and the Xabclib

We consider a form of matrix auto-tuning that includes not only solver performance and stability but the energy performance as well. A simple power model is adopted that infers power consumption from runtime resource usage. The model is combined with conventional matrix solution performance measures into an aggregate objective score that is then used by a Genetic Algorithm, deployed together with the Xabclib matrix autotuning library, to search for the optimal solver parameters.

Nordin Zakaria

HPC service center
Universiti Teknologi PETRONAS, Malaysia
nordinzakaria@petronas.com.my

Ken Naono
R&D department
Hitachi Asia Malaysia, Malaysia
knaono@has.hitachi.com.my

Anindya Jyoti Pal
HPC service center
Universiti Teknologi PETRONAS
anind.jyotipal@petronas.com.my

Takao Sakurai
Central Research Laboratory
Hitachi Ltd.
takao.sakurai.ju@hitachi.com

MS154

IMEX Schemes for Hyperbolic Systems and Kinetic Equations with Diffusion Relaxation

We consider IMEX schemes for hyperbolic systems with stiffff relaxation in the so-called diffusion limit. In such regime the system relaxes towards a convection-diffusion equation. The first objective is to show that traditional partitioned IMEX schemes will relax to an explicit scheme for the limit equation with no need of modication of the original system. Of course the explicit scheme obtained in the limit suffers from the classical parabolic stability restriction on the time step. The main goal is to present an approach, based on IMEX schemes, that in the diffusion limit relaxes to an IMEX scheme for the convection-di ffusion equation, in which the diffusion is treated implicitly. This is achieved by a novel reformulation of the problem, and subsequent application of IMEX schemes to it. Several numerical examples including neutron transport equations confirm the theoretical analysis.

Sebastiano Boscarino
University of Catania
boscarino@dmi.unict.it

MS154

Explicit Time Stepping for Radiative Transfer based on Mixed Variational Formulations

Based on a mixed variational formulation, we derive an explicit time stepping scheme of leap-frog type for PN-FEM discretizations of radiative transfer. We discuss stability aspects, the appropriate choice of the time-step, and the efficient implementation avoiding inversion of mass matrices. Numerical tests will be presented that illustrate the performance of the new algorithm on a few benchmark problems.

Herbert Egger
Numerical Analysis and Scientific Computing
Darmstadt University of Technology
herbert.eggerematma.tum.de

MS154

A New Spherical-Harmonics Scheme for Multi-Dimensional Radiation Transport

Recent work by [McClarren & Hauck, Robust and accurate filtered spherical harmonics expansions for radiative transfer. J. Comput. Phys., 229(16):5597−5614, Aug. 2010] suggests that the filtered spherical harmonics method represents an efficient, robust and accurate method for radiation transport, at least in the two-dimensional (2D) case. We extend their work to the three-dimensional (3D) case and find that all of the advantages of the filtering approach identified in 2D are present also in the 3D case. We reformulate the filter operation in a way that is independent of the timestep and of the spatial discretization. We also explore different second- and fourth-order filters and find that the second-order ones yield significantly better results. Overall, our findings suggest that the filtered spherical harmonics approach represents a very promising method for 3D radiation transport calculations.

David Radice
Max-Planck Institut für Gravitationsphysik
david.radice@aei.mpg.de

David Radice
Max-Planck Institut fuer Gravitationsphysik
david.radice@aei.mpg.de

Ernazar Abdikamalov
California Institute of Technology
abdik@tapir.caltech.edu

Luciano Rezzolla
Max-Planck Institut für Gravitationsphysik
rezzolla@aei.mpg.de

Christian Ott
Caltech
cott@tapir.caltech.edu

MS154

StaRMAP – A Staggered Grid Approach for Arbitrary Order Linear Moment Methods of Radiative Transfer

A simple method to solve the P_N equations of radiative transfer is presented. A specific coupling between the moments in the equations naturally induces a second-order accurate finite difference scheme on staggered grids. While the scheme does not possess limiters, its simplicity gives rise to a very efficient implementation in MATLAB. The code, which is available for download, can solve problems with ten million degrees of freedom in space, angle, and time within a few seconds.

Benjamin Seibold
Temple University
seibold@temple.edu

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

MS155

An Adaptive High Order Finite Element Scheme for the Monodomain Equation

Despite extensive research, obtaining accurate numerical solutions to the cardiac electrical propagation equations remains very time-consuming. Attempts have been made to reduce the computational burden using adaptive algorithms, but success has been limited due to the high costs associated with finite element matrix re-assembly. We present an adaptive algorithm based on high-order finite elements and controlled using an a posteriori error indicator which avoids this re-assembly problem and provides the potential for significant efficiency improvements.

Christopher Arthurs
Computing Laboratory
University of Oxford
chris.arthurs@comlab.ox.ac.uk

MS155

Efficiency Considerations for High-performance Computing with Applications to Cardiac Electrophysiology

We discuss techniques for improving performance of scientific codes with a particular emphasis on solving the reaction-diffusion equations associated with cardiac electrophysiology. Some of these methods involve programming choices that decrease memory usage and/or processing time. Others methods are architecture-dependent, including traditional parallelization and graphics processing unit-based implementations, where decisions regarding memory access patterns and problem decomposition are critical. We discuss the implementation of these various approaches and quantify the performance improvements they generate.

Elizabeth M. Cherry
Rochester Institute of Technology
School of Mathematical Sciences
excsma@rit.edu

MS155

Strongly Scalable Numerical Approaches for Modeling Coupled Cardiac Electro-mechanics at High Spatial Resolution

Biophysically detailed multiscale computer models of the heart are increasingly important in advancing our understanding of integrated cardiac function in health and disease. However, such detailed multiphysics simulations are computationally vastly demanding. In previous studies, using up to 16k cores we demonstrated strong scalability for solving the bidomain equations. More recently, we implemented a proper domain decomposition algebraic multigrid bidomain solver which can be compiled for execution on both CPUs and GPUs in distributed memory environments. In this study first results are presented on extending this solver framework for electro-mechanically coupled multi-physics simulations.

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

MS155

Efficient Computational Techniques for Modeling Electro-mechanical Interactions in the Heart

Computer simulations stand out as a promising path towards increased understanding of electro-mechanical interactions in the heart, and in particular to fully uncover the role of mechano-electric feedback in post infarct arrhythmias. However, the simulations are challenging to perform because of the complexity and strong non-linearity of the relevant mathematical models. In this talk we will address these computational issues, and propose a set of solution methods based on operator splitting techniques. The derived computational methods will then be applied to a model of infarct injured sheep hearts, to study the effect of MEF in the border zone surrounding the infarcted region.

Joakim Sundnes
Simula Research Laboratory
sundnes@simula.no

MS156

A Parallel and Dynamically Adaptive 3D Cubed-sphere Grid Framework for Hyperbolic Conservation Laws

A parallel block-adaptive simulation framework is described for hyperbolic conservation laws on 3D cubed-sphere grids. We use a multi-dimensional finite-volume approach, which naturally maintains uniform second-order accuracy at the sector boundaries and corners of the cubed-sphere grid without requiring special interpolation or reconstruction procedures. This simplifies implementation of parallelism and adaptivity in our hierarchical multi-block framework. Numerical tests demonstrate accuracy and efficiency of the approach and show excellent parallel scalability on thousands of computing cores.

Hans De Sterck, Lucian Ivan
University of Waterloo
Applied Mathematics
hdesterck@uwaterloo.ca, livan@uwaterloo.ca

Scott Northrup
University of Toronto
Institute for Aerospace Studies
northrup@scinet.utoronto.ca

Clinton P. Groth
University of Toronto Institute for Aerospace Studies
Canada
groth@utias.utoronto.ca

MS156

Grid Refinement in the GFDL High Resolution Atmosphere Model (HiRAM): Stretched and Nested Grid

Abstract not available at time of publication.

Lucas Harris
NOAA/GFDL
lucas.harris@noaa.gov

MS156

Variable Resolution Capabilities of the Community Atmosphere Model's Cubed-sphere Spectral-element Configuration

We will present results from high-resolution global atmospheric simulations using CAM-SE: The Community Atmosphere Model), running with the Spectral finite Element dynamical core from NCAR's High-Order Method Modeling Environment (HOMME). CAM-SE uses fully unstructured conforming quadrilateral grids, including cubed-sphere grids for global quasi-uniform resolution of the Earth's atmosphere. For global 1/4 and 1/8 degree resolutions, CAM-SE runs efficiently on hundreds of thousands of processors on modern Cray and IBM supercomputers and obtains excellent simulation throughput.

Mark A. Taylor
Sandia National Laboratories, Albuquerque, NM
mataylo@sandia.gov

Katherine J. Evans
Oak Ridge National Laboratory
evanskj@ornl.gov

Oksana Guba
Sandia National Laboratories
NM, USA.
onguba@sandia.gov

Peter H. Lauritzen
NCAR
CGD
pel@ucar.edu

Mike Levy, Rich Neale
NCAR
mlevy@ucar.edu, neale@ucar.edu

James Overfelt
Sandia National Laboratories
jroverf@sandia.gov

MS156

Utilizing Grid Refinement in the Cubed-sphere Spectral Element Option of CAM to Model Tropical Cyclones

We utilize the variable-resolution Spectral Element dynamical core of the NCAR Community Atmosphere Model (CAM-SE) on a cubed-sphere grid to forecast tropical cyclones. This setup permits high resolution in low-latitude ocean basins where tropical cyclones are prevalent while maintaining continuity within the remainder of the global domain. We present short-term forecasts of selected storms and compare performance to globally-uniform and limited area models currently utilized operationally for hurricane prediction.

Colin M. Zarzycki
University of Michigan
zarzycki@umich.edu

Christiane Jablonowski
University of Michigan
Ann Arbor MI 48109-2143
cjablono@umich.edu

MS157

The SORD Code for Rupture Dynamics

Simulations provide an important tool for investigating high-frequency seismic energy generation. We model spontaneous rupture within a 3D isotropic viscoelastic solid using a mimetic generalized finite-difference method. The relevant phenomena span many orders of magnitude in spatial scale, requiring high discretization resolution. Emerging petascale computational resources are making possible new applications of this method, such as a physics-based probabilistic seismic hazard analysis (PSHA) map for the entire state of California.

Geoffrey Ely
Argonne National Laboratory
gely@anl.gov

MS157

Earthquake Shaking from Rupture Dynamics and Seismic Wave Scattering

Predicting earthquake ground-motions in the frequency band of engineering interest (0-10 Hz) represents an important challenge in computational science. Realistic simulations rely on fundamental earthquake source physics, describing fault geometry and how frictional strength evolves during rupture, and seismic wave propagation through heterogeneous Earth structure. The generation of synthetic broadband seismograms requires state-of-the art computational resources and highly scalable and efficient numerical codes. The mini-symposium invites contributions that address this subject.

Walter Imperatori
KAUST
walter.imperatori@kaust.edu.sa

Christian Pelties
LMU Muenchen
pelties@geophysik.uni-muenchen.de

Martin Galis, Martin Mai
KAUST
martin.galis@kaust.edu.sa, martin.mai@kaust.edu.sa

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

MS157

The SeisSol Code: Efficient Implementation of the ADER-DG Method for Large-Scale Dynamic Earthquake Simulations

We will present a complete software solution for earthquake simulations based on the arbitrary high-order derivative Discontinuous Galerkin (ADER-DG) method. The implementation achieves excellent scalability on large-scale high-performance computing infrastructure using unstructured tetrahedral meshes. The presentation focuses on recent code optimizations, in particular the introduction of multi-threading and the efficient implementation of small matrix-matrix multiplications, as well as a large-scale dynamic earthquake rupture case study.

Christian Pelties
LMU Muenchen
pelties@geophysik.uni-muenchen.de

Alex Breuer
TU Muenchen
breuera@in.tum.de

Vipin Sachdeva
IBM Research, Austin
vsachde@us.ibm.com

Walter Imperatori
KAUST
walter.imperatori@kaust.edu.sa

Alice Gabriel

LMU Muenchen
gabriel@geophysik.uni-muenchen.de

Michael Bader
Technische Universität München
Department of Informatics
bader@in.tum.de

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

Martin Mai
KAUST
martin.mai@kaust.edu.sa

MS157

Construction of Models and Meshes for Large-Scale Earth Science Applications

The construction of domain definitions and meshes are critical steps in the execution of earth science simulations. This presentation will discuss on-going efforts on the extension of geometric model construction, meshing, and partitioning technologies, originally developed to work with CAD representations, to meet the needs of large-scale earth science simulations using well controlled unstructured meshes. Emphasis will be placed on issues associated with the construction of the non-manifold geometric model needed by the mesh generators. high-level commands.

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

Cameron Smith
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

MS158

Extracting Novel Insight from Probabilistic Machine-learned Classification Catalogs

As the size of astronomical light curve datasets grows, astronomers must become further removed from the inference workflow. But instead of labeling a certain source as definitely a member of one class, machine-learned classification produces class probability vectors. I discuss the calibration of probabilistic classification, using taxonometic information, with an eye towards the ultimate goal: producing novel scientific insight as a result the classification catalogs.

Joshua S. Bloom
University of California, Berkeley
jbloom@astro.berkeley.edu

MS158

Never Look Twice: Prefiltering Approaches for Dealing with Pesky Amounts of Biological Sequence Data

Raw sequencing data is noisy and redundant, but it's also annoyingly large and sometimes difficult to process efficiently. Rather than improving the performance of downstream applications, we have focused on efficient preprocessing and prefiltering algorithms that lay a foundation for streaming online compressive computing in sequence analysis. It turns out such algorithms work really well.

C. Titus Brown
Michigan State University
ctb@msu.edu

MS158

Electrical Signals in the Human Brain: Computational Challenges

Electrophysiological signals present a unique opportunity to record real-time activity from behaving human brains. While methods for recording large amounts of data have improved drastically, computational techniques for modeling and analyzing this data still leave much to be desired. Such improvements would have implications in studying mechanisms of neural communication and applications such as brain-computer interfaces. This talk will detail some leading questions in this area, and discuss challenges facing computational cognitive neuroscientists.

Chris Holdgraf
University of California, Berkeley
choldgraf@berkeley.edu

MS158

Managing Large Datasets and Computation Workflows with Python

Data collection and generation during the last decade has become increasingly easier – so much so that almost all scientific communities need new solutions for processing and managing large datasets. In particular, there exist two main problems: 1) the compute power of a single machine is not enough to complete tasks in a reasonable amount of time; and 2) the memory of a single machine is not large enough to load and process all the data. The recent progress of Python-based tools for Big Data offers users new and novel methods of managing data workflows – tools which I believe elevate the novice and give power to the domain expert. From data storage and management to modeling and data analysis, these new tools and concepts can help individuals, small teams, and large organizations manage the Big Data deluge.

Benjamin L. Zaitlen
Continuum Analytics
ben.zaitlen@continuum.io

MS159

Robust and Scalable Strategies for Coupling Biogeochemical Reaction and Solute Transport in Earth System Modeling Frameworks

The Earth fosters a diverse set of physical, chemical, and biological processes that operate over a range of spatial and temporal scales from nanometers to kilometers and microseconds to millennia. The coupling of these complex and often uncertain processes can be challenging and imprecise within numerical models. This presentation illustrates coupling strategies within the petascale reactive flow and biogeochemical transport code PFLOTRAN and demonstrates the accuracy and parallel performance on real-world Earth system problems.

Glenn Hammond
Pacific Northwest National Laboratory
glenn.hammond@pnnl.gov

MS159

Moment-based Scale-bridging Algorithms for Multi-physics Kinetic Systems on Emerging Architectures

The advent of modern parallel architectures, and the promise of exascale computing resources, brings new challenges to computational science. Dependably taking advantage of billion-way parallelism and distinct levels of parallelism will stress solver algorithms. We are developing a moment-based multi-physics scale-bridging solver algorithm to help meet some of this challenge. Development of similar moment-based acceleration methods can be found in a variety of application areas which are described by multi-physics kinetic systems. We will provide algorithm fundamentals, specific application results and open research questions. We will provide some indication of algorithm performance on emerging architectures.

Dana Knoll, Luis Chacon
Los Alamos National Laboratory
nol@lanl.gov, chacon@lanl.gov

G. Chen, C. Newman, H. Park, J. Payne, R. Rauenzahn
LANL
gcc@ornl.gov, cnewman@lanl.gov, hkpark@lanl.gov, payne@lanl.gov, rick@lanl.gov

W. Taitano
L
taitano@lanl.gov

J Willert, C.T. Kelley
NCSU
jawiller@ncsu.edu, ctk@ncsu.edu

MS159

Multiphysics Coupling Tools Applied to Large-scale Simulations of a Light Water Nuclear Reactor Core

Obtaining robust and efficient solutions for multiple-time-scale multi-physics systems is challenging. This talk will describe a general toolkit under development for use in coupling software codes on HPC platforms. We will discuss solution algorithms (including Picard iteration, Jacobian-free Newton-Krylov, and Fully coupled Newton methods), application interfaces, and data layout/transfer requirements in the context of HPC platforms. We will show results for a neutronics/CFD coupling for simulating an assembly in a light water nuclear reactor.

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Roscoe Bartlett
Oak Ridge National Laboratory
bartlettra@ornl.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

MS159

Managing Complexity in Multi-physics Calculations on Modern and Emerging HPC Architectures

Complexity in HPC software stems from two primary sources: complexity in the physics and deployment on heterogeneous and changing architectures. . Different modeling choices imply different coupling between components in software; a reflection of the different physical processes being described. Trying to support large, evolving code bases on evolving architectures can lead to major software maintenance challenges. We describe two abstractions that allow these challenges to be addressed in a scalable manner without imposing abstraction penalties. First, we employ task graphs to represent complex multiphysics problems, enabling the programmer to focus on writing individual pieces of software by removing challenges associated with integrating that into the code base by scheduling computation, memory management, etc. automatically. Secondly, we introduce a domain-specific language embedded in C++ that facilitates rapid code development and alleviates the need to maintain versions of code for execution on serial, multithreaded, or GPU architectures.

James C. Sutherland
University of Utah
James.Sutherland@utah.edu

MS160

A Semismooth Newton-CG Method for Full Waveform Seismic Inversion with Parameter Constraints

We present a semismooth Newton-PCG method for full waveform seismic inversion that uses a Moreau-Yosida regularization to handle box constraints on the material parameters and a trust-region globalization. The matrix-free implementation relies on adjoint-based gradient and Hessian-vector computations and parallelization with MPI. We analyze the proposed optimization method in a function space setting. Numerical results are shown for an application to marine geophysical exploration.

Christian Boehm
Technische Universitaet Muenchen
boehm@ma.tum.de

Michael Ulbrich
Technische Universitaet Muenchen
Chair of Mathematical Optimization
mulbrich@ma.tum.de

MS160

Performance and Real Data Application of 3D Frequency-domain Full-waveform Implementations: Frequency-domain Direct-solver versus Time-domain-based Modeling

3D frequency-domain solutions of the wave-equation for full waveform inversion (FWI) can be computed with several methods. We implemented two massively-parallel algorithms including the direct-solver and the time-domain modeling coupled with a discrete Fourier transform for 3D acoustic FWI. We applied the two approaches to ocean-bottom-cable data from the Valhall field. We discuss the performances, pros and cons of both methods, which depend on the wave physics, the acquisition geometry and the computing architectures.

Romain Brossier
ISTerre
University Joseph Fourier
romain.brosssier@obs.ujf-grenoble.fr

Vincent Etienne
GEOAZUR
University Nice - CNRS
etienne@geoazur.unice.fr

Guanghui Hu
ISTerre
University Joseph Fourier
guanghui.hu@ujf-grenoble.fr

Stephane Operto
GEOAZUR
CNRS
operto@geoazur.obs-vlfr.fr

Jean Virieux
ISTerre
University Joseph Fourier
jean.virieux@ujf-grenoble.fr

MS160

Multi-level Multi-frequency Full Waveform Inversion: Convergence and Computation

Abstract not available at time of publication.

Maarten de Hoop
Center for Computational & Applied Mathematics
Purdue University
mdehoop@math.purdue.edu

MS161

Efficient Methods for Computing and Estimating Probability of Failure in PDE Systems

We present an algorithm for computing probability of failure in systems governing by partial differential equations. Since in many practical systems, simulations of the PDE systems can be highly computationally intensive, it is essential to avoid sampling the PDE many times. Our approach is based on the availability of low-fidelity model of the system. The low-fidelity model is of low accuracy but can be simulated fast. One example of low-fidelity model is coarse grid computation for the PDE system. We present a method to use a large number of samples of the low-fidelity model as a predictor of the failure probability, and then use a small number of samples of the original PDE system as a corrector. The final result of the failure probability can be of high accuracy, and induces much reduced simulation cost.

Jing Li, Dongbin Xiu
Purdue University
li176@purdue.edu, dxiu@purdue.edu

MS161

Nonintrusive Polynomial Chaos on Nested Unstructured Meshes

We present a novel method for construction of polynomial interpolation grids on arbitrary Euclidean geometries on which there is a probability density. These grids are eminently suitable for uncertainty quantification: they are constructed using the standard polynomial Chaos basis, they are nested so that refinement strategies can be employed, they are applicable for high-dimensional spaces. These grids do not suffer computational curse-of-dimensionality-type restrictions: the number of nodes in the grid is arbitrary and user-specified.

Akil Narayan
University of Massachusetts Dartmouth
akil.narayan@umassd.edu

MS161

QMC Lattice Methods for PDE with Random Coefficients

This talk presents recent developments in applying Quasi-Monte Carlo methods (specifically lattice methods) to elliptic PDE with random coefficients. A guiding example is the flow of liquid through a porous material, with the permeability modeled as a high-dimensional random field with log- normal covariance. In this work we apply the theory of lattice methods in weighted spaces to design lattice rules with good convergence properties for expected values of functionals of the solution of the PDE.

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

James Nichols
University of New South Wales, Sydney
james.ashton.nichols@gmail.com

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

Ivan G. Graham
University of Bath
I.G.Graham@bath.ac.uk

Robert Scheichl
University of Bath
Department of Mathematical Sciences
masrs@maths.bath.ac.uk

MS161

Sufficient Model Reduction on High-dimensional Stochastic Input in Uncertainty Quantification

In order to overcome the curse of dimensionality in high-dimensional stochastic systems, we combine conventional model reduction techniques, such as principle component analysis (PCA) and manifold learning, that focus on representing stochastic input in a lower-dimensional space with sufficient model reduction (SMR) methods, such as sliced inverse regression and kernel dimension reduction for supervised learning, that seek sufficient reduction of predictors while preserving the information on responses. In this way, a high-dimensional stochastic input could be reduced to the minimum number of random variables that can accurately predict particular responses of interest.

Nicholas Zabaras, Jiang Wan
Cornell University
zabaras@cornell.edu, jw688@cornell.edu

MS162
Design Issues for Manycore Application-library Interfaces

Although there has been much focus on algorithms and solvers for scalable manycore systems, less effort has been applied to understanding the requirements and possibilities for the application-library interface. This prioritization was reasonable, given the critical need for new algorithms and the relatively large amount of time spent in the solvers. However, eventually we need to have a scalable manycore interface as well. In this presentation we discuss some of the issues and possibilities for application-library interfaces on scalable manycore systems, focusing on different ways to think about the roles and responsibilities of each component.

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS162
Porting and Optimizing a Large-Scale CFD Application with CUDA and OpenACC

Computational Fluid Dynamics (CFD) applications are one of the most important applications executed with high-speed supercomputers. Especially, GPU-based supercomputers have been showing remarkable performance of CFD applications. However, GPU-programing is still difficult to obtain high performance, which prevents legacy applications from being ported to GPU environments. We apply CUDA and OpenACC, which are programming language for GPU with different features, to a Large-Scale CFD application UPACS and evaluate these performance.

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

MS162
Multithreading API in the PLASMA Library

The PLASMA numerical library relies on a variety of multithreading mechanisms, including static and dynamic thread scheduling. PLASMA's superscalar scheduler, QUARK, offers powerfull tools for parallel task composition, such as support for nested parallelism and provisions for task aggregation. The dynamic nature of PLASMA's internal mode of operation, exposes its user, the application developer, to an array of new capabilities, such as asynchronous mode of operation, where library functions calls act like non-blocking MPI communications, i.e. return before work completion. This talk will discuss new opportunities and difficulties steming from the asynchronous design paradigm.

Jakub Kurzak
University of Tennessee Knoxville
kurzak@eecs.utk.edu

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

Jack Dongarra
University of Tennessee
dongarra@eecs.utk.edu

MS162
Implementation of FEM Application on GPU with StarPU

We are aiming at accelerating FEM application on multicore CPUs with GPU. While it isn't easy to obtain very good performance because FEM application requires many non-sequential memory accesses, our GPU implementation has obtained better performance than multicore CPUs. We are now trying to accelerate FEM with StarPU runtime system. In this talk, we talk about the implementation and performance of it.

Satoshi Ohshima
The University of Tokyo
Information Technology Center
ohshima@cc.u-tokyo.ac.jp

Takahiro Katagiri
Information Technology Center, The University of Tokyo
katagiri@cc.u-tokyo.ac.jp

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

Samuel Thibault, Raymond Namyst
University of Bordeaux 1, Inria
samuel.thibault@labri.fr, raymond.namyst@labri.fr

MS163
Approximate Preconditioners for the Unsteady Navier-Stokes Equations and Applications to Hemodynamics Simulations

We present scalable preconditioners based on approximate versions of efficient preconditioners for the Navier–Stokes equations, namely the Pressure Convection-Diffusion preconditioner, Yosida, and SIMPLE. We exploit factorizations of the linearized system where inverses are handled using specific embedded preconditioners. Weak and strong scalability results illustrate this approach using benchmarks relevant to hemodynamics simulations. All the computations are carried out using the open source finite element library LifeV based on Trilinos.

Gwenol Grandperrin
EPFL Lausanne
CMCS - MATHICSE
gwenol.grandperrin@epfl.ch

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

Simone Deparis
Ecole Politechnique Federale de Lausanne
simone.deparis@epfl.ch

MS163

An Incompressible Viscous Flow Finite Element Solver for 1D-3D Coupled Fluid Models

The talk adresses an application of conforming finite element method (FEM) for a 1D–3D coupled incompressible flow problem. Coupling conditions are introduced to ensure a suitable bound for the cumulative energy of the model. Motivated by the simulation of flow over an inferior vena cava filter, we consider the coupling of a 1D graph and a 3D flow domain with highly anisotropic inclusions. The tetrahedra grid is locally refined and highly anisotropic. We study the stability and the accuracy of the discretization method and the performance of some state-of-the-art linear algebraic solvers for such flow configurations.

Maxim A. Olshanskii
Department of Mathematics
University of Houston
molshan@math.uh.edu

Tatiana Dobroserdova
Moscow State University M.V.Lomonosov
zdrapa@mail.ru

MS163

PALADINS: Scalable Time-adaptive Algebraic Splitting and Preconditioners for the Navier-Stokes Equations

We consider a class of second and third order time-accurate algebaric splitting schemes for the incompressible Navier-Stokes equations featuring a hierarchical structure prone to time-adaptivity. These schemes can be used both as solver or as preconditioner. We discuss some properties and technical details on the implementation of these scheme, called "PALADINS" (Parallel ALgebraic ADaptive Incompressible Navier-Stokes). We present scalability results and 3D applications to computational hemodynamics.

Umberto E. Villa
Dept. of Mathematics and Computer Science
Emory University
uvilla@emory.edu

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

MS163

Computational Algorithms for Stability Analysis of Incompressible Flows

Linear stability analysis of large-scale dynamical systems requires the computation of the rightmost eigenvalues of a series of large generalized eigenvalue problems. Using the incompressible Navier-Stokes as an example, we show that this can be done efficiently using a new eigenvalue iteration constructed using an idea of Meerbergen and Spence that reformulates the problem into one with a Lyapunov structure. This method requires solution of Lyapunov equations, which in turn entail solutions of subproblems consisting of linearized Navier-Stokes systems. We describe the methodology and demonstrate efficient iterative solution of the Navier-Stokes systems.

Howard C. Elman
University of Maryland, College Park
elman@cs.umd.edu

Minghao Wu
Applied Mathematics Program
University of Maryland
mwu@math.umd.edu

MS164

POD-Based Reduced-Order Models for Boussinesq Equations

Model reduction promises to have a significant impact on the design and control of energy efficient buildings. The standard proper-orthogonal decomposition (POD) combined with Galerkin projection is not sufficient to produce accurate representations of the airflow inside buildings at realistic parameter values (modeled by the Boussinesq equations). Therefore, we propose an LES-based closure model to reproduce the dissipation effects of the neglected modes in nonlinear terms appearing in the momentum and energy equations. Three-dimensional airflow simulations will be presented.

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

MS164

Window Proper Orthogonal Decomposition for Continuum and Atomistic Flow Simulations

Proper Orthogonal Decomposition (POD) is a spectral analysis tool often employed for data postprocessing and predictive modeling. Here we present a window POD (WPOD, Grinberg et al. ABME 2009) and demonstrate its utility in analysis of flow fields. We review the methodology and demonstrate application of the WPOD for analysis of transient and space-time intermittent flow regimes. We demonstrate that the time-varying POD eigenvalue spectrum provides an indication of a turbulent or laminar regime.

Leopold Grinberg
Brown University
lgrinb@dam.brown.edu

Alexander Yakhot
Ben Gurion University
yakhot@bgu.ac.il

George E. Karniadakis

Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS164

Reduced Basis Methods for Viscous Flows: Application to Inverse Problems in Haemodynamics

We review the current state of the art of the Reduced Basis method for Stokes and Navier-Stokes equations in parametrized geometries, and related a posteriori error estimation. We present a fully decoupled Offline/Online procedure for both reduced spaces construction and error bounds evaluation, as well as a general framework for the efficient solution of inverse problems. Furthermore, we show some numerical results related to applications of interest in haemodynamics.

Andrea Manzoni
International School for Advanced Studies, Trieste, Italy
amanzoni@sissa.it

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

MS164

Space-time Error Bounds for Reduced-order Approximations of Parametrized Boussinesq Equations

We present a space-time certified reduced basis method for the multi-parameter unsteady Boussinesq equations. We combine a space-time Petrov-Galerkin variational formulation, discontinuous Galerkin temporal discretization, Brezzi-Rappaz-Raviart *a posteriori* error bounds, and *hp*-adaptive greedy sampling to provide rapid and certified characterization of the flow over a wide range of parameters as well as long integration times. The space-time stability constant constitutes a single quantifiable measure of flow stability which may serve for example in uncertainty quantification.

Masayuki Yano
Massachusetts Institute of Technology, Cambridge, MA, US
myano@mit.edu

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

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

MS165

Stable Parareal in Time Method for First and Second Order Hyperbolic System

The parareal in time algorithm has been implemented on many types of time dependent problems with some success. However, some difficulty occurs for problems where the solution lacks regularity, either initially or during the evolution, as it is the case for hyperbolic system. In this talk we propose some way to cure it. We use the new method to solve a linear wave equation and a nonlinear Burger's equation, the results illustrate the stability of this variant of the parareal in time algorithm.

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

Yvon Maday
Université Pierre et Marie Curie
and Brown university
maday@ann.jussieu.fr

MS165

Coarse Grid Correction for the Neumann-Neumann Waveform Relaxation Method

The Neumann-Neumann waveform relaxation (NNWR) method for the heat equation converges superlinearly for finite time windows; however, the number of iterations required for convergence to a fixed tolerance increases quadratically in the number of subdomains. We show that by adding a coarse grid correction step, the modified method converges in two iterations for 1D problems, independent of the number of subdomains. We also analyze its convergence rate for the 2D case.

Felix Kwok
University of Geneva
Section de Mathématiques
felixk80@gmail.com

MS165

Comparing Implementation Strategies for Parareal with Spatial Parallelization

We explore and compare different implementation strategies for a combination of Parareal with spatial parallelism. One is a hybrid approach, using shared memory in time with distributed memory in space, the other a fully MPI-based approach, using distributed memory and message passing in space as well as time. Shared memory in time avoids messaging volume data and reduces the memory footprint of the code, but combining it with message passing in space is not straightforward.

Daniel Ruprecht, Rolf Krause
Institute of Computational Science
University of Lugano
daniel.ruprecht@usi.ch, rolf.krause@usi.ch

MS165

A Large-Scale Space-Time Multilevel Solver for the 3D Heat Equation

We present results on the coupling of a space-parallel parabolic multigrid solver with the Parallel Full Approximation Scheme in Space and Time (PFASST). By intertwining sweeps of spectral deferred corrections with Parareal iterations, PFASST allows to obtain parallelized time propagators with arbitrary order. We compare the performance of on- and inter-node parallelization in time and demonstrate the capabilities of these approaches using large-scale benchmarks on a recent IBM Blue Gene/Q installation.

Robert Speck
Jülich Supercomputing Centre
Forschungszentrum Jülich
r.speck@fz-juelich.de

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

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

MS166

Solving the Non-linear Eigenvector Problem - FEAST-based Alternative to Self-consistent Field Methods

The non-linear eigenvector problem of the form $\mathbf{A}[\mathbf{X}]\mathbf{X} = \lambda\mathbf{B}\mathbf{X}$, such as the one arising in electronic structure calculations, is commonly addressed using a self-consistent field method (SCF) where a series of linear eigenvalue problems need to be solved iteratively until convergence. Here, we present an extension of the symmetric FEAST algorithm for directly addressing the nonlinear eigenvector problem. We show that the new strategy offers a robust and efficient alternative to the traditional SCF methods.

Brendan Gavin
ECE Department
University of Massachusetts, Amherst
bgavin@ecs.umass.edu

Eric Polizzi
University of Massachusetts, Amherst, USA
polizzi@ecs.umass.edu

MS166

FEAST for the Non-Symmetric Eigenvalue Problem

The FEAST algorithm is extended to address the non-symmetric eigenvalue problem $\mathbf{AX} = \lambda\mathbf{BX}$ where A and B are real non-symmetric or complex non-Hermitian matrices. In order to obtain left and right eigenvectors, we present a dual subspace strategy for performing the contour integration on the Green's function in a given region of the complex plane. All desirable features of the original algorithm including robustness, scalability and parallelism are retained.

James Kestyn
ECE Department
University of Massachusetts, Amherst
jkestyn@student.umass.edu

Eric Polizzi
University of Massachusetts, Amherst, USA
polizzi@ecs.umass.edu

MS166

Parallel Solution of Sparse Rank-deficient Linear Systems

The talk is devoted to a highly parallel method for solving rank-deficient sparse linear least-square problems on multi-core platforms. The proposed method is based on the truncated SVD approach. The algorithm computes an orthonormal basis of the kernel using the FEAST Eigenvalue Solver first and then solves an extended system of linear equations. Intel® MKL PARDISO is used to solve resulting well-conditioned systems.

Sergey Kuznetsov
Intel Corporation
sergey.v.kuznetsov@intel.com

MS166

Subspace Iteration + Approximate Spectral Projection = FEAST

The FEAST algorithm for Hermitian generalized eigenproblem is based on a "density-matrix' approach. The FEAST software is well received as it exploits multiple levels of parallelism that are available on modern architectures. This presentation provides some theoretical analysis of the FEAST algorithm that is hitherto missing. The point of view is that FEAST is carrying out subspace iteration with Rayleigh-Ritz procedure on an approximate spectral projector.

Ping T. Tang
Intel Corporation
Peter.Tang@intel.com

MS167

Tree Decompositions in Logical and Probabilistic Inference

Inference is the problem of answering questions from knowledge and observations represented in a formal language or model. Research about efficient computational inference procedures applies insights and assumptions about logical languages and probabilistic models. The tree-decomposition abstraction enables structured and fast inference by detecting an underlying tree structure in the knowledge or model. In this talk I will present results, insights, applications, and future promise of tree decomposition in inference and decision making.

Eyal Amir
University of Illinois, Urbana-Champaign
eyal@illinois.edu

MS167

Supernodal and Multifrontal Sparse Matrix Factorization

For peak performance, sparse direct methods rely on supernodal and multifrontal techniques, which exploit the cliques that form during factorization. These cliques are related via the supernodal elimination tree which has the same properties as trees used in tree decomposition methods for dynamic programming problems on graphs. This talk presents supernodal trees, how their treewidth can be computed efficiently, and the role they play in sparse direct methods, including multifrontal methods on the GPU.

Timothy A. Davis
University of Florida
Computer and Information Science and Engineering

DrTimothyAldenDavis@gmail.com

MS167

Efficient Algorithms from Graph Structure Theory: Minors, Bidimensionality, and Decomposition

Graph minors, treewidth, and other graph structure theory enable new approximation and fixed-parameter algorithms for graph optimization problems. This talk describes two recent approaches. Bidimensionality theory provides general tools for designing fast (constructive, often subexponential) fixed-parameter algorithms, kernelizations, and approximation algorithms (often PTASs), for a wide variety of NP-hard graph problems for graphs excluding a fixed minor. Simplifying decompositions split such a graph into few small-treewidth graphs via deletions or contractions, also leading to PTASs.

Erik Demaine
MIT
edemaine@mit.edu

MS167

Tree Decompositions: Adapting Algorithms for Parallel Computation

Although many NP-hard graph optimization problems can be solved in polynomial time on graphs of bounded treewidth, the adoption of these techniques into mainstream scientific computation has been limited due to the high memory requirements of required dynamic programming tables and excessive running times of sequential codes. We discuss how INDDGO, an open source OpenMP/MPI software suite, addresses some of these challenges through algorithms and implementations, and what impediments still remain.

Blair D. Sullivan
Oak Ridge National Laboratory
sullivanb@ornl.gov

MS168

QR Factorizations and SVDs for Tall-and-skinny Matrices in MapReduce Architectures

The QR factorization and the SVD are two fundamental matrix decompositions with applications throughout scientific computing and data analysis. Matrices with many more rows than columns, so-called "tall-and-skinny matrices," are of particular interest. We describe how to compute a tall-and-skinny QR factorization on MapReduce architectures with four different algorithms, and our direct TSQR method is the only fast and numerically stable option. We extend the algorithms to compute the SVD with no performance changes. Advisors: James Demmel, UC Berkeley and David Gleich, Purdue University

Austin Benson
UC Berkeley
arbenson@stanford.edu

MS168

Restoration and Analysis of Apollo Lunar Data

The Lunar Ejecta and Meteorites Experiment was designed to measure particles that collide with the surface of the Moon. Original analyses revealed more activity during the passage of the terminator. Recently, there has been speculation that noise from other instruments may have influenced the results. The current analysis examines the data to see if there are correlations between impacts and electrical events. Analyses have revealed unusual patterns in all three sensors of the instrument. Advisor: David Williams, NASA Goddard Space Flight Center

Missy Gaddy
Wofford College
gaddymr@email.wofford.edu

MS168

Visualization of Cardiac Simulations Using Amira

As the incidence of heart disease in many developed countries continues to rise, research involving computational models of the heart is becoming increasingly important. Various software can be used to visualize and run simulations on cardiac models, with the purpose of better understanding, diagnosing, and treating heart problems. This research looked at software for segmenting CT and MRI images to generate 3D cardiac meshes, which were used to run electrical and other simulations. Advisors: Rodrigo Weber dos Santos, Angela Shiflet

Jonathan Hanson
Federal University of Juiz de Fora, Brazil
hansonjt@email.wofford.edu

MS168

Water Quality Monitoring of Maryland's Tidal Waterways

The Maryland Department of Natural Resources monitors the Chesapeake Bay and its tributaries with monitoring stations located throughout the tidal waterways. The status of the stations is assessed using the Wilcoxon Signed Rank Test. Our simulations demonstrated that log-transforming the data helped to reduce the Type I Error. We ranked the stations using a set of multiple comparison methods, including the Benjamini-Hochberg rejection method. An interactive GUI is created to facilitate this data analysis. Advisor: Nagaraj K. Neerchal

Rosemary Le
Brown University
rosemary_le@brown.edu

MS169

Shock Capturing for High-Order Discontinuous Galerkin Simulation of Transient Flow Problems

We present some recent result for shock capturing using high-order discontinuous Galerkin schemes. We extend the sensor-based artificial viscosity approach to transient problems, where we demonstrate how weakly coupled sensors give a simpler scheme, higher performance, and a more robust behavior. We also show how different levels of smoothness in the sensor affect the solutions, for transonic and supersonic flow problems in 2D and 3D.

Per-Olof Persson
University of California Berkeley
Dept. of Mathematics
persson@berkeley.edu

MS169

H-to-P Efficiently: Achieving Scalable Perfor-

mance using Hybrid Parallelism and Matrix Coalescence

Spectral/hp element methods utilise a high-order piecewise discretisation which permits both geometric flexibility and exponential convergence properties to be attained simultaneously. Some fluid flow problems have one or more geometrically homogeneous coordinate directions and computation can be significantly reduced by using a spectral discretisation in these directions. In this talk we explore how hybrid parallelisation of spectral-spectral/hp element discretisations can be optimised for a given problem and parallel environment. Additionally, elemental matrix operations typically achieve poor performance on tuned BLAS implementations due to their small size. We therefore explore how the performance of each process can be further improved through intelligent matrix coalescence.

Spencer Sherwin, Chris Cantwell
Imperial College London
s.sherwin@imperial.ac.uk, c.cantwell@imperial.ac.uk

David Moxey
Department of Aeronautics
Imperial College London
d.moxey@imperial.ac.uk

Robert Kirby
University of Utah
kirby@sci.utah.edu

MS169

Recent Development of a High Order Riemann-Solver-Free Space-Time Discontinuous Galerkin Method

In this talk, new development of a high-order Riemann-solver-free space-time discontinuous Galerkin method will be reported. The alternate cell-face solution updating strategy is employed to replace the original cell-vertex scheme for the convenience of the boundary condition treatment. The resulting discontinuous Galerkin cell-face scheme (DG-CFS) preserves the original methods most important features including (i) high-order in both space and time; (ii) Riemann-solver-free approach and (iii) working for both advection and diffusion equations.

Shuangzhang Tu
Jackson State University
shuangzhang.tu@jsums.edu

MS169

Adjoint-Based Mesh Adaptation for a Class of High-Order Hybridized Finite-Element Schemes for Convection-Diffusion Problems

We present a goal-oriented mesh-adaptation methodology, building on a class of hybridized finite-element schemes for (nonlinear) convection-diffusion problems. Using a discrete-adjoint approach, sensitivities with respect to output functionals of interest are computed to drive the adaptation. The theoretical framework is embedded in a unified formulation of a large class of hybridized, adjoint consistent schemes. Furthermore, a shock-capturing methodology is incorporated, enabling applications in high-speed compressible flow simulation.

Michael Woopen, Aravind Balan, Georg May
AICES
RWTH Aachen
woopen@aices.rwth-aachen.de,
balan@aices.rwth-aachen.de, may@aices.rwth-aachen.de

Jochen Schütz
IGPM
RWTH Aachen
schuetz@igpm.rwth-aachen.de

MS170

Asymptotically Exact DG Error Estimates for Convection Problems on Tetrahedral Meshes

We present several superconvergence results and asymptotically exact a posteriori estimates for discontinuous Galerkin methods applied to convection and convection-diffusion problems. We derive new superconvergence results with accurate error estimates for three-dimensional hyperbolic problems on tetrahedral meshes. On each element, the asymptotic behavior of DG errors depends on the mesh orientation with respect to the problem characteristics. Thus, elements are classified according to the number of inflow and outflow faces and in all cases enriched finite elements spaces are needed to show pointwise superconvergence. Numerical results are presented to validate the theory and show the behavior of our error estimates when the theory does not apply near shocks and discontinuities.

Slimane Adjerid
Department of Mathematics
Virginia Tech
adjerids@math.vt.edu

Idir Mechai
Department of Mathematics, Virginia Tech
mechaii@vt.edu

MS170

A Posteriori Error Estimates for an LDG Method Applied to Transient Convection-diffusion Problems in One Space Dimension

In this talk, new *a posteriori* error estimates for a local discontinuous Galerkin (LDG) formulation applied to transient convection-diffusion problems in one space dimension are presented and analyzed. These *a posteriori* LDG error estimates are computationally simple and are computed by solving a local steady problem with no boundary conditions on each element of the mesh. We first show that the leading error term on each element for the solution is proportional to a $(p+1)$-degree right Radau polynomial while the leading error term for the solution's derivative is proportional to a $(p+1)$-degree left Radau polynomial. These results are used to prove that, for smooth solutions, these error estimates at a fixed time converge to the true spatial errors in the L^2-norm under mesh refinement. More precisely, we show that our LDG error estimates converge to the true spatial errors at $\mathcal{O}(\langle\surd^{+\triangledown/\triangle})$ rate. Finally, we prove that the global effectivity indices in the L^2-norm converge to unity at $\mathcal{O}(\langle^{\infty/\in})$ rate. Our computational results indicate that the observed numerical convergence rates are higher than the theoretical rates.

