IP1

Extreme-Scale Computing: Accelerating Discovery and Innovation

With petascale systems now being applied to a wide variety of challenging problems, attention is turning to exascale computing capability and the promise of simulation-based science and engineering. The challenges of making the transition to exascale are formidable, but the expected benefits, including opportunities to accelerate scientific discovery and enable new technologies, provide a powerful argument for a substantial investment in delivering new capabilities that will enable high-fidelity simulations of real-world systems. * Managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract DE-AC05-00OR22725.

Thomas Zacharia
ORNL
zachariat@ornl.gov

IP2

Biomolecular Modeling and Simulation: A Field Coming of Age

I will present a recent assessment of the progress in biomolecular modeling and simulation, focusing on structure prediction and dynamics, by presenting the field's history, metrics for its rise in popularity, early expressed expectations, and current significant applications. Despite early unrealistic expectations and the realization that computer technology alone will not quickly bridge the gap between experimental and theoretical time frames, increases in computational power and improvements in algorithms and force fields are propelling the field onto a productive trajectory to become full partner with experiment. Research research on RNA and chromatin folding will also be presented.

Tamar Schlick
Howard Hughes Medical Inst
Department of Mathematics and Chemistry
schlick@nyu.edu

IP3

Kinetic Methods for CFD

Computational fluid dynamics is based on direct discretizations of the Navier-Stokes equations. The traditional approach of CFD is now being challenged as new multi-scale and multi-physics problems have begun to emerge in many fields – in nanoscale systems, the scale separation assumption does not hold; macroscopic theory is therefore inadequate, yet microscopic theory may be impractical because it requires computational capabilities far beyond our present reach. Methods based on mesoscopic theories, which connect the microscopic and macroscopic descriptions of the dynamics, provide a promising approach. We will present two mesoscopic methods: the lattice Boltzmann equation and the gas-kinetic scheme, and their applications to simulate various complex flows.

Li-Shi Luo
Old Dominion University
lluo@odu.edu

IP4

Application Performance in the Multi-Core Era Lessons to Be Learned for Exascale Computing

Tremendous efforts are put into hardware and software technologies to enable Exaflop computing before the end of this decade. Technological limitations will enforce architectural disruptions in the design of this new class of supercomputers. Still most of these changes are not clearly identified. However, they will challenge the sustainability of most existing software infrastructures and programming principles used today in computational science and engineering. Focusing on technological trends already implemented at moderate scale in modern parallel computers and clusters, we evaluate their impact on programming and use of parallel computers for scientific applications. Recognizing some long known golden performance rules and adapting them to the latest widespread compute resources is the basis for getting good performance today. This will also be a prerequisite for coping with the same problems at extreme scale in Exaflop computers.

Gerhard Wellein
Erlangen Regional Computing Center (RRZE), Germany
Gerhard.Wellein@rrze.uni-erlangen.de

IP5

Reverse Engineering The Face

Creating animated computer generated faces which can withstand scrutiny on the large screen is a daunting task. How does the face move? How does it reflect light? What information is relevant? How can it be captured and then transformed to convincingly breath life into a digital human or fantastic creature? The talk will give examples of new technologies and methodologies developed to achieve this in blockbuster films including "Avatar" and will point the way to the next generation of computer generated characters by showing the increasing importance of computational simulation and discovering what is really going on underneath the skin.

Mark Sagar
Weta Digital
marks@wetafx.co.nz

IP6

Discovery of Patterns in Global Earth Science Data using Data Mining

The climate and earth sciences have recently undergone a rapid transformation from a data-poor to a data-rich environment. In particular, climate and ecosystem related observations from remote sensors on satellites, as well as outputs of climate or earth system models from large-scale computational platforms, provide terabytes of temporal, spatial and spatio-temporal data. These massive and information-rich datasets offer huge potential for understanding and predicting the behavior of the Earth's ecosystem and for advancing the science of climate change. This talk will present our recent research in discovering interesting relationships among climate variables from various parts of the Earth.

Vipin Kumar
University of Minnesota
kumar@cs.umn.edu

IP7

Interoperable Unstructured Mesh Technologies for

Petascale Computations

The advent of petascale computing will enable increasingly complex, realistic simulations of PDE-based applications. Numerous software tools are used to help manage the complexity of these simulations, including computer-aided design systems used to represent the geometry of the computational domain, advanced mesh generation tools to discretize those domains, solution adaptive methods to improve the accuracy and efficiency