Mahboub Baccouch
Department of Mathematics
University of Nebraska at Omaha
baccouch@gmail.com

MS170

Adjoint Based Error Estimation for the Lax-Wendroff Method

Hyperbolic problems are of interest in many research areas. A popular class of methods for solving them are finite difference methods. While these methods have been shown to be quite effective, goal oriented error analysis is difficult to perform due to the fact that the method is not in variational form. In this work we show equivalency between a finite element method and the Lax-Wendroff method, and present an adjoint based error representation formula..

James Collins, Don Estep, Simon Tavener
Colorado State University
jbcolli2@gmail.com, estep@stat.colostate.edu, tavener@math.colostate.edu

MS170

Finite Volume Adjoint Error Estimates for Weak Solutions

Many codes exist for hyperbolic equations that employ finite volume methods, which are often nonlinear. A method is presented for adjoint-based *a posteriori* error calculations with finite volume methods. We demonstrate the flexibility in implementation of this post-processing technique and demonstrate the accuracy with smooth and discontinuous solutions.

Jeffrey M. Connors
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
connors4@llnl.gov

MS171

An IDE Integrated Cross-Platform Build System for Scientific Applications

The goal of our work is to support the development of scientific and engineering applications while keeping its performance portable. To this end, we need to estimate performances of individual applications on various systems and detect their non-portable parts based on the estimation. Therefore, in this work, we have developed an IDE integrated cross-platform tool to automatically build applications on remote HPC systems. By enabling the applications to run on various systems, we can improve their functional portabilities, and also expect to improve their performance portabilities.

Shoichi Hirasawa
Tohoku Univ./JST CREST
hirasawa@sc.isc.tohoku.ac.jp

Hiroyuki Takizawa
Tohoku University/JST CREST
tacky@isc.tohoku.ac.jp

Hiroaki Kobayashi
Tohoku University, Japan
koba@isc.tohoku.ac.jp

MS171

An Autotuning Framework for Adapting OpenCL Kernels to Diverse Architectures

OpenCL promises to enable source code portability across vendor architectures. However, compiling OpenCL kernels for high-performance execution on diverse hardware is challenging. The widely varying forms of parallelism and memory organizations create a highly non-linear search space for compiler transformations. In particular, numerous combinations of transformations need to be evaluated either explicitly or implicitly in order to settle on the final configuration for a hardware target. In order to address this challenge, we have developed an autotuning framework for OpenCL programs. It takes an OpenCL kernel as an input and evaluates combinations of compiler transformations. We will report the cost and benefit of using this autotuning framework in compiling a wide variety of OpenCL kernels for different OpenCL devices.

Hee-Seok Kim
University of Illinois at Urbana-Champaign and MulticoreWare
kim868@illinois.edu

Matthieu Delahaye
MulticoreWare
matthieu@multicorewareinc.com

Wen-Mei Hwu
University of Illinois at Urbana-Champaign and MulticoreWare
w-hwu@illinois.edu

MS171

Assessing Library Performance with TAU

Autotuning technologies have been successful in concealing the complex decision-making process and restructuring for obtaining near optimal instances of kernels or libraries on a given platform. Yet, assessing the performance portability of kernels or libraries on existing and upcoming platforms may require the identification of additional parameters, and testing the performance characteristics of those. In this presentation we will discuss how performance analysis utilities can be employed to search for complementary opportunities for improving performance.

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

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

MS171

A Cost-Efficient Approach for Automatic Algorithm Selection of Collective Communications

As the size and the complexity of computers increase, selection and tuning the technologies for implementing communication libraries have become important issues. This talk introduces a method that selects a suitable algorithm of collective communications at runtime. In addition to the static information such as the number of processes and the size of message, this method also examines the runtime information such as relative locations of processes to choose an appropriate algorithm efficiently.

Takeshi Nanri
Kyushu University, Japan

JST CREST, Japan
nanri@cc.kyushu-u.ac.jp

Hironobu Sugiyama
Kyushu University
hirosugi3@i.softbank.jp

Keiichiro Fukazawa
Kyushu University, Japan
JST CREST
fukazawa@cc.kyushu-u.ac.jp

MS172

Scalable Algorithms for Real-Space and Real-Time First-Principle Simulations

The FEAST eigenvalue algorithm is presented beyond the "black-box" solver as a fundamental modeling framework for the electronic structure problem. Within this framework, the domain decomposition muffin-tin strategy can now be performed exactly for the all-electron DFT problem without resorting to traditional approximations such as as linearization and pseudopotential techniques. Additionally, the FEAST framework is also ideally suited for performing real-time TDDFT calculations using a direct spectral decomposition of time-ordered evolution operators.

Eric Polizzi
University of Massachusetts, Amherst, USA
polizzi@ecs.umass.edu

MS172

Polynomial Filtered Lanczos Algorithms and Spectrum Slicing

We present a polynomial filtering technique for extracting extreme or interior eigenvalues of large sparse matrices. This general approach can be quite efficient in the situation when a large number of eigenvalues is sought, as is the case in electronic structure calculations for example. However, its competitiveness depends critically on a good implementation. The method presented relies on a combination of the Lanczos algorithm with partial reorthogonalization and polynomial filtering based on least-squares polynomials. Details on implementation and a few numerical experiments will be presented.

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

Haw-ren Fang
Department of Computer Science & Engineering
University of Minnesota
hrfang@cs.umn.edu

MS172

Fast Inversion Methods for NEGF-based Simulation of Nanoelectronics Devices with Scattering

We will present numerical methods that allow fast and robust simulations of coupled electro-thermal transport in ultra-scaled nanostructures. For that purpose the simulation capabilities of the existing quantum transport solver OMEN will be improved by adding electron-phonon and phonon-phonon scattering. We will present numerical methods for computing all diagonal elements of both the retarded and lesser Green's function; As part of this research, the selected inversion approach will be extended to complex and general matrices as needed in the NEGF formalism.

Olaf Schenk
Department of Mathematics and Computer Science
Institute of Computational Science
olaf.schenk@usi.ch

MS172

Acceleration Techniques for Electronic Structure Calculation

We present a number of techniques for accelerating density functional theory based electronic structure calculation. One of these techniques aims at reducing the complexity for evaluating the electron density. In particular, we compute the electron density without diagonalizing the Kohn-Sham Hamiltonian. Another technique involves improving the convergence of the self-consistent field iteration by developing effective preconditioners.

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

MS173

Towards an Ultra Efficient Kinetic Scheme for the Boltzmann Equation

In this work we present a new ultra efficient numerical method for solving kinetic equations. The key idea, on which the method relies, is to solve the collision part on a grid and then to solve exactly the transport linear part by following the characteristics backward in time. The main difference between the method proposed and semi-Lagrangian methods is that here we do not need to reconstruct the distribution function at each time step. This allows to tremendously reduce the computational cost of the method and it permits for the first time to compute solutions of full six dimensional kinetic equations on a single processor laptop machine. Numerical examples, up to the full three dimensional case, are presented.

Giacomo Dimarco
Institut de Mathématiques de Toulouse
giacomo.dimarco@math.univ-toulouse.fr

MS173

Inverse Lax-Wendroff Method for Boundary Conditions of Boltzmann Type Models

In this talk we present a new algorithm based on Cartesian mesh for the numerical approximation of the kinetic models on complex geometry boundary. Due to the high dimensional property, numerical algorithms based on unstructured meshes for a complex geometry are not appropriate. Here we propose to adapt the inverse Lax-Wendroff procedure, which was recently introduced for conservation laws [S. Tan and C.W. Shu], to the kinetic equations. Applications in 1D3D and 2D3D of this algorithm for Boltzmann type operators (BGK, ES-BGK models) is then presented and numerical results illustrate the accuracy properties of this algorithm.

Francis Filbet
University of Lyon
filbet@math.univ-lyon1.fr

Chang Yang
University Claude Bernard, Lyon
yang@math.univ-lyon1.fr

MS173

Conservative Spectral Method for Collision Operators with Anisotropic Scattering

I will present the extension to anisotropic scattering mechanisms for the conservative spectral method developed by Gamba and Tharkabhushnanam, This method can admit a wide variety of collisional models, and we study the results obtained in the grazing collisions (Landau) limit.

Jeff Haack
Department of Mathematics
University of Texas at Austin
haack@math.utexas.edu

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

MS173

Averaged Kinetic Equations on Graphs

We derive a kinetic equation for flows on directed graphs lines with applications to production lines. We collect data from a company to derive transition probabilities and present a kinetic model that allows for a homogenization procedure, yielding a macroscopic transport model for large networks on large time scales.

Michael Herty
RWTH Aachen Universtiy
Department of Mathematics
herty@mathc.rwth-aachen.de

MS174

A Mass and Momentum Conserving Discontinuous Galerkin Shallow-water Model on the Cubed-sphere

The momentum equations used for the spherical shallow-water (SW) model equations are either in non-conservative form or in the vector-invariant flux-form with prognostic variables $(u, v, h)^T$. However, these formulations lack the formal conservation of momentum. A rigorous form of the mass and momentum conserving flux-form SW equations with prognostic variables $(uh, vh, h)^T$ on the cubed-sphere has been formulated, which leads to strong form of hyperbolic SW system. This system has more technical advantages than the usual vector-invariant form. Solving this system in non-orthogonal curvilinear geometry (such as the cubed-sphere) is a challenge due to tensorial representation with numerous metric terms. Both set of equations discretized using the discontinuous Galerkin method. A variety of standard SW tests have been performed for a rigorous comparison of conservation of integral invariants such as momentum, energy and potential enstrophy. The results will be presented in the seminar.

Lei Bao
University of Colorado, Boulder
lei.bao@colorado.edu

MS174

Numerical Framework for Atmospheric Modeling on Cubed Sphere by Multi-Moment Scheme

Multi-moment method defines more than one DOFs within each element to build local high-order schemes. On cubed sphere, compact stencil is beneficial not only to efficiently exchanging information over different patches, but also to effectively reducing the excessive errors due to the discontinuous coordinate systems along patch boundaries. In this talk, we will report a multi-moment SWE model up to fifth-order accuracy and a global AMR technique using multi-moment finite volume formulation on cubed sphere.

Chungang Chen
Xi'an Jiaotong Unversity
cgchen@mail.xjtu.edu.cn

Feng Xiao
Tokyo Institute of Technology
xiao@es.titech.ac.jp

Xingliang Li, Xueshun Shen
Center of Numerical Weather Predication,
China Meteorological Administration
lixl@cams.cma.gov.cn, shenxs@cams.cma.gov.cn

MS174

A Multi-dimensional Fourth-order Accurate Finite-volume Scheme on 3D Cubed-sphere Grids

A fourth-order accurate central essentially non-oscillatory (CENO) finite-volume scheme is presented for hyperbolic conservation laws on 3D cubed-sphere grids. The multi-dimensional approach uses k-exact reconstruction together with a monotonicity procedure that switches between high-order and low-order reconstruction based on a smoothness indicator. Hexahedral cells with trilinear faces are employed to handle nonplanar cell faces, achieving uniform high-order accuracy throughout the cubed-sphere grid. The 3D CENO scheme is implemented in a parallel dynamically adaptive simulation framework.

Lucian Ivan, Hans De Sterck, Andree Susanto
University of Waterloo
Applied Mathematics
livan@uwaterloo.ca, hdesterck@uwaterloo.ca,
asusanto@uwaterloo.ca

Clinton P. Groth
University of Toronto Institute for Aerospace Studies
Canada
groth@utias.utoronto.ca

MS174

Adaptive Fourth-Order Cubed Sphere Discretization for Non-Hydrostatic Atmosphere Simulations

We present an adaptive, conservative finite volume approach for non-hydrostatic atmospheric dynamics on a 3D cubed-sphere mapping with thin shells. Our method is fourth-order accurate in space, and uses higher-order least squares-based interpolation to compute stencil operations near block and refinement boundaries. Our discretization is adaptive in both time and horizontal spatial directions, while the radial direction is treated implicitly (using a fourth-order RK IMEX scheme) to eliminate time step constraints from vertical acoustic waves.

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

Phillip Colella, Peter Mccorquodale
Lawrence Berkeley National Laboratory
PColella@lbl.gov, PWMcCorquodale@lbl.gov

Paul Ulrich
UC Davis - LAWR
paullrich@ucdavis.edu

MS175

Immersed Boundary Method Simulations of Red Blood Cells

We have developed a variable-viscosity and variable-density Immersed Boundary method that makes effective use of the Fast Fourier transform despite the variable coefficients in the equations of motion. We have applied this method to study red blood cells, which rely on their remarkable flexibility to squeeze through capillaries. We demonstrate that our computations recover the physiological equilibrium shapes. Further, we describe our simulations of shear flow and single-file motion within capillaries. More recent work that includes a realistic model of the spectrin cytoskeleton will also be discussed.

Thomas Fai
Courant Institute
tfai@cims.nyu.edu

Boyce E. Griffith
Leon H. Charney Division of Cardiology
New York University School of Medicine
boyce.griffith@nyumc.org

Yoichiro Mori
School of Mathematics
University of Minnesota
ymori@umn.edu

Charles S. Peskin
Courant Institute of Mathematical Sciences
New York University
peskin@cims.nyu.edu

MS175

A Narrow-band Gradient-augmented Level Set Method for Incompressible Two-phase Flow

We have incorporated the gradient-augmented level set method (GALSM) for use in two-phase incompressible flow simulations by interpolating velocity values and introducing a re-initialization procedure. The method is conducted on a narrow band around the interface, reducing computational effort, while maintaining an optimally local advection scheme and providing sub-grid resolution. Ocean wave simulation is the primary motivation, and numerous benchmarks have been conducted, including a new comparison with wave tank data.

Curtis Lee, John Dolbow
Duke University
calee181@gmail.com, jdolbow@duke.edu

Peter J. Mucha
University of North Carolina
mucha@unc.edu

MS175

Adjoint-Based Error and Sensitivity Analysis in Transport-Depletion Problems

We describe the implementation and testing of an adjoint solver for the neutron/nuclide depletion equations. We do this in a flexible, multi-physics friendly framework for computing uncertainty and sensitivity information via the adjoint variable and describe our implementation in a massively-parallel transport solver. We then discuss performance trade-offs, algorithmic challenges, and scaling results as we push the problem towards large-scale simulations.

Hayes Stripling, Marvin Adams
Texas A&M University
h.stripling@tamu.edu, mladams@ne.tamu.edu

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

MS175

Efficient, Parametrically-Robust Nonlinear Model Reduction using Local Reduced-Order Bases

The large computational cost associated with high-fidelity physics-based simulations has limited their use in many practical situations. Model Order Reduction (MOR) is an attempt to reduce the computational cost of such simulations by searching for an approximate solution in well-chosen, low-dimensional affine subspaces. Here, I present recent developments in nonlinear MOR theory that achieve increased parametric robustness and additional speedup. These new techniques will be applied to fluid mechanical test problems.

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

David Amsallem, Charbel Farhat
Stanford University
amsallem@stanford.edu, CFarhat@stanford.edu

MS176

Bootstrapping Big Data in the Cloud

Bag of Little Bootstraps (BLB) is a recently developed and easily parallelizable algorithm which assesses the quality of a statistical estimator function on a large data set. Via Selective Embedded Just-in-Time Specialization (SEJITS), we dynamically convert such input high level estimator functions to low level, efficient Scala code to then deploy BLB to the cloud. This allows domain experts to quickly code BLB applications, thus providing for scalable and swifter manipulation of big data.

Peter Birsinger
University of California, Berkeley
pbirsinger@berkeley.edu

MS176

Big Data? Never Ask The Same Question Twice

The title refers to a holonic approach to the design of algorithms so that the same record never has to be read twice. With examples of application and their economic impact founded in LEGOs planning & logistic operations, this paper illustrates the relevance for real-time decision-making in an irreversible economic context, where the data-set evolves in discrete iterations. As aggregate run-time with discretely evolving data-sets appears to be absent in the debate on algorithmic design, the paper encourages further research in this area.

Bjorn Madsen
LEGO Group
bjorn.madsen@operationsresearchgroup.com

MS176

Spatial Analysis and Big Data: Challenges and Opportunities

We share our experiences as developers of new spatial analytical methods with a heavy social science perspective and identify new challenges that arise from "big data" in these realms. Much of the current stack of spatial analysis, statistics, and econometric software has not taken full advantage of new HPC environments due to a general lack of consensus on best practices. By engaging with the wider community of computational scientists we hope to identify paths forward.

Sergio Rey, Luc Anselin
GeoDa Center for Geospatial Analysis and Computation
Arizona State University
srey@asu.edu, luc.anselin@asu.edu

MS176

yt: Massively Parallel Astrophysical Simulation Analysis Made Easy

Astrophysical simulations have reached the point where the challenge of running ever-larger simulations is becoming overshadowed by the task of analyzing their data. We present recent enhancements to the yt analysis toolkit aimed at the analysis of large datasets, including multi-level parallelism and in-situ analysis capable of communicating back to running simulations. The goal is to provide a toolkit that allows the user to craft functionality that can be easily made to work in parallel.

Britton Smith
Michigan State University
smit1685@msu.edu

Matthew Turk
Columbia University
matthewturk@gmail.com

MS177

Experiences with Mini-Apps on Intel's Xeon Phi Architecture

Intel's recently announced Xeon Phi architecture represents a blend of features designed for high performance computing including increased length of vector processing units, increased thread counts and greater core density. Since 2011 Sandia National Laboratories has been working with Intel and several Office of Science Co-Design centers to port, analyze and optimize the performance of important mini-applications and libraries on the Xeon Phi architecture. In this talk we present the experiences gained from these exercises and thoughts of the next steps we will take to improve performance still further.

Simon D. Hammond
Scalable Computer Architectures
Sandia National Laboratories
sdhammo@sandia.gov

Richard Barrett
Sandia National Laboratories
rfbarre@sandia.gov

MS177

Introduction to Stampede and the Intel MIC Architecture

In this talk, we will introduced the Stampede supercomputing system deployed at TACC over the course of the summer and fall of 2012. Stampede is a large-scale Linux cluster which is first to deploy a significant number of Intel Xeon Phi co-processors using the Intel Many Integrated Core Architecture (MIC). We will present the Stampede architecture, programing models, and some early results enabled by the Xeon Phi coprocessors.

Bill Barth
Texas Advanced Computing Center
University of Texas at Austin
bbarth@tacc.utexas.edu

MS177

A Unified Approach to Heterogeneous Architectures using the Uintah Framework

Uintah is a large-scale, parallel multi-physics framework to simulate fluid-structure interaction problems on structured AMR grids using the ICE flow solver and Material Point Method. Uintah now scales to 256k cores on Jaguar by using hybrid MPI/multi-threaded parallelism. A recently developed heterogeneous runtime system further enables computational tasks to be offloaded to accelerators. In this work we demonstrate and analyze the performance of Uintah using both native and offload modes on the Intel MIC co-processor.

Alan Humphrey
Scientific Computing and Imaging Institute
University of Utah
ahumphrey@sci.utah.edu

Qingyu Meng
SCI Institute
Univeristy of Utah
qymeng@cs.utah.edu

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

MS177

MPI Communication on Stampede with MIC using

MVAPICH2: Early Experience

TACC Stampede is the first large-scale deployment of Intel Many Integrated Core (MIC) architecture. MVAPICH2 is one of the most widely-used open-sourceMPI libraries on high performance computing clusters with InfiniBand. This talk will present our early experience in enhancing MVAPICH2 for TACC Stampede with Intel MIC. An overview of the design enhancements (point-to-point and collective communication) and their impact on performance of end applications will be presented.

Dhabaleswar K. Panda, Sreeram Potluri, Devendar Bureddy
The Ohio State University
panda@cse.ohio-state.edu, potluri@cse.ohio-state.edu, bureddy@cse.ohio-state.edu

MS178

A Flexible and Extensible Multi-process Simulation Capability for the Terrestrial Arctic

The frozen soils of the Arctic contain vast amounts of stored organic carbon which is vulnerable to release as temperatures warm and permafrost degrades. The critical process models required for simulating degradation and release include subsurface thermal hydrology of freezing/thawing soils, thermal processes within ice wedges, mechanical deformation processes, overland flow, and surface energy balances including snow dynamics. The Arctic Terrestrial Simulator, based upon Amanzi, is being developed to enable flexible experimentation in coupling these processes.

Ethan T. Coon, Gianmarco Manzini, Markus Berndt, Rao V. Garimella
Los Alamos National Laboratory
ecoon@lanl.gov, gm.manzini@gmail.com, berndt@lanl.gov, rao@lanl.gov

David Moulton
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
moulton@lanl.gov

Scott Painter
Los Alamos National Laboratory
painter@lanl.gov

MS178

Sundance: High-level Components for Automation of PDE Simulation Development

We describe Sundance, a component toolkit for automating the assembly of high-performance parallel PDE simulations from high-level problem descriptions. We show how compact, high-level representations of a problem's symbolic form and geometry can be interpreted to coordinate the action of efficient low-level computational kernels.

Kevin Long, Robert C. Kirby
Texas Tech University
kevin.long@ttu.edu, Robert_Kirby@baylor.edu

MS178

Multilevel Preconditioner Components for Multi-Physics Problems on Adaptively Refined Meshes

Newton-Krylov and nonlinear Krylov methods for multi-physics problems require effective preconditioners for efficiency. Many multi-physics problems discretized on structured AMR grids require the solution of elliptic subcomponents as part of the preconditioner. Multilevel solution methods for elliptic problems that can exploit the natural multilevel hierarchy of structured AMR grids are discussed. We will describe the performance of both synchronous and asynchronous multilevel methods particularly in the context of nonlinear radiation diffusion problems.

Bobby Philip, Zhen Wang
Oak Ridge National Laboratory
philipb@ornl.gov, wangz@ornl.gov

Manuel Rodriguez Rodriguez, Mark Berrill
ORNL
mrs@ornl.gov, mbt@ornl.gov

MS178

Cache-aligned Data Structures for Unstructured Meshes

Abstract not available at time of publication.

Daniel Sunderland
Sandia National Laboratories
dsunder@sandia.gov

MS179

High-Performance Filtered Queries in Attributed Semantic Graphs

An analytic query views an attributed semantic graph through a filter that passes only edges of interest. In our Knowledge Discovery Toolbox, an open-source system for high-perfomance parallel graph computation, the user can define a Python filter that applies to every graph operation. We address the performance challenge of per-edge filtering with selective embedded specialization, automatically translating filters into an efficiency language, CombBLAS. This greatly accelerates Python KDT graph analytics on clusters and multicore CPUs.

John R. Gilbert
Dept of Computer Science
University of California, Santa Barbara
gilbert@cs.ucsb.edu

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

Armando Fox
University of California, Berkeley
fox@cs.berkeley.edu

Shoaib Kamil
Massachusetts Institute of Technology
skamil@cs.berkeley.edu

Adam Lugowski
UC Santa Barbara
alugowski@cs.ucsb.edu

Leonid Oliker, Samuel Williams

Lawrence Berkeley National Laboratory
loliker@lbl.gov, swwilliams@lbl.gov

MS179

Large-Scale Graph-Structured Machine Learning: GraphLab in the Cloud and GraphChi in your PC

Abstract not available at time of publication.

Carlos Guestrin, Joseph Gonzalez
Carnegie-Mellon University
guestrin@cs.cmu.edu, jegonzal@cs.cmu.edu

MS179

Are We There Yet? When to Stop a Markov Chain while Generating Random Graphs?

Markov chains are commonly used to generate random networks through rewiring of its edges. However, provable bounds on the number of iterations required for an independent instance are impractical. Practitioners end up using ad hoc bounds for experiments. In this work, we will present our methods for computing practical bounds on the number of iterations, and present experimental results for generating graphs with a given degree distributions and joint degree distribution.

Ali Pinar
Sandia National Labs
apinar@sandia.gov

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

C. Seshadhri
Sandia National Labs
scomand@sandia.gov

MS179

Analyzing Graph Structure in Streaming Data with STINGER

Analyzing static snapshots of massive, graph-structured data cannot keep pace with the growth of social networks, financial transactions, and other valuable data sources. Our software framework, STING (Spatio-Temporal Interaction Networks and Graphs), uses a scalable, high-performance graph data structure to enable these applications. STING supports fast insertions, deletions, and updates on graphs with semantic information and skewed degree distributions. STING achieves large speed-ups over parallel, static recomputation on both common multicore and specialized multithreaded platforms.

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

David A. Bader, Robert C. Mccoll
Georgia Institute of Technology
bader@cc.gatech.edu, rmccoll3@gatech.edu

David Ediger
Georgia Institute of Technology
School of Computational Science and Engineering
dediger@gatech.edu

MS180

Open Access for Models in Molecular Biophysics – Progress and Challenges

Abstract not available at time of publication.

Jaydeep P. Bardhan
Department of Electrical and Computer Engineering
Northeastern University
jaydeep_bardhan@rush.edu

MS180

Open Science in Molecular Simulations

Abstract not available at time of publication.

Ahmed Ismail
Department of Mechanical Engineering
RWTH Aachen University
ahmed.ismail@avt.rwth-aachen.de

MS180

How to Succeed in Reproducible Research without Really Trying

Abstract not available at time of publication.

Geoffrey M. Oxberry
Lawrence Livermore National Laboratory
goxberry@gmail.com

MS180

Exorcising Numerical Ghosts from ab initio Calculations of Electron Transport

The marriage of electron transport models with quantum chemistry codes has enabled the use of computation for exploring electron transport. Unfortunately, "better" basis sets often unphysically increase the calculated properties. The cause of this "ghost transmission" has proven elusive, in part due to the numerous, sometimes implicit, approximations invoked in the models. In this talk I diagnose ghost transmission, emphasize the importance of unit testing, and develop an open-source, ghost-busted model for electron transport.

Matthew Reuter
Eugene P. Wigner Fellow
Oak Ridge National Laboratory
reutermg@ornl.gov

MS181

Full Waveform Inversion Across the Scales

We present a newly-developed full waveform inversion scheme that incorporates seismic data on multiple spatial scales. Based on a decomposition of a multi-scale Earth model into various single-scale sub-models via 3D non-periodic homogenisation, the method allows us to simultaneously resolve the details of both the Earths crust and mantle. We demonstrate the applicability and efficiency of our technique in a multi-scale full waveform inversion for Europe and western Asia.

Andreas Fichtner
Utrecht University
fichtner@geo.uu.nl

Paul Cupillard
IPG Paris
paulcupillard@gmail.com

Erdinc Saygin
The Australian National University
u2509812@anu.edu.au

Yann Capdeville
Universite de Nantes
yann.capdeville@univ-nantes.fr

Tuncay Taymaz
Istanbul Technical University
ttaymaz@gmail.com

Antonio Villasenor
Barcelona Center for Subsurface Imaging
antonio@ictja.csic.es

Trampert Jeannot
Utrecht University
j.a.trampert@uu.nl

MS181

Dimensionality Reduction in FWI

Abstract not available at time of publication.

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

MS181

Projected-gradient Schemes and Sharp Interfaces in 3D Seismic Inversions

Abstract not available at time of publication.

Tarje Nissen-Meyer
ETH Zurich
Switzerland
tarjen@ethz.ch

Loredana Gaudio
University of Basel
loredana.gaudio@unibas.ch

Max Rietmann
University of Lugano
max.rietmann@usi.ch

Olaf Schenk
Department of Mathematics and Computer Science
Institute of Computational Science
olaf.schenk@usi.ch

Piero Basini
University of Toronto
piero.basini@utoronto.ca

MS181

Elastic and Anelastic Structure of the European Upper Mantle based on Adjoint Tomography

Harnessing high-performance computers and numerical methods to constrain physical properties of Earth's interior is becoming one of the most important topics in structural seismology. We use spectral-element and adjoint methods to iteratively improve 3D elastic and anelastic structure of the European upper mantle. This study involves two stages: in stage one, only phase measurements are employed to investigate radial anisotropic elastic structure. In stage two, both phase and amplitude measurements are applied to simultaneously constrain elastic and anelastic structure.

Hejun Zhu
Department of Geosciences
Princeton University
hejunzhu@princeton.edu

MS182

Error Analysis for Galerkin POD Approximation of the Nonstationary Boussinesq Equations

Abstract not available at time of publication.

S.S. Ravindran
University of Alabama
Huntsville, Alabama
ravinds@uah.edu

MS182

Reduced Order Modeling of Buoyancy Driven Incompressible Flows

In this talk, we apply proper orthogonal decomposition based reduced order modeling (ROM) to the time-dependent Boussinesq equations:

$$\begin{aligned} \partial_t \mathbf{u} - \nu_\mathbf{e} \boldsymbol{\Delta} \mathbf{u} + (\mathbf{u}\cdot\nabla)\mathbf{u} + \nabla \mathbf{p} &= -\beta\theta\mathbf{g}, \\ nabla \cdot \mathbf{u} &= 0, \\ partial_t \theta - a_e \Delta\theta + (\mathbf{u}\cdot\nabla)\theta &= c_p^{-1}\dot{q}^v. \end{aligned}$$

We discuss the function space setting and variational solution of the equations and how it relates to the ROM problem. We offer computational illustrations which indicate that for buoyancy driven flows, the ability of the ROM to reproduce the flow depends upon both the space utilized for the velocity and that for the temperature. We aim to utilize ROM to reproduce the solution from a parallel Boussinesq solver which uses a Chorin projection method.

John P. Roop
Department of Mathematics
North Carolina A&T
jproop@ncat.edu

MS182

New POD Error Expressions, Error Bounds, and Asymptotic Results for Model Reduction of Parabolic PDEs

The derivations of existing error bounds for POD-based reduced order models of time varying partial differential equations have relied on bounding the error between the POD data and various POD projections of that data. We prove that these data approximation errors can be computed exactly, using only the POD eigenvalues and modes. We apply our results to derive new POD model reduction error bounds and convergence results for the two dimensional Navier-Stokes equations.

John Singler

Missouri S & T
Mathematics and Statistics
singlerj@mst.edu

MS182

Proper Orthogonal Decomposition Reduced-Order Models of Complex Flows

In numerical simulations of complex flows, model reduction techniques are frequently used to make the computation feasible. Proper orthogonal decomposition is one of the most commonly used methods to generate reduced-order models for turbulent flows dominated by coherent structures. To balance the low computational cost required by a reduced-order model and the complexity of the targeted turbulent flows, appropriate closure modeling strategies need to be employed. In this talk, we will present new nonlinear closure methods for proper orthogonal decomposition reduced-order models, develop rigorous error estimates, and design efficient algorithms for them. Applications of the closure models in realistic engineering problems such as energy efficient building design and control will also be discussed.

Zhu Wang
University of Minnesota
wangzhu@ima.umn.edu

MS183

An Overview of Computational Methods for the Dynamics of Schrdinger Equations

The Schrödinger equation is a Partial Differential Equation that can be met in different areas of physics and engineering. Its most well-known use concerns quantum mechanics. Other fundamental applications are related to the study of Bose-Einstein condensates where the understanding of the Gross-Pitaevskii equation plays a crucial role. The numerical approximation of such equations is therefore fundamental. In this talk, we will review various solutions to build efficient numerical schemes, describe their properties and show their efficiency in various conditions.

Christophe Besse
Universite des Sciences et Technologies de Lille
christophe.besse@math.univ-lille1.fr

Xavier L. Antoine
Institut Elie Cartan Nancy (IECN)
Université de Lorraine
xavier.antoine@univ-lorraine.fr

Weizhu Bao
National University of Singapore
Department of Mathematics
bao@math.nus.edu.sg

MS183

A Posteriori Error Control and Adaptivity for Schrodinger Equations

We derive optimal order a posteriori error bounds for Crank-Nicolson fully discrete approximations for linear Schrodinger equations. For the spatial discretization we use finite element spaces that are allowed to change in time. The a posteriori error estimates are established using appropriate time-space reconstructions and a modified elliptic reconstruction that leads to estimators, which reflect correctly the physical properties of the continuous problem. Our theoretical results are validated by numerical experiments on various model problems.

Irene Kyza
IACM-FORTH
kyza@iacm.forth.gr

Theodoros Katsaounis
Associate Professor, Department of Applied Mathematics
University of Crete
thodoros@tem.uoc.gr

MS183

An Asymptotic Preserving Numerical Method for the Nonlinear Schrödinger Equation in the Semiclassical Regime

This is a joint work with Rémi Carles and Christophe Besse We present a new decomposition "à la Grenier" for the defocusing nonlinear Schrödinger equation in the semiclassical regime. We prove the local well-posedness of the system (and its global well-posedness in dimension one) and its convergence to the semiclassical limit before shocks appear in the limit system. Moreover, we construct a numerical scheme which is Asymptotic-Preserving, i.e. which is consistent and of order two when the semiclassical parameter ε is fixed, and which degenerates when $\varepsilon \to 0$ into a consistent numerical scheme for the limit system.

Florian Mehats
IRMAR - Université de Rennes 1
florian.mehats@univ-rennes1.fr

Remi Carles
CNRS & Universite Montpellier 2
Remi.Carles@math.cnrs.fr

Christophe Besse
Universite des Sciences et Technologies de Lille
christophe.besse@math.univ-lille1.fr

MS183

Recent Developments of Fast Algorithms for High Frequency Waves in the Semi-classical Regime

I will review some recent developments on fast algorithms for computing high frequency waves in the semi-classical regime.

Jianliang Qian
Department of Mathematics
Michigan State University
qian@math.msu.edu

MS184

A Greedy Strategy for Sparse Approximation of PDEs with High-dimensional Random Inputs

This talk is concerned with the sparse approximation of PDEs with high-dimensional random inputs where the quantity of interest (QOI) admits a sparse polynomial chaos (PC) expansion. We discuss a greedy strategy to select the PC basis functions whose coefficients are "large" from realizations of the QOI. We demonstrate the convergence of the proposed approach when the realizations of QOI are generated via Monte Carlo sampling as well as a more efficient sequential sampling design.

Jerrad Hampton
University of Colorado, Boulder
jerrad.hampton@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
doostan@colorado.edu

MS184
Dimension Reduction in Nonlinear Statistical Inverse Problems

The Bayesian approach to inverse problems in principle requires posterior sampling in high or infinite-dimensional parameter spaces. However, the intrinsic dimensionality of such problems is affected by prior information, limited data, and the smoothing properties of the forward operator. Often only a few directions are needed to capture the change from prior to posterior. We describe a method for identifying these directions through the solution of a generalized eigenvalue problem, and extend it to nonlinear problems where the data misfit Hessian varies over parameter space. This scheme leads to more efficient Rao-Blackwellized posterior sampling schemes.

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

Tiangang Cui, Tarek Moselhy
MIT
tcui@mit.edu, tmoselhy@mit.edu

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

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

MS184
High-dimensional Polynomial Chaos Basis Selection with Bayesian Compressive Sensing

For a complex model with a large number of input parameters, building Polynomial Chaos (PC) surrogate models is challenged by insufficient model simulation data as well as by a prohibitively large number of spectral basis terms. Bayesian sparse learning approaches are implemented in order to detect a sparse polynomial basis set that best captures the model outputs. We enhanced the Bayesian compressive sensing approach with adaptive basis growth and with a data-driven, piecewise-PC surrogate construction.

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

Dan Ricciuto, Peter Thornton
Oak Ridge National Laboratory
ricciutodm@ornl.gov, thorntonpe@ornl.gov

MS184
Mori-Zwanzig Approach to Nonlinear Stochastic Differential Equations with High-dimensional Parametric-type Uncertainty

We propose a new approach to obtain exact PDF equations for low-dimensional nonlinear functionals of the solution to high-dimensional stochastic differential equations, including SODEs and SPDEs. The new method does not suffer from the curse of dimensionality and, in principle, it allows us to avoid the integration of the full stochastic system, and solve directly for the PDF of the low-dimensional functional we are interested in.

Daniele Venturi
Brown University
daniele_venturi@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS185
Global Modes for Controller Selection and Placement in Compressible Turbulent Flows

A method for estimating the optimal location and type of flow control to use in compressible, turbulent flows is developed. Through linearizing the compressible Navier-Stokes equations about an unstable equilibrium point, the forward and adjoint equations of motion are used to estimate the structural sensitivity of the flow using a generalized “wavemaker’ concept. Matrix optimization is used to enhance the structural sensitivity over possible controller types (e.g., mass or energy sources) and locations, with optimal solutions identifying advantageous control strategies. The algorithms and theory are applied to a high-subsonic separating boundary layer appearing in an S-duct and compared with more traditional optimal methods through adjoint-based gradient information. It is found that optimum locations for mass and energy sources are typically located upstream of those momentum sources, and use different physical mechanisms for affecting the flow.

Daniel J. Bodony
University of Illinois at Urbana-Champaign
bodony@illinois.edu

MS185
Optimal Active Separation Control on Airfoils using Discrete Adjoint Approach

Blowing and suction type of active flow control techniques are being widely used to prevent or delay the flow separation in order to enhance the performance of aerodynamic configurations. Typically, this type of flow control is achieved by varying the actuation parameters such as amplitude, frequency, position and direction of blowing and suction. An efficient way of finding the optimal set of actuation parameters is by using the adjoint based optimisation methods. In the present work, a discrete adjoint solver has been developed for the optimal active control of unsteady incompressible viscous flows by using the Automatic/Algorithmic Differentiation (AD) techniques. An advantage of this approach is that the discrete realisations of the turbulence models are algorithmically differentiable and hence the frozen turbulence assumption is not required as it is often necessary for continuous adjoint methods. The adjoint code is applied to the optimisation problem of lift maximisation of a NACA 4412 airfoil using sinusoidal blowing and suction. Numerical results based on discrete adjoints are compared with the continuous adjoint approach.

Nicolas R. Gauger, Anil Nemili
Department of Mathematics and CCES
RWTH Aachen University
gauger@mathcces.rwth-aachen.de,
nemili@mathcces.rwth-aachen.de

Emre Oezkaya, Stefanie Guenther
RWTH Aachen University
ozkaya@mathcces.rwth-aachen.de,
guenther@mathcces.rwth-aachen.de

MS185

Shape Calculus and Unsteady PDE Control

Optimization and optimal control is considered for problems where the shape is the actual unknown to be found. Exploiting the structure of such problems leads to surface formulations of the gradient that circumvent most of the problems usually arising in shape optimization, such as computing sensitivities of the mesh deformation. The talk features applications of this technique in areas of wave propagation problems and computational fluid dynamics.

Stephan Schmidt
Imperial College London
s.schmidt@imperial.ac.uk

MS185

Towards Design and Optimization in Periodic and Chaotic Unsteady Aerodynamics

This talk focus on challenges in gradient based optimization for both periodic unsteady flows and chaotic, aperiodic unsteady flows. Checkpointing-based adjoint method for sensitivity analysis will first be introduced, followed by demonstration of the sensitivity divergence problem that aeries when the objective function is a long time averaged quantity. The fundamental cause of this problem is analyzed, and solutions to this problem, including the recently developed windowing method and Least Squares Sensitivity method, are introduced.

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

MS186

Sharing Thread Pools and Caches for Inter-library Composition and Multicore Performance

PETSc's multicore approach involves a thread communicator that can run on the user's choice of threading models and provides weak structured synchronization primitives. This is necessary to support applications and other libraries that were implemented using various programming models such as raw pthreads, OpenMP, and TBB. For performance, affinity must be managed and data structures organized so that cores can effectively utilize bandwidth and share caches while avoiding contention.

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory
jedbrown@mcs.anl.gov

MS186

Manycore Performance Portability through Mapped Multidimensional Arrays

Performance on manycore-accelerators is dependent upon architecture-specific data access patterns. The pervasive question of whether to use arrays-of-structures or structures-of-arrays is a fundamentally wrong question. The correct question is "what data structure abstraction is required to achieve manycore performance-portability?' Our answer is "device-mapped multidimensional arrays.' The KokkosArray library transparently inserts, at compile time, device-specific multidimensional array maps via conventional C++ template meta-programming. CPU/GPU performance-portability is demonstrated through MPI + KokkosArray hybrid parallel FEM proxy-applications.

H. Carter Edwards
Sandia National Laboratories
hcedwar@sandia.gov

MS186

Multilevel Programming Paradigms for Smart-tuned Exascale Computational Science

Exascale "hypercomputers" are expected to have highly hierarchical architectures with nodes composed by "lot of-core" processors and accelerators. The different programming levels (from clusters of processors loosely connected to tightly connected "lot-of-core" processors and/or accelerators) will generate new difficult algorithm issues. New language and framework should be defined and evaluated. In this talk, we present multilevel programming paradigms for exascale computing and propose example based on YML.

Serge G. Petiton
CNRS/LIFL and INRIA
serge.petiton@lifl.fr

MS186

A Hierarchical Parallel Implementation of a Contour Integral-based Eigensolver on Trilinos

We consider solving large-scale sparse eigenvalue problems. An eigensolver based on contour integral has been proposed by Sakurai and Sugiura. This solver has a hierarchical structure and is suitable for massively parallel supercomputers. In this solver, linear systems with multiple right-hand sides need to be solved. Trilinos includes useful packages for solving such kind of linear systems. In this talk, we present an implementation of the eigensolver on Trilinos and show its performance.

Tetsuya Sakurai
Department of Computer Science

University of Tsukuba
sakurai@cs.tsukuba.ac.jp

Yasunori Futamura
University of Tsukuba
futamura@mma.cs.tsukuba.ac.jp

Lei Du
University of Tsukuba, Japan
JST-CREST
du@mma.cs.tsukuba.ac.jp

Hiroto Tadano
Department of Computer Science
University of Tsukuba
tadano@cs.tsukuba.ac.jp

MS187

Optimal Control of Variable Density Navier-Stokes Equations

Optimal control problems for partial differential equations arise in various applications,such as in engineering,tomography and finance. Here we are concerned with the Navier-Stokes equation as state equation and a source control function to obtain a desired solution,also referred to as a velocity tracking problem. Using Tikhonov regularization and a Lagrange multiplier,a saddle point operator system arises where we can use equal order discretization for the state variable,the control function and the Lagrange multiplier. This enables elimination of the control function. The reduced system takes a particular two-by-two block matrix form for which we can construct a special preconditioner, the inverse of which involves just the inverses of two matrices that are linear combinations of the block matrices in each row of the given matrix.In the N-S problem,the are a mass matrix and an advection perturbed diffusion matrix. There is no need to solve any Schur complement system.The resulting eigenvalues are real and positive with a small condition number, which holds uniformly with respect to both discretization and method parameters.

Owe Axelsson, He Xin, Maya Neytcheva
Department of Information Technology
Uppsala University
Owe.Axelsson@it.uu.se, he.xin@it.uu.se,
maya.neytcheva@it.uu.se

MS187

Efficient Augmented Lagrangian-type Preconditioning for the Oseen Problem using Grad-Div Stabilization

Grad-Div stabilization can be exploited in a preconditioner for the Oseen Problem. It turns out that it behaves similar to the classical augmented Lagrangian approach, but with the advantage of being able to easily construct the system matrix efficiently. This simplifies the construction of inner preconditioners. I will discuss the difficulty of the trade-off between solution accuracy from stabilization and solver efficiency. Finally I will present numerical results to demonstrate the preconditioner.

Timo Heister
Texas A&M University
Department of Mathematics
heister@math.tamu.edu

MS187

Performance of SIMPLE-type Preconditioners in CFD Applications for Maritime Industry

CFD applications in maritime industry, for example hull resistance prediction, involve high Reynolds number flows modelled by the incompressible Reynolds-averaged Navier-Stokes equations. The system of equations is discretized with a cell-centered finite-volume method with colocated variables. After linearization, various SIMPLE-type preconditioners can be applied to solve the discrete system. In this presentation, we discuss their performance for flows with Reynolds number up to 10^9 and cell aspect ratio up to 10^6. http://ta.twi.tudelft.nl/nw/users/vuik/papers/Kla12V.pdf

Chris Klaij
Maritime Research Institute Netherlands
c.klaij@marin.nl

C Vuik
Faculty of Electrical Engineering, Mathematics
and Computer Science, Delft University of Technology,
NL
c.vuik@tudelft.nl

MS188

Solution of Parabolic Inverse Coefficient Problem Via Reduced Order State Equations

We introduce a method for numerical solution of a one-dimensional parabolic inverse coefficient problem with time domain data. The problem is formulated as non-linear least squares optimization. Unlike the traditional output least squares approach, the proposed method employs non-linear preconditioning, which greatly improves the conditioning of the optimization functional and its convexity. As a result, the method attains high quality reconstructions in just a few Gauss-Newton iterations. The construction of the non-linear preconditioner is based on ideas from model order reduction and rational approximation. We study different choices of matching conditions for projection based model order reduction and osculatory rational interpolation. Among those choices we identify those that correspond to a non-linear preconditioner with desired properties. Performance of the proposed method is evaluated in a number of numerical experiments involving both smooth and discontinuous coefficients.

Alexander V. Mamonov
Rice University
mamonov@ices.utexas.edu

MS188

Stability-Corrected Extended Krylov Method for Wavefield Problems in Unbounded Domains

In this talk, we present a new Extended Krylov Subspace Method for exterior wavefield problems. We start by showing that the solution of such problems can be expressed in terms of stability-corrected operator exponents. Subsequently, the exponents are approximated by structure-preserving and unconditionally stable reduced-order models (ROMs) drawn from an extended Krylov space. The performance of the method is illustrated through a number of examples in which we simulate electromagnetic and

acoustic wavefield propagation.

Rob Remis
Circuits and Systems Group
Delft University of Technology
r.f.remis@tudelft.nl

Vladimir L. Druskin, Mikhail Zaslavsky
Schlumberger-Doll Research
druskin1@slb.com, mzaslavsky@slb.com

MS188

Krylov Subspace Methods for Large Scale Constrained Sylvester Equations

Constrained Sylvester equations arise in various applications, and in particular in control theory, in the design of reduced-order observers. In this talk we present a new formulation of the algebraic problem which possesses certain advantages in the case of large scale data. Projection solvers for the resulting matrix equation will be discussed, whose approximate solution exactly satisfies the given constraint. Numerical experiments will be reported to illustrate the new methodology.

Stephen D. Shank
Temple University
sshank@temple.edu

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

MS188

Inverse Problems for Large-Scale Dynamical Systems in the H2-Optimal Model Reduction Framework

The Rational Krylov subspace (RKS) projection method with application to the inverse problems was considered. We derive a representation for the reduced Jacobian as the product of a time-dependent and a stationary part. Then we show that the RKS satisfying the Meier-Luenberger necessary H_2 optimality condition completely annuls the influence of approximation error on the inversion result. We compare our inversion against other nearly optimal RKS's based on Zolotarev problem and adaptive pole selection algorithm.

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

Aria Abubakar
Schlumberger Doll Research
aabubakar@slb.com

Vladimir L. Druskin, Tarek Habashy
Schlumberger-Doll Research
druskin1@slb.com, habashy1@slb.com

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

MS189

Improved Stability Estimates for the *hp*-Raviart-Thomas Projection Operator on Quadrilaterals

In this talk we derive improved stability estimates for the *hp*-Raviart-Thomas projection operator on quadrilaterals. Such estimates may be useful per se, but also have important applications, e.g., in the inf-sup stability proofs and a posteriori error estimation for *hp*-DG methods. In particular, we show that the stability constant of the RT-projector as a mapping in H^1 grows not faster than $p^{3/2}$, where p is the polynomial degree, being an improvement of the bound p^2 reported by Schoetzau et. al in SIAM J. Numer. Anal., 40 (2003), 2171–2194.

Alexey Chernov
University of Bonn
Hausdorff Center for Mathematics
chernov@hcm.uni-bonn.de

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

MS189

Commuting Diagram of TNT Elements on Cubes

Abstract not available at time of publication.

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

Weifeng Qiu
City University of Hong Kong
qiuw78@gmail.com

MS189

The Discontinuous Petrov-Galerkin Method for the Stokes Problem

We discuss well-posedness and convergence theory for the discontinuous Petrov Galerkin (DPG) method applied to the classical Stokes problem. The Stokes problem is an iconic troublemaker for standard Bubnov Galerkin methods; if discretizations are not carefully designed, they may exhibit non-convergence or locking. By contrast, DPG does not require us to treat the Stokes problem in any special manner. We illustrate and confirm our theoretical convergence estimates with numerical experiments.

Nathan Roberts
University of Texas at Austin
nroberts@ices.utexas.edu

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

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

MS189

Biorthogonal Basis Functions in *hp*-Adaptive FEM

for Elliptic Obstacle Problems

The talk presents an hp-adaptive mixed finite element discretiziation for a non-symmetric elliptic obstacle problem where the dual space is discretized via biorthogonal basis functions. The resulting algebraic system only includes box constraints and componentwise complementarity conditions. This special structure is exploited to apply efficient semismooth Newton methods using a penalized Fischer-Burmeister NCP-function in each component. Adaptivity is accomplished via a posteriori error control which is also introduced. Several numerical experiments show the applicability of the constructed biorthogonal basis functions.

Andreas Schroeder
Humboldt-Universitaet zu Berlin
andreas.schroeder@sbg.ac.at

Lothar Banz
Leibniz University Hannover
banz@ifam.uni-hannover.de

MS190

A Portable OpenMP Runtime Library based on MCA APIs

Programming multicore embedded systems is a challenge. These systems typically consist of heterogeneous cores operating on different ISAs, OSes and dedicated memory systems. Are there adequate software toolsets that can exploit capabilities of such systems? We will discuss about the industry standards, MCA APIs and how they could be used with OpenMP to provide the light-weight runtime library for embedded platforms.

Sunita Chandrasekaran
University of Houston
sunita@cs.uh.edu

MS190

Toward Parallel Applications for the Year of Exascale: Requirements for Resilient, Scalable Programming Models

Presently we are on the threshold of mass deployment of multilevel parallelism across most application areas. There are many programming models, languages and architectures from which to pick, and the number of choices is growing. In this presentation we discuss some of the principles of parallel application development that have produced todays codes, how we can address these principles going forward, and what we need from programming models.

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS190

The GPU Revolution, What Computational Chemistry and Battlefield Earth have in Common

This talk will look at the evolution of hardware features in commodity accelerators and the drivers influencing the architecture. The application of accelerators in HPC has lead to a wide array of software tools and underlying programming standards. We will snapshot the role NVIDIA has played in this effort, and illustrate a few successes and the challenges ahead.

Duncan Poole
NVIDIA
dpoole@nvidia.com

MS190

Handling the Power, Performance and Reliability Battle in Programming Models

Energy consumption and resilience are primary issues for high supercomputer utilization at Exascale systems. These issues cannot be addressed in isolation as there are proven tradeoffs between them. In this talk, I will discuss the importance of programming models in designing fault tolerance solutions, and present a case study with computational chemistry. I will also present the methodologies for energy efficiency and a case of co-design here. I will also present food for thought for fault tolerance and energy efficiency based on upcoming architectural trends.

Abhinav Vishnu
Pacific National Northwest Laboratory
abhinav.vishnu@pnnl.gov

MS191

Automated Adjoints of Finite Element Discretizations

A new approach for automatically deriving adjoint models is presented. The discretised PDE is formulated in a high level language that resembles the mathematical notation. Our approach differentiates this high level specification and uses code generation to implement the adjoint model (using the FEniCS system). I demonstrate that this approach automatically and robustly generates adjoint models for a wide class of PDE models. Examples in the field of optimization and stability analysis are presented.

Simon W. Funke
Imperial College London
Department of Earth Science and Engineering
s.funke09@imperial.ac.uk

Patrick Farrell
Department of Earth Science and Engineering
Imperial College London
patrick.farrell@imperial.ac.uk

Marie E. Rognes
Simula Research Laboratory
meg@simula.no

David Ham
Imperial College
david.ham@imperial.ac.uk

MS191

Oneshot Design Optimisation with Bounded Retardation

To complete an optimization run in a small multiple of the number of iterations required for a simulation run, it must be ensured that the contractivity factor of the combined ‘one-shot’ scheme is only a certain fraction closer to 1 than that of the user supplied fixed-point solver. This property we call bounded retardation of the convergence speed of a one-shot optimization compared to just a simulation.

We provide a bound on the retardation factor in terms of problem characteristics, mostly related to the Hessian of the Lagrangian. A key question is how close the preconditioning matrix of the design step should be to the reduced Hessian.

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

MS191

Dynamically and Kinematically Consistent Global Ocean-ice State and Parameter Estimation with a General Circulation Model and its Adjoint

Over the last decade the consortium on "Estimating the Circulation and Climate of the Ocean" (ECCO) has been producing optimal estimates of the global time-evolving circulation of the ocean. These estimates form the basis for addressing various problems in climate research. At the heart of the effort is the state-of-the-art MIT ocean general circulation model (MITgcm) and its adjoint. The project has taken rigorous advantage of algorithmic differentiation (AD) to generate efficient, scalable adjoint code. Use of AD has enabled the maintenance of up-to-date adjoint model versions in an environment of vigorous code development. As a result, today adjoint versions exist for the ocean component itself, and for components simulating sea ice, sub-ice shelf cavity circulation, and ocean biogeochemical processes. In this talk, we provide an overview of the estimation infrastructure, highlight sample applications related decadal ocean climate variability and predictability, and discuss future directions.

Patrick Heimbach
Massachusetts Institute of Technology
heimbach@mit.edu

MS191

Nonlinear Adjoint Looping in Thermoacoustics

Thermoacoustic oscillations are currently one of the biggest problems facing aircraft engine manufacturers, particularly because these oscillations seem to be triggered by very little noise. In this paper, nonlinear adjoint looping is used to calculate the optimal starting perturbation of a simple non-linear thermo-acoustic system. This shows that the system exploits non-normal transient growth to reach a stable limit cycle, even when starting with less energy than the corresponding unstable limit cycle.

Matthew P. Juniper
University of Cambridge
Department of Engineering
mpj1001@eng.cam.ac.uk

MS192

Parallel I/O Optimizations for a Massively Parallel Electromagnetic System

Checkpointing is a very effective approach for failure restart and post processing. However, this approach could result in heavy I/O load and may cause an I/O bottleneck on a massively parallel systems, especially future exascale systems with extremely high concurrency. In this talk, we present our application-level checkpoint approaches using tuned collective I/O, application data aggregation I/O and a threaded data aggregation model for asynchronous I/O. We discuss some production performance improvement of a massively parallel electromagnetic solver (NekCEM) on the different supercomputer architectures including the IBM Blue Gene/Q.

Jing Fu
Department of Computer Science
Rensselaer Polytechnic Institute
jingfu4@gmail.com

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Robert Latham
Argonne National Laboratory
robl@mcs.anl.gov

Christopher Carothers
Computer Science
Rensselaer Polytechnic Institute
chrisc@cs.rpi.edu

MS192

Accelerating Performance of a Petascale Electromagnetic Solver NekCEM with MPI+CUDA

Modern supercomputing architectures are increasing the number of cores per node and adding GPU co-processors to increase the instruction throughput. Compute-intensive applications need to take advantage of higher throughput by employing both distributed and shared memory programming models. In this talk, we discuss our approach for accelerating computational electromagnetics application code NEKCEM using MPI and CUDA programming models and demonstrate its performance on the leadership-class computing systems such as Eureka/Gadzooks on the IBM BG/P and Cray Titan.

Azamat Mametjanov
Mathematics and Computer Science
Argonne National Laboratory
azamat@mcs.anl.gov

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Boyana Norris
Argonne National Laboratory
norris@mcs.anl.gov

MS192

A Scalable Electromagnetic Solver for Applications in Nanoscale Materials

We will present recent advances in algorithmic and software development for efficient and accurate electromagnetics modeling that can benefit a wide range of relevant research communities and industries involved in the production of plasmonic devices, photovoltaic cells, electronic and storage devices, and semiconductors. The core algorithms are implemented into a petascale electromagnetic code that is based on an efficient communication kernel and memory-reduced framework. Performance and scalability analysis on the advanced computing architectures such as the IBM BG/Q and Cray XK6 will be demonstrated as well as some preliminary results based on a hybrid MPI/shared-memory model.

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Jing Fu
Department of Computer Science
Rensselaer Polytechnic Institute
jingfu4@gmail.com

Azamat Mametjanov
Mathematics and Computer Science
Argonne National Laboratory
azamat@mcs.anl.gov

Ying He
Department of Mathematics
Purdue University
heyingxmu@gmail.com

Paul F. Fischer
Argonne National Laboratory
fischer@mcs.anl.gov

MS192

Investigation on using Different High-order bases for Some Electromagnetic Simulations

Electromagnetic simulations need highly efficient numerical methods, which depend on many factors, such as expansion bases, mesh, algorithms and parallel models. In this talk, we will investigate the effects of using different expansion bases in some electromagnetic simulations. Their performance will be compared in detail, which includes accuracy, speed, and parallel efficiency. Based on these, advantages and disadvantages of using different bases will be studied. Some simulation results will be presented using the suitable solvers.

Jin Xu, Xiaohe Zhufu, Ruifeng Zhao
Institute of Software, ISCACS
Chinese Academy of Science, China
xu_jin@iscas.ac.cn, xiaohe@iscas.ac.cn,
ruifeng@iscas.ac.cn

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

MS193

Scalable Physics-based Preconditioning for 3D Extended MHD

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 3D XMHD. The method is based on a conceptual Schur-complement decomposition, which 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

MS193

Block Preconditioners for Coupled Fluids Problems

Many important scientific systems require solution of extensions to standard incompressible flow models, whether by incorporating additional nonlinear effects or by coupling to other processes. We consider approximate block factorization preconditioners for the algebraic systems resulting from linearization and FE discretization of these systems. In particular, we extend existing block-structured preconditioners (such as those of Elman, et al.), combining effective preconditioners for Navier-Stokes with other elements to obtain preconditioners for magnetohydrodynamics and coupled fluids problems.

Victoria Howle
Texas Tech
victoria.howle@ttu.edu

Robert C. Kirby, Geoffrey Dillon
Texas Tech University
Robert_Kirby@baylor.edu, geoffrey.dillon@ttu.edu

MS193

Block-Oriented Preconditioners for the Solution of the Semiconductor Drift-Diffusion Equations

We apply block-oriented preconditioners to the implicit solution of the drift-diffusion equations for semiconductor device modeling. The equations are discretized by a stabilized finite element method to produce the nonlinear coupled system, then solved with a parallel preconditioned Newton-Krylov method. The subblocks are solved by algebraic multigrid methods. The performance of these preconditioners will be compared to preconditioners that solve the fully-coupled algebraic system with fully-coupled AMG technique.

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

MS193

The Fast Adaptive Composite-grid Method for a 3-temperature Radiation Diffusion System with Adaptive Mesh Refinement

We describe the fast adaptive composite-grid (FAC) preconditioner applied to a non-equilibrium 3-temperature radiation diffusion problem. Multiple temporal and spatial scales make the associated initial value problem very challenging to solve. These challenges are addressed by fully implicit time integration and dynamic adaptive mesh refinement. At every timestep, a large scale nonlinear system is solved by the Jacobian-free Newton Krylov approach preconditioned by FAC. We will demonstrate the performance with accuracy, efficiency and scalability results.

Zhen Wang, Bobby Philip, Manuel Rodriguez Rodriguez, Mark Berrill
Oak Ridge National Laboratory
wangz@ornl.gov, philipb@ornl.gov, rodriguezrom@ornl.gov, berrillma@ornl.gov

MS194

Phase Response Theory Reveals Roles of Central and Sensory Inputs in Cockroach Locomotion

Abstract not available at time of publication.

Einat Fuchs
Princeton University
einat@princeton.edu

MS194

Design of Experiments of Parametric Manifolds with Application to Machine Vision and Material Identification

Abstract not available at time of publication.

James Penn
Massachusetts Institute of Technology
pennjam@mit.edu

MS194

On the Usefulness of Model Reduction Techniques in the Quest to Eradicate Infectious Diseases

Abstract not available at time of publication.

Joshua Proctor
Intellectual Ventures
JoshLProctor@gmail.com

MS195

Analysis of Glassy Potential Energy Landscapes

Glasses are an ill-understood class of materials. Although many potential topologies have been proposed for the underlying energy landscape, little is known about the true form of the potential energy surface that governs glassy dynamics. We explore techniques to sample the underlying energy landscape. We then use dimensionality reduction techniques to extract patterns in these landscapes, searching for underlying trends or coarse-grained structure in the detailed potential energy landscape of glasses.

Carmeline Dsilva
Princeton University
cdsilva@princeton.edu

MS195

Mapping Sugars Along Catalytic Itineraries: A Case Study in Exploring Multi-dimensional Landscapes

Accurately mapping the substrate free energy landscape is computationally demanding and of particular importance to understand enzyme function. We compared a suite of methods to examine how enzymes break down polysaccharides. The results highlight how the molecular and electronic structure of carbohydrates change as a function of perturbation from their solution-stable structures, lending them amenable to catalysis. These new insights will have wide-ranging applications from biomedical sciences to cellulose conversion for biofuel production.

Heather Mayes
Northwestern University
hmayes@u.northwestern.edu

MS195

Tensor Hypercontraction Theory: A Physically-Motivated Rank Reduction Method for Electronic Structure Theory

The Tensor Hypercontraction representation is a new, physically-motivated compression scheme which reduces the ubiquitous electron repulsion operator from a fourth-order tensor to a product of five second-order tensors. Advantages of this representation include substantial formal scaling reductions, decreased memory and/or communications requirements, and enhanced reliance on matrix-multiplication kernels. This talk discusses the mathematics of the compression scheme, as well as ongoing initiatives to apply the representation within many areas of electronic structure theory.

Robert M. Parrish
Georgia Tech
robparrish@gatech.edu

MS195

Absorption Spectra and Photoexcitation Dynamics in Phenylacetylene Dendrimers using TDDFT

Photochemical conversion of solar energy involving light harvesting molecules has been identified as a promising pathway to practical alternative energy technologies. Design of higher efficiency artificial photosynthetic systems necessitates a detailed understanding of photoexcitation and energy transfer processes. Here, we use time-dependent density functional theory to describe the electronic structure and photoexcitation in dendrimeric systems. Absorption spectra and electronic relaxation mechanisms of a wide range of dendrimers are computed and agreement with recent experiments is discussed.

Aaron Sisto
Stanford
asisto@stanford.edu

Todd Martinez
Department of Chemistry
Stanford University
toddjmartinez@gmail.com

MS196

The Community Earth System Model: Enabling High-resolution Climate Simulations

The Community Earth System Model (CESM) is a widely-used global climate model that enables climate simulations over meaningful periods of time and at high resolutions. CESM is composed of separate component models that

simulate the ocean, atmosphere, ocean, sea-ice, and land surface. A central coupler provides interpolation and regridding of the boundary conditions between the component models. We will discuss some examples of the computational challenges involved with climate modeling software.

Allison H. Baker, John M. Dennis
National Center for Atmospheric Research
abaker@ucar.edu, dennis@ucar.edu

MS196

A Space-Time Domain Decomposition Method for Stochastic Parabolic Problems

We consider a domain decomposition based implicit space-time approach for solving stochastic parabolic PDEs. The equation is first discretized in space and time using a stochastic Galerkin method and then decoupled into a sequence of deterministic systems with a Karhunen-Loeve expansion and double orthogonal polynomials. A Schwarz preconditioned recycling GMRES method is employed to solve the systems with similar structures. We report experiments obtained on a parallel computer with a large number of processors.

Cui Cong
University of Colorado at Boulder
cui.cong@colorado.edu

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

MS196

Resilience at Extreme Scale: System Level, Algorithmic Level or Both?

Resilience is a critical problem for extreme scale numerical simulations. The most credible solution is still based on checkpoint/restart with its high overheads or hardware cost. It has been shown recently that some algorithmic approaches and some code characteristics can help reducing these costs through combined system-algorithmic/application approaches. However, we are still looking for a right solution to this simple question: how to reduce simultaneously and significantly state saving and recovery times?

Franck Cappello
INRIA and University of Illinois at Champaign-Urbana
fci@lri.fr

MS196

Exploiting Hierarchies in Algorithms, Software, and Applications

This presentation discusses exploiting algorithmic and software hierarchies in scientific computing. These hierarchies present opportunities to manage complexity, exploit the changing landscape of high-performance computing, and reduce the time required to simulate physical phenomena. In this talk, we focus on the impact of hierarchy-aware Krylov methods in subsurface flow to overcome bottlenecks in global synchronization, and we discuss related issues in optimization software.

Lois Curfman McInnes, Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
curfman@mcs.anl.gov, tmunson@mcs.anl.gov

Jie Chen
Argonne National Laboratory
jiechen@mcs.anl.gov

Hong Zhang
Argonne National Lab
hzhang@mcs.anl.gov

MS197

Block Preconditioners for Implicit Atmospheric Climate Simulation in CAM-SE

We discuss the development of block preconditioners in an effort to reduce computational costs associated with implicit time integration of atmospheric climate models within CAM-SE. We construct a fully implicit framework based on the shallow water equations and view the subsidiary linear system as a block matrix. Preconditioners are derived based on approximate block factorization.

Aaron Lott
Lawrence Livermore Nat'l Lab
lott3@llnl.gov

Katherine J. Evans
Oak Ridge National Laboratory
evanskj@ornl.gov

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

MS197

Physics-based Preconditioners for Ocean Simulation

We examine physics-based preconditioners for free-surface, fully implicit, fully coupled time integration of the momentum and continuity equations of ocean dynamics; thus reducing errors and increasing stability due to traditional operator splitting. The nonlinear system is solved via Jacobian-free Newton-Krylov, where we reformulate semi-implicit barotropic-baroclinic splitting as a preconditioner. Thus the desired solution is timestep converged with timesteps on the order of the dynamical timescale. We provide numerical examples and compare to explicit methods.

Christopher K. Newman, Dana A. Knoll
Los Alamos National Laboratory
cnewman@lanl.gov, nol@lanl.gov

MS197

Implicit Solvers for Coupled Overland and Subsurface Flow

Subsurface and overland flow models constitute a significant portion of fresh water resource simulations, and coupling these models can produce insights into efficient management of these renewable resources. We consider a subsurface flow model coupled to kinematic and diffusive wave approximations of overland flow. We will discuss an implicit solution approach to these coupled models and overview the formulation and effectiveness of a block preconditioning approach.

Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
dan.osei@gmail.com

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

Reed M. Maxwell
Department of Geology and Geologic Engineering
Colorado School of Mines
rmaxwell@mines.edu

Laura Condon
Colorado School of Mines
lcondon@mymail.mines.edu

MS197

A Domain Decomposition based Implicit Method for Compressible Euler Equations in Atmospheric Modeling

We discuss some multilevel domain decomposition based fully implicit methods for solving the nonlinear system of hyperbolic equations arising from global climate modeling. With the fully implicit approach, the time step size is no longer limited by the stability condition, and with multilevel overlapping Schwarz preconditioners, good scalabilities are obtained on computers with a large number of processors. Numerical results are provided to show the conservation properties and the parallel scalability of the methods.

Chao Yang
Dept. of Computer Science
University of Colorado at Boulder
chao.yang@colorado.edu

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

MS198

Multi-level Monte Carlo for Continuous Time Markov Chain Models of Intracellular Biochemical Processes

Multi-level Monte Carlo is a relatively new method that computes expectations of stochastic processes to a desired tolerance significantly faster than previous methods. I will detail the basic idea of multi-level Monte Carlo and show how to implement it in the continuous time Markov chain setting, which is a modeling choice commonly used in the biosciences.

David F. Anderson
Department of Mathematics
University of Wisconsin Madison
anderson@math.wisc.edu

MS198

Efficient Simulation of Mesoscopic Reaction-diffusion Kinetics via Operator Splitting

A fractional step method offers many advantages for spatial stochastic simulation. It makes efficient simulation of large systems possible by using approximate methods to update the system due to diffusive transfers on the grid. It also simplifies more efficient parallel simulation since decoupling the operators enables more parallelism over the splitting time step. We propose a computable strategy to adaptively select that time step and illustrate the efficiency of our approach in parallel implementations.

Andreas Hellander
Department of Computer Science
University of California, Santa Barbara
andreash@cs.ucsb.edu

Brian Drawert
University of California Santa Barbara
bdrawert@cs.ucsb.edu

Michael Lawson
Univerity of Califorina Santa Barbara
mjl@cs.ucsb.edu

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

MS198

First-Passage Kinetic Monte Carlo Methods for Reaction-Drift-Diffusion Processes

Earlier versions of First-Passage Kinetic Monte Carlo, a stochastic algorithm for simulating reaction-diffusion processes, rely on analytic solutions of the diffusion equation and do not allow for drift. We have developed a variation of the algorithm using a discretization of the Fokker-Planck equation to incorporate drift due to arbitrary potentials. In this talk, we will demonstrate accuracy and convergence of our algorithm, compare various implementation approaches, and discuss applications to reaction-drift-diffusion processes in cell biology.

Ava J. Mauro
Department of Mathematics and Statistics
Boston University
avamauro@bu.edu

MS198

Computational Analysis of Stochastic Reaction-diffusion Equations

How to choose the computational compartment or cell size for the stochastic simulation of a reaction-diffusion system is still an open problem. We discuss a new criterion based on a global measure of the sensitivity of the reaction network that predicts a grid size that assures that the concentrations of all species converge to a spatially-uniform solution. This criterion applies for all orders of reactions and encompasses both diffusing and non-diffusing species

Hans G. Othmer
University of Minnesota
Department of Mathematics
othmer@math.umn.edu

MS199

Scalable Algorithms for Function Approximation

and Error Estimation on Arbitrary Sparse Samples

Stochastic collocation methods are an attractive choice to characterize uncertainty because of their non-intrusive nature. High dimensional stochastic spaces can be approximated well for smooth functions with sparse grids. There has been a focus in extending this approach to non-smooth functions using adaptive sparse grids. We have developed 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 compare these methods to more traditional statistical approaches.

Rick Archibald
Computational Mathematics Group
Oak Ridge National Labratory
archibaldrk@ornl.gov

MS199

Tensor-based Algorithms for the Optimal Model Reduction of Stochastic Problems

Tensor-based methods are receiving a growing attention for their use in high dimensional applications in scientific computing where functions of multiple parameters have to be approximated. Here, we present algorithms that are able to directly construct an approximation of optimal tensor decompositions of the solution of stochastic equations, without a priori information on the solution. Optimality can be achieved with respect to a desired metric.

Anthony Nouy, Marie Billaud-Friess, Loic Giraldi, Olivier Zahm
LUNAM Universite, Ecole Centrale Nantes, CNRS, GeM
anthony.nouy@ec-nantes.fr,
marie.billaud-friess@ec-nantes.fr, loic.giraldi@ec-nantes.fr,
olivier.zahm@ec-nantes.fr

MS199

Sliced Cross-validation for Surrogate Models

Multi-fold cross-validation is widely used to assess the accuracy of a surrogate model in uncertainty quantification. Despite its popularity, this method is known to have high variability. We propose a method, called sliced cross-validation, to mitigate this drawback. It uses sliced space-filling designs to construct structured cross-validation samples such that the data for each fold are space-filling. Extensions of the method to situations with high-dimensional inputs will be discussed. Numerical examples and theoretical results will be given to illustrate the proposed method.

Peter Qian
University of Wisconsin - Madison
peterq@stat.wisc.edu

MS199

On the Consistency of Calibration Parameter Estimation in Deterministic Computer Experiments

Calibration parameters in deterministic computer experiments are those attributes that cannot be measured or available in physical experiments. Kennedy-OHagan suggested an approach to estimate them by using data from physical as well as computer experiments. We develop an asymptotic theory for kernel interpolations to study the calibration consistency problem and show that the KO method leads to asymptotically inconsistent calibration due to overfitting. This calibration inconsistency can be remedied by modifying the original estimation procedure. (Based on joint work with Rui Tuo, Chinese Academy of Sciences)

Jeff Wu
Georgia Institute of Technology
jeffwu@isye.gatech.edu

MS200

Differential Geometric MCMC Methods and Applications

I will discuss the latest advances in Markov chain Monte Carlo methodology that exploit the natural underlying Riemannian geometry of many statistical models. Such algorithms automatically adapt to the local correlation structure of the model, providing highly efficient means of performing Bayesian inference for inverse problems. I will provide examples of Bayesian inference using these methods on a variety of challenging statistical models, including dynamical systems described by nonlinear differential equations.

Ben Calderhead, Mark Girolami
University College London
b.calderhead@ucl.ac.uk, m.girolami@ucl.ac.uk

MS200

Derivation and Low-rank Computation of the Bayesian Filter

We derive the non-linear Bayesian update (NLBU) and develop low-rank numerical algorithms for its evaluation. NLBU in a non-sampling functional approximation setting will be derived from the variational problem associated with conditional expectation. Whereas the linear BU is a linear function of the prediction mismatch, here we will use higher order polynomials. An important intermediate subproblem which appears during the BU is the increase of the stochastic dimension after each update. The reason is the new random variables which come from the random measurement noise.

Alexander Litvinenko
TU Braunschweig, Germany
litvinen@tu-bs.de

Hermann G. Matthies
Institute of Scientific Computing
Technische Universität Braunschweig
wire@tu-bs.de

MS200

Implicit Particle Methods for Data Assimilation

Many applications in science and engineering require that an uncertain model be updated by a stream of incomplete and noisy data. I will present a sequential Monte Carlo method for this problem that can avoid many of the pitfall that arise from a large problem dimension. The basic idea is to first identify regions of large probability and then focus attention on these regions. I will illustrate the theory with examples from applications in geophysics.

Matthias Morzfeld
Department of Mathematics

Lawrence Berkeley National Laboratory
mmo@math.lbl.gov

Alexander J. Chorin
University of California, Berkeley
Mathematics Department
chorin@math.berkeley.edu

MS200
Bayesian Data Assimilation with Optimal Maps

We develop novel map-based schemes for sequential data assimilation, i.e., nonlinear filtering and smoothing. One scheme involves pushing forward a fixed reference measure to each filtered state distribution, while an alternative scheme computes maps that push forward the filtering distribution from one stage to the next. A key advantage of the map approach is that it inherently avoids issues of sample impoverishment, since the posterior is explicitly represented as a transformation of a reference measure, rather than with a particular set of samples. The computational complexity of our algorithm is comparable to state-of-the-art particle filters.

Tarek Moselhy
MIT
tmoselhy@mit.edu

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

MS201
Control Volume Finite Element Method for Drift-diffusion Equations on General Unstructured Grids

We present a new Control Volume Finite Element Method for the drift-diffusion equations, which uses edge elements to define an exponentially fitted elemental current. This current obviates the need for the control volumes to be topologically dual to the finite elements and results in a method that is stable and accurate on general unstructured finite element grids. Simulations of a silicon PN diode and a MOSFET device demonstrate the performance of the new scheme.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Kara Peterson, Xujiao Gao
Sandia Natl. Labs
kjpeter@sandia.gov, xngao@sandia.gov

MS201
Lagrangian Hydrodynamics for Compressible Fluids

The entropy viscosity technique is extended to the Lagrangian framework for solving the compressible Euler equations. Various types of parabolic regularizations are discussed and a minimum entropy principle is proved. The method is illustrated numerically on 2D and 3D test cases.

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

Bojan Popov
Department of Mathematics
Texas A&M University
popov@math.tamu.edu

Vladimir Tomov
Department of Mathematics Texas A&M, USA
tomov@math.tamu.edu

MS201
High-order Curvilinear ALE Hydrodynamics

The Arbitrary Lagrangian-Eulerian (ALE) framework, which forms the basis of many shock hydrodynamics codes, consists of alternating Lagrange and advection phases. In this talk we will discuss our work on high-order extensions of the ALE advection phase, including curvilinear mesh optimization based on appropriately defined high-order topological "mesh Laplacian" and smoothing operators, as well as new DG advection algorithms for conservative and accurate remap of high-order fields. We will report results from single-material advection tests in a parallel research code.

Robert Anderson
LLNL
Ctr for Appl Scientific Comp
anderson110@llnl.gov

Veselin Dobrev
Lawrence Livermore National Laboratory
dobrev1@llnl.gov

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

Robert Rieben
Lawrence Livermore National Laboratory
rieben1@llnl.gov

MS201
Adaptive Material Interface Capturing Methods

A method used to capture a material interface carries with it a number of considerations that manifest themselves in the character of the solution. The nature of a physically admissible solution is key. Unfortunately, these considerations are often only implicitly manifested in the methods chosen to propagate the interface. We examine these considerations and how to bring them into the light of day. This provides a path forward toward better more physically motivated methods.

William J. Rider
Sandia National Laboratory
wjrider@sandia.gov

MS202
Performances of Krylov Solvers for Reactor Physics Simulation on Petascale Architectures

The governing equation in neutronic simulation for reactor physics application is the Boltzmann neutron transport equation. In order to solve this equation, one has to solve an eigenproblem. Efficient parallel eigensolvers are compulsory to reach high performances simulations for reactor physics. We will present in this talk some performances results of Arnoldi eigensolvers applied to reactor physics problem on the petascale heterogeneous CURIE machine on both CPUs and GPUs configurations.

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

Jérôme Dubois
CEA-Saclay
CEA/DEN/DANS/DM2S/SERMA
jerome.dubois@cea.fr

MS202

Performance Evaluation of Multi-threaded Iterative Solver on Recent Processors

Performance evaluation of multi-threaded sparse triangular solver is conducted on recent multi-core / many-core processors such as Intel Sandy Bridge processor. A sparse triangular solver is involved in IC(0) preconditioning, SOR method, Gauss-Seidel smoother etc., and it is utilized in many practical applications. Our multi-threaded solver is based on block multi-color ordering, which is one of parallel ordering techniques. The effect of blocking of unknowns is examined in recent processors.

Takeshi Iwashita, Akihiro Ida
Academic Center for Computing and Media Studies
Kyoto University
iwashita@media.kyoto-u.ac.jp, ida@media.kyoto-u.ac.jp

Masatoshi Kawai
Graduate School of Informatics
Kyoto University
kawai@ais.sys.i.kyoto-u.ac.jp

Hiroshi Nakashima
ACCMS
Kyoto University
h.nakashima@media.kyoto-u.ac.jp

MS202

Modeling of Epidemic Spread and Eigenvalue Computation

We present epidemiological modeling techniques of the spread of infectious diseases and show that below a threshold number of infected individuals, the spread of the epidemic will stop. We highlight that this epidemic threshold, beyond which infections become endemic, can be represented by the largest eigenvalue of the adjacency matrix representing the network of individuals. We present an efficient method to calculate the major eigenvector and reach a fast reaction facing the outbreak of epidemics.

Nahid Emad
University of Versailles, PRiSM Laboratory
Nahid.Emad@prism.uvsq.fr

Zifan Liu
University of Versailles
zifan.liu@prism.uvsq.fr

Michel Lamure
University of Lyon 1
France
michel.lamure@univ-lyon1.fr

Sofian ben Amor
Versailles University
France
s.benamor@gmail.com

MS202

Parallel CFD Code using ppOpen-HPC for Post-peta-scale Systems

ppOpen-HPC is an open source infrastructure for development and execution of large-scale scientific applications on various types of post-peta-scale systems with automatic tuning. This talk provides an example of development of 3D parallel CFD code on ppOpen-HPC, and overviews data structures, API and strategy for automatic tuning of ppOpen-HPC. Target code is developed on multicore clusters with OpenMP/MPI hybrid parallel programming models, and on CPU-GPU heterogeneous environment. Performance of the developed code is also demonstrated.

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

MS203

A Locally Conservative Eulerian-Lagrangian Method for a Two-Phase Flow Problem

We develop an Eulerian-Lagrangian numerical method for a system of two conservation laws in one space dimension modeling a simplified two-phase flow problem in a porous medium. We approximate tracing along the tracelines by imposing local mass conservation principles for both phases and optimizing self-consistency. Numerical results demonstrate that the method can handle problems with shocks and rarefactions on coarse spatial grids using time steps larger than the CFL limit.

Todd Arbogast
Dept of Math; C1200
University of Texas, Austin
arbogast@ices.utexas.edu

Chieh-Sen Huang
ational Sun Yat-sen University
Kaohsiung, Taiwan
huangcs@math.nsysu.edu.tw

Thomas F. Russell
NSF
trussell@nsf.gov

MS203

Finite Element Methods for the Fully Nonlinear Monge-Ampere Equation using a Local Discrete Hessian

In this talk, we will discuss a family of numerical methods for the Monge-Ampere equation, a fully nonlinear second order PDE. The approach is based on the concept of a discrete Hessian recently introduced by Aguilera & Morin (2009), Huang et al. (2010) and Lakkis & Pryer

(2011). However, the definition of our discrete Hessian is entirely local, making the resulting linear system within the Newton iteration much easier to solve. Replacing the Hessian in the PDE by its discrete counterpart, we obtain schemes that converge even when the exact solution possesses strong singularities.

Michael J. Neilan
University of Pittsburgh
Department of Mathematics
neilan@pitt.edu

MS203

New Phase-field Models and Energy Stable Numerical Schemes for Multiphase Flows with Different Densities

I shall present two new phase field models, one incompressible and the other quasi-incompressible, for multiphase flows with different densities. I shall also present efficient and energy stable numerical schemes, as well as some numerical results which validate the flexibility and robustness of these phase-field models.

Jie Shen
Purdue University
Department of Mathematics
shen@math.purdue.edu

MS203

Numerical Approximation of Oldroyd-B Fluids

The Oldroyd–B equations model the flow of fluids containing rod–like elastic molecules. This model couples the momentum equation with an equation governing the evolution of the elastic components, and numerical simulation is notoriously difficult (the high Weisenberg problem). This system of partial differential equations can be derived from Hamiltonian's principle which reveals a subtle balance between inertia, transport, and dissipation effects. This talk will focus on the structural properties of these equations which provide insight into why naive numerical schemes may fail, and the ingredients required to construct stable numerical schemes.

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

MS204

Multilevel Approximate Inversion

Computing the diagonal entries of the inverse of a sparse matrix arises in several computational applications such as covariance matrix analysis in uncertainty quantification, or when evaluating Green's functions in computational nanoelectronics. We discuss the approximate matrix inversion using a multilevel incomplete factorization. Its approximate inverse is represented recursively as a sum of matrices with increasing rank but decreasing norm. Taking the levels backwards we can successively update the diagonal entries of the matrix inverse.

Matthias Bollhöfer
TU Braunschweig
m.bollhoefer@tu-bs.de

MS204

Usage of Domain Decomposition Smoothers in Multigrid Methods

Local block solves are a building block of domain decomposition methods. Furthermore they can be used as smoothers for multigrid methods. In this context they have some advantages over point smoother, e.g., a higher arithmetic intensity that is beneficial on nowadays computer architectures. We analyzed different block smoothers and we will point out their advantages and disadvantages when used in a multigrid setting.

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

Karsten Kahl
Bergische Universität Wuppertal
Department of Mathematics
kkahl@math.uni-wuppertal.de

MS204

Bootstrap AMG

We present in this talk a Bootstrap approach to adaptive AMG introducing the so-called "Least Squares Interpolation" (LSI) and a "Bootstrap Setup" which enables us to compute accurate LSI operators using a multigrid approach in the setup to efficiently compute prototypes of algebraically smooth error. We demonstrate the potential of the Bootstrap AMG approach in the application to a variety of problems, each illustrating a certain aspect of the method.

Karsten Kahl
Bergische Universität Wuppertal
Department of Mathematics
kkahl@math.uni-wuppertal.de

MS204

Robust Solution of Singularly Perturbed Problems using Multigrid Methods

In order to resolve important features of solutions to singularly perturbed differential equations, specialized discretization techniques are frequently used. Here, we consider classical finite difference schemes on meshes adapted to resolve boundary layers in reaction-diffusion equations. We show that standard direct solvers exhibit poor scaling behaviour when solving the resulting linear systems. We consider, instead, standard robust multigrid preconditioners for these linear systems, and we propose and prove optimality of a new block-structured preconditioning approach.

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

Niall Madden
National University of Ireland, Galway
Niall.Madden@NUIGalway.ie

MS205

Reproducibility and Computationally Intensive,

Data-driven Research

Since Jon Claerbout adopted and started promoting reproducible research practices much has changed. While the problems for reproducibility of computational results has grown in conjunction with increases in computing power and storage densities, there has also been a steady growth in awareness of these problems and strategies to address them. In this minisymposium, we will discuss several recent attempts to come to terms with reproducibility in computational research. Topics will include education, publication, forensics and scientific integrity, as well as new technologies for provenance tracking and literate programming.

Kenneth J. Millman
University of California, Berkeley
millman@berkeley.edu

Vincent J. Carey
Channing Division of Network Medicine
Harvard Medical School
stvjc@channing.harvard.edu

MS205

Disseminating Reproducible Computational Research: Tools, Innovations, and Best Practices

Computation is now widely recognized as central to the scientific enterprise, and numerous efforts are emerging to incorporate code and data sharing into standards of research dissemination. This goal is challenging from a number of perspectives, including effective research practices. In this talk I discuss novel innovations and best practices for facilitating code and data sharing, both at the time of publication and during the research itself, that support the underlying rational of reproducible research.

Victoria Stodden
Columbia University
Statistics
vcs@stodden.net

MS205

Rethinking How we Work with Documents

Reproducible research tools require capturing all of the different avenues and lines of exploration in research. The document should be a database that can be rendered in different ways for different audiences, allowing dynamic results replacing code, enabling reader interactivity to explore different approaches and what-ifs. We also want to be able to programmatically query, update and verify this document database. All of this leads us to a different structure and approach for authoring documents.

Duncan W. Temple Lang
Statistics
UC Davis
duncan@wal.ucdavis.edu

MS205

Reproducible Research on the Web: From Homework, Blogging to Open Journals

Reproducible research used to be tied to LaTeX (e.g. Sweave in R) for statisticians, which has a steep learning curve and lacks many features of the web. The underlying idea of literate programming, however, is language agnostic. In this talk, we introduce the R package **knitr** as a general-purpose tool for reproducible research, with an emphasis on dynamic report generation on the web with Markdown, including reproducible homework, blog posts and online journals in statistics.

Yihui Xie
Department of Statistics
Iowa State University
xie@yihui.name

MS206

Numerical Solution of the Bloch-Torrey Equation Applied to the DMRI of Biological Tissue

We propose a numerical method to solve the Bloch-Torrey partial differential equation in multiple diffusion compartments to simulate the bulk magnetization of a sample under the influence of a diffusion gradient. We couple a mass-conserving finite volume discretization in space with a stable time discretization using an explicit Runge-Kutta-Chebyshev method. We are able to solve the Bloch-Torrey PDE in multiple compartments for an arbitrary diffusion sequence with reasonable accuracy for moderately complicated geometries in computational time that is on the order of tens of minutes per bvalue on a laptop computer. We show simulation results for nearly isotropic as well as anistropic diffusion, for the PGSE as well cosine OGSE sequences.

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

Jing-Rebecca Li
INRIA Saclay
jingrebecca.li@inria.fr

MS206

Image based Simulations towards Understanding Tissue Microstructure with MRI

Diffusion-weighted MRI and other quantitative MRI sequences like multi-exponential T2 are known to report on the microanatomical integrity of tissue. However, given the microanatomical complexity of tissue, mathematical and computational models are important to aid in the interpretation of these experiments. We present on histology-based simulations of white matter mircoanatomy for comparison with quantitative MRI, as well as the computational techniques required to make these simulations feasible.

Kevin Harkins
Vanderbilt University
kevin.harkins@vanderbilt.edu

MS206

Reduced Models of Multiple-compartment Diffusion MRI in the Intermediate Exchange Regime

We model the magnetization in biological tissue due to a diffusion gradient by a two compartment Bloch- Torrey partial differential equation with infinitely thin permeable membranes. We formulate a ODE model for the magnetization and show the simpler ODE model is a good approximation to the Bloch-Torrey PDE model for a variety of gradient shapes. Using the ODE model we determine of the change in the cellular volume fraction from the signal attenuation obtained before and after cell swelling. This method requires only the ADC and Kurtosis of the two signal attenuations and the numerical solution of an ODE system.

Jing-Rebecca Li
INRIA Saclay
jingrebecca.li@inria.fr

MS206

Time-dependent Diffusion: From Microstructure Classification to Biomedical Applications

We show how a bulk diffusion measurement can distinguish between different classes of microgeometry. Based on the specific values of the dynamical exponent of a velocity autocorrelator measured with diffusion MRI, we identify the relevant tissue microanatomy in muscles and in brain, and the microstructural changes in ischemic stroke. Our framework presents a systematic way to identify the most relevant part of structural complexity with diffusion.

Dmitry S. Novikov
New York University School of Medicine
dmitry.novikov@nyumc.org

Els Fieremans
New York University
els.fieremans@nyumc.org

Jens Jensen, Joseph Helpern
Medical University of South Carolina
jense@musc.edu, helpern@musc.edu

MS207

A General Framework for Stable Reconstructions from Non-uniform Fourier Samples

Abstract not available at time of publication.

Ben Adcock
Purdue University
adcock@purdue.edu

MS207

Finite Fourier Frame Approximation using the Inverse Polynomial Reconstruction Method

The inverse polynomial reconstruction method (IPRM) was developed to resolve the Gibbs phenomenon in the spectral reconstruction of piecewise analytic functions. We demonstrate that the IPRM is suitable for approximating the finite inverse Fourier frame operator as a projection onto the polynomial space. The IPRM can also remove the Gibbs phenomenon from the Fourier frame approximation of piecewise smooth functions. Numerical results show that the IPRM is robust, stable, and accurate for non-uniform Fourier data.

Jae-Hun Jung, Xinjuan Chen
SUNY at Buffalo
jaehun@buffalo.edu, xc5@buffalo.edu

Anne Gelb
Arizona State University
annegelb@asu.edu

MS207

Convolutional Gridding and Frame Approximation

The technique of covolutional gridding (CG) has been widely used in applications with non-uniform (Fourier) data such as magnetic resonance imaging (MRI). On the other hand, its error analysis is not fully understood. We consider it as a frame approximation and present an error analysis accordingly. Moreover, we propose a generalized convolutional gridding (GCG) method as an improved frame approximation.

Guohui Song
Clarkson University
gsong@clarkson.edu

Anne Gelb
Arizona State University
annegelb@asu.edu

MS207

Robust Sub-Linear Time Fourier Algorithms

We present a new deterministic algorithm for the sparse Fourier transform problem, in which we seek to identify $k \ll N$ significant Fourier coefficients from a signal of bandwidth N. Previous deterministic algorithms exhibit quadratic runtime scaling, while our algorithm scales linearly with k in the average case. Via a multiscale approach our algorithm is extremely robust against noise.

Yang Wang
Michigan State University
ywang@math.msu.edu

MS208

Non-Gaussian Test Models for Prediction and State Estimation with Model Errors

A class of statistically exactly solvable non-Gaussian test models are introduced where a generalized Feynman-Kac formulation reduces the exact behavior of conditional statistical moments to the solution of inhomogeneous Fokker-Planck equations modified by linear lower order coupling and source terms. This procedure is applied to a test model with hidden instabilities and combined with information theory to address the coarse-grained ensemble prediction in a perfect model and improving long range forecasting in imperfect models.

Nan Chen
New York University
nan.chen@nyu.edu

MS208

Stability and Convergence of a Fully Discrete Fourier Pseudo-spectral Method for Boussinesq Equation

In this paper, we discuss the nonlinear stability and convergence of a fully discrete Fourier pseudo-spectral method coupled with specially designed time-stepping of second order for the numerical solution of the “good” Boussinesq equation. Our results improve the known results obtained by Frutos et al. In particular, this stability condition does not impose a restriction on time-step that is dependent on the spatial grid size.

Wenqiang Feng

Missouri University of Science and Technology
fw253@mst.edu

MS208

First and Second Order Schemes for Applications of Dynamic Density Functional Theory

In this talk I will present the first and second order (in time) unconditional energy stable schemes for nonlocal Cahn-Hilliard (CH) equation, nonlocal Allen-Cahn (AC) equation and some models derived from dynamic density functional theory (DDFT). I will briefly derive these nonlocal models and discuss the relation between nonlocal CH equation and DDFT. Also the relation between DDFT and classical CH equation and phase field crystal equation will be discussed. I will briefly introduce properties of schemes such as stability and convergence. Numerical simulations of nonlocal CH equation with anisotropic correlation functions will be given. Also numerical simulations for implications of DDFT, such as hard sphere model, will be presented.

Zhen Guan
The University of Tennessee
zguan@utk.edu

MS208

Operator-splitting for Convection-reaction-diffusion Equations

For reaction-diffusion systems with both stiff reaction and diffusion terms, implicit integration factor (IIF) method and its high dimensional analog compact form (cIIF) serve as an efficient class of time-stepping methods. For nonlinear hyperbolic equations, front tracking method is one of the most powerful tools to dynamically track the sharp interfaces. Meanwhile, weighted essentially non-oscillatory (WENO) methods are a class of start-of-the-art schemes with uniform high order of accuracy in smooth regions of the solution. In this talk, IIF/cIIF is coupled with front tracking or WENO by the second-order symmetric operator splitting approach to solve advection-reaction-diffusion equations. In the methods, IIF/cIIF methods treat the stiff reaction-diffusion equations, and front tracking/WENO methods handle hyperbolic equations that arise from the advection part.

Xingfeng Liu
University of South Carolina
xfliu@math.sc.edu

MS209

Beyond Treewidth in Graphical Model Inference

While exact probabilistic inference in graphical models typically has cost exponential in the treewidth, we discuss cases where this breaks down. For example, when the distribution contains determinism, graph triangulations implying higher than optimal treewidth can have unboundedly faster inference. Also, when the nature, rather than the degree, of interaction is limited, inference cost can become polynomial even for unboundedly large treewidth. Examples include submodular interaction functions or those that indirectly utilize submodularity.

Jeff A. Bilmes
University of Washington, Seattle
bilmes@u.washington.edu

MS209

Graph Width Metrics, Well-Quasi Ordered Sets and Fixed Parameter Tractability: History, Applications and Scalable Implementations

Parameterized computation has evolved from a mere complexity theoretic aberration to a powerful and highly respected technique for solving difficult combinatorial problems. Papers on fixed parameter tractability (FPT) now appear with regularity in top computer science conferences and journals. It has not always been an easy journey. We will survey FPT's history, its current known range of applications, the central role played by graph width metrics, and the significance of effective parallel implementations.

Michael A. Langston
Department of Electrical Engineering and Computer Science
University of Tennessee
langston@eecs.utk.edu

MS209

Toward Tree-like Structure in Large Informatics Graphs

Large informatics graphs such as large social and information networks are often thought of as having some sort of tree-like or hierarchical structure. We describe recent empirical and theoretical results aimed at extracting meaningful tree-like structure from large informatics graphs in a scalable and robust way. In particular, empirical properties of cut-based methods such as tree decompositions and metric-based methods such as delta-hyperbolicity, as well as their similarities and differences, will be reviewed.

Michael Mahoney
Stanford University
Applied and Computational Mathematics
mmahoney@cs.stanford.edu

MS209

Chordal Graphs and Clique Trees in Sparse Matrix Algorithms

We consider two problems in parallel sparse matrix computation where efficient algorithms are enabled by the clique tree. Simple greedy algorithms that eliminate certain leaves of the clique tree solve these problems. Proving the correctness of the algorithms provides fresh insight into the collection of vertex separators in a chordal graph. Hence, in this context, sparse matrix algorithms provde new results in chordal graph theory.

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

MS210

Fully Computable a posterior Error Estimators for Stabilized Conforming Finite Element Approximations

Abstract not available at time of publication.

Alejandro Allendes
Universidad Técnica Federico Santa Mara
allendes.flores.alejandro@gmail.com

MS210

Error Estimation for VMS-stabilized Acoustic Wave Propagation

Abstract not available at time of publication.

Brian Carnes
Sandia National Laboratories
bcarnes@sandia.gov

MS210

Output-based hp-adaptive Simulations of High-Reynolds Number Compressible Flows

We present a method for concurrent mesh and polynomial-order adaptation with the objective of direct minimization of output error using a selection process for choosing the optimal refinement option from a discrete set of choices that includes directional spatial resolution and approximation order increment. The scheme is geared towards compressible viscous aerodynamic flows, in which solution features make certain refinement options more efficient compared to others. We present results for 2D and 3D turbulent flows.

Marco Ceze, Krzysztof Fidkowski
University of Michigan
mceze@umich.edu, kfid@umich.edu

MS210

Gradient-norm Error Estimation for PDE-constrained Optimization

When solving a PDE-constrained optimization problem numerically, we would like to know if the obtained design accurately reflects the true optimum. One measure of accuracy would be to compute an error estimate for the objective or the Lagrangrian; however, a more natural measure of accuracy for an optimum is the discrepancy in the first-order optimality conditions. To this end, we show how the adjoint-weighted residual method can be used to construct a posteriori error estimates for the norm of the gradient. The error estimate requires two additional adjoint variables, but only at the beginning and end of each optimization cycle. Moreover, the adjoint systems can be formed and solved with limited additional infrastructure.

Jason E. Hicken
postdoctoral fellow, Stanford University
hickej2@rpi.edu

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

MS211

Auto-tuning and Smart-tuning Approaches for Efficient Krylov Solvers on Petascale Architectures

For decades, supercomputers have carried on delivering more and more computational power using more and more complex architecture. It reflects in the parameterization of the basic algorithms used by the computer codes to achieve good performances. We experiment a methodology based on statistical approach for autotuning. The objective is to help the code users define the best set of parameters on a key algorithm (Matrix.Vector multiplication) depending on the numerical problem and the computing environment.

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

Anthony Leroy Drummond
Lawrence Berkeley National Laboratory
ladrummond@lbl.gov

France Boillod-Cerneux
CNRS/LIFL
france.boillod-cerneux@cea.fr

Jerome Dubois
CEA-DEN
jerome.dubois@cea.fr

Gisele Ndongo Eboum
University of Paris XIII, France
andongobecker@aol.com

MS211

Iterative Method for Sparse Linear Systems using Quadruple Precision Operations on GPUs

The convergence of iterative methods such as a Krylov subspace method may be affected by round-off errors and extended precision operations may reduce the iterations to get the convergence. We implemented Krylov subspace methods for sparse linear systems on an NVIDIA GPU using quadruple precision operations and evaluated their performance. We used double-double arithmetics to perform quadruple precision arithmetic operations. On the GPUs, the computation time of one iteration on quadruple precision is up to approximately twice as that on double precision. In some cases by utilizing quadruple precision operations we can reduce the computational time to get the convergence as compared with the computation using full double precision operations.

Daichi Mukunoki
Graduate School of Systems and Information Engineering
University of Tsukuba
mukunoki@hpcs.cs.tsukuba.ac.jp

Daisuke Takahashi
University of Tsukuba
Graduate School of Systems and Information Engineering
daisuke@cs.tsukuba.ac.jp

MS211

Toward Tunable Multi-Scheme Parallelization

High performance computers will have more SIMD widths and a larger number of cores. System scale ranges from a single processor (1TF) to a supercomputer (1EF). We need more flexible parallelization methodologies, which must be based on composite extension of Amdahl's law and hierarchical extension of Gustafson's law. Gradual parallelization and deparallelization, reusable parallel components, and autotuning will be a part of key concepts, and a vision on programming tools for them will be discussed.

Reiji Suda
Department of Computer Science, The University of Tokyo

reiji@is.s.u-tokyo.ac.jp

MS212
Communicating in Science and not being Afraid of Tenacious Self-promotion

Publish-or-perish is no news to anyone, but is that all there is in science communication? Far from it, there are a myriad considerations: some political and even legal in nature (copyright, anyone?), but most are social. Twitter for tenure? Absurd as it may sound, there are advantages to using social media, which some scientists are already benefitting from. At the very least, scientists should invest in their web presence and identity. And while we mention presence, how about honing those presentation skills? Zen for slides, Tufte for plots, and did you know that your font selection can help the review of your proposal?

Lorena A. Barba
Department of Mechanical Engineering
Boston University
labarba@bu.edu

MS212
The Two Body Problem

The two body problem in physics can be solved by dealing with a few equations, while solving the two body problem in academic career is much harder, maybe no correct solution at all. Here, I just want to share my own experience of two body problem–my struggle, attempts, failures and result, as well as the experience of some dual-career couples that I know. I hope this will be helpful for people in similar situation.

Bo Dong
University of Massachusetts Dartmouth
bdong@umassd.edu

MS212
Preparing for Tenure and Promotion

Preparing for tenure and/or promotion is arguably one of the most stressful times in one's academic career. When preparing your case, it is worth bearing in mind that most academic institutions have T&P processes which require each case be reviewed at different levels within the institution before a final decision is made. I will give pointers for making a strong case for tenure and/or promotion based on my recent past experience making recommendations at three of these levels within my home institution.

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

MS212
Educating Undergraduate Women in Mathematics

I will present my experiences at Spelman College, educating undergraduate women in mathematics.

Monica Stephens
Spelman College
Atlanta, GA
mstephens@spelman.edu

MS213
Spatial Stochastic Modelling of the Hes1 Pathway

Individual mouse embryonic stem cells have been found to exhibit highly variable difffferentiation responses under the same environmental conditions. The noisy cyclic expression of Hes1 and its downstream genes are known to be responsible for this, but the mechanism underlying this variability in expression is not well understood. We show that the observed experimental data and diverse difffferentiation responses can be explained by a spatial stochastic model of the Hes1 gene regulatory network.

Marc Sturrock
University of Dundee
msturrock@maths.dundee.ac.uk

Andreas Hellander
Department of Computer Science
University of California, Santa Barbara
andreash@cs.ucsb.edu

Anastasios Matzavinos
Iowa State University
tsasos@iastate.edu

Mark Chaplain
Division of Mathematics
University of Dundee
chaplain@maths.dundee.ac.uk

MS213
Modeling of Stochastic Diffusion and Reaction Processes in Mixed Dimensions in Systems Biology

Abstract not available at time of publication.

Per Lotstedt
Department of Information Technology
Uppsala University, Sweden
Per.Lotstedt@it.uu.se

MS213
Intrinsic and Extrinsic Noise in Genetic Oscillations

Abstract not available at time of publication.

Shev MacNamara
University of Oxford
Oxford, UK
shev.mac@gmail.com

MS213
Linear Algebra for Difference Equations, Networks and Master Equations

Abstract not available at time of publication.

Gilbert Strang
Massachusetts Institute of Technology
gs@math.mit.edu

MS214
Comparison of Uncertainty Quantification Meth-

ods for Nonlinear Parabolic Systems

We compare exiting probabilistic methods for propagation of epistemic uncertainty in parabolic models such as the nonlinear Richards' equation with spatially distributed, uncertain input parameters. For the estimation of first and second order moments, we show that Monte Carlo methods outperform global and locally adaptive Stochastic Collocation methods for input with short correlation lengths and high variances, due to an increase in problem dimensionality and the loss of regularity of state variables in probability space.

David A. Barajas-Solano
Dept. of Mechanical and Aerospace Engineering
University of California, San Diego
d2baraja@ucsd.edu

Daniel M. Tartakovsky
University of California, San Diego
dmt@ucsd.edu

MS214

An Implementation of Polynomial Chaos Expansion on Discontinously Dependent Model Paramters

One of the main advantageous of polynomial chaos expansions and stochastic collocation method is it's exponential convergence rate. However, this rate is lost because of Gibb's phenomena if there is a discontinuity in the dependency structure between model parameters. In this talk a new all-purpose method for setting up an expansion with dependent model parameters will be presented, and fast convergence property will be shown irrespectively of probability structure.

Jonathan Feinberg
Simula Research Laboratory
University of Oslo
jonathfe@simula.no

MS214

Visualizing Gaussian Process Uncertainty using Smooth Animations

Animations can visualize probability distributions: each frame shows a random draw from the distribution, and is correlated with its neighbors so the motion is continuous. I focus on Gaussian distributions, including Gaussian Processes which are used in kriging. Existing work interpolates between "keyframes' (*mutually independent* random draws from the distribution), but motion becomes kinky at these keyframes. My approach treats all frames on equal footing, yielding smooth, natural-looking timetraces. Code is included.

Charles Hogg
Google
charles.r.hogg@gmail.com

MS214

Assessment of Numerical Uncertainty in the Solution of Inverse Problems, for Different Observations Conditions

Abstract The Bayesian Inference method is applied to calibrate parameters of the differential equation used to model diffusion processes. The forward model is estimated using both analytical and numerical solution of the ODE. The numerical uncertainty induced by time discretization is studied through comparing results associated with the analytical and numerical forward models. Also, effect of time step size on numerical uncertainty and computational cost is investigated.

Negin Yousefpour
Texas A&M University
negin.yousefpour@tamu.edu

Zenon Medina-Cetina
Zachry Department of Civil Engineering
Texas A&M University
zenon@tamu.edu

Hans Petter Langtangen
Center for Biomedical Computing
Simula Research Laboratory and University of Oslo
hpl@simula.no

Are Magnus Bruaset, Stuart Clark
Simula Research Laboratory
arem@simula.no, stuart@simula.no

MS215

Sensitivity Analysis in Weak-Constraint 4D-Var: Theoretical Aspects and Applications

Suboptimal weighting of the information provided by models and measurements poses a fundamental limitation on the performance of atmospheric data assimilation systems (DAS). Theoretical aspects of adjoint sensitivity analysis are presented in the context of weak-constraint four-dimensional variational data assimilation. Evaluation of the model forecast sensitivity to the information vector and to the DAS representation of the information error statistics, covariance parameter optimization, and a priori performance assessment are discussed.

Dacian N. Daescu
Portland State University
Department of Mathematics and Statistics
daescu@pdx.edu

MS215

Efficient Implementations of the Ensemble Kalman Filter

This research presents efficient implementations of the ensemble Kalman filter based on Singular Value (SVD) and Cholesky decompositions. In addition, a novel implementation is derived from the Sherman Morrison formula. This direct method exploits the special structure of the data error covariance matrix which, in practice, is often diagonal. The complexity of the proposal is equivalent to well-known, efficient implementations of the EnKF. The proposed method is tested using realistic atmospheric and oceanic models. In terms of accuracy, not significant difference is shown in the results for the compared methods. Moreover, the elapsed time of the simulation is reduced when the proposed implementation of the EnKF is used.

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

Elias Nino-Ruiz

Virginia Polytechnic Institute and State University
enino@vt.edu

MS215

Trust Region Adaptive POD/DEIM 4D-Var for a Finite-Element Shallow Water Equations Model

A proper orthogonal decomposition (POD) coupled with discrete empirical interpolation method (DEIM) constructs efficient reduced-order inverse problem four - dimensional variational (4D-Var) data assimilation for a nonlinear finite element (FEM) shallow water equations(SWE) model. Different approaches of POD/DEIM trust-region 4D-Var data assimilation problem are compared, including a dual-weighed method for snapshots selection. Using DEIM we reduce the computational complexity of the Trust-Region POD-4D-Var FEM-SWE model by decreasing the CPU time required to calculate the solutions of the forward and adjoint models and indirectly by reducing the condition number of Hessian of cost functional, thus accelerating convergence of minimization process.

Razvan Stefanescu
Florida State University
rstefanescu@fsu.edu

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

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

MS215

A Hybrid Variational-ensemble Data Assimilation Method

In this presentation we will discuss discus a possibility to selectively use methodological advantages from ensemble and variational data assimilation systems while avoiding sub-optimal features of the original algorithms. We will present a new development of static error covariance model for use in hybrid system that explores the structure of circulant matrices, while maintaining the complexity of existing variational error covariance models.

Milija Zupanski
Cooperative Institute for Research in the Atmosphere
Colorado State University
Milija.Zupanski@colostate.edu

MS216

Data Fusion based on Coupled Matrix and Tensor Factorizations

Data fusion enhances knowledge discovery, in particular, in complex data mining problems. The task of fusing data, however, is challenging since data are often incomplete, heterogeneous, i.e., in the form of higher-order tensors and matrices, and have both overlapping and non-overlapping components. We formulate data fusion as a coupled matrix and tensor factorization problem and propose an all-at-once optimization algorithm, which easily extends to coupled analysis of incomplete data sets. We demonstrate the usefulness of our approach in metabolomics applications.

Evrim Acar
University of Copenhagen
evrim@life.ku.dk

MS216

Data-Driven Analysis and Fusion of Medical Imaging Data

Data-driven methods such as independent component analysis (ICA) have proven quite effective for the analysis of functional magnetic resonance (fMRI) data and for discovering associations between fMRI and other medical imaging data types such as electroencephalography (EEG) and structural MRI data. Without imposing strong modeling assumptions, these methods efficiently take advantage of the multivariate nature of fMRI data and are particularly attractive for use in cognitive paradigms where detailed a priori models of brain activity are not available. This talk reviews major data-driven methods that have been successfully applied to fMRI analysis and fusion, and presents examples of their successful application for studying brain function in both healthy individuals and those suffering from mental disorders such as schizophrenia.

Tulay Adali
University of Maryland
Baltimore County
adali@umbc.edu

MS216

MetaFac: Community Discovery via Relational Hypergraph Factorization

This work aims at discovering community structure in rich media social networks, through analysis of time-varying, multi-relational data. Community structure represents the latent social context of user actions. It has important applications in information tasks such as search and recommendation. Social media has several unique challenges. (a) In social media, the context of user actions is constantly changing and co-evolving; hence the social context contains time-evolving multi-dimensional relations. (b) The social context is determined by the available system features and is unique in each social media website. In this work we propose MetaFac (MetaGraph Factorization), a framework that extracts community structures from various social contexts and interactions. Our work has three key contributions: (1) metagraph, a novel relational hypergraph representation for modeling multi- relational and multi-dimensional social data; (2) an efficient factorization method for community extraction on a given metagraph; (3) an on-line method to handle time-varying relations through incremental metagraph factorization. Extensive experiments on real-world social data collected from the Digg social media website suggest that our technique is scalable and is able to extract meaningful communities based on the social media contexts. We illustrate the usefulness of our framework through prediction tasks. We outperform baseline methods (including aspect model and tensor analysis) by an order of magnitude.

Jimeng Sun
IBM T.J. Watson Research Center
jimeng@us.ibm.com

MS216

Looking for Common Features Across a Collection of Matrices Using the Higher-Order GSVD

The higher-order GSVD is a way of simultaneously reducing each matrix in a collection $\{D_1, \ldots, D_N\}$ to a form that permits one to identify common features. Each D_i has the same number of columns. The invariant subspace associated with the minimum eigenvalue of the very nasty matrix $(S_1 + \ldots + S_n)(inv(S_1) + \ldots + inv(S_N))$ is involved where $S_i = D_i^{'} D_i$. However, we are able to get at this subspace safely via the minimization of a very interesting quadratic form. Everything reverts to the ordinary generalized SVD when $N = 2$.

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

Orly Alter
Sci Comp & Imaging (SCI) Inst, Bioeng & Human Genetics Depts
University of Utah
orly@sci.utah.edu

Sri Priya Ponnapalli
Bloomberg
sripriyaponnapalli@gmail.com

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

MS217

Optimal Experimental Design under Uncertainty

We focus on large-scale optimal experimental design for statistical inverse problems. We consider a model involving recovery of initial concentration field in an advection-diffusion problem. We use a Bayesian inference framework with Gaussian prior and likelihood. Due to linearity of the parameter-to-observable map, the posterior is also Gaussian, with mean and covariance given by solution of a deterministic inverse problem. The goal is to optimize the location of observation points to maximize information gain.

Alen Alexanderian
University of Texas at Austin
alen@ices.utexas.edu

Noemi Petra
Institute for Computational Engineering and Sciences (ICES)
The University of Texas at Austin
noemi@ices.utexas.edu

Georg Stadler, Omar Ghattas
University of Texas at Austin
georgst@ices.utexas.edu, omar@ices.utexas.edu

MS217

Joint Inversion

Inverse problems are inheritable non-unique and regularization is needed to obtain stable and reasonable solutions. The regularization adds information to the problem and determines which solution, out of the infinitely many, is obtained. In this talk we discuss the case when a-priori information exists in the form of either known structure or in the form of another inverse problem for a different property. The challenge is to include such information in the inversion process.

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

MS217

Approximate Dynamic Programming for Sequential Bayesian Experimental Design

Optimal design maximizes the value of costly experimental data. Popular approaches for designing multiple experiments are suboptimal: these include *open-loop* approaches that choose all experiments simultaneously, or *greedy* methods that optimally design the next experiment without accounting for the future. We instead formulate experimental design in a *closed-loop dynamic programming* (DP) framework that yields the true optimal sequence of experiments. We solve the DP problem using methods of approximate dynamic programming, coupled with polynomial surrogates and stochastic optimization. Preliminary results demonstrate the superiority of the closed-loop approach.

Xun Huan, Youssef M. Marzouk
Massachusetts Institute of Technology
xunhuan@mit.edu, ymarz@mit.edu

MS217

Estimating and using Information in Inverse Problems

Information theory provides convenient mechanisms to manage and unify various components of the inversion process especially when uncertainty has to be addressed. Here we calculate information densities from adjoints to guide mesh adaptivity, regularization parameters, and sensor location identification. These aspects have historically been handled in an ad-hoc and incohesive manner. Our information density approach not only provides a unifying mechanism, but also demonstrates improved performance on a convection-diffusion problem in comparison to standard methods.

Bart Vanbloemenwaanders
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
bartv@sandia.gov

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

MS218

SciHadoop: Array-based Query Processing in Hadoop

Hadoop has become the de facto platform for large-scale data analysis in commercial applications and increasingly in scientific applications. Scientific data, typically structured but stored in binary file formats, enables different assumptions and optimizations than unstructured data. Our system, SciHadoop, utilizes knowledge of that structure to yield big performance gains by creating efficient input splits; enabling holistic combiners; minimizing and localizing communications; producing early results; and producing balanced, contiguous output.

Joe Buck
UC Santa Cruz
Personal
buck@soe.ucsc.edu

MS218

What is MapReduce and How can it Help with my Simulation Data?

Abstract not available at time of publication.

Paul Constantine
Stanford University
Mechanical Engineering
paul.constantine@stanford.edu

David F. Gleich
Purdue University
dgleich@purdue.edu

MS218

Benchmarking MapReduce Implementations for Scientific Applications

With data production increasing rapdily due to the ever growing application needs, the MapReduce model and its implementations need to be further evaluated and optimized. Several MapReduce frameworks with various degrees of conformance to the key tenets of the model are available today, each, optimized for specific features. We will discuss the design of a standard benchmark suite for quantifying, comparing, and contrasting the performance of MapReduce platforms under a wide range of representative use cases found in scientific applications. The aim of the benchmark is to allow scientists to choose the MapReduce implementation that best suits their application's needs.

Madhusudhan Govindaraju
SUNY Binghamton
mgovinda@binghamton.edu

MS218

Dynamic Mode Decomposition with MapReduce

Dynamic mode decomposition (DMD) is a spectral analysis technique that identifies coherent features of a fluid flow, based on a series of snapshots of the flow field (Schmid, 2010). In order to treat databases exceeding 1 Tb in size, we develop an implementation of the DMD algorithm based on MapReduce tall-and-skinny QR factorization (Constantine & Gleich, 2011). The method is demonstrated on a three dimensional simulation database of a screeching supersonic jet.

Joseph W. Nichols
Stanford University
jwn@stanford.edu

MS219

Certified Parameter Optimization with Reduced basis Surrogate Models

PDE-constrained parameter optimization problems with arbitrary output functionals may be quite expensive, when using standard PDE-solvers. Instead we use RB surrogate models for approximately solving the optimization problem. Ingredients of our scheme comprise RB-spaces for the solution, its sensitivity derivatives and rigorous a-posteriori error bounds for the solution, derivatives, output functional and the suboptimal parameters. Experiments on an instationary convection-diffusion problem demonstrate the benefits of the approach.

Markus Dihlmann, Bernard Haasdonk
University of Stuttgart
dihlmann@mathematik.uni-stuttgart.de,
haasdonk@mathematik.uni-stuttgart.de

MS219

A Certified Reduced basis Approach for Parametrized Linear-quadratic Optimal Control Problems

In this talk, we discuss effective model reduction of linear-quadratic optimal control problems governed by parametrized elliptic and parabolic partial differential equations. To this end, we employ the reduced basis method as a surrogate model for the solution of the optimal control problem and develop rigorous *and* efficiently evaluable *a posteriori* error bounds for the optimal control *and* the associated cost functional. Numerical results are presented to confirm the validity of our approach.

Mark Kaercher, Martin Grepl
RWTH Aachen University
kaercher@aices.rwth-aachen.de,
grepl@igpm.rwth-aachen.de

MS219

Goal-oriented Inference for Nonlinear PDE-constrained Inverse Problems

High-dimensional parameter spaces pose a significant challenge to existing inference methods, particularly in the context of limited data. When the quantity of interest depends on the parameter estimate, computational effort may be wasted resolving components of the parameter not required to make accurate predictions. We present a goal-oriented inference method that exploits the context of required predictions. Previous work on linear problems is extended to the nonlinear setting.

Chad E. Lieberman
MIT
celieber@mit.edu

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

MS219

Use of Reduced Order Models in Iterative Optimization Solvers

When solving nonlinear optimization problems with partial differential equations, information from the previous iteration is sometimes useful to accelerate the convergence of the nonlinear optimization solver. We show how to use information obtained from a reduced order model in the

previous iterations in the solver of a quadratic subproblem, for example. This leads to deflated MINRES or CG algorithms, where we consider the convergence of these methods. We give numerical results for an optimization problem with a nonlinear parabolic partial differential equation in the 3-dimensional case.

Ekkehard W. Sachs
University of Trier
Virginia Tech
sachs@uni-trier.de

Xuancan Ye
University of Trier
ye@uni-trier.de

MS220

Using Shannon Entropy to Estimate Convergence of Cmfd-Accelerated Monte Carlo

Coarse Mesh Finite Difference (CMFD)[Smith, TANS, v44, 1984] is a common moment based acceleration technique that can be used to solve transport equations. The statistical error that arises from CMFD-accelerated Monte Carlo can mask the convergence of the CMFD correction values. This work demonstrates how Shannon entropy[Ueki and Brown, ANS, 2003] can be used to determine when a CMFD iteration is near a convergence limit created by the statistical noise of the Monte Carlo method.

Mathew Cleveland
Lawrence Livermore National Laboratory
cleveland7@llnl.gov

Todd Palmer
Oregon State University
palmerts@engr.orst.edu

Nick Genitle
Lawrence Livermore National Laboratory
gentile1@llnl.gov

MS220

Monte Carlo Simulation Methods in Moment-based Scale-bridging Algorithm for Neutral Particle Transport Problems

Recently, we have extended a moment-based scale-bridging algorithm to thermal radiative transfer problems. The algorithm accelerates a solution of the high-order (HO) radiation transport equation using a "discretely consistent" low-order (LO) continuum system. In this talk, we discuss a recent progress of the use of Monte Carlo simulation methods, which are specifically tailored to the scale-bridging algorithm, such as an energy conserving tally and asymptotic approximation for optically thick spatial cells.

Hyeongkae Park
Los Alamos National Laboratory
hkpark@lanl.gov

Jeff Densmore, Allan Wollaber
LANL
jdd@lanl.gov, wollaber@lanl.gov

Dana Knoll
Los Alamos National Laboratory
nol@lanl.gov

Rick Rauenzahn
LANL
rick@lanl.gov

MS220

Acceleration of Nuclear Reactor Neutronics Eigenvalue Problems with Non-Linear Low-Order Projection Operators

Nuclear reactor core analyses using continuous-energy Monte Carlo neutron transport models have been accelerated using low-order operator projections onto coarse spatial and energy mesh diffusion theory models. A discontinuity factor formulation of the diffusion equations (which allow exact preservation of Monte Carlo solution on the coarse mesh) is used to produce deterministic eigenvalue equations at each stage of the Monte Carlo solution. Converged eigenvectors obtained by solving the low-order diffusion equations are then used to scale and prolong the Monte Carlo source distributions, resulting in significant reductions in the number of inactive and active histories needed to compute reliable reactor spatial power distributions and eigenvalues.

Kord Smith, Benoit Forget, Bryan Herman, Paul Romano
Massachusetts Institute of Technology
kord@mit.edu, bforget@mit.edu, bherman@mit.edu, paul.k.romano@mit.edu

MS220

A Hybrid Approach to Nonlinear Acceleration of Transport Criticality Computations

In nuclear engineering, the computation of the dominant eigenvalue of the neutron transport equation is highly important for the analysis of nuclear reactors. Recent work in the field has led to advancements in efficient deterministic numerical techniques for computing this eigenvalue, including Nonlinear Diffusion Accelerated Power Iteration (NDA-PI). In an effort to obtain higher accuracy solutions, it is becoming more popular to use Monte Carlo simulations to compute the eigenvalue and associated eigenvector, as these simulations allow for a continuous treatment of space, angle and energy. We look to build hybrid solvers that combine the efficiency of NDA-PI with the high-accuracy solutions provided by otherwise expensive Monte Carlo simulations.

Jeffrey A. Willert
North Carolina State University
jawiller@ncsu.edu

MS221

DPG Methods for Transport and the Inviscid Euler Equations

The DPG method, introduced by Demkowicz and Gopalakrishnan, is a Petrov-Galerkin method derived from the minimization of the operator residual associated with a variational formulation. For transport under the ultra-weak variational formulation, DPG demonstrates optimal convergence for cases in which the upwind DG method is suboptimal. Unfortunately, the DPG solution for systems of transport equations manifests problems related to solution regularity in the crosswind direction; we discuss these issues in detail and propose regularization solutions through the vanishing viscosity method.

Jesse L. Chan
University of Texas at Austin
jchan985@gmail.com

MS221

Comparison of Multigrid Performance for Stabilized and Algebraic Flux Correction FEM for Convection Dominated Transport

In this empirical study we will compare the performance of Algebraic Multigrid (AMG) applied to stabilized and algebraic flux corrected discretizations of convection dominated problems. The Algebraic Flux Corrected transport algorithm considered in this study results in an operator with the M-matrix property which has theoretical properties more amenable to AMG as compared to stabilized methods. The performance of these methods on benchmark convection-diffusion and fluid flow problems will be demonstrated.

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

John N. Shadid
Sandia National Laboratories
jnshadi@sandia.gov

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Dmitri Kuzmin
University Erlangen-Nuremberg
kuzmin@am.uni-erlangen.de

MS221

Characteristic Discontinuous Galerkin for Tracer Advection in Climate Modeling

We will describe our development of Characteristic Discontinuous Galerkin (CDG) for tracer advection, on arbitrary convex polygon meshes, for use in climate modeling. CDG can be viewed as an extension of Prather's 1986 moment method to unsplit advection on general meshes and higher orders of accuracy. Methods for bounds preservation will be discussed, and comparisons made with Runge-Kutta DG (Cockburn & Shu), semi-Lagrangian DG (Restelli, Bonaventura, & Sacco), and Hancock-Huynh DG (Lo & Van Leer).

Robert B. Lowrie
Los Alamos National Laboratory
Los Alamos, NM 87545 USA
lowrie@lanl.gov

Todd Ringler
Los Alamos National Laboratory
ringler@lanl.gov

MS221

A Semi-Lagrangian Discontinuos Galerkin Transport Scheme on the Cubed Sphere

Because of its geometric flexibility and high parallel efficiency, DG (discontinues Galerkin) method is becoming increasingly popular in atmospheric and ocean modeling. However, a major drawback of DG method is its stringent CFL stability restriction associated with explicit time-stepping. We adopt a dimension-splitting approach where a regular semi-Lagrangian (SL) scheme is combined with the DG method, permitting longer time steps. The SLDG scheme is inherently conservative and has the option to incorporate a local positivity-preserving filter for tracers. The SLDG scheme is tested for various benchmark advection test-suites on the cubed-sphere and results will be presented in the seminar.

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

Wei Guo
University of Houston
wguo126@gmail.com

MS222

Mesoscale Investigations of the Influence of Capillary Heterogeneity on Multiphase Flow of Fluids in Rocks

We are focused on understanding the influence of mesoscale heterogeneity on multiphase flow of CO2 and brine. This largely ignored spatial scale provides the opportunity to develop a rigorous understanding of major controls on CO2 migration when a continuum description of the multiphase physics is applicable. Methods have been developed to characterize subcore scale capillary heterogeneity, to perform high-resolution simulations that that accurately replicate experiments, and accurately predict multiphase flow in complex rocks.

Sally M. Benson
Energy Resources Engineering
Stanford University
smbenson@stanford.edu

MS222

Multiscale Algorithms for Reactive Transport in Porous Media

I will present two multiscale methods including a Langevin approach and a dimension reduction method based on a computational closure. The purpose of these methods is to provide an accurate description of the system averages while retaining critical pore-scale information. The advantages, range of applicability and limitations of the mentioned above multiscale methods will be discussed.

Alexandre Tartakovsky
Pacific Northwest National Laboratory
alexandre.tartakovsky@pnl.gov

MS222

Pore Scale Reactive Transport Modeling using Adaptive, Finite Volume Methods with a Look toward Upscaling

We have developed a high performance simulation capability to model pore scale multi-component reactive transport in geologic subsurface media, especially that obtained from image data. Our approach is conservative, accurate and robust, based on adaptive, finite volume methods. As such, the adaptive functionality can be used to provide communication between models at different spatial scales, and, thus, provide the ability to directly upscale local pore scale data to the Darcy scale through flux matching techniques.

David Trebotich
Lawrence Berkeley National Laboratory
treb@lbl.gov

MS222

Geometric Comparisons in Porous Media Simulation

We extract geometric descriptors that characterize flow properties and the pore space of a material. By measuring distances between the descriptors, we compare multiple materials.

Gunther H. Weber
Lawrence Berkeley Lab
ghweber@lbl.gov

Dmitriy Morozov
Lawrence Berkeley National Lab
dmitriy@mrzv.org

MS223

Rational Krylov Subspace Methods for Transient Electromagnetic Geophysical Forward Modeling

In recent work we explored how the initial value problem for the quasi-static Maxwell's equations arising in transient electromagnetic modeling (TEM) can be solved efficiently in the frequency domain by solving a small projected model for each relevant frequency using simple shift-and-invert type Krylov subspace projection much in the style of classical model order reduction for linear time-invariant control problems. In this work we present more advanced rational Krylov subspace methods employing more elaborate pole selection techniques. We also compare this frequency domain approach with solving the problem in the time domain using the same rational Krylov subspace methods to approximate the action of the matrix exponential.

Oliver G. Ernst
TU Bergakademie Freiberg
Fakultaet Mathematik und Informatik
ernst@math.tu-freiberg.de

Stefan Güttel
School of Mathematics
The University of Manchester
stefan.guettel@manchester.ac.uk

Ralph-Uwe Börner
Institute of Geophysics
TU Bergakademie Freiberg
rub@geophysik.tu-freiberg.de

MS223

Rational Approximations through Finite Element Discretization

Abstract not available at time of publication.

Murthy N. Guddati
North Carolina State University
mnguddat@ncsu.edu

MS223

Spectrally Adaptive Rational Quadrature of Markov Functions

Rational Arnoldi is a powerful method for approximating expressions of the form $f(A)b$ or $b^H f(A)b$, where A is a large sparse matrix, $f(A)$ is a matrix function, and b is a vector. The selection of asymptotically optimal parameters for this method is crucial for its fast convergence. We present and investigate a novel strategy for the automated parameter selection when the function f is of Markov type, such as the square root or the logarithm. The performance of this approach is demonstrated by numerical examples involving symmetric and nonsymmetric matrices. This is joint work with Leonid Knizhnerman.

Stefan Guettel
University of Oxford
stefan.guettel@manchester.ac.uk

MS223

Recursive Relations for Rational Krylov Methods

Abstract not available at time of publication.

Lothar Reichel
Kent State University
reichel@math.kent.edu

MS224

Reproducible Research and Omics: Thoughts from the IOM Review

Between 2007 and 2010, several genomic signatures were used to guide patient therapy in clinical trials in cancer. Unfortunately, the signatures were wrong, and trials proceeded despite warnings to this effect. The Institute of Medicine (IOM) subsequently reviewed the level of evidence that should be required in such situations. Many recommendations focus on reproducibility and data integrity, including directives to funders, journals, and regulatory agencies. We briefly review the report and implications for reproducible research.

Keith Baggerly
Department of Bioinformatics and Computational Biology
UT MD Anderson Cancer Center
kabagger@mdanderson.org

MS224

A Portrait of One Scientist as a Graduate Student

In this talk, I will focus on the *how* of reproducible research. I will focus on specific tools and techniques I have found invaluable in doing research in a reproducible manner. In particular, I will cover the following general topics (with specific examples in parentheses): version control and code provenance (git), code verification (test driven development, nosetests), data integrity (sha1, md5, git-annex), seed saving (random seed retention) distribution of datasets (mirroring, git-annex, metalinks), light-weight analysis capture (ttyrec, ipython notebook)

Paul Ivanov
Redwood Center for Theoretical Neuroscience
University of California, Berkeley
pi@berkeley.edu

MS224

Reproducible Research in Graduate Education in the Computational Sciences

Instilling good habits of reproducible computational research can be woven throughout graduate student education. Current software tools allow drafting of live documents that contain theoretical derivations, generation of computational code, verifiable execution of code, and preparation of reports. The talk presents experience in this approach in graduate courses at UNC.

Sorin Mitran
University of North Carolina Chapel Hill
mitran@unc.edu

MS224

Publishing Reproducible Research: Thoughts on Journal Policy

I will also discuss the role that journals can play in encouraging reproducible research and will review the recent reproducibility policy at the journal Biostatistics.

Roger D. Peng
Department of Biostatistics
Johns Hopkins Bloomberg School of Public Health
rpeng@jhsph.edu

MS225

Modeling Gene Regulatory Networks through Bayesian Structure Learning

Gene regulatory networks modeled as Bayesian networks are known to provide high quality inference. However, exact Bayesian inference is NP-hard. In this talk, we present scalable parallel exact and heuristic methods for Bayesian inference. The exact method achieves both work and space optimality compared to the serial algorithm. The heuristic method operates by predicting and subsequently refining parent sets of the genes in the network. Experimental results are presented using synthetic and model pathways.

Olga Nikolova
Graduate student
Iowa State University
olia@iastate.edu

Jaroslaw Zola, Srinivas Aluru
Iowa State University
zola@iastate.edu, aluru@iastate.edu

MS225

PASQUAL: Parallel Techniques for Next Generation Genome Sequence Assembly

Short-read assembly of a genome from sequence data gathered from Next Generation Sequencing (NGS) platforms has become a highly challenging and fundamental problem in bioinformatics. In this talk we discuss algorithms and compressed string graph-based data structures for shared memory parallel architectures that address several computational challenges of short-read assembly. Experiments on up to 40 cores demonstrate that our method delivers the fastest time-to-solution and the best trade-off between speed, memory consumption, and solution quality.

Xing Liu, Pushkar Pande
School of Computational Science and Engineering
Georgia Institute of Technology
xing.liu@gatech.edu, pushkar.pande@gatech.edu

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

David A. Bader
Georgia Institute of Technology
bader@cc.gatech.edu

MS225

Theory, Application and Challenges for Graph-theoretic Models in Computational Biology

With the combinatorics inherent in the field of computational biology alongside a recent emergence of high-throughput instruments for data generation, the role of graph-theoretic modeling has been elevated to take a center stage in biological discovery. In this overview talk, we discuss the application potential and the imminent challenges of graph-theoretic analytics in modern-day computational biology viz. sheer volume of data, diversity in type, and a quest to reveal hidden data interconnectivity at different levels.

Ananth Kalyanaraman
Associate Professor
School of EECS, Washington State University
ananth@eecs.wsu.edu

MS225

Graph Algorithms in Flow Cytometry

Flow cytometry (FC) measures fluorescence of proteins in single cells, enabling us to identify cellular subpopulations relevant to immunology, cell signaling, and oncology. Recent experimental developments lead to large-scale, high dimensional FC data, necessitating novel combinatorial algorithms for analyzing the data. We describe algorithms for registering cell populations across multiple samples, experimental conditions and times. The combinatorics employed include edge covers of minimum weight and several concepts from phylogenetic combinatorics.

Ariful Azad
Purdue University
aazad@purdue.edu

Arif Khan
Department of Computer Science
Purdue University
khan58@cs.purdue.edu

Bartek Rajwa
Bindley Bioscience Center
Purdue University
brajwa@purdue.edu

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

MS226

Approximation of the Fokker-Planck Equation of FENE Dumbbell Model

We propose a new weighted weak formulation for the Fokker-Planck equation of FENE dumbbell model, and prove its well-posedness in weighted Sobolev spaces. We also propose simple and efficient semi-implicit time-discretization schemes and prove that they are unconditionally stable. We then construct two Fourier-Jacobi spectral-Galerkin algorithms which enjoy the following properties: (i) it is unconditionally stable, spectrally accurate and of optimal computational complexity; (ii) it conserves the volume naturally, and provide accurate approximation to higher-order moments of the distribution function; and (iii) it can be easily extended to coupled non-homogeneous systems. Extension to full Navier-Stokes Fokker-Planck system by using sparse spectral methods will also be discussed.

Jie Shen
Dept. of Mathematics
Purdue University
shen7@purdue.edu

MS226

Modeling Tissue Self-assembly in Bio-fabrication using Kinetic Monte Carlo Simulations

We present a three-dimensional lattice model to study self-assembly and fusion of multicellular aggregate systems by using kinetic Monte Carlo (KMC) simulations. This model is developed to describe and predict the time evolution of postprinting morphological structure formation during tissue or organ maturation in a novel biofabrication process (or technology) known as bioprinting. In this new technology, live multicellular aggregates as bio-ink are used to make tissue or organ constructs via the layer-by-layer deposition technique in biocompatible hydrogels; the printed bio-constructs embedded in the hydrogels are then placed in bioreactors to undergo the self-assembly process to form the desired functional tissue or organ products. Here we implement our model with an efficient KMC algorithm to simulate the making of a set of tissues/organs in several designer's geometries like a ring, a sheet and a tube, which can involve a large number of cells and various other support materials like agarose constructs etc. We also study the process of cell sorting/migration within the cellular aggregates formed by multiple types of cells with different adhesivities.

Yi Sun
University of South Carolina
yisun@math.sc.edu

MS226

Theoretical and Computational Advances in Modeling Active Nematic Liquid Crystal Polymers and its Applications

Active liquid crystal polymers can be found in many man-made and natural material systems. Their features include spontaneous local molecular orientation and self-propelled molecular motion. The F-actin in cytoplasm is an outstanding example of the active material system. The traditional (passive) liquid crystal theory has been modified to account for the active (nonequilibrium) forcing. In this talk, I will discuss a mathematical analysis of a class of polar nematic active liquid crystal models in simple shear, Poiseuille flows, and other simple geometries under various stress balance conditions. 2-D numerical simulation with respect to two types of boundary conditions will be discussed as well. A rich set of spatial-temporal patterns in collective molecular orientation and flows will be revealed and so will be their relation to the active parameters in the model.

Qi Wang
University of South Carolina
qwang@math.sc.edu

MS226

Numerical Stability for Incompressible Euler Equation

Fully discrete pseudo spectral numerical schemes to 2-D and 3-D incompressible Euler equation are presented in this talk. To ensure the numerical stability for this inviscid equation, an artificial viscosity is added with a required numerical accuracy. A local in time convergence analysis is established for smooth solution and a global in time energy stability is assured with a suitable choice of the artificial viscosity term.

Cheng Wang
University of Massachusetts
Dartmouth, MA
cwang1@umassd.edu

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

MS227

Model Form Issues in Uncertainty Quantification

Abstract not available at time of publication.

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS227

Extreme-Scale Stochastic Inversion

We present a computational framework for solution of discretized infinite-dimensional Bayesian inverse problems. We address several computational issues related to the appropriate choice of prior, consistent discretizations, tractable treatment of the Hessian of log posterior, and scalable parallel MCMC algorithms for sampling the posterior. We apply the framework to the problem of global seismic inversion, for which we demonstrate scalability to 1M earth parameters, 630M wave propagation unknowns, and 100K cores on the Jaguar supercomputer.

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

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

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

Georg Stadler
University of Texas at Austin
georgst@ices.utexas.edu

Lucas Wilcox
HyPerComp
me@lucaswilcox.com

MS227

Bayesian Inference of Wind Drag using AXBT Data

An adapative, sparse, pseudospectral sampling algorithm is applied to efficiently construct a polynomial chaos surrogate of the response of the ocean circulation to uncertain drag parameters. The surrogate is then exploited to infer uncertain drag parameters, using AXBT data collected during typhoon Fanapi. The analysis leads to sharp estimates for the saturation of the wind drag coefficient and the corresponding wind speed; however, the data was not informative regarding the drag behavior at higher speeds.

Ihab Sraj
Duke University
ihab.sraj@duke.edu

Mohamed Iskandarani
Rosenstiel School of Marine and Atmospheric Sciences
University of Miami
MIskandarani@rsmas.miami.edu

Ashwanth Srinivasan, Carlisle Thacker
University of Miami
asrinivasan@rsmas.miami.edu, carlisle.thacker@noaa.gov

Justin Winokur
Duke University
justin.winokur@duke.edu

Alen Alexanderian
Johns Hopkins University
aalexa20@jhu.edu

Chia-Ying Lee, Shuyi Chen
University of Miami
chiaying@orca.rsmas.miami.edu, schen@rsmas.miami.edu

Omar M. Knio
Duke University
omar.knio@duke.edu

MS227

Multiscale Methods for Large-Scale Bayesian Inversion

In many ill-posed inverse problems, especially those involving transport in porous media, the forward model smooths the input parameters. In this context, we can exploit conditional independence of scales, identified through multiscale forward solvers, to efficiently sample the Bayesian posterior. We sample a low-dimensional coarse-scale problem with MCMC and "project' these samples to the fine scale with approximate iterative techniques. This approach is well-suited to parallel computation. We will describe not only the algorithm but its implementation via a new high-level inversion toolbox.

Matthew Parno, Youssef M. Marzouk
Massachusetts Institute of Technology
mparno@mit.edu, ymarz@mit.edu

MS228

Iterative Methods and Spectral Approximation of Fast Rotating Gross-Pitaevskii Equations

The aim of this talk is to get fast converging pseudospectral methods for computing stationary solutions to the Gross Pitaevskii equation with large rotations and nonlinearites. The pseudospectral method is standard and based on FFTs. Concerning the numerical solution of the resulting non explicit linear system, we show that usual methods used in the literature makes the solver diverge. We propose to use Krylov solvers and build some physics based preconditioners that provide substantial gains in term of computational times to reach the ground state with very high accuracy. Numerical validations for 2D and 3D case will support our approach.

Xavier L. Antoine
Institut Elie Cartan Nancy (IECN)
Université de Lorraine
xavier.antoine@univ-lorraine.fr

Romain Duboscq
IECN, Nancy, France
romain.duboscq@iecn.u-nancy.fr

MS228

Improved Sobolev Gradient Methods for Solving the Stationary Gross-Pitaevskii Equation with Rotation

We compute vortex states of a rotating Bose-Einstein condensate by direct minimization of the Gross-Pitaevskii energy functional. We extensively compare different minimization algorithms with improved steepest descent methods based on Sobolev gradients. In particular, we show that mixed Newton-Sobolev gradient methods offer appealing convergence properties. Advantages and drawbacks of each method are summarized. We present numerical setups using 6th order finite difference schemes and finite elements with mesh adaptivity that were successfully used to compute a rich variety of difficult cases with quantized vortices.

Ionut DANAILA
Laboratoire de Mathématiques Raphaël Salem, Université
de
Rouen, FRANCE
ionut.danaila@univ-rouen.fr

Parimah Kazemi

Ripon College, Ripon, Wisconsin, USA
parimah.kazemi@gmail.com

MS228

Dimension Reduction of the Nonlinear Schrdinger Equation with Coulomb Interaction under Anisotropic Potentials

We present a rigorous dimension reduction analysis from the 3D (i.e. in 3 spatial dimensions) Schrödinger-Poisson system (SPS) to lower dimensional models, arising in the limit of infinitely strong confinement in two or one space dimensions, namely the Surface Adiabatic Model (SAM) and the Surface Density Model (SDM) in 2D or Line Adiabatic Model (LAM) in 1D. In particular, we explain and demonstrate that the 2D Schrödinger-Poisson Model (SPM) is not appropriate to simulate a "2D electron gas" of point particles confined into a plane (or more general 2D manifold), whereas SDM as the correct model can be successfully applied by utilizing the inversion of Square root of Laplacian to describe the Coulomb interaction. Finally, we study the ground state and dynamics of SDM in different setups.

Yong Zhang
Wolfgang Pauli Institute, Wien, Austria
sunny5zhang@gmail.com

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

Huaiyu Jian
Dept. of Mathematical Sciences, Tsinghua Univ.
hjian@math.tsinghua.edu.cn

Norbert Mauser
Wolfgang Pauli Institute c/o Fak. Mathematik
mauser@courant.nyu.edu

MS228

Numerical Methods for Rotating Dipolar BEC based on a Rotating Lagrange Coordinate

In the talk, a simple and efficient numerical method for studying the dynamics of rotating dipolar BEC is presented. Using coordinate transformation, we eliminate the angular rotational term from the Gross-Pitaevskii equation (GPE). Then we develop pseudospectral type methods to solve the GPE. The accuracy of the method is discussed. In addition, the dynamics of quantized vortex lattice in rotating dipolar BEC are studied.

Yanzhi Zhang
Missouri University of Science and Technology
zhangyanz@mst.edu

Weizhu Bao
National University of Singapore
Department of Mathematics
bao@math.nus.edu.sg

Qinglin Tang
National University of Singapore
g0800880@nus.edu.sg

MS229

A New Spectral Method for Numerical Solution of the Unbounded Rough Surface Scattering Problem

I shall present a new spectral method for solving the unbounded rough surface scattering problem. The method uses a transformed field expansion to reduce the boundary value problem with a complex scattering surface into a successive sequence of transmission problems of the Helmholtz equation with a plane surface. Hermite orthonormal basis functions are used to further simply the transmission problems to fully decoupled one-dimensional two-point boundary value problems with piecewise constant wavenumbers, which can be solved efficiently by a Legendre-Galerkin method. Ample numerical results will be presented to demonstrate the new spectral method is efficient, accurate, and well suited to solve the scattering problem by unbounded rough surfaces.

Ying He, Peijun Li
Department of Mathematics
Purdue University
he14@math.purdue.edu, lipeijun@math.purdue.edu

Jie Shen
Purdue University
Department of Mathematics
shen@math.purdue.edu

MS229

Spectral Element Discontinuous Galerkin Lattice Boltzmann Method for Convection Heat Transfer

Spectral Element Discontinuous Galerkin (SEDG) Lattice Boltzmann Method (LBM) employs body-fitted unstructured mesh and is capable of dealing with complex geometry with high-order accuracy. Explicit time marching, diagonal mass matrix and minimum communication cost for numerical flux make SEDG-LBM very efficient flow solver. We present SEDG-LBM algorithms based on two-distribution function and hybrid approaches and compare them for convection heat transfer problems. Various types of boundary conditions will be described in detail.

Saumil Patel, Kalu Chibueze Uga
Department of Mechanical Engineering
The City College of New York
saumil.patel134@gmail.com, uga.kalu@gmail.com

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Taehun Lee
Department of Mechanical Engineering
The City College of New York
thlee@ccny.cuny.edu

MS229

Boundary Perturbation Methods for Surface Plasmon Polaritons

Surface Plamson Polaritons (SPPs) arise in a number of technologically important applications including enhancement of signals, highly sensitive sensing, and the design of metamaterials. In this talk we discuss several new classes of Boundary Perturbation Methods for simulating SPPs which are accurate, robust, and extremely fast. We will outline the numerical methods, display simulation results for SPP configurations, and discuss how these algorithms can be effectively utilized for optimization and design of devices.

David P. Nicholls
University of Illinois at Chicago
nicholls@math.uic.edu

MS229

Boundary Treaments for Second Order Wave Equations by Pseudospectral and Runge-Kutta-Nystrm Methods

Many wave phenomena are described by second order wave equations, for instance, general relativity, acoustics, and electromagnetic waves. For these wave problems, the domain may be large compared the wavelength, and the waves have to propagate long distances. Hence, numerically solving these problems requires long time integration and accumulation of numerical dispersion error may affect the simulation quality. It can be shown that in preserving low accumulation of dispersion error during long time integration, high-order accurate methods are more efficient than the low-order methods. However, high-order schemes are very sensitive to the imposition of boundary conditions, and great care must be exercised to ensure stable computations of high-order schemes. In this talk we present a high-order scheme based on the pseudospectral Legendre approximation in space and the Runge-Kutta-Nyström algorithm in time, which can be adopted in a multi-domain computational framework, to solve second order wave equations. The key toward to the success of constructing such a stable scheme hinges upon properly imposing penalty boundary conditions at every collocation equations. We shall use one dimensional space problems to illustrate the conceptual ideas of the methods, and special attention is paid upon analyzing the stability of the scheme subject to various types of boundary conditions, including Dirichlet, Neumann, Robin, and material interface boundary conditions. Through conducting energy estimates it is shown that the scheme can be made stable by properly choosing the penalty parameters. The present one-dimensional scheme can be generalized for multidimensional problems even defined on curvilinear coordinates. Numerical experiments for model problems defined on one and two dimensional spaces are conducted, and we observe the expected convergence results.

Chun-Hao Teng
Department of Applied Mathematics
National Chiao Tung University
chunhao.teng@gmail.com

MS230

Excelling: Transition from School to Aerospace and Defense Industry

There are commonalities between being a graduate researcher and working in the defense industry - both require well defined scope of the problem and increasingly the work in done in collaborative teams. There are, however, also significant differences. School environment is usually a very protective, productive, and creative one. You are taught to raises expectations from your own self. Excellence is defined by your capability to learn. The transition from school where the focus is on learning without constraints to the defense industry, which is governed by requirements set by the customer, can be a very turbulent one. In a generic defense industry, a mark of excellence is how well one learns to adapt to the hierarchy and conforms to a certain behavior pattern. The work in graduate school expects a very out of box thinking and a solution over longer time scales than the industry which expects a very process oriented thinking. Since the goals are so different, transitioning from one environment to another can be turbulent, especially, if you are a strong and intelligent woman engineer/scientist.

Roochi Chopra
Raytheon
roochi_chopra@raytheon.com

MS230

A Dual Career: Experiences as a Researcher and a Program Manager

In this talk I will present an overview of my experiences as a former academic (and still active researcher) and a program manager for applied and computational mathematics at the Air Force Office of Scientific Research (AFOSR), and will discuss the challenges and rewards of pursuing this dual career.

Fariba Fahroo
Air Force Office of Scientific Research
fariba.fahroo@afosr.af.mil

MS230

Experiences in Industry

In this talk, I will discuss my experience as an industrial post-doc at IMA UMN and as a researcher in industrial environment at IBM Watson research lab.

Raya Horesh
Institute of Mathematics and its Applications
University of Minnesota
rhoresh@ima.umn.edu

MS230

The Best of Both: Federally Funded Research and Development Center (FFRDCs) at the Juncture of Industry and Academia

The work performed at Federally Funded Research and Development Centers straddles academic research and industrial design and development. Such an environment provides insight into the advancement of new technologies and analysis along with insight into bringing these advances into real world designs. The panelist will focus on the benefits and drawbacks of working concurrently on an expansive array of projects as well as the characteristics necessary to be successful in such a career.

Julia Mullen
Computing & Communications Ctr
Worcester Polytechnic Inst
jsm@wpi.edu

MS231

Energy Aware Performance Metrics

Energy aware performance metrics are absolutely necessary in order to better appreciate the performance of algorithms on modern architectures. Although recent advances, such as the FLOPS/WATT metric, gave an important push to the right direction, we will show that deeper investigations are needed, if we are to overcome the power barrier for reaching Exaflop. We will showcase tools that allow accurate, on chip power measurements that shed new light on the energy requirements of important kernels.

Costas Bekas, Alessandro Curioni
IBM Research - Zurich
BEK@zurich.ibm.com, zur@zurich.ibm.com

MS231

Application-aware Energy Efficient High Performance Computing

The energy cost of running an HPC system can exceed the cost of the original hardware purchase. This has driven the community to attempt to understand and minimize energy costs wherever possible. We present an automated framework, Green Queue, for customized application-aware Dynamic Voltage-Frequency Scaling (DVFS) settings to reduce the energy consumption of large scale scientific applications. Green Queue supports making CPU clock frequency changes in response to intra-node and internode observations about application behavior. Our intra-node approach reduces CPU clock frequencies and therefore power consumption while CPUs lacks computational work due to inefficient data movement. Our inter-node approach reduces clock frequencies for MPI ranks that lack computational work. We investigated these techniques on a set of large scientific applications on 1024 cores of Gordon, an Intel Sandybridge based supercomputer at the San Diego Supercomputer Center. Our optimal intra-node technique showed an average measured energy savings of 10.6% and a maximum of 21.0% over regular application runs. Our optimal inter-node technique showed an average 17.4% and a maximum of 31.7% energy savings.

Laura Carrington
San Diego Supercomputing Center
lcarring@sdsc.edu

MS231

A 'Roofline' Model of Energy and What it Implies for Algorithm Design

Abstract not available at time of publication.

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

MS231

Power Bounds and Large Scale Computing

Energy and power are widely recognized as significant challenges for future large scale systems. Current processors, in particular the Intel Sandy Bridge family, already include mechanisms to limit power levels dynamically. However, these mechanisms apply without consideration of the power levels of other nodes, which may be lower and allow a "hot" node to consume more power. This talk will discuss options and techniques for limiting power and energy consumption that better suit large-scale systems. It will also detail initial experiences with Intel's Running Average Power Limit (RAPL) on a large Linux system and discuss possible extensions.

Bronis R. de Supinski
Livermore Computing
Lawrence Livermore National Laboratory
bronis@llnl.gov

MS232

Convex Collective Matrix Factorization

In many realistic applications, multiple interlinked sources of data are available and they cannot be easily represented in the form of a single matrix. Collective matrix factorization has recently been introduced to improve generalization performances by jointly factorizing multiple relations or matrices. In this paper, we extend the trace norm for matrix factorization to the collective matrix factorization case. This norm defined on the space of relations is used to regularize the empirical loss, leading to a convex formulation of the problem. Similarly to the trace norm on matrices, we show that the collective-matrix completion problem admits a fast iterative singular-value thresholding algorithm. The collective trace norm is also characterized as a decomposition norm, useful to find an optimal solution thanks to an unconstrained minimization procedure. Empirically we show that stochastic gradient descent suits well for solving the convex collective factorization even for large scale problems. We also show that the proposed algorithm directly solving the convex problem is much faster than unconstrained gradient minimization optimizing in the space of low-rank matrices.

Guillaume Bouchard
Xerox Research Centre Europe
guillaume.bouchard@xrce.xerox.com

MS232

Generalised Coupled Tensor Factorisation and Graphical Models

We discuss coupled tensor factorisation from a statistical perspective, building on generalised linear models and probabilistic graphical models. We express a model using a factor graph where the factorisation is achieved via a message passing algorithm. This provides a practical approach that enables development of application specific custom models, Bayesian model selection as well as algorithms for joint factorisations where several tensors are factorised simultaneously. We illustrate the approach on signal processing applications.

A. Taylan Cemgil
Bogazici University
taylan.cemgil@boun.edu.tr

MS232

Linked Multilinear Component Analysis, Multiway Canonical Correlation Analysis and Partial Least Squares

We will present models and algorithms for multiway component analysis, tensor canonical correlation analysis (MCCA) and higher-order partial least squares (HO-PLS). We also discuss emerging models and approaches for multi-block constrained matrix/tensor decompositions in applications to group- and linked-multiway component analysis, feature extraction, classification and clustering. We also overview our recent related results: http://www.bsp.brain.riken.jp/ cia/recent.html#tensor

Andrzej Cichocki
RIKEN Brain Science Institute
cia@brain.riken.jp

MS232

Exact Line and Plane Search for Tensor Optimization by Global Minimization of Bivariate Polynomials and Rational Functions

Line search (LS) and plane search (PS) are an integral component of many optimization algorithms. We pose independent component analysis (ICA) as a data fusion problem in which a PS subproblem naturally arises. In tensor optimization problems LS and PS often amount to minimizing a polynomial. We introduce a scaled LS and PS and show they are equivalent to minimizing a rational function. Lastly, we show how to compute the global minimizer of real and complex (scaled) LS and PS problems accurately and efficiently by means of a generalized eigenvalue decomposition.

Laurent Sorber
KU Leuven
Laurent.Sorber@cs.kuleuven.be

Ignat Domanov
KU Leuven Kulak
ignat.domanov@kuleuven-kulak.be

Marc Van Barel
Katholieke Universiteit Leuven
Department of Computer Science
marc.vanbarel@cs.kuleuven.be

Lieven De Lathauwer
KU Leuven Kulak
lieven.delathauwer@kuleuven-kulak.be

MS233

NWChem Quantum Many-body Methods on the Intel MIC Architecture

The NWChem quantum many-body methods found in the Tensor Contraction Engine module are widely used for scientific applications but are both computation and memory intensive, hence require large parallel computers. We present our initial findings concerning the porting of these methods to the Intel MIC architecture using multiple programming models. We apply inspector-executor schemes to eliminate the overhead of dynamic load-balancing and allow for greater data reuse, which reduces both intranode and internode communication.

Jeff R. Hammond
Argonne National Laboratory
Leadership Computing Facility
jhammond@alcf.anl.gov

David Ozog
University of Oregon
david.m.ozog@gmail.com

MS233

Optimizing Numerical Weather Prediction Performance and Scaling on the Intel MIC

Further increases in weather and climate simulation capability will depend on exploiting all levels of application parallelism. What are the opportunities for exposing and exploiting parallelism in atmospheric models despite relatively large memory footprints and low computational intensities? What are the challenges and prospects for programmability? We present status and preliminary results of efforts to employ the Intel Many Integrated Core processor in the Weather Research and Forecast (WRF) model.

John Michalakes
National Renewable Energy Laboratory
john.michalakes@nrel.gov

MS233

Parallelization Strategies for High-order Discretized Hyperbolic PDEs

We target high-performance implementations for the solution of hyperbolic PDEs on Stampede using a high-order discontinuous Galerkin finite element discretization. One key challenge in achieving high throughput for these applications is that these are heavily memory-bound; additionally the flop/mop ratio changes with the polynomial order of the elemental basis. The other challenge arises from the asymmetry between the CPU and the co-processor making load-balance difficult. We consider parallelization strategies addressing these challenges.

Hari Sundar
Institute for Computational and Engineering Sciences
University of Texas at Austin
hari@ices.utexas.edu

Jesse Kelly, Omar Ghattas
University of Texas at Austin
jkelly@ices.utexas.edu, omar@ices.utexas.edu

MS233

Evaluating Intel's Many Integrated Core Architecture for Climate Science

Abstract not available at time of publication.

Henry Tufo
NCAR
University of Colorado at Boulder
tufo@cs.colorado.edu

MS234

Solving Sequences of Linear Systems for Stochastic Collocation based Uncertainty Quantification

Using stochastic collocation for uncertainty quantification in partial differential equations (PDEs) with random data requires solving sequences of linear systems. We use Krylov subspace recycling to obtain efficient solution of such sequences. We propose a new recycle subspace selection criterion, describe a new recycling algorithm, and adapt the existing recycling solvers. The underlying PDEs include an elliptic diffusion equation, Maxwell's equations, and a heat equation. The results show speed ups of up-to 55% in time.

Kapil Ahuja
Max Planck Institute
for Dynamics of Complex Technical Systems
ahuja@mpi-magdeburg.mpg.de

Michael Parks
Sandia National Laboratories
parks@sandia.gov

Eric De Sturler

Virginia Tech
sturler@vt.edu

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

MS234

On using AMG Preconditioners to Accelerate the XFEM-Monte Carlo Approach for Uncertainty Quantification in Homogenization

An AMG approach is proposed for solution of linear systems arising from unit cell problems within the XFEM-Monte Carlo framework to quantify uncertainties in homogenization. The XFEM linear systems can be partitioned to contain an invariant sub-matrix corresponding to the standard DOFs along with a sub-matrix corresponding to the enrichment DOFs that change for each Monte Carlo realization. This property is harnessed to optimally reuse parts of the preconditioner thereby reducing the AMG setup costs.

Badri Hiriyur
Weidlinger Associates Inc., USA
badri.hiriyur@wai.com

MS234

An Algebraic Multigrid Solver for Elliptic PDEs with Random Coefficients

In this talk we present a study of performance for algebraic multigrid (AMG) method on elliptic partial differential equations (PDEs) with random coefficients arising from discretization of groundwater flow problem. We will further discuss the effects of use of fixed coarse grids and computational cost savings by this approach.

Minho Park
University of Nottingham
min.park@nottingham.ac.uk

MS234

Preconditioning Stochastic Collocation Saddle-Point Systems

We discuss linear algebra issues and preconditioning for saddle point systems arising from stochastic collocation discretisations of PDEs with random data. Model problems include: a mixed formulation of a second-order elliptic problem (with uncertain diffusion coefficient), the steady-state Navier-Stokes equations (with uncertain Reynolds number) and a second-order elliptic PDE on a random domain, solved via the fictitious domain method.

Catherine Powell
School of Mathematics
University of Manchester
c.powell@manchester.ac.uk

MS235

Implicit Timesteping of the Community Atmosphere Model Spectral Element Code Utilize GPU Accelerators

With the commissioning of hybrid multicore HPC systems, such as the Cray XK6 Titan supercomputer at the Oak Ridge Leadership Computing Facility (OLCF), we are bridging the way to the exascale era of high performance computing which will extend the resolution of regional climate predictions with realistic clouds and chemistry. In order to realize the full potential of Titan and similar hardware, we are enhancing the Community Earth System Model (CESM) to take advantage of the GPU accelerators. Here, we discuss the adaptation of the spectral element dycore in the Community Atmospheric Model(CAM) to run on hybrid architecture and the step towards implicit timesteping.

Rick Archibald
Computational Mathematics Group
Oak Ridge National Labratory
archibaldrk@ornl.gov

MS235

Development of an IMEX Integration Method for Sea Ice Dynamics

Most sea ice model numerical schemes involve a splitting in time between the momentum and the continuity equations. It has been shown that this approach can lead to a numerical instability or failure of the momentum equation implicit solver for large advective time steps. To cure this problem, an IMplicit EXplicit (IMEX) time integration technique is being developed. Results show that the IMEX technique improves the accuracy of the solution and reduces the number of failures of a JFNK solver. These less frequent failures also lead to a decrease in the CPU time required.

Jean-François Lemieux
Environment Canada
jean-francois.lemieux@ec.gc.ca

MS235

Fast Offline Ocean Tracer Transport Model for the Community Earth System Model

Abstract not available at time of publication.

Francois Primeau
Earth System Science Dept.
University of California, Irvine
fprimeau@uci.edu

MS235

Software and Algorithms for Implementing Scalable Solvers in Climate Codes

We have had success in using the Trilinos software from the FASTMath SciDAC institute to implement implicit and steady-state solvers in climate codes. Since each application has different requirements and capabilities, having a toolbox of available software and algorithms has led to flexibility and extensible implementations that can be tailored to the application. We will show results from Ice Sheet simulations (CISM) and progress in atmosphere, ocean, sea ice, and tracer transport applications.

Andrew Salinger
Applied Computational Methods Dept, Sandia National Labs
Sandia National Labs
agsalin@sandia.gov

Katherine J. Evans

Oak Ridge National Laboratory
evanskj@ornl.gov

MS236

An Efficient Resampling Algorithm for Estimation of Fisher Information Matrix using Prior Information

The Fisher information matrix (FIM) is a critical quantity in several aspects of mathematical modeling, e.g., confidence interval calculation, optimal input selection in experimental design. Analytical determination of the FIM in a general setting may be difficult or almost impossible due to intractable modeling requirements or/and intractable high-dimensional integration. We will present an efficient resampling algorithm for estimation of the FIM, particularly, for the cases when some elements of the FIM are analytically known a priori, and the rest of the elements are unknown. Such an interesting structure of the FIM is reported to be observed in the context of many engineering applications.

Sonjoy Das
University at Buffalo
sonjoy@buffalo.edu

James C. Spall
Johns Hopkins University
james.spall@jhuapl.edu

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

MS236

Design or Experiments for Multi-phyiscs Problems

We demonstrate the use of information theory on inversion for multi physics problems. In particular we evaluate the use of different measurements on a multi-physics prototype that emulates glacier flow, atmospheric transport and heat. Information metrics are explored to help manage different measurements. We have developed a flexible and extensible computational framework based on various components from Trilinos. This enables large scale optimization through adjoints, which is used in reduced and full space optimization algorithms.

Andrew Davis
MIT
davisad@mit.edu

Bart Vanbloemenwaanders
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
bartv@sandia.gov

MS236

Data-driven Model Mis-specification Mitigation

In the context of inverse problems, often the ability to simulate data is limited to the extent that only a discrete, incomplete observation operator can be specified. Such misspecified observation operator may describe approximated physics, approximated numerics, linearization of non-linear processes, etc. In this study, we developed a nuclear norm stochastic optimization technique to supplement the model with rank restriction given a training set of models and data.

Lior Horesh
Business Analytics and Mathematical Sciences
IBM TJ Watson Research Center
lhoresh@us.ibm.com

Ning Hao
Tufts University
Department of Mathematics
ning.hao@tufts.edu

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

MS236

Coherence Metric for Optimal Compressive Sensing

The efficiency of compressive sensing depends on random projections of the sampled signal. However, in practical encoders/decoders, one uses deterministic matrices, which have non-zero coherence between sampling matrices and the bases used to represent the signal sparsely. This results in sub-optimal sampling. We investigate how the degree of coherence correlates with the accuracy of reconstruction (i.e., completeness of the estimated sparse respresentation) as well as the uncertainty in the estimated signal.

Sean A. Mckenna
Sandia National Laboratories
samcken@sandia.gov

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

MS237

Coarse-grained Modeling of Supported and Tethered Bilayers

Biomembranes compartmentalize cells and are fundamental for any living organism. The fundamental building block for biomembranes are lipid bilayers. Here we present a multiscale molecular modeling investigation of biomembranes. Especially the interplay between membranes and supports will be discussed. Supported Lipid Bilayers are an abundant research platform for understanding the behavior of cell membranes as they allow for additional mechanical stability and enable characterization techniques not reachable otherwise. However, in computer simulations these systems have been studied only rarely up to now. We present systematic studies on different length scales of the changes that a support and tethering to it inflicts on a phospholipid bilayer using molecular modeling. We characterize the density and pressure profiles as well as the density imbalance induced by the support. We determine the diffusion coefficients and characterize the influence of corrugation of the support. We also measure the free energy of transfer of phospholipids between leaflets using the coarse-grained Martini model.

Roland Faller
Department of Chemical Engineering and Materials Science
University of California, Davis

rfaller@ucdavis.edu

MS237

Stochastic Reaction-diffusion Simulation on Multiple Scales

Abstract not available at time of publication.

Mark Flegg
Mathematical Institute
University of Oxford
flegg@maths.ox.ac.uk

MS237

Matrix Calculations in Diffusion Approximations for Molecular Motors

A discussion of matrix methods to bridge nano-scale molecular motor model, which includes both diffusive and kinetic components, to a meso-scale stepping model. Such a method serves as an alternative to monte carlo methods, facilitating rapid sensitivity analysis over a large parameter space.

John Fricks
Dept of Statistics
Pennsylvania State University
fricks@stat.psu.edu

MS237

Fluctuating Lipid Bilayer Membranes with Diffusing Protein Inclusions: Hybrid Continuum-Particle Numerical Methods

Many proteins through their geometry and specific interactions with lipids induce changes in local membrane material properties. To study such phenomena we introduce a new hybrid continuum-particle description for the membrane-protein system that incorporates protein interactions, hydrodynamic coupling, and thermal fluctuations. We investigate how collective protein effects influence membrane mechanical properties. We discuss interesting numerical aspects that are required to obtain good translation invariance. Finally, we discuss a coarse grained model that incorporates important hydrodynamics.

Jon Karl Sigurdsson
Department of Mathematics
University of California, Santa Barbara
jonksig@gmail.com

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

MS238

Multi-level Communication Avoiding LU and QR Factorizations for Hierarchical Platforms

This study focuses on LU and QR factorizations of dense matrices on multi-level hierarchical platforms. We first introduce a new platform model called HCP. The focus is set on reducing the communication requirements of the studied algorithms at each level of the hierarchy. We extend lower bounds on communications to the HCP model and introduce two multi-level LU and QR algorithms tailored for those platforms. Finally, we provide a detailed performance analysis.

Mathias Jacquelin
LRI - INRIA Saclay, France
mathias.jacquelin@inria.fr

Amal Khabou
NRIA Saclay, 4 Rue Jacques Monod 91893
amal.khabou@inria.fr

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

MS238

A Communication Optimal N-Body Algorithm for Long-Range Direct Interactions

Algorithms for long-range, direct particle interactions traditionally work by dividing the particles among processors and synchronously exchanging particles to compute interactions. This approach works well when computation dominates communication, but it breaks down at scale when communication dominates. We present a communication-optimal algorithm that uses extra memory to replicate the particles. In practice, the algorithm yields modest speedups on communication-bound problem instances and provides significantly better strong scaling to tens of thousands of processors.

Michael Driscoll
UC Berkeley
mbdriscoll@eecs.berkeley.edu

Evangelos Georganas
EECS
UC Berkeley
egeor@eecs.berkeley.edu

Penporn Koanantakool
UC Berkeley
penpornk@eecs.berkeley.edu

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

Katherine Yelick
UC Berkeley
Lawrence Berkeley National Laboratory
kayelick@lbl.gov

MS238

Algorithmic Adventures in 3D

We describe an extension of the Scalable Universal Matrix Multiplication Algorithms (SUMMA) from 2D to 3D process grids; the underlying idea is to lower the communication volume through storing redundant copies of one or more matrices. This talk focuses on element-wise matrix distributions, which lead to allgather-based algorithms. We begin by describing an allgather-based 2D SUMMA, describe its generalization to 3D process grids, and then discuss theoretical and experimental performance benefits of the new algorithms.

Martin D. Schatz

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

MS238

Shape-morphing in LU Factorizations

We present a sequential communication-avoiding LU factorization algorithm. By communication avoiding, we mean that the volume of data transferred in cache misses or page faults (or explicit input-output operations) is close to optimal, and that the size of each data transfer can be proportional to the size of the fast memory (cache). The key idea in the algorithm is to morph the data layout of the reduced matrix and the factors repeatedly throughout the algorithm. The goal is to use a column-major when columns are eliminated, but to use block layout during matrix multiplication and triangular solves. Our key finding is that the shape-morphing steps do not contribute significantly to the asymptotic communication complexity of the computation.

Grey Ballard
UC Berkeley
ballard@cs.berkeley.edu

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

Benjamin Lipshitz
UC Berkeley
lipshitz@berkeley.edu

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

Oded Schwartz
UC Berkeley
oded.schwartz@gmail.com

MS239

Randomized Asynchronous Iterative Linear Solvers

Asynchronous methods for solving linear equations have been researched since Chazan and Miranker published their pioneering paper on *chaotic relaxation* in 1969. While these methods were shown to converge, no bounds were proven on the rate of convergence, and in practice their convergence was sometimes shown to be very slow. In this talk we will present new asynchronous iterative linear solvers. A key component in the new algorithms is the use of randomization. We show that unlike previous algorithms, the rate of convergence for new algorithms can be bounded (in expectancy), and that they can be competitive in practice.

Haim Avron
Business Analytics & Mathematical Sciences
IBM T.J. Watson Research Center
haimav@us.ibm.com

Alex Druinsky
Tel-Aviv University
alexdrui@post.tau.ac.il

Anshul Gupta, Prabhanjan Kambadur
IBM T J Watson Research Center
anshul@us.ibm.com, pkambadu@us.ibm.com

MS239

Investigating the Convergence of Asynchronous Iterative Methods

When solving large sparse systems of linear equations in parallel with synchronous iterative techniques, it is necessary to exchange newly-computed vector elements between processors at every iteration. This talk describes our current work on developing and analysing asynchronous iterative algorithms that avoid time-consuming synchronous communication operations and thus allow execution to proceed without having to wait for new data to arrive. We present theoretical results supported by practical experimentation.

Iain Bethune
The University of Edinburgh
ibethune@epcc.ed.ac.uk

J. Mark Bull
Edinburgh Parallel Computing Centre
m.bull@epcc.ed.ac.uk

Nicholas Dingle
School of Mathematics, University of Manchester
nicholas.dingle@manchester.ac.uk

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

James Hook
School of Mathematics
University of Manchester
james.hook@manchester.ac.uk

MS239

Asynchronous Multilevel Methods on Adaptively Refined Grids

Emerging architectures and the drive towards exascale machines has seen a renewed interest in asynchronous solution methods for PDE systems. In this talk we will consider the AFACx method which is an asynchronous version of the Fast Adaptive Composite Grid (FAC) method. AFACx is a multilevel solution method most often used on structured adaptively refined grids for the solution of elliptic problems. AFACx has been demonstrated for both scalar and system elliptic PDEs. We will describe recent numerical work and the performance of AFACx as a solver and as a preconditioner. Progress permitting we will also describe the use of chaotic smoothers within AFACx for both multicore and GPU architectures.

Bobby Philip, Zhen Wang
Oak Ridge National Laboratory
philipb@ornl.gov, wangz@ornl.gov

Manuel Rodriguez Rodriguez, Mark Berrill
ORNL
mrs@ornl.gov, mbt@ornl.gov

MS239

Asynchronous Stochastic Optimization Methods

Stochastic gradient descent (SGD) has become a popular method for many machine learning tasks, particular when dealing with massive data sets. Some attempts to parallelize SGD require performance-destroying memory locking and synchronization. We describe an approach in which a centrally stored decision variable is updated asynchronously by the different cores. We present convergence theory and computational results. We also discuss Lagrangian-based extensions that potentially allow a higher degree of parallelism.

Christopher Re, Ben Recht
Computer Sciences Department
University of Wisconsin-Madison
chrisre@cs.wisc.edu, brecht@cs.wisc.edu

Stephen Wright
University of Wisconsin
Dept. of Computer Sciences
swright@cs.wisc.edu

MS240

Solution of Inverse Problems with Limited Forward Solver Evaluations: a Fully Bayesian Framework

Solving inverse problems based on computationally demanding forward solvers is ubiquitously difficult since one is necessarily limited to just a few observations of the response surface. This limited information induces additional uncertainties on the predicted design distributions. In this work we quantify this epistemic uncertainty by employing our recently developed fully Bayesian surrogates based on Gaussian processes.

Ilias Bilionis
Center for Applied Mathematics
Cornell University, Ithaca
ib227@cornell.edu

Nicholas Zabaras
Cornell University
zabaras@cornell.edu

MS240

An Analysis of Infinite-dimensional Bayesian Inverse Shape Acoustic Scattering and its Numerical Approximation

We present and analyze an infinite dimensional Bayesian inference formulation, and its numerical approximation, for the inverse problem of inferring the shape of an obstacle from scattered acoustic waves. Given a Gaussian prior measure on the shape space, whose covariance operator is the inverse of the Laplacian, the Bayesian solution of the inverse problem is the posterior measure given by the Radon-Nikodym derivative with respect to the Gaussian prior measure. The well-posedness of the Bayesian formulation in infinite dimensions is proved, including the justification of the Radon-Nikodym derivative and the continuous dependence of the posterior measure on the observation data via the Hellinger distance. Next, a finite dimensional approximation to the Bayesian posterior is proposed and the corresponding approximation error is quantified. The approximation strategy involves a Nyström scheme for approximating a boundary integral formulation of the forward Helmholtz problem and a Karhunen-Loève truncation for approximating the prior measure. The main result of this work is that the convergence rate for approximating the Bayesian inverse problem is spectral, and this directly inherits the spectral convergence rates of the approximations of both the prior and the forward problem.

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

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

MS240

Evaluation of Gaussian Approximations to Bayesian Inverse Problems in Subsurface Models

We discuss numerical evaluations of Gaussian approximations to the posterior distribution that arises in Bayesian data assimilation for reservoir characterization. In particular, we evaluate (i) maximum a posterior estimate, (ii) randomized maximum likelihood, (iii) EnKF (and variants). For our evaluation, we use a state-of-the art MCMC to accurately characterize the posterior distribution of the log-permeability conditioned to dynamic data. Therefore, MCMC provides a gold standard against which to compare the performance of the Gaussian approximations.

Marco Iglesias
Civil and Environmental Engineering
MIT
m.a.iglesias-hernandez@warwick.ac.uk

Kody Law
Mathematics Institute
University of Warwick
kodylaw@gmail.com

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

MS240

An Adaptive Sparse-Grid High-Order Stochastic Collocation Method for Bayesian Inference with Computationally Expensive Simulations

We develop an adaptive sparse-grid high-order stochastic collocation (aSG-hSC) method to quantify parametric and predictive uncertainty of computationally expensive physical systems with Bayesian approach. In our method, a surrogate system for the logarithm of the posterior probability density function (PPDF) is constructed using the local adaptive sparse-grid approximation. Moreover, we incorporate the high-order local hierarchical polynomial basis (e.g. quadratic, cubic basis) to further improve the accuracy of the surrogate system and reduce the computational expense.

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

Max Gunzburger
Florida State University
gunzburg@fsu.edu

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

MS241

Methods for Long-time Lagrangian Transport on the Sphere

We solve the advection equation in its Lagrangian form using a collection of disjoint particles and panels covering the sphere, for applications of climate modeling. For short time scales, the advection equation is satisfied exactly by each Lagrangian computational element. For longer simulations, we encounter the problem of mesh distortion. We introduce a remeshing scheme that uses adaptive refinement and interpolation of the Lagrangian parameter to preserve the range of the advected tracer without introducing any numerical oscillations.

Peter A. Bosler
University of Michigan
Applied and Interdisciplinary Mathematics
pbosler@umich.edu

Christiane Jablonowski
University of Michigan
Ann Arbor MI 48109-2143
cjablono@umich.edu

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

MS241

A Conservative Semi-Lagrangian Transport Scheme on Spectral Element Cubed-sphere Grids

A conservative semi-Lagrangian scheme for the spectral-element (SE) spherical grids (SPELT) has been developed. Each SE is overlaid by finite-volume (FV) cells, where we build a biquadratic mulit-moment reconstruction procedure. Conservation is guaranteed by a flux-based characteristic semi-Lagrangian approach. However, the velocity is specified on the host (SE) grid system. SPELT allows us to adopt arbitrary unstructured SE discretization and can therefore be used for future (unstructured) climate simulations (CAM-SE). The new scheme is tested on cubed-sphere geometry using several challenging tests.

Christoph Erath
NCAR
University of Colorado, Boulder
erath@ucar.edu

Ramachandran Nair
National Center for Atmospheric Research
rnair@ucar.edu

Henry Tufo
NCAR
University of Colorado at Boulder
tufo@cs.colorado.edu

MS241

Advection Scheme in CAM-SE

CAM-SE is the Community Atmosphere Model (CAM) configured with the spectral element dynamical core, which uses conforming quadrilateral grids. The spectral element method requires stabilization. If grids are quasi-uniform, adding a constant amount of hyper-viscosity is a very effective form of stabilization. Here we focus on a tensor-based hyper-viscosity operator for grids with high variability of scales.

Oksana Guba
Sandia National Laboratories
NM, USA.
onguba@sandia.gov

Michael Levy
NCAR
mlevy@ucar.edu

James Overfelt
Sandia National Laboratories
jroverf@sandia.gov

Mark A. Taylor
Sandia National Laboratories, Albuquerque, NM
mataylo@sandia.gov

MS241

Optimization-based Remap and Transport: A Divide and Conquer Strategy for Feature-preserving Discretizations

We present an optimization framework for the preservation of physical features in PDE discretizations and relate the recently introduced optimization-based remap of flux and mass to this framework. Remap is cast as a quadratic program whose optimal solution minimizes the distance to a suitable target quantity, responsible for the method's accuracy, subject to a system of inequality constraints, which separately maintain physical features. We use optimization-based remap to demonstrate shape-adaptable, conservative and bound-preserving transport algorithms.

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Kara Peterson
Sandia Natl. Labs
kjpeter@sandia.gov

MS242

Fast Multigrid Solvers Long Range Potentials

This talk will present multigrid solvers for computing electrostatic potentials. These parallel solvers can be used as building blocks in coupled multiphysics applications, such as particle transport in electro-osmotic or electro-rheological flows.

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ruede@informatik.uni-erlangen.de

Daniel Ritter
University of Erlangen-Nuremberg
daniel.ritter@informatik.uni-erlangen.de

Dominik Bartuschat
Lehrstuhl f. Simulation
Uni Erlangen-Nuremberg
bartuschat@i10.informatik.uni-erlangen.de

MS242

A Finite Element Method for the Total Variation Flow Without Regularization

The TV flow, that is the subgradient flow of the energy generated by the BV-norm, and related equations are called very singular diffusion equations, since in flat regions ($|\nabla u| = 0$) the diffusion is so strong that becomes a nonlocal effect. We propose a method for the solution of this class of equations, which involves no regularization and is unconditionally stable and convergent. To deal with the fact that the underlying nonlinear problems are solved only approximately, we devise an a posteriori error estimator. Applications to materials science are currently under investigation.

Abner J. Salgado
Department of Mathematics
University of Maryland
abnersg@math.umd.edu

MS242

Conforming vs. Nonconforming Finite Element Methods for Fluid and Solid Mechanics

The use of nonconforming elements to solve fluid and fluid mechanical problems has shown several advantages over the use of conforming counterparts. Replacing conforming elements by nonconforming ones is simple and guarantees numerical stability in many cases. We give a brief review on the recent development in quadrilateral nonconforming finite elements. We will then discuss several new interesting observations about these quadrilateral nonconforming elements. Error estimates and numerical results will be presented.

Dongwoo Sheen
Seoul National University
Department of Mathematics
sheen@snu.ac.kr

MS242

Superconvergence of Polynomial Spectral Collocation Methods

Abstract not available at time of publication.

Zhimin Zhang
Wayne State University
Department of Mathematics
ag7761@wayne.edu

MS243

Hierarchical Matrix Preconditioners and the Preservation of Vectors

Approximate $\mathcal{H}$-matrix preconditioners provide a robust and efficient way to solve large systems of linear equations resulting from the FE discretization of elliptic pdes. Due to ill-conditioning, the accuracy of the $\mathcal{H}$-matrix approximation has to be raised with n, which spoils the efficiency. We present a new approach that is based on the preservation of certain vectors during the $\mathcal{H}$-matrix approximation. The new technique significantly improves the quality of the preconditioner.

Mario Bebendorf
University of Bonn
bebendorf@ins.uni-bonn.de

MS243

Block Filtering Factorizations

In this talk we describe a family of preconditioners that are suitable for matrices arising from the discretization of a system of PDEs on unstructured grids. To address the scalability problem of many existing preconditioners, these preconditioners preserve directions of interest from the input matrix and are able for example to filter low frequency modes and alleviate their effect on convergence. The input matrix can have an arbitrary sparse structure, and can be reordered using nested dissection to allow a parallel implementation. We present a set of numerical experiments on matrices with anisotropies and jumps in the coefficients that show the efficiency of our preconditioners.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Frederic Nataf
Laboratoire J.L. Lions
nataf@ann.jussieu.fr

Riadh Fezzani
INRIA, France
riadh.fezzani@inria.fr

MS243

Elliptic Preconditioner for Accelerating the Self Consistent Field Iteration of Kohn-Sham Density Functional Theory

Kohn-Sham density functional theory (KSDFT) is the most widely used electronic structure theory for condensed matter systems. We present a new preconditioner for accelerating the self consistent field (SCF) iteration for solving the Kohn-Sham equations. The new preconditioner, which is called the elliptic preconditioner, is constructed by solving an elliptic partial differential equation. The elliptic preconditioner is shown to be effective for large inhomogeneous systems at low temperature.

Lin Lin
Lawrence Berkeley National Laboratory
linlin@lbl.gov

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

MS243

Multilevel Low-rank Approximation Preconditioners

A new class of methods based on low-rank approximations which has some appealing features will be introduced. The methods handle indefiniteness quite well and are more amenable to SIMD compuations, which makes them attractive for GPUs. The method is easily defined for Symmetric Positive Definite model problems arising from Finite Difference discretizations of PDEs. We will show how to extend to general situations using domain decomposition concepts.

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

Ruipeng Li
Department of Computer Science & Engineering
University of Minnesota
rli@cs.umn.edu

MS244

Application-specific Reduced Order Quadratures for Parameterized Problems

We present an algorithm to generate an application-specific reduced order quadrature (ROQ) for products of parameterized functions. Such integrals need to be computed many times, for example, while doing matched filtering to search for gravitational waves (GWs). If a reduced basis (RB) or any other projection based model reduction technique is applied, the dimensionality of integrands is reduced dramatically, however the cost of evaluating the reduced integrals still scales as the size of the original problem. Using discrete empirical interpolation (DEIM) points as ROQ nodes leads to a computational cost that only depends linearly on the number of reduced basis. Generation of a RB via a greedy procedure requires defining a training set, which for the product of functions can be very large. We notice that this direct approach can be impractical in many applications, and propose instead a two-step greedy targeted towards approximation of such products. The accuracy and efficiency of the two-step greedy and ROQ are verified by the error estimates and operation counts presented. In addition, we find that for the particular application in gravitational wave physics the two-step greedy speeds up the computations in the offline stage by two order of magnitude and two order of magnitude savings are observed in actual evaluations of integrals when ROQs are used as a downsampling strategy for equidistant samples assimilated from data. While the primary focus is on quadrature rules for products of parameterized functions, our method is easily adapted for integrations of single parameterized functions and some examples are considered. Joint work with: R. H. Nochetto, S. Field, M. Tiglio, F. Herrmann.

Harbir Antil
The University of Maryland
hantil@gmu.edu

MS244

Model Reduction for Fluid Flows Based on Nonlinear Balanced Truncation

For a vast variety of fluid flows the dynamics are governed by the Navier-Stokes equations which are highly nonlinear. For this class of systems linear model reduction techniques fail. In this paper after showing existence of solutions to the required equations, a computational algorithm for nonlinear balanced truncation is proposed for Galerkin models corresponding to the Navier-Stokes equations. The Galerkin models present a quadratic type nonlinearity that can be exploited in solving the Hamilton-Jacobi equations. The method is based on Taylor series expansion in which the computation of successively higher order terms reduces to solving consecutively higher order systems of linear algebraic equations.

Seddik Djouadi
Department of Electrical Engineering and Computer Science
University of Tenneessee, Knoxville
djouadi@eecs.utk.edu

Samir Sahyoun, Jin Dong
University of Tennessee
ssahyoun@utk.edu, jdong@ut.edu

MS244

Combined Reduced Basis Approximations and Statistical Approaches for Datat Assimulation

Reduced basis approximations allow to simulate the solution to systems modeled by parameter dependent PDE's. These methods allow to provide certified rapid on-line approximation of the solutions. We explain in this presentation how these complexity reduction methods can be used to improve data assimilation by adding a large amount of knowledge based of the model, constructed thanks to prior understanding by the specialists that have built the models.

Yvon Maday
Université Pierre et Marie Curie
and Brown university
maday@ann.jussieu.fr

MS244

Certified Reduced Order Methods for Optimal Flow Control Problems

We propose a reduced basis framework for the numerical solution of optimal control problems for parametrized viscous incompressible flows. We mainly focus on control problems for the (Navier-) Stokes equations involving infinite-dimensional control functions, thus requiring a suitable reduction of the whole optimization problem, rather than of the sole state equation. We provide rigorous and efficiently evaluable a posteriori error estimates, as well as preliminary examples of application arising from haemodynamics.

Gianluigi Rozza
SISSA, International School for Advanced Studies
Trieste, Italy
gianluigi.rozza@sissa.it

Federico Negri
Mathematics Institute of Computational Science and Eng, EPFL
federico.negri@epfl.ch

Andrea Manzoni
International School for Advanced Studies, Trieste, Italy
amanzoni@sissa.it

MS245

Monte-Carlo Simulation of Diffusion in Fractal Do-

mains

Diffusion-weighted MRI has became an important source of information about the dynamics in and the structure of natural or artificial materials (e.g., rocks, cements, human organs). In spite of intensive research, the relation between the microstructure and the signal formed by diffusing nuclei remains poorly understood, mainly due to lack of efficient algorithms and models for multi-scale porous media. To overcome this limitation, we developed a fast random walk (FRW) algorithm with gradient encoding which exploits the multi-scale character of the medium. In this talk, we present an application of this algorithm to a Menger sponge which is formed by multiple channels of broadly distributed sizes and often used as a model for soils and materials. Using this model, we investigate the role of multiple scales onto diffusion-weighted signals.

Denis Grebenkov
Ecole Polytechnique, France
Laboratoire de Physique de la Matiere Condensee
denis.grebenkov@polytechnique.edu

Hang Tuan Nguyen
CEA Neurospin, France
nguyentuanhang@yahoo.com

MS245
Hermite Functions in Modeling Diffusion MRI Data: From Applications to Fundamentals

Properties of Hermite functions make them ideally-suited to problems of modeling diffusion-weighted (DW) magnetic resonance (MR) signals and estimating propagators. We illustrate the utility of Hermite functions to characterize the anisotropy and diffusion-time dependence of the diffusion process. The ability of the model to represent different signal profiles suggests that the approach could be used even when no information regarding the underlying microstructure is availablea desirable characteristic for imaging applications.

Evren Ozarslan
Department of Radiology, Brigham and Women's Hospital
Harvard Medical School
evren@bwh.harvard.edu

Cheng Guan Koay, Peter J. Basser
NIH
guankoac@mail.nih.gov, pjbasser@helix.nih.gov

MS245
Diffusion Dynamics in Porous Media

Diffusion is known to cause multiple relaxation peaks in NMR relaxation behavior. However, in real porous materials, pore size distribution also results in multiple or broad relaxation times. We have devised a two-dimensional NMR method in order to distinguish these two scenarios. Using analytical and numerical solutions of the diffusion dynamics and numerical simulation we have demonstrated this effect and found it consistent with experimental results.

Yi-Qiao Song
Massachusetts General Hospital and Harvard Medical school, R
ysong@nmr.mgh.harvard.edu

Giovanna Carneiro
Massachusetts General Hospital and Harvard Medical school
ysong@nmr.mgh.harvard.edu

Larry Schwartz, Michael Prange
Schlumberger
ysong@nmr.mgh.harvard.edu,
ysong@nmr.mgh.harvard.edu

MS246
Evaluating Noise in Complex Networks

As in all computations involving real-world systems, the results of network analysis are affected by experimental, subjection and computational choices. However, network analysis algorithms are primarily based on graph theory and therefore assume the inputs to be exact. In this talk, we will demonstrate how user choices and computational limitations can affect network analysis results and discuss how concepts from numerical analysis such as conditioning and stability can be extended to evaluate the effect of noise in this domain.

Sanjukta Bhowmick, Sriram Srinivasan
Department of Computer Science
University of Nebraska, Omaha
sanjukta.bhowmick1@gmail.com,
sriramsrinivas@unomaha.edu

Vladimir Ufimtsev
University of Nebraska, Omaha
vufimtsev@unomaha.edu

MS246
On the Resilience of Graph Clusterings

Are clusters found by popular graph clustering algorithms significant? One way to answer this is by measuring cluster resilience as follows: repeatedly perturb the input graph by adding one edge, and for each new edge recluster the vertices and calculate the distance between the original and modified clustering. We hypothesize that the distribution of distances contains information about the stability of clusters in the original graph and present applications of this technique to synthetic and real-world graphs.

Evan Fields, Tzu-Yi Chen
Pomona College
ejf12009@mymail.pomona.edu, tzuyi@cs.pomona.edu

MS246
Quantification of Uncertainty in Network Summary Statistics

It is common in the analysis of network data to present a network graph and then cite various summary statistics thereof (e.g., degree distribution, clustering coefficient, conductance, etc.). However, most real-world networks derive from low-level measurements that are noisy. Surprisingly, there has been little effort to date aimed at quantifying the uncertainty induced in network summary statistics by noise in the underlying measurements. I will discuss work being done in my group towards this goal.

Eric D. Kolaczyk
Boston University
Dept of Mathematics & Stats
kolaczyk@math.bu.edu

MS246

Impact of Graph Perturbations on Structural and Dynamical Properties

Much of the research on networks arising in online social media and bioinformatics involves networks that are inferred by sampling or indirect measurements, making them noisy and incomplete. We study the impact of perturbations on structural and dynamical properties of networks. Specifically, find that threshold models are sensitive to perturbations, which can significantly alter the fixed point properties. Further, the effects depend on the nature of the perturbation.

Abhijin Adiga
Virginia Tech
abhijin@vbi.vt.edu

Henning Mortveit
Dept. of Mathematics and Virginia Bioinformatics Institute
Virginia Tech
hmortvei@vbi.vt.edu

Chris Kuhlman
Virginia Tech
ckuhlman@vbi.vt.edu

Anil Vullikanti
Dept. of Computer Science, and Virginia Bioinformatics Inst.
Virginia Tech
akumar@vbi.vt.edu

MS247

Numerical Stability of Vortex Soliton under Optical Lattice and Harmonic Potential

With only parabolic potential, 2D vortex soliton collapses (for the focusing case) when its norm exceeds certain threshold (which is much smaller than the norm that will cause its intrinsic instability). After including an optical lattice, there exists a second stability region (up to T=2000) when the lattice strength is moderate to large. However, some of these new stable solutions collapse after T¿5000. We investigate whether it is a property of the solution or simply a numerical artifact.

Qian-Yong Chen
Department of Mathematics & Statistics
U. of Massachusetts at Amherst
qchen@math.umass.edu

MS247

Analysis of Formal Order of Accuracy of WENO Finite Difference Scheme

In the presence of critical points, the formal order of accuracy of WENO schemes is reduced if the sensitivity parameter is chosen to be a small constant value or as a function of grid spacing. We proved a much weaker sufficient condition for the WENO-Z schemes by analyzing the nonlinear weights using some interesting properties of the smoothness indicators and illustrated with one dimensional Euler equations.

Wai-Sun Don
San Diego State University
wsdon@dam.brown.edu

MS247

Unconditionally Stable Numerical Scheme for Two Phase Models in Karst Aquifers

We present numerical methods that are unconditionally stable for certain phase field models of two phase flows in Karst aquifers.

Xiaoming Wang
FSU
wxm@math.fsu.edu

MS247

Adaptive Multigrid, Discontinuous Galerkin Methods for Cahn-Hilliard Type Equations

I will define and analyze a mixed DG, convex splitting scheme for a modified Cahn-Hilliard equation. The equation represents a diffuse interface model for phase separation in diblock coploymer blends and permits rather exotic solutions compared to the classical Cahn-Hilliard equation, including solutions like those from the phase field crystal model. The talk will cover theoretical energy stability and convergence results and also the practical, efficient solution of the algebraic equations via an adaptive nonlinear multigrid method. This is joint work with A. Aristotelous (SAMSI) and O. Karakashian (UTK).

Steven M. Wise
Mathematics Department
University of Tennessee
swise@math.utk.edu

MS248

Splitting Schemes for Incompressible Fluid-structure Interaction with Unfitted Interface Formulations

The stability and accuracy of the numerical approximations of incompressible fluid-structure interaction problems is very sensitive to the way the interface coupling conditions (kinematic and kinetic continuity) are treated at the discrete level. In this talk we will provide an overview of the existing loosely coupled (or explicit coupling) schemes within a fitted mesh framework and then discuss their extension to the unfitted case.

Miguel A. Fernandez
INRIA
miguel.fernandez@inria.fr

MS248

Immersed Finite Element Methods for Parabolic Equations with Moving Interface

Three Crank-Nicolson-type immersed finite element (IFE) methods are presented for solving parabolic equations whose diffusion coefficient is discontinuous across a time dependent interface. Instead of the body-fitting mesh needed by the traditional finite elements for solving interface problems, these IFE methods can use a structured mesh because IFEs can handle interface jump conditions without requiring the mesh to be aligned with the interface. Several disadvantages of the body-fitting mesh for time-dependent interface problems will be discussed. And then a fixed structured mesh for IFEs will be utilized to resolve these problems. Numerical examples are provided to demonstrate features of the three IFE methods.

Xiaoming He
Department of Mathematics and Statistics
Missouri University of Science and Technology
hex@mst.edu

Tao Lin
Virginia Tech
tlin@vt.edu

Yanping Lin
Department of Mathematical and Statistics Science, Universit
yanlin@ualberta.ca

Xu Zhang
Virginia Tech
xuz@vt.edu

MS248

A High Order Discontinuous Galerkin Nitsche Method for Elliptic Problems with Fictitious Boundary

We present a DG method, based on the method of Nitsche, for elliptic problems with an immersed boundary representation on a structured grid. In such methods small elements typically occur at the boundary, leading to breakdown of the discrete coercivity as well as numerical instabilities. In this work we propose a method that avoids using small elements on the boundary by associating them to a neighboring element with a sufficiently large intersection with the domain.

August Johansson
University of California, Berkeley
Dept. of Mathematics
august@math.berkeley.edu

MS248

Analysis and Implementation of a Nitsche-based Domain-bridging Method for Fluid Problems

We propose a Nitsche-based domain-bridging formulation for a general class of stabilized finite element methods for the Stokes problem. To ensure the stability of the method and to avoid ill-conditioned linear algebra systems, a ghost-penalty is added by extending the least-squares stabilization terms to the overlap region. We explain how general overlapping domain methods can be implemented efficiently by employing sophisticated algorithms and data structures from the field of computational geometry.

Andre Massing
Simula Research Laboratory
massing@simula.no

Mats G. Larson
Umeå University
mats.larson@math.umu.se

Anders Logg, Marie E. Rognes
Simula Research Laboratory
logg@simula.no, meg@simula.no

MS248

A Cut Finite Element Method for a Stokes Interface Problem

We present a finite element method for the Stokes equations involving two immiscible incompressible fluids with different viscosities and with surface tension. The interface separating the two fluids does not need to align with the mesh. We propose a Nitsche formulation which allows for discontinuities along the interface with optimal a priori error estimates. A stabilization procedure is included which ensures that the method produces a well conditioned stiffness matrix.

Peter Hansbo
Jönköping University
peter.hansbo@jth.hj.se

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

Sara Zahedi
Information Technology
Uppsala University
sara.zahedi@it.uu.se

MS249

Relaxation Models for Two-phase Flows

Dynamical two-phase flow models represent flows in various states of disequilibrium, where the driving forces towards full equilibrium are represented as *relaxation terms*. There is a deep connection between the entropy production of the relaxation process and the stability and wave speeds of the resulting equilibrium model. For various explicit models, we will discuss the issues of hyperbolicity and the *subcharactistic condition*; imposed equilibrium conditions will always reduce the mixture speed of sound.

Tore Flatten
SINTEF Energy Research
tore.flatten@sintef.no

Halvor Lund
Department of Energy and Process Engineering
Norwegian University of Science and Technology
halvor.lund@ntnu.no

MS249

A Mixture-energy-consistent 6-equation Two-phase Numerical Model for Cavitating Flows

We model cavitating flows by a variant of the 6-equation single-velocity two-phase model with stiff pressure relaxation of Saurel–Petitpas–Berry. Our novel formulation employs phasic total energy equations instead of the internal energy equations of the classical system. This allows us to design a simple numerical model that ensures consistency with mixture total energy conservation and agreement of the relaxed pressure with the mixture equation of state. Heat and mass transfer terms are also considered.

Marica Pelanti
ENSTA ParisTech
marica.pelanti@ensta-paristech.fr

Keh-Ming Shyue

National Taiwan University
shyue@ntu.edu.tw

MS249
Liquid-gas Mixtures and Diffuse Interfaces Computations at All Speeds

All speed flows and in particular the low Mach number asymptotics algorithms are addressed for the numerical approximation of a mechanical equilibrium multiphase flow model. During the computation, the interface is considered as a numerically diffused zone, captured as well as all present waves (shocks, expansion waves). Many applications with liquid-gas interfaces involve a very wide range of Mach number variations. Therefore, it is important to address numerical methods free of restrictions regarding the Mach number.

Richard Saurel, Sébastien Le Martelot
Polytech Marseille
Richard.Saurel@polytech.univ-mrs.fr,
sebastien.lemartelot@polytech.univ-mrs.fr

Boniface Nkonga
INRIA PUMAS - Université de Nice
boniface.nkonga@unice.fr

MS249
Multi-model Simulation of Compressible Two-phase Flows

The simulation of water flows in Pressurized Water Reactor requires the use of different models according to the local characteristics of the flow. These models may be formally linked by asymptotic limits but the dedicated codes are developed independently. We present how to couple these different models and codes and how to optimize the location of the coupling interfaces, in order to use as much as possible the simplest models without deteriorating the accuracy.

Nicolas Seguin
Laboratoire Jacques-Louis Lions
University Paris 6, France
nicolas.seguin@upmc.fr

MS250
Popularity versus Similarity in Growing Networks

Popularity is attractive – this is the formula underlying preferential attachment, a popular explanation for the emergence of scaling in growing networks. If new connections are made preferentially to more popular nodes, then the resulting distribution of the number of connections that nodes have follows power laws observed in many real networks. Preferential attachment has been directly validated for some real networks, including the Internet. Preferential attachment can also be a consequence of different underlying processes based on node fitness, ranking, optimization, random walks, or duplication. Here we show that popularity is just one dimension of attractiveness. Another dimension is similarity. We develop a framework where new connections, instead of preferring popular nodes, optimize certain trade-offs between popularity and similarity. The framework admits a geometric interpretation, in which popularity preference emerges from local optimization. As opposed to preferential attachment, the optimization framework accurately describes large-scale evolution of technological (Internet), social (web of trust), and biological (E.coli metabolic) networks, predicting the probability of new links in them with a remarkable precision. The developed framework can thus be used for predicting new links in evolving networks, and provides a different perspective on preferential attachment as an emergent phenomenon.

Fragkiskos Papadopoulos
Department of Electrical Engineering
Cyprus University of Technology
f.papadopoulos@cut.ac.cy

Maksim Kitsak
Cooperative Association for Internet Data Analysis
University of California, San Diego
mkitsak@caida.org

M. Angeles Serrano, Marian Boguna
Departament de Fisica Fonamental
University of Barcelona
marian.serrano@ub.edu, marian.boguna@ub.edu

Dmitri Krioukov
Cooperative Association for Internet Data Analysis
University of California, San Diego
dima@caida.org

MS250
Testing Model Fit with Algebraic Statistics

We address the problem of model validation or goodness-of-fit testing: to assess whether a given model can be considered as a satisfactory generative model for the data at hand. As argued in Diaconis-Sturmfels '98, a stochastic search of the space of networks with given values of sufficient statistics by Markov Chain Monte Carlo algorithms yields bonafide tests for goodness of fit. Markov bases guarantee to connect all networks with the same sufficient statistics, thus enabling this stochastic search. They are the only set of moves that guarantee that the random walk will give the real distribution. We show the algebraic tools used to derive Markov bases for several network models.

Sonja Petrovic
Department of Statistics
Pennsylvania State University
petrovic@psu.edu

Alessandro Rinaldo, Stephen Fienberg
Carnegie Mellon University
arinaldo@stat.cmu.edu, fienberg@stat.cmu.edu

MS250
Utilizing Spectral Methods for Uncued Anomaly Detection in Large-scale, Dynamic Networks

Abstract not available at time of publication.

Matt Schmidt, Nadya Bliss
MIT Lincoln Laboratory
matthew.schmidt@ll.mit.edu, nt@ll.mit.edu

MS250

Recent Results in Statistical Network Analysis

Abstract not available at time of publication.

Patrick Wolfe
University College London
patrick.j.wolfe@gmail.com

MS251

A Scalable Data Centric Model for Security and Provenance

Enthusiasm for Big Data is spurring innovation in new technologies that seek to transform information into knowledge. Using the foundation of Apache Accumulo we introduce a data centric security model that compliments the advances in Big Data technology and provides adaptability for data scientists. Our approach re-conceptualizes the relationship among data, users and applications, unlocking the potential for innovation and serving as a mechanism for the integration of disparate data sets. We address data provenance, integration of disparate data sets, and the interaction of a diverse community to realize the potential of Big Data to combat the spread of infectious disease, climate change, and other issues of scientific importance.

Oren J. Falkowitz
Sqrrl.com
oren@sqrrl.co

MS251

Large Data Analysis using the Dynamic Distributed Dimensional Data Model (D4M)

The growth of bioinformatics, social analysis, and network science is forcing computational scientists to handle unstructured data in the form of genetic sequences, text, and graphs. Triple store databases are a key enabling technology for this data and are used by many large Internet companies (e.g., Google Big Table and Amazon Dynamo). Triple stores are highly scalable and run on commodity clusters, but lack interfaces to support efficient development of the mathematically based algorithms of the typed used by computational scientists. D4M (Dynamic Distributed Dimensional Data Model) provides a parallel linear algebraic interface to triple stores. Using D4M, it is possible to create composable analytics with significantly less effort than using traditional approaches. The central mathematical concept of D4M is the associative array that combines spreadsheets, triple stores, and sparse linear algebra. Associative arrays are group theoretic constructs that use fuzzy algebra to extend linear algebra to words and strings. This talk describes the D4M technology, its mathematical foundations, application, and performance.

Jeremy Kepner
MIT Lincoln Laboratory
kepner@ll.mit.edu

MS251

How to Achieve Scalable Complex Analytics

In this talk we distinguish simple (SQL) analytics which are popular in business data warehouses from complex analytics, which most scientists are interested in. These are usually domain-specific codes which have matrix operations as inner loops. We first discuss why RDBMSs are not likely to work well for complex analytics. We then continue with a collection of tactics to make such complex analytics perform well and scale to arbitrary size data sets.

Michael Stonebraker
MIT CSAIL
stonebraker@csail.mit.edu

MS251

Streaming Algorithms and Large Astronomical Data Sets

As scientific data sets get larger, random access patterns to data become increasingly untenable. As a result, sequential streaming is the primary way for scalable data analysis. Also, the main statistical challenges are shifting, as uncertainties are no longer statistical but due to systematic errors. We discuss various astronomical analysis challenges related to and streaming algorithms, e.g. incremental robust PCA and how these can be implemented as part of database access patterns.

Alex Szalay
Johns Hopkins University
szalay@jhu.edu

MS252

Graphlets: A Scalable, Multi-scale Decomposition for Large Social and Information Networks

Abstract not available at time of publication.

Edo Airoldi
Harvard University
airoldi@fas.harvard.edu

MS252

From Cells to Tissue: Coping with Heterogeneity when Modelling the Electrophysiology of the Human Heart

Abstract not available at time of publication.

Kevin Burrage
University of Oxford, & Queensland University of Technology
& Institute for Molecular Bioscience, UQ
kevin.burrage@gmail.com

MS252

Stochastic Simulation Service: Towards an Integrated Development Environment for Modeling and Simulation of Stochastic Biochemical Systems

Abstract not available at time of publication.

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

MS252

Exploiting Stiffness for Efficient Discrete Stochastic Biochemical Simulation

Gillespie's stochastic simulation algorithm (SSA) has become an invaluable tool for simulating biochemical models in a way that captures the inherent randomness of these systems. However, the SSA is inefficient for models featuring stochastic stiffness arising from the presence of multiple timescales. In this talk, we present some counterintuitive implications of stochastic stiffness and describe methods for exploiting stiffness to improve simulation efficiency. Techniques for automating and accelerating existing algorithms will also be discussed.

Kevin Sanft
St. Olaf College
kevin@kevinsanft.com

MS253

The Powers that be in HPC

The power consumption of supercomputers ultimately limits their performance. The current challenge is not whether we will can build an exaflop system by 2018, but whether we can do it in less than 20 megawatts. The SCAPE Laboratory at Virginia Tech has been studying the tradeoffs between performance and power for over a decade. We've developed an extensive tool chain for monitoring and managing power and performance in supercomputers. We will discuss the implications of our findings for exascale systems and some research directions ripe for innovation.

Kirk Cameron
Department of Computer Science
Virginia Tech
cameron@cs.vt.edu

MS253

Lower Bounds on Algorithm Energy Consumption: Current Work and Future Directions

By extending communication lower bounds on algorithms via linear models of machine energy consumption, we have derived theoretical bounds on the minimal amount of energy required to run an algorithm on a given machine. To use these bounds for HW/SW cotuning, efficient parameters for energy models must be calculated. We discuss initial approaches to parameter calculation and further development of energy bounds into broader classes of codes.

Andrew Gearhart
UC Berkeley
agearh@cs.berkeley.edu

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

MS253

Energy-Aware Dense and Sparse Linear Algebra

Power is becoming a crucial challenge that the high performance computing community will have to face to efficiently leverage the Exascale systems that will be available at the end of this decade. In this talk we will address several aspects related to this issue on multicore and many-core (GPU) processors. Specifically, we introduce a simple model of power dissipation, and a tools to compose a power tracing framework. Then, using experimental data, we evaluate both the potential and real energy reduction that can be attained for the task-parallel execution of dense and sparse linear algebra operations on multi-core and many-core processors, when idle periods are leveraged by promoting CPU cores to a power-saving C-state.

Enrique S. Quintana-Ortí
Universidad Jaume I
quintana@icc.uji.es

MS253

Locality Aware Scheduling of Sparse Computations for Energy and Performance Efficiencies

We consider the problem of increasing the performance and energy efficiencies of sparse matrix and graph computations on multicore processors. Such systems have complex non-uniform memory access (NUMA) cache and memory that exhibit significant variations in data access latencies. We describe a scheme for fine-grain task scheduling to cores that takes into account the probabilities of hits in cache. We present results indicating that our scheme leads to near ideal speed-ups and large improvements in performance and energy efficiencies compared to traditional methods.

Padma Raghavan
The Pennsylvania State Univ.
Dept of Computer Science Engr.
raghavan@cse.psu.edu

MS254

Testing of Management Objective Hypotheses by Solution of Constrained Inverse Problems

Often, in the environmental sciences, too much faith is placed in model predictions when there is no scientific basis for doing so. We assert that model-based decision-making should implement the scientific method, and present an approach best described as model-based hypothesis testing. The method seeks to reject the hypothesized occurrence of a future event by demonstrating that it is incompatible with historical measurements of system state and expert knowledge.

Rachel S. Blakers
NCGRT, iCAM, Fenner School of Environment and Society
Australian National University
rachel.blakers@anu.edu.au

MS254

Employing Imprecise Knowledge for Model Parameter Inference and Prediction

In Bayesian analysis, the presence of ambiguity in subjective or intersubjective knowledge often makes it difficult to formulate precise prior probability distributions. The concept of imprecise probability provides a framework for characterizing such ambiguity. When imprecise probabilities are used to represent ambiguous prior knowledge about model parameters, there are consequences for posteriors and model results. By extending earlier studies, we show that, under weak regularity conditions, the Density Ratio Class of imprecise probability measures is invariant under Bayesian inference, marginalization, and propagation through deterministic models. These invariance properties are desirable because they make it possible to describe sequential Bayesian learning and prediction using imprecise probabilities within a unified framework. We proceed to describe algorithms that exploit these invariance properties to minimize the additional computational burden compared to the use of precise priors. To demonstrate the feasibility of the proposed approach, the concepts and numerical methods developed are applied to a simple empirical ecological model.

Simon Rinderknecht
Eawag
Swiss Federal Institute of Aquatic Science and Technology
simon-lukas.rinderknecht@gmx.ch

Carlo Albert
Eawag
carlo.albert@eawag.ch

Mark Borsuk
Thayer School of Engineering
Dartmouth College
mark.borsuk@dartmouth.edu

Nele Schuwirth
Eawag
Swiss Federal Institute of Aquatic Science and Technology
nele.schuwirth@eawag.ch

Hans-Rudolf Kuensch
Seminar for Statistics
ETH-Zurich
hans-ruedi.kuensch@stat.math.ethz.ch

Peter Reichert
Eawag: Swiss Federal Inst. of Aquatic Science and Technology
peter.reichert@eawag.ch

MS254

Uncertainty Assessment by Algebraic Plausible Worst-case Falsification of Conclusions

The minisymposium topic is introduced, using a typology to highlight key aspects of deep uncertainty. Key techniques are summarised, including imprecise probabilities, measures of robustness and various approaches to falsifying hypotheses. Using groundwater management examples, we show that algebraic techniques to falsification can efficiently identify plausible worst case scenarios, effectively stress testing model conclusions. Accounting for uncertainty in this way helps surface assumptions and can simplify the application of other techniques.

Joseph H. Guillaume, Anthony J. Jakeman
NCGRT, iCAM, Fenner School of Environment and Society
Australian National University
joseph.guillaume@anu.edu.au, tony.jakeman@anu.edu.au

MS254

Many Objective Robust Decision Making for Complex Environmental Systems Undergoing Change

We present a many objective robust decision making framework, combining evolutionary optimization, robust decision making (RDM), and interactive visual analytics. Many objective evolutionary search enables the discovery of tradeoffs between objectives. Subsequently, RDM determines the robustness of the component tradeoff solutions to deeply uncertain future conditions. Results suggest that including robustness as a decision criterion can dramatically change the formulation of environmental management problems as well as the negotiated selection of candidate alternatives.

Joseph R. Kasprzyk
Department of Civil and Environmental Engineering
The Pennsylvania State University
jrk301@psu.edu

Shanthi Nataraj
RAND Corporation
snataraj@rand.org

Patrick Reed
The Pennsylvania State University
preed@engr.psu.edu

Robert Lempert
RAND Corporation
lempert@rand.org

MS255

A Composed and Multilevel Solution Framework for Nonlinear PDEs

Nonlinear PDE solvers enjoy a fairly large design space that is underexplored in practical application. Hierarchies of solution techniques, such as local or multilevel nonlinear solvers accelerated by nonlinear Krylov or quasi-Newton methods, provide a flexible alternative to the nested Newton-Krylov iteration. We present a framework, available in the PETSc library, for large-scale composed and multilevel nonlinear solvers. We demonstrate the benefits and drawbacks of these methods on problems of interest.

Peter R. Brune
Mathematics and Computer Science Division
Argonne National Laboratories
brune@mcs.anl.gov

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

MS255

A Multilevel Stochastic Collocation Algorithm for Optimization of PDEs with Uncertain Coefficients

Optimization of PDEs with uncertain coefficients is expensive because problem size depends on both spatial and sample space discretizations. I analyze the use of MG/OPT to formulate a multilevel stochastic collocation algorithm by exploiting the hierarchical structure of sparse grids. For a class of problems and methods, I present convegence analysis as well as explicit cost and error bounds for one V-cycle of MG/OPT. I provide numerical illustrations confirming these bounds.

Drew P. Kouri
Mathematics and Computer Science Division
Argonne National Laboratory
dpkouri@mcs.anl.gov

MS255

Using Inexact Gradients in a Multilevel Optimization Algorithm

Many optimization algorithms require gradients of the

model functions, but computing accurate gradients can be expensive. We study the impact of inexact gradients on the multilevel optimization algorithm MG/Opt. MG/Opt recursively uses coarse models to obtain search directions for fine models. However, MG/Opt requires fine-level gradients to define the recursion. We consider various possible sources of error in the gradients, and demonstrate that in many cases the effect of the errors is benign.

Stephen G. Nash
George Mason University
Systems Engineering & Operations Research Dept.
snash@gmu.edu

Robert Michael Lewis
College of William & Mary
Williamsburg, VA 23187
rmlewi@wm.edu

MS255

Parallel Software for the Optimization Algorithm DIRECT with High-cost Function Evaluations

This talk concerns derivative-free global optimization using a large number of extreme-scale function evaluations, illustrated by parameter estimation for basis set selection in the ab initio nuclear physics code MFDn (Many Fermion Dynamics, nuclear). A hierarchical three-tier scheme for integrating a massively parallel DIRECT implementation, pVTdirect, with MFDn is considered. Additional pVTdirect modifications to support a dynamic number of processes to accommodate unpredictable runtime memory demands are discussed.

Masha Sosonkina
Ames Laboratory/DOE
Iowa State University
masha@scl.ameslab.gov

Layne T. Watson
Virginia Polytechnic Institute and State University
Departments of Computer Science and Mathematics
ltwatson@computer.org

MS256

A(nother) Randomized Rank-revealing Decomposition: Generalized RURV

Fast, randomized rank-revealing decompositions have emerged in the last decade, in particular, as a means to construct low-rank or sparse approximations of matrices. Various types of algorithms emerged, some involving less randomization (through the use of FFT-like matrices), and some more. We present here one that uses more randomization, but the results are guaranteed for a wider range of numerical ranks, and as a bonus, it is straightforwardly generalizable to a product of matrices and inverses. This turns out to be a very desirable feature, given that it was developed for a communication-minimizing non-symmetric eigenvalue algorithm using divide-and-conquer.

Ioana Dumitriu
University of Washington, Seattle
dumitriu@math.washington.edu

MS256

Beating MKL and Scalapack at Rectangular Matrix Multiplication Using the BFS/DFS Approach

We present CARMA, a communication-avoiding parallel rectangular matrix multiplication algorithm, attaining significant speedups over both MKL and ScaLAPACK. Combining the recursive BFS/DFS approach of Ballard, Demmel, Holtz, Lipshitz and Schwartz (SPAA '12) with the dimension splitting technique of Frigo, Leiserson, Prokop and Ramachandran (FOCS '99), CARMA is communication-optimal, cache- and network-oblivious, and simple to implement (*e.g.*, 60 lines of code for the shared-memory version), performing well on both shared- and distributed-memory machines.

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

David Eliahu
UC Berkeley
deliahu@berkeley.edu

Armando Fox
University of California, Berkeley
fox@cs.berkeley.edu

Shoaib Kamil
Massachusetts Institute of Technology
skamil@cs.berkeley.edu

Benjamin Lipshitz, Oded Schwartz, Omer Spillinger
UC Berkeley
lipshitz@berkeley.edu, oded.schwartz@gmail.com, omers88@berkeley.edu

MS256

Communication Avoiding ILU(0) Preconditioner (CA-ILU(0))

We present a parallel communication-avoiding ILU(0) preconditioner for systems $Ax = b$, where A is sparse. First, A is reordered using nested dissection, then the blocks and separators are reordered to avoid communication at the expense of performing redundant computations. Our special reordering doesn't affect the convergence rate of the ILU(0) preconditioned systems as compared to nested dissection reordering of A, while it reduces data movement and redundant flops. Thus, CA-ILU(0) preconditioner should have good performance.

Sophie Moufawad
LRI - INRIA Saclay, France
sophie.moufawad@lri.fr

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

MS256

Scalable Numerical Algorithms for Electronic Structure Calculations

This talk will focus on dense linear algebra and tensor computations in application to electronic calculations. I will introduce 2.5D algorithms, which are designed to minimize communication between processors.

The algorithms employ limited data-replication to theoretically lower communication costs with respect to standard (ScaLAPACK/Elemental) algorithms. I will detail implementations of 2.5D matrix multiplication and LU factorization, which employ topology-aware mapping on supercomputers such as BG/P and BG/Q. 2.5D algorithms have been employed on BG/Q by QBox, a parallel DFT code for electronic structure calculations. The communication minimization and topology-aware mapping schemes will also be extended to tensor contractions as embodied by Cyclops Tensor Framework. The Coupled Cluster method for systems with high electronic correlation has already been implemented on top of this framework and is achieving strong performance results on BG/Q.

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

MS257

Reduced Collocation Methods: Reduced basis Methods in the Collocation Framework

In this talk, we present a reduced basis method well-suited for the collocation framework. Two fundamentally different algorithms will be presented. This work provides a reduced basis strategy to practitioners who prefer a collocation, rather than Galerkin, approach. Furthermore, one of these two algorithms *eliminates* a potentially costly online procedure that is needed for non-affine problems with Galerkin approach. Numerical results demonstrate the high efficiency and accuracy of the reduced collocation methods, which match or exceed that of the traditional reduced basis method in the Galerkin framework.

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

MS257

A Component-based Reduced basis Method for Many-parameter Systems

We present a reduced basis method based that combines standard reduced basis approximations with a static condensation formulation. This allows us to develop standalone parametrized reduced basis components that can be connected together to form large systems with many parameters. Also, this enables an appealing "system assembly" approach, which we demonstrate in an interactive GUI. We show numerical results from applications drawn from structural analysis and acoustics.

David Knezevic
Harvard University
dknezevic@seas.harvard.edu

MS257

Non-intrusive Reduced basis Approximations for Parameter Dependent PDEs

Reduced basis methods take advantage of the small Kolmogorov dimension of the manifold of all solutions to some parameter dependent partial differential equation (when the parameter vary). These allow to propose rapid approximation methods for real time on-line approximations based on reduction of complexity. Numerical analysis based on a posteriori error estimation allow to provide additional tools that allow to validate the computations with reduced basis approximations. These approximations methods are ad'hoc and require a preliminary off-line analysis that uses all the previously tools together a fine classical approximation of e.g. finite element type. The reduced basis will be an alternative and complement to these classical methods and will allow to mimic it. The rapid on-line implementation and off-line learning process require to have a full access to the original classical solver. When this is not the case we have to imagine alternative tools that allow to circumvent this lack of full access (as can be the case when the code is an industrial one). We shall present some recent advances in the direction of the non intrusive implementation.

Yvon Maday
Université Pierre et Marie Curie
and Brown university
maday@ann.jussieu.fr

MS257

Component-based Reduced basis Simulations: Conjugate Heat Transfer and Transient Problems

We present some applications of the Static Condensation Reduced Basis Element (SCRBE) method: a domain decomposition with reduced basis at the intradomain level to populate a Schur complement at the interdomain level. We present numerical results for a conjugate heat transfer problem (2D car radiator model) as well as transient heat transfer problems on 3D configurations, which demonstrate the flexibility, accuracy and computational efficiency of our approach.

Sylvain Vallaghe, Anthony T. Patera
Massachusetts Institute of Technology
svallagh@mit.edu, patera@mit.edu

MS258

A Fluid-Based Preconditioner for Fully Implicit Kinetic Electrostatic Plasma Simulations

We introduce a fluid-based preconditioner for a recently proposed energy- and charge-conserving fully implicit, particle-in-cell electrostatic algorithm in 1D [Chen, Chacón, Barnes. *J.Comp.Phys*, **230**, 7018 (2011)]. The approach employs a linearized form of the first two moment equations, with suitable kinetic closures, to find approximate updates of the self-consistent electric field from the linearized Ampere's law. The effectiveness of the resulting preconditioner will be demonstrated on a challenging multiscale ion acoustic wave problem.

Guangye Chen
ORNL
cheng2@ornl.gov

Luis Chacón
Los Alamos National Laboratory
chacon@lanl.gov

MS258

Coupled Simulation of Continuum Material Point Method with Molecular Dynamics Numerical Statistics

Material point method uses both Lagrangian and Eulerian descriptions for material motion, avoids mesh distortion and numerical diffusion issues. When combined with Molecular Dynamics in a scale bridging algorithms, we have the advantage of simulating whole domain with minimal communication between molecular dynamic systems to achieve parallel computation efficiency close to the embarrassing parallelism. We present examples to show that this algorithm can be implemented in combined GPU/CPU platform.

Xia Ma
LANL
xia@lanl.gov

Duan Z. Zhang, Dana Knoll
Los Alamos National Laboratory
dzhang@lanl.gov, nol@lanl.gov

MS258

Moment Acceleration of Fokker-Planck Collision Operator for Electrostatic Plasma Physics Simulation

For time step sizes significantly larger than the collision time scale, a fully implicit solution to the Fokker-Planck operator requires many fixed point iterations to converge. The convergence of the fixed point iteration for the Fokker-Planck collision operator will be accelerated using an emerging class of moment based accelerators. It is demonstrated that upon convergence, we obtain a fully consistent kinetic and moment solution.

William T. Taitano
University of New Mwxico
Los Alamos National Laboratory
taitae25@unm.edu

Dana Knoll
Los Alamos National Laboratory
nol@lanl.gov

Anil K. Prinja
University of New Mexico
prinja@unm.edu

MS258

An Accelerated Free Surface, Z-Level Ocean Model using a Moment-Based Approach and Trilinos

We study a moment based method for a free-surface z-level ocean model. In this approach, the full three dimensional momentum and continuity equations are coupled to a set of two dimensional moment equations, obtained by vertical integration. This formulation allows relaxation of the timestep size by isolating the stiff physics to the reduced system. We provide numerical examples to support our study and compare to traditional implementations.

Geoff Womeldorff, Christopher K. Newman, Dana Knoll
Los Alamos National Laboratory
womeld@lanl.gov, cnewman@lanl.gov, nol@lanl.gov

MS259

Problem Formulation for Multi-source Optimization in Complex Systems

The operation of a single complex artifact is governed by many models that are combined at various stages of design, together with some statement of constrained optimization to form a model used for designing the artifact a design problem formulation. The use of models within various formulations influences the tractability and robustness of the solution. We discuss the synthesis of models originating from a variety of sources into problem formulations and the resulting computational properties.

Natalia Alexandrov
NASA Langley Research Center
n.alexandrov@nasa.gov

MS259

Sensitivity Analysis and Information Fusion for Multi-source Optimization

A multifidelity approach to design seeks to exploit optimally all available information. Existing methods generally calibrate or replace low-fidelity results using higher fidelity information. Here we propose an approach based on estimation theory to fuse information from multifidelity sources. Mathematical interrogation of the uncertainty in quantities of interest is achieved via global sensitivity analysis, which provides guidance for adaptation of model fidelity. The methodology is demonstrated on a wing-sizing problem for a high-altitude, long-endurance vehicle.

Doug Allaire, Karen E. Willcox
Massachusetts Institute of Technology
dallaire@mit.edu, kwillcox@MIT.EDU

MS259

On the Use of Self-organizing Maps for Data Management in a Multifidelity Analysis Context

The use of Self Organizing Maps (SOM), unsupervised neural networks also known as Kohonens Maps, for the management of data in a variable fidelity analysis environment is discussed. In particular the employment of SOM for screening purposes is introduced and applied to the multifidelity aerodynamic analysis of an aircraft wing. Moreover the possibility to use those networks to address local errors of low fidelity models for corrective purposes is finally proposed.

Laura Mainini
Polytechnic University of Turin
mainini.laura@gmail.com

MS259

Numerical Bouillabaisse: Combining Experiments, Simulations and Machine Learning to Guide Optimization

In conventional optimization one iteratively samples a single source of information concerning a function $f(.)$, e.g., one samples $f(x)$ itself. However in practice, often at each step one must choose which of several information sources about f to sample. Here I illustrate how to use semi-supervised machine learning to statistically combine samples from multiple sources. I also introduce an algorithm for deciding which information source to sample next and with what input values.

David Wolpert
NASA Ames Research Center
david.h.wolpert@gmail.com

MS260

Multiscale Methods and Analysis for the Nonlinear Klein-Gordon Equation in the Nonrelativistic Limit Regime

In this talk, I will review our recent works on numerical methods and analysis for solving the nonlinear Klein-Gordon (KG) equation in the nonrelativistic limit regime, involving a small dimensionless parameter which is inversely proportional to the speed of light. In this regime, the solution is highly oscillating in time and the energy becomes unbounded, which bring significant difficulty in analysis and heavy burden in numerical computation. We begin with four frequently used finite difference time domain (FDTD) methods and obtain their rigorous error estimates in the nonrelativistic limit regime by paying particularly attention to how error bounds depend explicitly on mesh size and time step as well as the small parameter. Then we consider a numerical method by using spectral method for spatial derivatives combined with an exponential wave integrator (EWI) in the Gautschi-type for temporal derivatives to discretize the KG equation. Rigorious error estimates show that the EWI spectral method show much better temporal resolution than the FDTD methods for the KG equation in the nonrelativistic limit regime. In order to design a multiscale method for the KG equation, we establish error estimates of FDTD and EWI spectral methods for the nonlinear Schrodinger equation perturbed with a wave operator. Finally, a multiscale method is presented for discretizing the nonlinear KG equation in the nonrelativistic limit regime based on large-small amplitude wave decompostion. This multiscale method converges uniformly in spatial/temporal discretization with respect to the small parameter for the nonlinear KG equation in the nonrelativistic limite regime. Finally, applications to several high oscillatory dispersive partial differential equations will be discussed.

Weizhu Bao
National University of Singapore
Department of Mathematics
bao@math.nus.edu.sg

MS260

Uniform Error Estimates of An Exponential Wave Integrator Sine Pseudospectral Method for Nonlinear Schrödinger Equation with Wave Operator

We propose an exponential wave integrator sine pseudospectral (EWI-SP) method for the nonlinear Schrödinger equation (NLS) with wave operator (NLSW), and carry out rigorous error analysis. NLSW is NLS perturbed by the wave operator with strength described by a dimensionless parameter $\varepsilon \in (0, 1]$. In this work, we show that the proposed EWI-SP possesses the optimal uniform error bounds at $O(\tau^2)$ and $O(\tau)$ in τ (time step) for well-prepared initial data and ill-prepared initial data, respectively, and spectral accuracy in h (mesh size) for the both cases, in the L^2 and semi-H^1 norms.

Yongyong Cai
National University of Singapore
yongyong.cai@gmail.com

Weizhu Bao
National University of Singapore
Department of Mathematics
bao@math.nus.edu.sg

MS260

Fast Computation of Time-Domain Electromagnetic Scattering Problems with Exact Transparent Boundary Conditions

Wave propagations in unbounded media arise from diverse applications. Intensive research has been devoted to frequency-domain simulation, e.g., for the time-harmonic Helmholtz problems and Maxwell's equations. Here, we are interested in time-domain computation, which is known to be more flexible in capturing wide-band signals and modeling more general material inhomogeneities and nonlinearities. In this talk, we shall show how to use tools in complex analysis to analytically evaluate circular and spherical nonreflecting boundary conditions (NRBCs), and how to efficiently deal with time-space globalness of such boundary conditions. Fast spectral-Galerkin solvers together with stable time integration and techniques for handling general irregular scatterers will be introduced for the simulation. We intend to demonstrate that the interplay between analytic tools, accurate numerical means and sometimes brute force hand calculations can lead to efficient methodologies for challenging simulations. This is a joint work with Bo Wang and Xiaodan Zhao.

Li-Lian Wang
School of Physical & Mathematical Sciences
Nanyang Technological University
lilian@ntu.edu.sg

MS260

Uniformly Correct Multiscale Time Integrators for Highly Oscillatory Second Order Differential Equations

In this talk, two multiscale time integrators (MTIs), motivated from two types of multiscale decomposition by either frequency or frequency and amplitude, are proposed and analyzed for solving highly oscillatory second order differential equations with a dimensionless parameter $0 < \epsilon \ll 1$. In fact, the solution to this equation propagates waves with wavelength at $O(\epsilon^2)$ when $0 < \epsilon \ll 1$. We rigorously establish two independent error bounds for the two MTIs as $O(\tau^2/\epsilon^2)$ and $O(\epsilon^2)$ for $\epsilon \in (0, 1]$ with $\tau > 0$ as step size, which imply that the two MTIs converge uniformly with linear convergence rate at $O(\tau)$ for $\epsilon \in (0, 1]$ and optimally with quadratic convergence rate at $O(\tau^2)$ in the regimes when either $\epsilon = O(1)$ or $0 < \epsilon \ll \tau$. Thus the meshing strategy requirement (or ϵ-scalability) of the two MTIs is $\tau = O(1)$ for $0 < \epsilon \ll 1$, which is significantly improved from $\tau = O(\epsilon^3)$ and $\tau = O(\epsilon^2)$ requested by finite difference methods and exponential wave integrators to the equation, respectively. At last numerical results are reported to support the two error bounds.

Xiaofei Zhao
National University of Singapore
Department of Mathematics
a0068978@nus.edu.sg

Bao Weizhu, Dong Xuanchun
Department of Mathematics
National University of Singapore
bao@math.nus.edu.sg, matdx@nus.edu.sg

MS261
Title Not Available at Time of Publication

Abstract not available at time of publication.

Yalchin Efendiev
Dept of Mathematics
Texas A&M University
efendiev@math.tamu.edu

MS261
Estimation of Anthropogenic CO2 Emissions from Sparse Observations using a Multiresolution Random Field Model

We present a method to estimate fossil-fuel CO2 emissions from observed CO2 concentrations at a sparse set of locations. The emission field is represented using wavelets, with nightlight images providing the spatial model. Sparsity-enforced estimation is used to reconstruct the emission field; wavelets unconstrained by the observations are set to zero. The method is tested with synthetic data. Emissions from North-eastern US are estimated most accurately, and the South-west is the least constrained by observations.

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

Vineet Yadav
Carnegie Institution for Science
Stanford, CA
vineety@stanford.edu

Sean A. Mckenna
Sandia National Laboratories
samcken@sandia.gov

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

Bart van Bloeman Waanders
Sandia National Laboratories
bartv@sandia.gov

MS261
Quantification of Parameter Uncertainties in Nonlinear Distributed Material Models

Piezoelectric, magnetic and shape memory alloy (SMA) materials offer unique capabilities for energy harvesting and reduced energy requirements in aerospace, aeronautic, automotive, industrial and biomedical applications. However, all of these materials exhibit creep, rate-dependent hysteresis, and constitutive nonlinearities that must be incorporated in models and model-based robust control designs to achieve their full potential. Furthermore, models and control designs must be constructed in a manner that incorporates parameter and model uncertainties and permits predictions with quantified uncertainties. In this presentation, we will discuss Bayesian techniques to quantify parameter uncertainties in nonlinear distributed models arising in the context of smart systems. We will also discuss highly efficient techniques to propagate these uncertainties through models to quantify uncertainties associated with quantities of interest.

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

Nathanial Burch
SAMSI
nathanial.burch@gmail.com

Zhengzheng Hu
Department of Mathematics
North carolina State University
zhu4@ncsu.edu

Michael Hays
Florida State University
mrh04c@fsu.edu

William Oates
Department of Mechanical Engineering, Florida State Universi
woates@eng.fsu.edu

MS261
Calibration of Uncertain Parameters in Stochastic Turbulence Models

Abstract not available at time of publication.

Catalin S. Trenchea
Department of Mathematics
University of Pittsburgh
trenchea@pitt.edu

William Layton
University of Pittsburgh
wjl@pitt.edu

Clayton G. Webster
Oak Ridge National Laboratory
webstercg@ornl.gov

MS262
Multiscale Modeling of Energy Storage Device

An efficient implementation of the Poisson-Nernst-Plank - classical Density Functional Theory (PNP-cDFT), a multiscale method for the simulations of charge transport in nanocomposite materials, is designed and evaluated. Spatial decomposition of the multi particle system is employed in the parallelization of classical density functional theory (cDFT) algorithm. Furthermore, a truncated strategy is used to reduce the computational complexity of cDFT algorithm. The simulation results show that the parallel implementation has close to linear scalability in parallel computing environments for both 1D and 3D systems. Additionally, we develop robust numerical algorithms to find steady state fluxes in 2D circular nano-particles made up of rutile TiO2. The equations we solve come from density functional theory coupled with the Poisson-Nernst-Plank formalism. In solving steady state fluxes, we come to a better understanding of how battery properties are influenced by nano-particle arrangement which influence the flux of Lithium ions.

Lin Guang

PNNL
guang.lin@pnl.gov

MS262

Multiscale Computation of Heterogeneous Surface Catalytic Reactions

Kinetic Monte Carlo simulations are typically performed to determine steady state values in kinetic reactions for catalysis. We present a full non-linear master equation capable of resolving these dynamics; we also present a linearized approximation whose solutions and eigenstructure can be fully classified. The equations are simple, yet have many degrees of freedom. We use symbolic packages to generate code to solve the resulting ODEs for the master equation and its linear approximation.

Gregory Herschlag
Department of Mathematics
University of North Carolina at Chapel Hill
gregoryh@unc.edu

Sorin Mitran
University of North Carolina Chapel Hill
mitran@unc.edu

MS262

Numerical Simulation of Reactive Particle Compacts

A computational model is developed for the simulation of transient reactions in heterogeneous porous compacts of reactive multilayered particles. The evolution of the reaction is described in terms of a reduced model formalism that efficiently accounts for thermal transport at the macroscale, and for atomic mixing and chemical heat release at the micro or nanoscale. Computations are used to analyze the effects of microstructural parameters on the mean properties of the reaction front.

Leen Alawieh
Johns Hopkins University
leen@jhu.edu

Ihab Sraj
Duke University
ihab.sraj@duke.edu

Timothy Weihs
Johns Hopkins University
weihs@jhu.edu

Omar M. Knio
Duke University
omar.knio@duke.edu

MS262

Quantifying Uncertainty in Ionic Flow through a Silica Nanopore

We discuss uncertainty quantification in MD simulations of concentration-driven flow through a silica nanopore. We model all components, namely water, silica, sodium and chloride ions, and explore the sensitivity of the flow to the pore diameter and gating charge. A Bayesian inference formalism, coupled with polynomial chaos representations, is used for this purpose. In addition to sensitivity analysis, the study focuses on complex effects arising due to system heterogeneity as well as intrinsic thermal fluctuations.

Francesco Rizzi
Department of Mechanical Engineering
Johns Hopkins University
francesco.rizzi@duke.edu

Reese Jones
Co-author
rjones@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

MS263

A Global Jacobian Method for Simultaneous Solution of Mortar and Subdomain Variables in Nonlinear Porous Media Flow

We describe a new algorithm to perform non-overlapping domain decomposition with nonlinear model problems, called the Global Jacobian (GJ) method. The are two main ideas: (1) we linearize the global system in both subdomain and interface variables simultaneously to yield a single Newton iteration; and (2) we algebraically eliminate subdomain velocities (and optionally, subdomain pressures) to solve either the 1st or 2nd Schur complement systems. The GJ method improves upon the previous nonlinear mortar algorithm, which required two nested Newton iterations and a Forward Difference (FD) approximation.

Benjamin Ganis
The University of Texas at Austin
Center for Subsurface Modeling
bganis@ices.utexas.edu

Mika Juntunen
University of Texas at Austin, ICES / CSM
mojuntun@gmail.com

Gergina Pencheva
University of Texas at Austin
gergina@ices.utexas.edu

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

Ivan Yotov
Univeristy of Pittsburgh
Department of Mathematics
yotov@math.pitt.edu

MS263

Feedback Control of the Boussinesq Equations with Application to Control of Energy Efficient Building Systems

In this talk, we present theoretical and numerical results for feedback control of the Boussinesq Equations. The problem is motivated by design and control of energy efficient building systems. In particular, new low energy concepts such as chilled beams and radiant heating lead to problems with Dirichlet, Neumann and Robin type boundary conditions. It is natural to consider control formulations that account for minimizing energy consumption and providing reasonable performance. We discuss a LQR type control problem for this system with Robin/Neumann boundary control inputs and apply the results to a 2D problem to illustrate the ideas and demonstrate the computational algorithms.

John A. Burns
Virginia Tech
Interdisciplinary Center for Applied Mathematics
jaburns@vt.edu

Xiaoming He
Department of Mathematics and Statistics
Missouri University of Science and Technology
hex@mst.edu

Weiwei Hu
University of Southern California
weiweihu@usc.edu

MS263

Model Reduction of Nonlinear PDEs Using Group POD

We propose a new method to reduce the cost of computing nonlinear terms in projection based reduced order models with global basis functions. We develop this method by extending ideas from the group finite element method to proper orthogonal decomposition (POD). Numerical results for a scalar 2D Burgers' equation show that the proposed group POD reduced order models are as accurate and are computationally more efficient than standard POD models of the Burgers' equation.

John Singler
Missouri S & T
Mathematics and Statistics
singlerj@mst.edu

Benjamin Dickinson
Munitions Directorate, Air Force Research Laboratory
Eglin AFB
benjamin.dickinson.1@us.af.mil

MS263

New Efficient Splitting Methods for Flow Problems

In this talk, we develop new reliable and efficient splitting methods for simulating fluid flows in two different problems: magnetohydrodynamics and Navier-Stokes equations with Coriolis force. For the former, we study partitioned methods which allow us to uncouple the problem by solving one each of the subphysics problems per time step (without iteration). For the latter, we present a fast-slow wave splitting method with Crank-Nicolson Leap-Frog scheme, which is unconditionally stable. We will discussed the stability and accuracy of our methods and give numerical experiments that support the theory. This is the joint work with Nan Jiang, William Layton and Catalin Trenchea (all from University of Pittsburgh).

Hoang A. Tran
University of Pittsburg
Department of Mathematics
hat25@pitt.edu

MS264

Quasi-Newton Update of Preconditioners for the Linearized Newton System Arising from 3d Discretizations of Groundwater Flow Models

We discuss quasi-Newton updates of preconditioners for linear systems arising in 3D groundwater flow problems.

Luca Bergamaschi
Universita di Padova
Italy
berga@dmsa.unipd.it

MS264

Update Preconditioners and Incomplete Decompositions using Approximate Inverses

In the past ten years we propose a framework for updating preconditioners in factorized form for sequences of general linear systems based on the use of inversion and sparsification of the incomplete factorizations. Examples of successful applications to sequences of systems arising in linear and nonlinear PDEs, image restoration, inexact Newton solvers, optimization, approximation of functions of large matrices and some issues of the proposed strategies will be considered.

S. Bellavia, D. Bertaccini, and B. Morini, Nonsymmetric preconditioner updates in Newton-Krylov methods for nonlinear systems, SIAM J. Sci. Comput., 33 (2011), pp. 2595-2619.
M. Benzi and D. Bertaccini, Approximate inverse preconditioning for shifted linear systems, BIT Numerical Mathematics, 43 (2003), pp. 231-244.
D. Bertaccini, Efficient preconditioning for sequences of parametric complex symmetric linear systems, Electronic Trans. on Num. Anal., 18 (2004), pp. 49-64.
D. Bertaccini and F. Sgallari, Updating preconditioners for nonlinear deblurring and denoising image restoration, Applied Numerical Mathematics, 60 (2010), pp. 994-1006.

Daniele Bertaccini
Department of Mathematics
Universita' di Roma Tor Vergata
bertaccini@mat.uniroma2.it

MS264

Incremental Approximate LU Factorizations and Applications

A problem that is common to many applications is to solve a sequence of linear systems whose matrices change only slightly from one step to the next. We address the issue of solving these systems using an ILU preconditioner without recomputing entirely the ILU factorizations but by updating them instead. Mathematical properties and implementation aspects of these incremental methods are discussed and numerical tests on a collection of linear systems and on the 2D Navier-Stokes equation with variable density are presented.

Caterina Calgaro
Laboratoire Paul Painleve - UMR 8524
Universite Lille 1, Sciences et Techno.
caterina.calgaro@univ-lille1.fr

Jean-Paul Chebab

Laboratoire Amienois de Mathématiques Fondamentales et Appl.
Universite de Picardie Jules Verne
jean-paul.chehab@u-picardie.fr

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

MS264
Updating Preconditioners for Model Reduction and Other Parameterized Systems

After a brief introduction on various approaches to updating preconditioners for sequences of linear systems, we will discuss our recent efforts for updating preconditioners for parameterized linear systems. Such systems arise in a range of applications such as acoustics, (parametric) model reduction, and inverse problems. We will show a very general approach to updating preconditioners for modest changes in parameters, thereby significantly reducing the cost of computing preconditioners while maintaining good convergence.

Eric De Sturler
Virginia Tech
sturler@vt.edu

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

MS265
Placing Communicating Tasks Apart to Maximize Effective Bandwidth

Topology-aware mapping algorithms should take the application behavior (latency versus bandwidth-bound) and target platform characteristics (routing schemes, implementation of collectives, etc.) into account. When mapping on n-dimensional tori, most previous work has focused on bringing communicating tasks closer on the network. We present a non-intuitive idea of moving tasks farther apart to increase the number of available routes and effective bandwidth on torus networks. Bandwidth-bound applications can benefit significantly from heuristics that implement such mappings.LLNL-ABS-571694

Abhinav Bhatele, Todd Gamblin
Lawrence Livermore National Laboratory
bhatele@llnl.gov, gamblin2@llnl.gov

MS265
Improving MPI Process Mapping for Cartesian Topologies on Multicore Nodes

We consider the problem of MPI process mapping for modern architectures that contain many cores on a single CPU with complex NUMA architectures. Although the MPI standard provides routines for process mapping in Cartesian topologies, most implementations do nothing to reduce off-node communications. We evaluate methods to improve process mapping on multicore nodes. We evaluate performance for hybrid nodes with accelerators; multiple MPI processes sharing an accelerator are used to overlap host-device communication with computation.

William M. Brown
Computational Biology
Sandia National Laboratories
brownw@ornl.gov

MS265
Task Mapping for Noncontiguous Allocations

This talk presents task mapping algorithms for non-contiguously allocated parallel jobs. Previous work on task mapping either uses a very general model that is hard to apply in practice or assumes that jobs are allocated to be completely isolated from each other. We pose the mapping problem for non-contiguous jobs utilizing a simple stencil communication pattern. We then devise several task mapping algorithms for this problem and evaluate their performance using simulations and experiments.

David Bunde
Knox College
dbunde@knox.edu

Vitus Leung
Sandia National Laboratories
vjleung@sandia.gov

MS265
Ensuring Continued Scalability of Mesh Based Hydrocodes

CTH is a widely used shock hydrodynamics code based on the bulk synchronous parallel programming model. At increasingly large processor counts, we noted that its nearest neighbor communication performance significantly degraded. MiniGhost, a miniapp in the Mantevo suite, was configured to mimic CTH's communication patterns, exposed the problem, and led to a solution that illustrates and informs a fundamental issue that must be considered in our preparations for exascale computing.

Courtenay T. Vaughan
Sandia National Laboratories
ctvaugh@sandia.gov

MS266
Energy Conserving Local Discontinuous Galerkin Methods for the Wave Propagation Problems

Wave propagation problems arise in a wide range of applications. The energy conserving property is one of the guiding principles for numerical algorithms, in order to minimize the phase or shape errors after long time integration. In this presentation, we develop and analyze a local discontinuous Galerkin (LDG) method for solving the wave equation. We prove optimal error estimates and the energy conserving property for the semi-discrete formulation. The analysis is extended to the fully discrete LDG scheme, with the energy conserving high order time discretization. Numerical experiments have been provided to demonstrate the optimal ratesof convergence. We also show that the shape of the solution, after long time integration, is well preserved due to the energy conserving property.

Yulong Xing
Department of Mathematics
Univeristy of Tennessee / Oak Ridge National Lab
xingy@math.utk.edu

MS266

Energy Stable Numerical Schemes and Simulations of Two Phase Complex Fluids on Phase Field Method

We present an energetic variational phase-field model for the two-phase Incompressible flow with one phase being the nematic liquid crystal. The model leads to a coupled nonlinear system satisfying an energy law. An efficient and easy-to-implement numerical scheme is presented for solving the coupled nonlinear system. We use this scheme to simulate two benchmark experiments: one is the formation of a bead-on-a-string phenomena, and the other is the dynamics of drop pinching-off. We investigate the detailed dynamical pinch-off behavior, as well as the formation of the consequent satellite droplets, by varying order parameters of liquid crystal bulk and interfacial anchoring energy constant. Qualitative agreements with experimental results are observed.

Xiaofeng Yang
University of South Carolina
xfyang@math.sc.edu

MS266

Krylov Implicit Integration Factor WENO Methods for Advection-diffusion-reaction Systems

Implicit integration factor (IIF) methods are originally a class of efficient “exactly linear part” time discretization methods for solving time-dependent partial differential equations with linear high order terms and stiff lower order nonlinear terms. In this paper, we developed new Krylov IIF-WENO methods to deal with both semilinear and fully nonlinear advection-diffusion-reaction equations. The methods are free of operator splitting error and can be designed for arbitrary order of accuracy. The stiffness of the system is resolved well and the methods are stable by using time step sizes which are just determined by the non-stiff hyerbolic part of the system.

Tian Jiang, Yongtao Zhang
Notre Dame University
tjiang@nd.edu, yzhang10@nd.edu

MS266

A Simple WENO Limiter for RKDG Methods

We investigate a simple limiter using weighted essentially non-oscillatory (WENO) methodology for the Runge-Kutta discontinuous Galerkin (RKDG) methods solving conservation laws, with the goal of obtaining a robust and high order limiting procedure to simultaneously achieve uniform high order accuracy and sharp, non-oscillatory shock transitions. The main advantage of this limiter is its simplicity in implementation, especially for multi-dimensional meshes.

Xinghui Zhong
Michigan State University
zhongxh1@gmail.com

MS267

Bayesian Inference with Processed Data Products

Abstract not available at time of publication.

Kenny Chowdhary
Sandia National Laboratories
kchowdh@sandia.gov

MS267

Probabilistic Schwarz Coupling for Fault-Tolerance and Scalability

In the drive towards exascale computing, simulation codes need to scale effectively to ever more cores, and be resilient against faults. We discuss a novel probabilistic Schwarz coupling approach, which employs a probabilistic representation of the knowledge about the solution at subdomain interfaces, and as such can naturally account for faults and other nondeterministic events. This probabilistic representation is updated asynchronously with information from independent subdomain computations for sampled values of the interface distribution.

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

Khachik Sargsyan, Cosmin Safta, Gilbert Hendry
Sandia National Laboratories
ksargsy@sandia.gov, csafta@sandia.gov,
ghendry@sandia.gov

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

MS267

Stochastic Reduced Models

Abstract not available at time of publication.

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

MS267

Statistically Robust and Parallel Load Balanced Sampling Algorithms for Bayesian Analysis and Multimodal Distributions

The Bayesian analysis of mathematical models often require the evaluation of multidimensional integrals related to posterior probability density functions (PDFs) of uncertain model parameters. Such integrals rarely can be computed analytically, which motivates the approximate calculation of their values with stochastic simulation methods that sample the corresponding posterior PDFs. In this talk we will discuss a recently proposed method that is statistically robust and parallel load-balanced. We will also present some numerical results.

Ernesto E. Prudencio
Institute for Computational Engineering and Sciences
University of Texas at Austin
prudenci@ices.utexas.edu

Sai Hung Cheung
School of Civil and Environmental Engineering
Nanyang Technological University, Singapore
shcheung@ntu.edu.sg

MS268

Adaptive Multi-resolution Solver for Multi-phase Flows with Sharp Interface Model and Efficient Data Structure

In this work, we present a block-based multi-resolution solver coupled with sharp interface model (MR-SIM) for multi-phase flows. While the solver updates the overlapped-block system according to two separate procedures, i.e. tracking interface position and MR analyzing for each individual phase, it uses a table-based storage-and-operation-splitting data structure. The data structure allows direct searches among blocks on the same or different resolution levels with indexes, and splits the block from its associated computational data. Since there is no extra memory required for the data associated with overlap regions of the blocks, highly efficient memory usage and inter-block communication are achieved.

Luhui Han, Xiangyu Hu, Nikolaus Adams
Technical University of Munich
Luhui.Han@aer.mw.tum.de, xiangyu.hu@aer.mw.tum.de, nikolaus.adams@tum.de

MS268

Interface Capturing with High-order Accurate Schemes: Pressure and Temperature Considerations

Compressible multifluid simulations are challenging due to the necessity to accurately represent discontinuities and transport processes. Direct application of shock capturing may produce spurious pressure oscillations and have been shown to generate temperature errors. The present focus is on high-order accurate interface capturing for multiphase flows. We show that, by appropriately transporting the relevant parameters of the equation of state, pressure, temperature and conservation errors can be prevented in gas/gas and gas/liquid flows.

Eric Johnsen
University of Michigan
ejohnsen@umich.edu

MS268

Simple Interface Sharpening Technique with Hyperbolic Tangent Function Applied to Compressible Two-Fluid Modeling

A very simple interface sharpening technique for compressible two fluids model is proposed. A main idea is using hyperbolic tangent function for reconstruction of volume fraction in the formulation of compressible two fluids model, which is similar to THINC or MTHINC scheme recently well-discussed in the incompressible multi-phase computation. This very simple technique sharpens the interface very much, although just the interpolation of volume fraction differs from conventional MUSCL scheme applied to two-fluids model. Several examples are shown in the presentation.

Taku Nonomura
JAXA/Institute of Space and Astronautical Science
nonomura@flab.isas.jaxa.jp

Keiichi Kitamura
Nagoya University
kitamura@fluid.nuae.nagoya-u.ac.jp

Kozo Fujii
JAXA/ISAS
fujii@flab.eng.isas.jaxa.jp

MS268

Eulerian Interface-sharpening Algorithms for Compressible Flow Problems

Our goal is to describe a novel Eulerian interface-sharpening approach for the efficient numerical resolution of interfaces arising from inviscid compressible flow in more than one space dimension. The algorithm uses the compressible Euler equations as the model system, and introduces auxiliary differential terms to the model so as to neutralize numerical diffusion that is inevitable when the original Euler system is solved by a diffused interface method. A standard fractional-step method is employed to solve the proposed model equations in two steps, yielding an easy implementation of the algorithm. Sample numerical results are shown to demonstrate the feasibility of the algorithm for sharpening compressible interfaces numerically.

Keh-Ming Shyue
Department of Mathematics, National Taiwan University
shyue@math.ntu.edu.tw

MS269

Variational Data Assimilation and Particle Filters

There is a subtle similarity between the 4D-Var and EnKF methods for linear, Gaussian models, that can be used to design a hybrid ensemble filter which performs better than the regular EnKF. This paper investigates the conjecture, that combining a variational method with particle filtering can deliver information useful for improving the performance of particle filters.

Haiyan Cheng
Willamette University
hcheng@willamette.edu

MS269

Ensemble Data Assimilation Using An Unstructured Adaptive Mesh Ocean Model

For an unstructured adaptive mesh ocean model, a supermesh technology is applied to the adapted meshes of the ensembles. Mesh adaptivity is also adopted around the observation locations to improve the efficiency of the data assimilation process and to address the key problem of the representativity of the observations in ocean data assimilation.

Juan Du
Institute of Atmospheric Physics, Chinese Academy of Science
dujuan10@mail.iap.ac.cn

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Jiang Zhu
Institute of Atmospheric Physics
Chinese Academy of Sciences,
jzhu@mail.iap.ac.cn

C.C. Pain, P.A. Allison
Department of Earth Science and Engineering
Imperial College London, U.K.
c.pain@imperial.ac.uk, p.a.allison@imperial.ac.uk

MS269

POD/DEIM 4-D Var Data Assimilation Applied to a Nonliner Adaptive Mesh Model

A novel POD 4-D Var model has been developed for a nolinear adaptive mesh model. For nonlinear problems, a perturbation approach is used to help accelerate the matrix equation assembly process. The discrete empirical interpolation method (DEIM) is further applied to the POD model and achieves a complexity reduction of the high order nonlinear term in the full model. A non-linear Petrov-Galerkin method is used for improving the stability of POD results without tuning parameters.

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Dunhui Xiao
Department of Earth Science and Engineering
Imperial Colleg London
dh.xiao@imperial.ac.uk

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

Juan Du
Institute of Atmospheric Physics, Chinese Academy of Science
dujuan10@mail.iap.ac.cn

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

MS269

Data Assimilation 2

Abstract not available at time of publication.

Ionel M. Navon
Florida State University
Department of Scientific Computing
inavon@fsu.edu

MS270

Landslides Thresholds for Early-Warning Systems: History, Challenges and Perspectives

The earliest thresholds for landslides were derived by visually fitting boundaries to a few triggering events using ground-based measured rainfall. Over the years, threshold estimation has faced increasing complexities: larger databases, real-time monitoring from both ground-based and remote sensing products, and demand for optimum performance for early-warning systems. The evolution of the state-of-the-art in landslide thresholds is presented, emphasizing recent developments in data mining and probabilistic classification with case studies in Europe and Central America.

Jose Cepeda
International Center for Geohazards
Norwegian Geotechnical Institute, Oslo Norway
jose.cepeda@ngi.no

Dalia Kirschbaum
Hydrological Sciences Laboratory
NASA Goddard Space Flight Center, MD, USA
dalia.b.kirschbaum@nasa.gov

Zenon Medina-Cetina
Zachry Department of Civil Engineering
Texas A&M University
zenon@tamu.edu

MS270

Groundwater Flow, Poroelastic Waves and Liquefaction

Liquefaction is one of the less well quantified consequences of earthquakes: engineering assessments of liquefaction risk often boil down to back of a cigarette packet calculations. Uncertainties are generally not considered. However such assessments provide a basis for land classification and foundation engineering. This talk summarises our attempts to build rational models of liquefaction and hence better understand the mechanisms that lead to liquefaction. As a backdrop we consider the case of Christchurch, New Zealand.

Nicholas Dudley Ward
Otago Computational Modelling Group
nick@ocmo.co.nz

Jari Kaipio
Department of Mathematics
University of Auckland
jari@math.auckland.ac.nz

Tiangang Cui
MIT
tcui@mit.edu

MS270

Accelerating Numerical Modeling of Waves Propagating Through 2-D Anisotropic Materials Using a Graphic Processing Card

We present an implementation and analysis of performance over multi-threaded executions using graphic cards as parallel computing devices to model numerically (using finite differences implemented in PyOpenCL) the behavior of waves propagating through 2-D anisotropic materials. This approach is inspired by laboratory experiments where a sample plate of different anisotropic materials inside a water tank is rotated and for every angle of rotation is subjected to the emission of an ultrasonic wave.

Ursula Iturraran
Universidad Nacional Autonoma de Mexico UNAM
ursula.iturraran@gmail.com

Miguel Molero
Centro de Acustica Aplicada y Evaluacion No Destructiva
CAEND (CSIC-UPM), Madrid, Spain
miguel.molero@csic.es

MS270

Diagnosis and Prognosis Analysis of Ecological Management under Varying Climate Change Scenarios

Unsustainable management of rangeland where energy development systems take place leads to the degradation of both the resource base, the value of the commodities, and the benefits these generate. Sustainable development and its inherent spatio-temporal complexities include interactions across social, ecological and economic paradigms, including climate change, considered a key triggering mechanism for sudden interventions across these processes. A Bayesian mapping and GIS tool is presented as applied to ecoregions of shale gas development.

Patricia Varela
Texas A&M University
varela_patricia@yahoo.com

Zenon Medina-Cetina
Zachry Department of Civil Engineering
Texas A&M University
zenon@tamu.edu

Bill Fox
Blackland Research & Extension Center
Texas AgriLife - Research, TX USA
bfox@brc.tamus.edu

Jay Angerer
Blackland Research & Extension Center
Texas AgriLife - Research, TX, USA
jangerer@brc.tamus.edu

Luis Alberto Munoz Ubando
Plenum Soft
Merida Yucatan, Mexico
luisalbertomunozubando@gmail.com

MS271

Exterior Calculus Stuff

Abstract not available at time of publication.

Anil N. Hirani
University of Illinois at Urbana-Champaign
Department of Computer Science
hirani@cs.illinois.edu

MS271

Mimetic Finite Difference Method for the Stokes Equations

We present two formulations of the mimetic finite difference (MFD) method for the Stokes equations on arbitrary polygonal meshes with convex and non-convex cells. The formulations correspond to C^0 and C^1 continuity of the velocity field. The numerical and theoretical analysis of the stability will be discussed in great details. We show how flexibility of the MFD method can be used to enforce stability for problematic mesh configurations.

Lourenco Beirao Da Veiga
Universita' degli Studi di Milano
lourenco.beirao@unimi.it

Konstantin Lipnikov, Gianmarco Manzini
Los Alamos National Laboratory
lipnikov@lanl.gov, gm.manzini@gmail.com

MS271

Is the Keller-Box Scheme Mimetic?

The Keller-Box scheme is a method for solving hyperbolic and parabolic PDEs that has been shown to have a number of attractive physical and mathematical properties. It can handle discontinuous material coefficients and works well on distorted or high aspect ratio meshes. Most recently the Keller-Box scheme has been shown by Reich to multi-symplectic (2000) and by Frank (2006) to always propagate waves in the correct direction. By deriving the scheme as a type of exact discretization we show that the differential operators in the Keller-Box scheme are mimetic. However, the mimetic gradient operator is shown to be fundamentally different from that found in mimetic FE (Nedelec) methods, or commonly used as the basis for SOM methods. And the method is not easily described by algebraic topology or discrete differential forms. The advantages and disadvantages of this unusual numerical method compared to other mimetic schemes are discussed.

Blair Perot
University of Massachusetts - Amherst
tba

MS271

Mimetic Methods Toolkit: An Object-Oriented Api Implementing Mimetic Discretization Methods with Application Examples in Oil Reservoir Simulation.

In this work, we introduce Mimetic Methods Toolkit (MTK), an object-oriented Application Programming Interface for the implementation of Mimetic Discretization Methods in developing computer applications of a scientific nature. We present examples on how can MTK be used to simulate pressure distribution in oil reservoirs. We present the fundamentals of the mathematical models for oil reservoirs as porous media, and we present an introduction to Mimetic Discretization Methods, while we discuss the advantages of using both Mimetic Discretization Methods and MTK. Finally, we explain how MTK implements these methods, and we compare the results against well-known discretization methods.

Eduardo Sanchez
San Diego State University
ejspeiro@gmail.com

MS272

Efficient Nonlinear Model Reduction Approach using Local Reduced bases and Hyper-reduction

A new model reduction approach for nonlinear dynamical systems developed. This method is based on three ingredients: 1) the definition of local reduced-order bases that can better capture the subspace containing the solution in a given regime than a global counterpart, 2) a local hyper-reduction technique resulting in an efficient algorithm and 3) an efficient basis update procedure. Applications to the reduction of CFD-based dynamical systems highlight the effectiveness of the method.

David Amsallem, Kyle Washabaugh
Stanford University
amsallem@stanford.edu, kwash@stanford.edu

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

Charbel Farhat
Stanford University
CFarhat@stanford.edu

MS272

Frequency-weighted $\mathcal{H}_2$-optimal Model Reduction

In this talk, we discuss different optimality conditions arising in the context of frequency weighted $\mathcal{H}_2$-model reduction for linear control systems. In particular, we establish a connection between a recently introduced distributed interpolation framework and a special class of linear matrix equations. We further provide a discussion on numerically efficient methods to construct locally optimal reduced models.

Tobias Breiten
Max-Planck-Institute for Dynamics of Complex Technical Systems, Magdeburg
breiten@mpi-magdeburg.mpg.de

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

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

MS272

Greedy Algorithms and Stable Variational Formulations

Greedy schemes for computing reduced bases for parameter dependent PDEs are typically based on efficiently computable surrogates for the actual current distance between the reduced space and the solution set S. It can be shown that the accuracy provided by the reduced spaces is "rate-optimal', when compared with the Kolmogorov n-widths of S, if the surrogate is tight. By this we mean that up to a constant the surrogate is also a lower bound for the distance of the reduced space from the solution set S. We show how to obtain such tight surrogates also for non-elliptic problems such as, for instance, pure transport equations or singularly perturbed unsymmetric problems by deriving variational formulations that are in a certain sense uniformly stable and by properly stabilizing the reduced spaces. The results are illustrated by some numerical experiments.

Wolfgang Dahmen, Gerrit Welper, Christian Plesken
RWTH Aachen
IGPM
dahmen@igpm.rwth-aachen.de, welper@igpm.rwth-aachen.de, plesken@igpm.rwth-aachen.de

MS272

Greedy Algorithms for Eigenvalue Problems

Greedy algorithms are promising algorithms to treat high-dimensional PDEs. Encouraging numerical results have already been obtained in a large number of applications. Recently, theoretical convergence results have been obtained for nonlinear convex optimization problems. The aim of this talk is to present new convergence results for this family of numerical methods regarding high-dimensional eigenvalue linear problems, along with some numerical illustrations.

Virginie Ehrlacher
CERMICS - Ecole des Ponts Paristech / INRIA
ehrlachv@cermics.enpc.fr

Tony Lelievre
Ecole des Ponts ParisTech
lelievre@cermics.enpc.fr

Eric Cances
Ecole des Ponts and INRIA, France
cances@cermics.enpc.fr

MS273

Multi-fidelity Analysis and Optimization for Low-boom Supersonic Aircraft

Abstract not available at time of publication.

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

MS273

Advanced Mathematical Techniques for New Systems-level Analysis and Optimization

We seek to advance the science of systems integration by systematically implementing sound scientific principles into the engineering process, a "physics-based approach.' Three areas of recent progress will be highlighted in this presentation. First, a generic dynamic stochastic model was developed that employs a methodology that predicts and corrects the time-varying parameters of a dynamic system. This procedure was employed to reduce the growth of the prediction interval for design predictions and is extensible to all dynamic stochastic mathematical models. Second, a formulation for calculating and measuring dynamic entropy generation rate was developed that is consistent with computational and experimental determinations. This formulation allows for direct calculation of losses and supports the physics-based modeling supported by experimental validations. The importance of calculating the entropy generation dynamically is illustrated in a numerical example. Third and last, the impact of sampling methodology on uncertainty propagation is detailed in a computational example. Psuedo-random sampling is compared to low-discrepancy-sequence sampling to illustrate numerically-induced uncertainties that influence uncertainty propagation and sensitivity analysis. Uncertainty analysis, uncertainty propagation and sensitivity analysis are all performed for a nonlinear system.

Jose A. Camberos
Air Force Research Laboratory
Multidisciplinary Science & Technology Center
jose.camberos@wpafb.af.mil

John Doty
Air Force Research Laboratory,
Multidisciplinary Science & Technology Center
jdoty1@udayton.edu

Ray Kolonay
Air Force Research Laboratory
Multidisciplinary Science & Technology Center
raymond.kolonay@wpafb.af.mil

MS273

Multifidelity Approaches for Parallel Multidisciplinary Optimization

This talk formulates two methods to parallelize the optimization of multidisciplinary systems. The first method decomposes the system optimization problem into multiple subsystem optimizations that are solved in parallel. The second method generates a list of designs at which computationally expensive simulations should be run, evaluates those designs in parallel, and then solves an inexpensive surrogate-based optimization problem. Both methods enable the use of multifidelity optimization and are high-fidelity gradient-optional, i.e. they exploit high-fidelity sensitivity information when available, but do not require gradients of the high-fidelity models.

Andrew I. March
Massachusetts Institute of Technology
amarch@ll.mit.edu

MS273

Optimization Under Uncertainty Using Control Variates

Accounting for uncertainty and variability is important to design robust and reliable systems, but it is often too computationally expensive to nest Monte Carlo simulation within design optimization. The control variate method can take advantage of the correlation between random model outputs at two design points to reduce estimator variance. By chaining control variates through the sequence of optimization design points, we can achieve significant computational savings. Numerical experiments demonstrate 80% reduction in model evaluations.

Leo Ng, Karen E. Willcox
Massachusetts Institute of Technology
leo_ng@mit.edu, kwillcox@MIT.EDU

MS274

Accelerating MCMC with Local Quadratic Models

In many inference problems, the cost of MCMC analyses is dominated by repeated evaluations of expensive forward models. Surrogate methods attempt to exploit model regularity by substituting an inexpensive approximation constructed from a limited number of model runs. The construction of globally accurate surrogates may be prohibitively expensive, however. We therefore propose to interleave MCMC with the construction of local quadratic surrogate models, borrowing from trust-region methods and approximation results in derivative free optimization.

Patrick R. Conrad
MIT
prconrad@mit.edu

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

MS274

Projections on Positive Random Variables in Finite Wiener Chaos Spaces

Abstract not available at time of publication.

Evangelia Kalligiannaki
University of Southern California
vageliwka@gmail.com

MS274

Preconditioned Bayesian Regression for Stochastic Chemical Kinetics

We develop a preconditioned Bayesian regression method that enables sparse polynomial chaos representations of noisy outputs for stochastic chemical systems with uncertain reaction rates. The approach is based on coupling a multiscale transformation of the state variables with a Bayesian regression formalism. Numerical experiments show that the approach accommodates large noise levels and large variability with uncertain parameters, and enables efficient and robust recovery of both the transient dynamics and the corresponding noise levels.

Olivier P. Le Maitre
LIMSI-CNRS
olm@limsi.fr

Alen Alexanderian
Johns Hopkins University
aalexa20@jhu.edu

Francesco Rizzi
Department of Mechanical Engineering
Johns Hopkins University
francesco.rizzi@duke.edu

Muruhan Rathinam
University of Maryland, Baltimore County
muruhan@umbc.edu

Omar M. Knio
Duke University
omar.knio@duke.edu

MS274

Hybrid Discrete/Continuum Algorithms for Stochastic Reaction Networks

Direct solutions Chemical Master Equation (CME) models governing Stochastic Reaction Networks (SRNs) are generally prohibitively expensive due to excessive numbers of possible discrete states in such systems. To enable efficiency gains we develop a hybrid approach where states with low molecule counts are treated with the discrete CME model while states with large molecule counts are modeled by the continuum Fokker-Planck equation. The performance of this novel hybrid approach is explored in canonical SRNs configurations.

Cosmin Safta, Khachik Sargsyan
Sandia National Laboratories
csafta@sandia.gov, ksargsy@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

PP1

Parameter Mesh Adaptivity for Ill-Posed Inverse Problems Based on the Dominant Modes of the Misfit Hessian

In ill-posed parameter identification problems, the observed data typically provide different levels of information about the parameter in different parts of the domain. In particular, the informativeness of the data is reflected in the eigenstructure of the misfit Hessian. Here, we present a mesh adaptivity scheme in which the parameter mesh is refined and coarsened to approximate the dominant (well informed) modes of the misfit Hessian.

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

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

PP1

Flexible Krylov Subspace Methods for Shifted Systems

We discuss methods for solving shifted systems of the form $K + \sigma_j M = b, \quad j = 1, \ldots, n_f$, where $\sigma_j \in C$ using an Arnoldi-based method, with flexible preconditioners of the form $K + \tau M$ that are inverted using iterative solvers. We demonstrate this method with error estimates and numerical examples from Oscillatory Hydraulic Tomography.

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

Arvind Saibaba
Stanford University
arvindks@stanford.edu

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

PP1

Spatiotemporal Dynamics of Cardiac Electrical Alternans

We study the behavior of cardiac alternans, a period-2 cardiac rhythm. During spatially discordant alternans, waves generated from long action potentials at the stimulus site become short as they propagate through the tissue and vice versa. Separating these regions are sites called nodes where action potentials remain identical from beat to beat. We analyze the dynamics of these nodes including their locations and movement with and without spatial gradients in electrophysiological properties.

Michael Bell
Rochester Institute of Technology
mmbsma@rit.edu

Elizabeth M. Cherry
Rochester Institute of Technology
School of Mathematical Sciences
excsma@rit.edu

PP1

Mass-Conservative Adaptive Mesh Unrefinement Computation for Shallow Water Flow Simulations

The h-adaptive unstructured mesh refinement (h-AMR) has been widely used to improve mesh resolution locally in order to achieve better quality of numerical simulations while sustaining limited computer power. In AMR, massconservation is often neglected when a node is simply deleted during coarsening. A local approach to conserve mass for unrefinement process is presented. Mathematical derivation for a higher-order time derivative and its detailed implementation are described. Experimental results also demonstrate our successful development.

Ruth Cheng
US Army Corps of Engineers
Ruth.C.Cheng@usace.army.mil

PP1

Overlapping Local/Global Iteration Scheme for Whole-Core Neutron Transport Calculation

In the overlapping local/global (OLG) iteration scheme described in this paper, local transport calculation and global diffusion-like calculation are coupled through interface boundary conditions. Local calculations are performed over half-assembly overlapping subregions since this overlapping takes into account the inter-assembly transport effect. The global calculation is performed via partial current-based coarse-mesh finite difference (p-CMFD) method to obtain the whole-core transport solution. This two-level iterative method is tested on several multi-slab and two-dimensional rectangular geometry problems, with encouraging results.

Seungsu Yuk, Nam Zin Cho
Korea Advanced Institute of Science and Technology
sixpig@kaist.ac.kr, nzcho@kaist.ac.kr

PP1

A GPU-Accelerated Method of Regularized Stokeslets

We present a GPU-accelerated computational method for simulating the mechanics of poroelastic materials. The elastic material is described by a collection of point forces, and we develop a parallel implementation of the method of regularized Stokeslets to solve for the fluid motion. Results from our method are compared with a naive serial algorithm for two and three spatial dimensions. We demonstrate the performance gain of our algorithm in a simulation of the dynamics of cellular bleb formation.

Calina A. Copos
University of California Davis
ccopos@math.ucdavis.edu

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

Wanda Strychalski
Department of Mathematics
University of California, Davis
wanda@math.ucdavis.edu

PP1

Sparse Matrix Operations on GPU Architectures

GPUs have been leveraged in numerous numerical applications to achieve tremendous gains in performance. Although operations on dense matrices have been studied extensively their sparse counterparts, with the exception of SpMV, have developed relatively slowly with few tangible implementations. In this work we present work concerning several data dependent and irregular operations on sparse matrices on GPU architectures: sparse matrix-matrix multiplication, graph contractions, and graph partitioning.

Steven Dalton
University of Illinois at Urbana-Champaign
dalton6@illinois.edu

PP1

Matrix Functions and the NAG Library

Functions of matrices are of growing interest in science and engineering due to the concise way they allow problems to be formulated and solutions to be expressed. The Numerical Linear Algebra Group at the University of Manchester has developed many algorithms for computing matrix functions. These are currently being implemented in the NAG Software Library. This poster will introduce some of the techniques used in the algorithms and discuss some of the implementation details.

Edvin Deadman
Numerical Algorithms Group and University of Manchester
edvin.deadman@nag.co.uk

PP1

Computational Tools for Digital Holographic Microscopy

Digital holographic microscopy is in principle a fast 3D imaging technique, but extracting precise 3D information from the 2D holograms is challenging, especially for complex samples. A promising approach used for colloidal systems involves fitting scattering models to the holograms. We demonstrate a related approach to imaging biological samples with much more complex shapes and structures. Our method makes use of the discrete dipole approximation and a new mathematical approach to solving inverse problems.

Thomas G. Dimiduk
Harvard University
Manoharan Lab
tdimiduk@physics.harvard.edu

Jerome Fung, Rebecca Perry, Vinothan Manoharan
Harvard University
fung@physics.harvard.edu, rperry@seas.harvard.edu, vnm@seas.harvard.edu

PP1

Group Steiner Problem: An Ant Colony Optimization Based Solution and Practical Applications

The Group Steiner Problem (GSP) is an important generalization of some basic NP-hard problems. Many complex real-world applications require solving the GSP in graphs modeling the topology of the given problem, such as: the design of Very Large Scale Integration (VLSI), the design of a minimal length irrigation network, and routing problem for wireless sensor networks. In this talk, we first describe topology models of these applications solved by searching the minimum group Steiner tree. Next, we show our design of some new algorithms based on Ant Colony Optimization model to solve the GSP in general graphs. Our experimental results show that our method outperforms the best other heuristic methods for GSP.

Thuan P. Do
Dr.
thuan-do@uiowa.edu

Duong Nguyen
School of Information and Communication Technology
Hanoi University of Science and Technology
thaiduongnguyen91@gmail.com

PP1

Block Conjugate Gradient Type Methods for the Approximation of Bilinear form $C^H A^{-1} B$

We consider approximating the bilinear form $C^H A^{-1} B$, where matrices $A \in C^{n\times n}$, C and $B \in C^{n\times m}$, usually $m \ll n$. The common way is first to solve $AX = B$, linear systems with multiple right-hand sides, by block Krylov subspace methods. Then, $\eta := C^H X$ can be obtained. Inspired by methods in [Strakos and Tichý, SISC, 2011], we propose block conjugate gradient type methods for $C^H A^{-1} B$. Numerical results will be presented to show the efficiency of our methods.

Lei Du
University of Tsukuba, Japan
JST-CREST
du@mma.cs.tsukuba.ac.jp

Yasunori Futamura
University of Tsukuba
futamura@mma.cs.tsukuba.ac.jp

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

PP1

An Asymptotic-Based Numerical Algorithm for Analyzing Nonlinear Waves

Resolving the complicated fluid-solid interactions of the mammalian cochlea requires solving a nonlinear wave problem; however, the geometric and material properties of the ear make this difficult to do with numerics alone. In this work a hybrid analytic-numeric approach is used. Asymptotic methods are utilized in conjunction with an iterative numerical method to achieve an approximate solution to

the nonlinear wave problem.

Kimberly Fessel, Mark Holmes
Rensselaer Polytechnic Institute
fessek@rpi.edu, holmes@rpi.edu

PP1

Scalable Methods for Large-Scale Bayesian Inverse Problems

We address the challenge of large-scale nonlinear statistical inverse problems by developing an adaptive Hessian-based Gaussian process response surface method to approximate the posterior pdf solution. We employ an adaptive sampling strategy for exploring the parameter space efficiently to find interpolation points and build a global analytical response surface far less expensive to evaluate than the original. The accuracy and efficiency of the response surface is demonstrated with examples, including a subsurface flow problem.

H. Pearl Flath
Institute for Computational Engineering and Sciences
The University of Texas at Austin
pflath@ices.utexas.edu

PP1

Parallel in Time Using Multigrid

One of the main challenges facing computational science with future architectures is that faster compute speeds will be achieved through increased concurrency, since clock speeds are no longer increasing. As a consequence, traditional time marching will eventually become a huge sequential bottleneck, and parallel in time integration methods are necessary. This poster discusses optimality and parallelization properties of one approach for parallelization in time, namely a multigrid-in-time method that is fairly unintrusive on existing codes.

Stephanie Friedhoff
Tufts University
Stephanie.Friedhoff@tufts.edu

Robert Falgout, Tzanio Kolev
Lawrence Livermore National Laboratory
falgout2@llnl.gov, kolev1@llnl.gov

Scott Maclachlan
Department of Mathematics
Tufts University
scott.maclachlan@tufts.edu

Jacob Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

PP1

A Fully Implicit Newton-Krylov Method for Euler Equations

Fully implicit methods for hyperbolic PDEs are relatively rare. We present a fully nonlinearly implicit in time method for the Euler equations of gas dynamics based on a Newton-Krylov approach. Numerical examples are presented for one-dimensional smooth linear wave propogation, and for the Riemann shock-tube problems with discontinuities. Numerical tests with CFL number as large as Order(10^3) demonstrate that the computed solutions are in good agreement with exact solutions. Preconditioning strategies will be presented.

Wei Gao
King Abdullah University of Science and Technology
wei.gao@kaust.edu.sa

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

PP1

Lighthouse Taxonomy: Delivering Linear Algebra Solutions

While scientists use linear algebra in a wide variety of applications, they typically lack the training required to develop high-performance implementations. We narrow this gap with a taxonomy entitled Lighthouse, which steers practitioners from algorithmic descriptions to high-performance implementations. Lighthouse unifies a routine database, search capabilities, and code generation and tuning. For tuning, Lighthouse employs the Build To Order (BTO) compiler that reliably achieves high performance for linear algebra via a variety of optimization techniques.

Paul L. Givens, David Johnson, Javed Hossain
CU - Boulder
pgivens4915@gmail.com, david.l.johnson@colorado.edu, javed.hossain@colorado.edu

Sa-Lin Bernstein
Computation Institute, University of Chicago
salin@mcs.anl.gov

Elizabeth Jessup
CU - Boulder
erjessup@gmail.com

Boyanna Norris
Argonne National Laboratory
norris@mcs.anl.gov

PP1

On the Relationship between Polynomial Chaos Expansions and Gaussian Process Regression

Gaussian process regression (GPR) and polynomial chaos expansions (PCE) are often used as surrogates to approximate the outputs of complex computational simulations—for statistical inference, uncertainty propagation, and other tasks. We use Mercer's theorem to relate the choice of kernel and training points in GPR to the choice of polynomial basis and quadrature scheme in PCE. This relationship is used to show the conditions under which the mean prediction of a Gaussian process is equivalent to a polynomial chaos expansion.

Alex A. Gorodetsky
Massachussets Institute of Technology
goroda@mit.edu

Tarek Moselhy
MIT
tmoselhy@mit.edu

Youssef M. Marzouk

Massachusetts Institute of Technology
ymarz@mit.edu

PP1

Effects of Abruptly Reversing Shear Flow on Red Blood Cells

Blood pumps and other cardiovascular devices subject red blood cells to fluctuating shear flow, which leads to shear stresses well beyond their normal physiological values and may cause hemolysis. Using the lattice Boltzmann and Immersed Boundary methods, we model the red blood cell as a viscoelastic biconcave capsule. The capsule dynamics and shape change under abruptly reversing shear flow are considered, particularly with respect to dependence on peak shear stress and capsule characteristics (e.g., bending stiffness and membrane viscosity).

John Gounley
Old Dominion University
jgounley@odu.edu

Yan Peng
Dept of Math. & Stat.
Old Dominion University
ypeng@odu.edu

PP1

Inverse Sensitivity Analysis for Storm Surge Models

We describe a framework for quantifying uncertainty in spatially varying parameters critical to the accuracy of model forecasts of storm surge. We characterize uncertainty in Manning's n (a bottom friction parameter) using an inverse sensitivity analysis applied to the ADCIRC storm surge model. We model bottom friction as a random field, and use a Karhunen-Loève expansion to obtain a finite dimensional approximation. The inverse sensitivity analysis is based on a measure-theoretic approach to inversion.

Lindley Graham
The University of Austin at Texas
Institute for Computational Engineering and Sciences (ICES)
lgraham@ices.utexas.edu

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

Troy Butler, Don Estep
Colorado State University
butler@stat.colostate.edu, estep@stat.colostate.edu

Joannes Westerink
Department of Civil Engineering and Geological Sciences
University of Notre Dame
jjw@nd.edu

PP1

A Multigrid Algorithm for Elliptic Type Problems Using the Hierarchical Element Structure

We present a multigrid algorithm for elliptic type problems, where the projection / restriction operators are built using the hierarchical element structure, in which the basis of degree n+1 is obtained as a correction to the degree n basis. This with respect to a geometric multigrid avoids the complication of building recursive grid structures. The algorithm is implemented in 2 and 3 dimensions, and will be extended to more complicated physics, including Navier-Stokes and FSI.

Janitha Gunatilake
Texas Tech University
janitha.gunatilake@ttu.edu

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

PP1

A Bayesian Framework for Uncertainty Quantification in the Design of Complex Systems

This poster presents a Bayesian framework for the design of complex systems, in which uncertainty in various parameters is characterized probabilistically, and updated through successive design iterations as new estimates become available. Incorporated in the model are methods to quantify system complexity and risk, and reduce them through the allocation of resources for redesign and refinement. This approach enables the rigorous quantification and management of uncertainty, thereby serving to help mitigate technical and programmatic risk.

Qinxian He, Douglas Allaire, John Deyst, Karen E. Willcox
Massachusetts Institute of Technology
qche@mit.edu, dallaire@mit.edu, deyst@mit.edu, kwillcox@MIT.EDU

PP1

Practical Experience with Gaussian Processes in Quantification of Margins and Uncertainties

Quantification of margins and uncertainties is a process by which performance and safety of engineered systems are assessed. Computational simulation, combined with experimental testing, plays a key role in this type of analysis. However, full system models can be too computationally intensive to fully explore the relevant uncertainties, so emulators such as Gaussian processes are often used. We describe our successes and challenges using Gaussian processes in assessing margins and uncertainties for an engineered system.

Patricia D. Hough, Jeff Crowell
Sandia National Laboratories
pdhough@sandia.gov, jacrowe@sandia.gov

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

PP1

Adaptive Time Stepping in Moose

The INL's Multi-physics Object Oriented Simulation Environment (MOOSE) had the ability to integrate over time using backwards Euler, BDF2, and Crank-Nicolson. Variable time step selection was highly limted. I have added in an AB2 predictor according to "Philip Gresho, David Griffiths, and David Silvester, Adaptive time-stepping for incompressible flow. Part 1: scalar advection-diffusion" for both Crank-Nicolson and BDF2. I am also adding in variable step-variable order time integration according.

John T. Hutchins
Boise State University'
Hutchinsjh@hotmail.com

Michael Pernice
Idaho National Laboratory
Michael.Pernice@inl.gov

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

PP1

A 3D Pharmacophore-Based Scoring Method for DOCK: Application to HIVgp41

Pharmacophore modeling incorporates geometrical and chemical features of either know inhibitors(s) or the targeted binding site in order to rationally design new drug leads. This presentation describes our efforts to encode a three dimensional pharmacophore matching similarity (FMS) scoring function into the program DOCK for use in virtual screening and de novo design. Application to the drug target HIVgp41 using binding profiles for known peptide inhibitors to generate reference pharmacophores will be shown.

Lingling Jiang
Rizzo Research Group, Stony Brook University
jianglilian68@gmail.com

Robert Rizzo
Department of Applied Math. & Stat., Stony Brook University
Laufer Center for Physical & Quantitative Biology
rizzorc@gmail.com

PP1

Optimal Dirichlet Boundary Control for the Navier-Stokes Equations

We consider an optimal Dirichlet boundary control problem for the Navier–Stokes equations. The control is considered in the energy space where the related norm is realized by the so-called Steklov–Poincaré operator. We introduce a stabilized finite element method for the optimal control problem with focus on lowest order elements. At the end, we present some numerical results demonstrating the differences between the realization of the control in $L_2(\Gamma)$ and the energy space $H^{1/2}(\Gamma)$.

Lorenz John
Institute of Computational Mathematics
Graz University of Technology
john@tugraz.at

Olaf Steinbach
TU Graz, Institute of Computational Mathematics
o.steinbach@tugraz.at

PP1

A Deformed Spectral Quadrilateral Multi-Domain Penalty Model for the Incompressible Navier-Stokes Equations

A penalty method is a variant of a spectral element method that weakly enforces continuity between adjacent elements and weakly enforces continuity at physical boundaries. Furthermore, at the boundaries, the PDE is also partially satisfied. The spirit of such a formulation is that in theory, a PDE operates arbitrarily close to any measure-zero boundary. Here, a previous spectral multi-domain penalty model for the incompressible Navier-Stokes equations is extended to include deformed boundaries for shoaling-type problems encountered in Environmental Fluid Mechanics. Some difficulties addressed include satisfying compatibility conditions in a (pseudo-)pressure Poisson equation that arises. A previous strategy to satisfy compatibility by use of a null singular vector is presented, and strategies to enforce compatibility for the deformed problem are discussed. Results are shown for standard incompressible flow benchmarks. The primary goal of this work is to model nonlinear internal wave propagation along a shallow, sloping bathymetry, as may be characteristic of a continental shelf region.

Sumedh Joshi
Center for Applied Mathematics
Cornell University
smj96@cornell.edu

Peter Diamessis
Department of Civil and Environmental Engineering
Cornell University
pjd38@cornell.edu

PP1

A Bayesian Approach to Feed Reconstruction in Chemical Processes

We develop a fully Bayesian hierarchical model to estimate the detailed chemical composition of a stream in a petroleum refinery from limited measurements (e.g., of bulk properties and elemental composition). Complex prior information, coupled with positivity and normalization constraints on the composition parameters, make it difficult to sample from the resulting high-dimensional posterior distribution. We propose a method to efficiently generate posterior samples and validate the method on example data sets.

Naveen Kartik, Youssef M. Marzouk
Massachusetts Institute of Technology
nkartik@mit.edu, ymarz@mit.edu

PP1

Nonlinear Dynamic State Reconstruction With Applications to Exposure-Based (Bio)Chemical Hazard Assessment

A new approach to exposure-based (bio)chemical hazard assessment is proposed through a nonlinear dynamic state reconstruction method. The formulation of the nonlinear state estimator problem is realized via a system of invariance nonhomogeneous functional equations, a general set of necessary and sufficient conditions for solvability is derived and an easily programmable series solution method is developed. Finally, the performance of the proposed state

estimator is evaluated in an illustrative case study.

Nikolaos Kazantzis
Department of Chemical Engineering
Worcester Polytechnic Institute
nikolas@wpi.edu

PP1

A Hybrid CPU-GPU Approach to Fourier-Based Image Stitching

We present a hybrid CPU-GPU system for the Fourier-based stitching of 2D optical microscopy images. The system uses CPU and GPU resources to achieve interactive stitching rates on large problems. It stitches a grid of 59×42 tiles in 29.5 s. This is a 42x speedup over an optimized sequential implementation which takes 20.5 min for the same workload. This is a major step towards computationally steerable experiments in Biology that rely on automated optical microscopes.

Walid Keyrouz, Bertrand Stivalet, Joe Chalfoun, Mary Brady
National Institute of Standards and Technology
walid.keyrouz@nist.gov, bertrand,stivalet@nist.gov, joe.chalfoun@nist.gov, mary.brady@nist.gov

Timothy Blattner, Shujia Zhou
Computer Science
University of Maryland Baltimore County
tblatt1@umbc.edu, szhou@umbc.edu

PP1

Non-Asymptotic Confidence Regions for Model Parameters of a Hybrid Dynamical System

In this paper, a non-asymptotic method is introduced for evaluating the uncertainties of model parameters of a class of hybrid dynamical systems. The new algorithm for this purpose is based on the Leave-out Sign-dominant Correlation Regions (LSCR) algorithm. This method has been known for generating non-conservative confidence regions of model parameters in system identification using only a finite number of data points. In order to extend the range of system types to which the LSCR method can be applied, a simple hybrid dynamical system with uncertain parameters is considered, and the effectiveness of the LSCR method is demonstrated using a simulated finite number of data points obtained from the system.

Sangho Ko
School of Aerospace and Mechanical Engineering
Korea Aerospace University
sanghoko@kau.ac.kr

PP1

Mathematical Analysis of Three-Dimensional Open Water Maneuverability by Mantas $MantaBirostris$

For aquatic animals, turning maneuvers may not be confined to a single coordinate plane, making analysis difficult particularly in the field. To measure turning performance for the manta ray, a large open-water swimmer, scaled stereo video recordings (30 fr/s) were collected around Yap, Micronesia. Movements of the cephalic lobes, eye and tail base were tracked to obtain three-dimensional coordinates. A mathematical analysis was performed on the coordinate data to calculate the turning rate and curvature (1/turning radius) by numerically estimating the derivative. Tikhonov regularization was used involving the minimization of a cost function with a parameter controlling the balance between data fidelity and regularity of the derivative. The approach gives a more accurate estimate of the derivative than conventional finite-difference schemes which can amplify noise leading to erroneous results. Data for 30 sequences of rays performing slow, steady turns showed a median turning rate of 46.48 deg/s with a median turning radius of 0.39 body lengths. Such turning maneuvers fall within the range of performance exhibited by swimmers with rigid bodies.

Allison Kolpas
Department of Mathematical Sciences
University of Delaware
akolpas@wcupa.edu

Frank Fish
Department of Biology
West Chester University of PA
ffish@wcupa.edu

Alex Meade
Department of Mathematics
West Chester University of PA
am688054@wcupa.edu

Michael Dudas
Dudas' Diving Duds
West Chester, PA
mail@dudasdiving.com

Keith Moored
Mechanical Engineering
Princeton University
kmoored@princeton.edu

PP1

Lqr Optimal Control for a Thermal-Fluid Dynamics

In recent years, considerable attention has been devoted to energy efficient buildings and controlling the thermal-fluid dynamics therein. In this work, we investigate the feasibility of POD for the Linear Quadratic Regulator Problem (LQR) of a coupled PDE. However, since the dynamics of a system (e.g. viscosity of air) is subject to change during a given time period, we are particularly interested in the sensitivity of the model to a change in the underlying Reynolds number.

Boris Kraemer
Virginia Tech
bokr@vt.edu

John A. Burns
Virginia Tech
Interdisciplinary Center for Applied Mathematics
jaburns@vt.edu

PP1

Matrix Interpolation Reduced Order Modeling For Microstructure Design

Matrix Interpolation Reduced Order Modeling (MIROM) uses radial basis functions to interpolate reduced matrices, rather than computing them directly as in proper orthogo-

nal decomposition. The resulting interpolation allows for a computationally cheaper reduced order model evaluation. This work will present results from applying MIROM to a microstructure design problem, where the design parameters are the microstructure dimensions and the outputs of interest are the coefficient of thermal expansion and the Young's modulus.

Kyle Lange, Dan White, Mark L. Stowell
Lawrence Livermore National Laboratory
lange9@llnl.gov, white37@llnl.gov, stowell1@llnl.gov

PP1

Large-Scale Stochastic Linear Inversion Using Hierarchical Matrices

Large scale inverse problems, which frequently arise in earth-science applications, involve estimating unknowns from sparse data. The goal is to evaluate the best estimate, quantify the uncertainty in the solution, and obtain conditional realizations, which are computationally intractable using conventional methods. In this talk, I will discuss the hierarchical matrix approach optimized for a realistic large scale stochastic inverse problem arising from a cross-well seismic tomography application.

Judith Yue Li
Stanford University
yuel@stanford.edu

Sivaram Ambikasaran
Institute for Computational and Mathematical Engineering
Stanford University
sivaambi@stanford.edu

Peter K Kitanidis
Stanford University
peterk@stanford.edu

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

PP1

Crosslinks: Connecting Topics Across Mathematics and Engineering Curricula

Crosslinks is a tool to discover and track the relationships between core concepts in mathematics and the engineering topics that follow. Its primary purpose is to improve knowledge transfer through use by students and faculty alike. The greatest learning moments occur when students discover connections between topics that once felt distinct; Crosslinks is intended to elicit and record those discoveries.

Chad E. Lieberman
MIT
celieber@mit.edu

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

Haynes Miller
MIT
hrm@math.mit.edu

PP1

Mathematical Modeling with Sensor Data

We present our work on the development of mathematical models used in conjunction with real-time sensor data to simulate certain physical processes, for example, air flow and heat transfer. In particular, we consider the use of both physics-based models (in the form of boundary value problems) and data-driven statistical models within a data center energy management system and present results from simulations conducted as part of case studies.

Vanessa Lopez-Marrero, Hendrik Hamann
IBM T. J. Watson Research Center
lopezva@us.ibm.com, hendrikh@us.ibm.com

Huijing Jiang
IBM T.J. Watson Research Center
huijiang@us.ibm.com

Xinwei Deng
Virginia Tech
xdeng@vt.edu

PP1

Bovine Lameness Detection Via Multidimensional Time-Series Force Data

Bovine lameness is a serious humanitarian concern and a costly problem for the US dairy industry. A mechanical lameness detection system developed at UMBC has been continuously collecting three dimensional force data from a herd of 750 dairy cows since 2012. To detect lameness, we reduce the raw data to a suitable heuristic representation and then train a hierarchical logistic model on scores from a veterinary examination. Our method achieves sensitivity and specificity above 85 percent.

Jonathan S. McHenry
Department of Mathematics and Statistics
UMBC
jon4@umbc.edu

Nagaraj Neerchal
Department of Mathematics and Statistics
University of Maryland, Baltimore County
nagaraj@umbc.edu

Uri Tasch, Jason Dunthorn
University of Maryland, Baltimore County
tasch@umbc.edu, jasond2@umbc.edu

Robert Dyer
University of Delaware
rdyer@udel.edu

PP1

Matrix Kronecker Products for Easy Numerical Implementation of PDEs and BCs in 2D

We illustrate the use of matrix Kronecker products to formulate 2D PDEs involving the laplacian or biharmonic operators on a rectangular domain with arbitrary boundary conditions. For Poisson-type equations, an iterative multigrid method is developed to solve the resulting Sylvester

system directly. For biharmonic equations that require two boundary conditions on each edge, we explain how to use matrix Kronecker products to implement the boundary conditions easily in the matrix formulation.

Ruben Glueck, Michael Franklin
Claremont Graduate University
ruben.glueck@cgu.edu, mbfranklin@gmail.com

Ali Nadim
Claremont Graduate University
Department of Mathematics
ali.nadim@cgu.edu

PP1

Random Ordinary Differential Equations for Multi-Storey Buildings

We investigate the numerical properties of modern methods for the simulation of random differential equations and their performance in non-tivial applied settings. In particular, the averaged Euler, the averaged Heun, and a K-RODE-Taylor scheme are used for ground-motion-induced excitation of multi-storey buildings subject to the Kanai-Tajimi earthquake model. The generalisation of our approach to random partial differential equations as well as an efficient vectorised implementation on different new CPU architectures are subject to current work.

Tobias Neckel, Alfredo Parra, Florian Rupp
Technische Universität München
neckel@in.tum.de, alfredo.parra@tum.de, rupp@ma.tum.de

PP1

Boundary Integral Equations for Waves in Linearly Graded Media

Boundary integral equations (BIEs) are a popular method for numerical solution of Helmholtz boundary value problems in a piecewise-uniform medium. We generalize BIEs for the first time to a continuously-graded medium in two dimensions, where the square of the wavenumber varies linearly with one coordinate. Applications include optics and acoustics in thermal/density/index gradients. We use contour integrals to rapidly evaluate the new fundamental solution, and give examples of interior and high-frequency scattering problems.

Brad Nelson, Alexander H. Barnett, Matt Mahoney
Dartmouth College
bjn@dartmouth.edu, ahb@math.dartmouth.edu, john.m.mahoney@dartmouth.edu

PP1

Multiphysical Coupled Problems for Vehicle Simulation

For engineering application in controller development for vehicles it is advantageous to be able to simulate efficiently coupled problems arising in a vehicle efficiently. Those problems are typically from multiphysical domains. We show problems, formulation and calculations for coupling of the mechanical system with 1D gasdynamic of the engine, cooling circuits, hybrid components and other domains needed for vehicle simulation.

Ralf U. Pfau
IMCC
MathConsult GmbH
ralf-uwe.pfau@avl.com

Roman Heinzle
IMCC
MathConsultGmbH
roman.heinzle@mathconsult.co.at

PP1

A Library of Tensors with Order-Oblivious Indexing: Fast Prototyping and Vectorization of Interpreted Scientific Codes

Interpreted languages with multi-array and vectorization support are widely used to fast prototype scientific codes. However, vectorized codes are difficult to write, debug, maintain, and optimize. First, we use labeled multi-array indices to help write, debug, and maintain scientific codes. Moreover, index positions are dynamic (order-oblivious) and can be optimized a posteriori to improve performance. Finally, we incorporate tensor products and contractions to enhance the applicability of the presented MATLAB and Python library.

Francisco J. Roca, David Moro, Ngoc Cuong Nguyen, Jaime Peraire
Massachusetts Institute of Technology
xeviroca@mit.edu, dmoro@mit.edu, cuongng@mit.edu, peraire@MIT.EDU

PP1

Efficient Solution of the Optimization Problem in Model-Reduced Gradient-Based History Matching

We present preliminary results of a performance evaluation study of several gradient-based state-of-the-art optimization methods for solving the nonlinear minimization problem arising in model-reduced gradient-based history matching. The issues discussed in this work also apply to other problems, such as production optimization in closed-loop reservoir management.

Marielba Rojas, Slawomir Szklarz
Delft University of Technology
marielba.rojas@tudelft.nl, s.p.szklarz@tudelft.nl

Malgorzata Kaleta
Shell Global Solutions International
malgorzata.kaleta@shell.com

PP1

Adaptive Simulation of Global Mantle Flows

I present adaptive discretization methods and efficient parallel solvers for the nonlinear Stokes systems arising in global mantle flow. The nonlinear Stokes equations are discretized using high-order elements that are discontinuous for the pressure. The block preconditioner for the Krylov iterations employs a BFBT approximation of the pressure Schur complement and an algebraic multigrid approximation of the viscous operator. Local mesh refinement enables resolution of localized features while keeping the size of the resulting algebraic systems amenable for solution on contemporary supercomputers.

Johann Rudi
ICES
The University of Texas at Austin

johann@ices.utexas.edu

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

Michael Gurnis
Caltech
gurnis@caltech.edu

Toby Isaac
ICES
The University of Texas at Austin
tisaac@ices.utexas.edu

Georg Stadler
University of Texas at Austin
georgst@ices.utexas.edu

Hari Sundar
Institute for Computational and Engineering Sciences
University of Texas at Austin
hari@ices.utexas.edu

PP1

A Reduced Basis Method for the Design of Metamaterials Through Optimization

Wave propagation through heterogeneous materials is often unintuitive and leads to unique phenomena. The design of such materials is therefore of big interest. The state-of-the-art numerical simulation tools are used for the simulation. Hybrid DG methods provide accurate and efficient solutions but design is only possible if NP-hard binary-optimization problems are formulated. A reduced-basis method with heuristic discrete programming techniques is introduced. Applications of this method for cloaking and other wave problems will be presented.

Joel Saa-Seoane
M.I.T.
jsaa@mit.edu

Cuong Ngoc Nguyen, Abby Men
MIT
cuongng@mit.edu, abby.men@gmail.com

Robert M. Freund
MIT
Sloan School of Management
rfreund@mit.edu

Jaume Peraire
MIT
peraire@mit.edu

PP1

High-Performance Computing in Simulating Carbon Dioxide Geologic Sequestration

In Carbon Sequestration, codes that simulate water-rock interaction and reactive transport sequentially solve mass balance equations for each control volume representing a lithology containing charged aqueous solutes. This is not well suited for execution on many-core computers. We present the theory and implementation of a numerical scheme whereby solute concentrations in all control volumes are solved simultaneously by constructing a large block-banded matrix. These matrices are factored with `SuperLU_DIST` (Berkeley Laboratory). Performance metrics are evaluated.

Eduardo Sanchez
San Diego State University
ejspeiro@gmail.com

PP1

Accelerator-Enabled Distributed-Memory Implementation of the Icon Dynamical Core

We present first results from the accelerator-enabled distributed-memory Icosahedral Non-hydrostatic (ICON) dynamical core, which will be part of the ICON climate model currently under development at the Max Planck Institute for Meteorology. The accelerator implementation utilizes the evolving OpenACC standard for directives augmenting existing Fortran compilers, such as Cray and PGI. An initial single-node prototype implementation halved time-to-solution on NVIDIA M2090 with respect to Intel Sandybridge for high resolution runs.

William Sawyer
Swiss Centre of Scientific Computing (CSCS)
Manno, Switzerland
wsawyer@cscs.ch

PP1

Blocking Symmetric Tensors

We introduce a new method of storing symmetric tensors (m-dimensional arrays) based on blocking in linear algebra. With this blocked storage scheme, we devise a blocked algorithm for the ttsm (tensor times same matrix) in all modes operation based on insights gained from linear algebra. The $C = BAB^T$ operation is the matrix-equivalent of the ttsm in all modes operation and is used as the basis for our insights.

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

PP1

Multi-Scale Atmospheric Chemical Transport Modeling with Wavelet-Based Adaptive Mesh Refinement (WAMR) Numerical Method

Application of non-adaptive numerical techniques for modeling of multi-scale atmospheric chemical transport often results in significant numerical errors due to poor spatial and temporal resolution. Here we present an adaptive multi-scale WAMR method applied to numerical simulation of transpacific pollution plume transport. It is shown that multilevel numerical grid efficiently adapts to the numerical solution development and the algorithm accurately

reproduces the plume dynamics at a reasonable computational cost unlike conventional non-adaptive numerical methods.

Yevgenii Rastigejev
Harvard University
ye_rast@yahoo.com

Artem N. Semakin
North Carolina A&T State University
arte-semaki@yandex.ru

PP1

Fluctuating Lipid Bilayer Membranes with Diffusing Protein Inclusions: Hybrid Continuum-Particle Numerical Methods

Many proteins through their geometry and specific interactions with lipids induce changes in local membrane material properties. To study such phenomena we introduce a new hybrid continuum-particle description for the membrane-protein system that incorporates protein interactions, hydrodynamic coupling, and thermal fluctuations. We investigate how collective protein effects influence membrane mechanical properties. We discuss interesting numerical aspects that are required to obtain good translation invariance. Finally, we discuss a coarse grained model that incorporates important hydrodynamics.

Jon Karl Sigurdsson
Department of Mathematics
University of California, Santa Barbara
jonksig@gmail.com

Paul J. Atzberger
University of California-Santa Barbara
atzberg@math.ucsb.edu

PP1

Preconditioning Techniques for Stochastic Conservation Laws

We derive a preconditioning technique for a class of stochastic conservation laws. The transformation, based on a space-time stretching, either pushes the solution field into a low rank manifold or expresses it as a low rank perturbation of a reference solution. Both cases allow efficient time integration through reduced basis methods or the generalized spectral decomposition, together with a substantial reduction in storage and computational burdens. The technique is particularly suited to long time integration problems where polynomial chaos approaches typically fail. An extension to a more general class of SPDEs is also provided.

Alessio Spantini
Massachusetts Institute of Technology
spantini@mit.edu

Lionel Mathelin
LIMSI - CNRS
mathelin@limsi.fr

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

PP1

Balanced Splitting Methods for Multidimensional Systems

Splitting methods are frequently used when solving large ODE systems such as those arising in reacting flow simulations. For some systems, standard splitting methods can introduce large steady-state errors. We introduce a new method, balanced splitting, which eliminates the steady state error by adding a balancing constant to each of the split terms. We analyze the method's stability and present examples based on typical reacting flow problems.

Raymond L. Speth
MIT
speth@mit.edu

William H. Green
MIT Chemical Engineering
whgreen@mit.edu

PP1

Referenceless Magnetic Resonance Temperature Imaging Approaches

Online temperature monitoring is mandatory for safe and successful thermal therapy treatments of patients. Two referenceless magnetic resonance temperature imaging approaches are compared: a l1-minimization edge detection approach and a PDE solution approach. Both methods have high computational costs due to real time requirements. The robustness of the PDE based approach is validated using polynomial chaos (DAKOTA framework) with additional high computational costs. Distributed computing and GPU implementations are employed to find solutions efficiently.

Wolfgang Stefan, Florian Maier
The University of Texas MD Anderson
wstefan@mdanderson.org, fmaier@mdanderson.org

Jason Stafford, John Hazle
MD Anderson Cancer Center
jstafford@mdanderson.org, jhazle@mdanderson.org

PP1

An Adaptive Simplex Cut-cell Method for High-order Discontinuous Galerkin Discretizations of Elliptic Interface Problems

We present a new approach for high-order discretizations of elliptic interface problems on unfitted meshes. The approach consists of a discontinuous Galerkin (DG) discretization and a simplex cut-cell technique. We show that no modification on DG bilinear form is needed for interface treatment. For irregularly shaped interfaces, we combine our strategy with an adaptive scheme to control the effect of geometry-induced singularities. High-order convergence is demonstrated for elliptic interface problems with regular and irregular interfaces.

Huafei Sun
MIT
sunhu@mit.edu

PP1

Stochastic Eulerian-Lagrangian Method with Thermal Fluctuations for Fluid-Structure Interactions

with Strong Coupling

We model a fluid-structure system subject to thermal fluctuations involving both Eulerian and Lagrangian reference frames. Although this description arises rather naturally, it presents both analytic and computational challenges. We treat the central issue of coupling between the two frames. Specifically, we simplify the viscous coupling between the immersed structures and the fluid in the regime of strong coupling by utilizing an asymptotic reduction on the infinitesimal generator of the stochastic differential equations.

Gil J. Tabak, Paul Atzberger
University of California at Santa Barbara
gtabak@umail.ucsb.edu, atzberger@math.ucsb.edu

PP1

Sport - An Effective Algorithm Towards Improving Bayesian Network Structure Learning

The Bayesian Network (BN) is a very powerful tool for causal relationship modeling and probabilistic reasoning. It facilitates deep understanding of very complex, high-dimensional problem domains. A key process of the building the BN is identifying its structure – a directed acyclic graph (DAG). In the literature, researchers proposed over fifty algorithms to discover the BN structure from data. The recent emerging of many BN-Structure learning algorithms (BN-SLAs) calls for generic techniques capable of improving BN structures learned by different algorithms. However, currently, these types of generic improvement techniques are lacking. This study proposes a novel three-phased algorithm called SPORT (Score-based Partial Order RefinemenT). Through three phases: Pruning, Thickening and Correcting, SPORT leverages Bayesian score function to iteratively simplify and improve BN structures. It is generic and pluggable to different BN-SLAs with marginal computational overhead. Empirical study applies SPORT to the BN structures learned by four major BN-SLAs: PC, TPDA, OR and MMHC. Improvements up to 64.4 % are observed among all the learned structures, confirming the effectiveness of SPORT towards improving BN structural learning.

Yan Tang
Hohai University
tangyan1981@gmail.com

Kendra Cooper
University of Texas at Dallas
kcooper@utdallas.edu

DeShan Tang
Hohai University
tds808@163.com

PP1

Simplifying Chemical Kinetic Systems under Uncertainty using Markov Chains

We propose a new approach to the simplification of chemical kinetic systems, particularly suited to systems with large uncertainties. The method measures the probability that elimination of a chemical species will influence the microstates of the system. We use a Markov process to model the transfer of atoms from one molecule to another via elementary reactions, and show that the absorption properties of the Markov chain identify chemical species to be eliminated. Preliminary results will be shown for the combustion of methane.

Luca Tosatto, Youssef M. Marzouk
Massachusetts Institute of Technology
ltosatto@mit.edu, ymarz@mit.edu

PP1

A Parallelized Model Reduction Library: Modred

We present modred, a Python library for model reduction, modal analysis, and system identification of large systems and datasets. It is parallelized for distributed-memory computing and has a comprehensive suite of automated tests. Its modular design allows it to interface with arbitrary data formats. The algorithms implemented include the Proper Orthogonal Decomposition (POD), Balanced POD (BPOD), Dynamic Mode Decomposition (DMD), Eigensystem Realization Algorithm (ERA), and Observer/Kalman Filter Identification (OKID).

Jonathan Tu, Brandt Belson
Princeton University
jhtu@princeton.edu, bbelson@princeton.edu

Clarence Rowley
Princeton University
Department of Mechanical and Aerospace Engineering
cwrowley@princeton.edu

PP1

An Estimation Theory Approach to Decision Under Uncertainty with Application to Wind Farm Siting

In this poster we present a methodology for the quantification and systematic reduction of risk in the design and development of complex systems. Our methodology exploits Bayesian estimation theory to track the evolution of probability distributions of critical parameters that characterize the risk in the development process, as well as sensitivity analysis techniques to reduce these uncertainties via efficient management of resources. The development of the methodology is demonstrated on the decision of whether or not to site a wind farm.

Fatma D. Ulker, Douglas L. Allaire, John J. Deyst, Karen E. Willcox
Massachusetts Institute of Technology
fdulker@mit.edu, dallaire@mit.edu, deyst@mit.edu, kwillcox@MIT.EDU

PP1

Updating Singular Subspaces for Latent Semantic Indexing

Latent Semantic Indexing (LSI) is a popular technique for intelligent information retrieval. Computationally, the approach is based on finding dominant singular vectors of term-document matrices, which are used to represent data in a lower-dimensional space. In practice, since the data are not static, the term-document matrices are frequently updated. The state-of-the-art algorithm for the corresponding updates of singular vectors is due to Zha and Simon [SIAM J. Sci. Comput., 21(2):782-791, 1999]. In this work, we propose a different updating approach that is based on the Rayleigh-Ritz method. The two updating schemes are compared for a few standard document collections.

Eugene Vecharynski
University of Colorado Denver
eugene.vecharynski@gmail.com

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

PP1

Application of Automatic Model Order Reduction to Electromagnetic Interactions

We are interested in electromagnetic coupling problems, an example problem is computing induced currents on printed circuit board traces given an incident electromagnetic wave. The geometry is parameterized, and we wish to solve for the induced currents for all values of the parameters. We use the frequency domain boundary element method. A new Model Order Reduction technique is developed; this technique combines hierarchical decomposition of parameter space, reduced basis method, and radial basis function interpolation.

Daniel White, Kyle Lange
Lawrence Livermore National Laboratory
white37@llnl.gov, lange9@llnl.gov

Matt Stephanson
Ohio State University
stephanson.2@buckeyemail.osu.edu

PP1

Inversion of Rheological Parameters of Mantle Flow Models from Observed Plate Motions

Modeling the dynamics of the Earth's mantle is critical for understanding the dynamics of the solid earth. Yet, there remain large uncertainties in the constitutive parameters employed within mantle convection models. Here we formulate an inverse problem to infer the rheological parameters that minimize the misfit between observed and modeled tangential surface velocity fields. The inverse problem is solved using a parallel scalable implementation of an adjoint-based quasi-Newton method.

Jennifer A. Worthen, Georg Stadler
University of Texas at Austin
jworthen@ices.utexas.edu, georgst@ices.utexas.edu

Noemi Petra
Institute for Computational Engineering and Sciences (ICES)
The University of Texas at Austin
noemi@ices.utexas.edu

Michael Gurnis
California Institute of Technology
gurnis@gps.caltech.edu

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu

PP1

Hardware-Aware Optimizations for Using Exafmm As a Preconditioner

The hardware landscape is changing and machine balance is converging, so that it is timely to revisit algorithms to adapt to this situation. In regards to the fast multipole method, there is growing interest in its application as a preconditioner. These two factors are an opportunity to implement optimizations that can make FMM competitive with mainstream preconditioners like multigrid. We will show evidence of this statement, using the exaFMM code framework.

Rio Yokota
King Abdullah University of Science and Technology
rio.yokota@kaust.edu.sa

Simon Layton
Boston University
slayton@bu.edu

Lorena A. Barba
Department of Mechanical Engineering
Boston University
labarba@bu.edu

PP1

A Method of Calculating Stress Intensity Factors at The Edges of a Crack Located Near a Welding Seam

We consider a crack located in the influence zone of a welding seam and parallel to it. We compute the stress intensity factors at the edges of the crack and present a method of calculating these factors, in which the seam is regarded as a periodic system of collinear cracks. We also showed that, with a high degree of accuracy, we can view a periodical system of cracks as a chain of three cracks.

Olga Zaydenvarg
National Aerospace University "Kharkiv Aviation Institute"
Department of Higher Mathematics
olga_z@inbox.ru

PP1

Population Size Effects in Genetic Algorithms for Auto Tuning

Scientific applications often rely on linear algebra computations. We describe the Build To Order (BTO) compiler which automates the optimization of those computations. BTO takes in a high level linear algebra specification, searches for the optimal combination of loop fusion, and shared memory parallelism, and outputs it in C. In this poster, we present the effect of population size on the overall performance of the genetic search algorithm, one of the search strategies in BTO.

Xing Jie Zhong
University of Colorado Boulder
xingjie.zhong@colorado.edu

Thomas Nelson
University of Colorado at Boulder
Argonne National Laboratory
thomas.nelson@colorado.edu

Elizabeth Jessup
University of Colorado Boulder
elizabeth.jessup@colorado.edu

Jeremy Siek
Department of Electrical and Computer Engineering
University of Colorado at Boulder
jeremy.siek@colorado.edu

PP1

Inverse Problems for Basal Boundary Conditions in A Thermomechanically Coupled Nonlinear Stokes Ice Sheet Model

Modeling the dynamics of polar ice sheets is critical for projection of future sea level rise. Yet, there remain large uncertainties in the boundary conditions at the base of the ice sheet. Here we study mathematical and computational issues in inversion of basal sliding coefficient and geothermal heat flux in a thermomechanically coupled nonlinear Stokes model using observations of surface flow velocities. We employ adjoint-based inexact Newton methods to solve the inverse problem.

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

Tobin Isaac
University of Texas at Austin
tisaac@ices.utexas.edu

Noemi Petra
Institute for Computational Engineering and Sciences (ICES)
The University of Texas at Austin
noemi@ices.utexas.edu

Georg Stadler
University of Texas at Austin
georgst@ices.utexas.edu

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

Omar Ghattas
University of Texas at Austin
omar@ices.utexas.edu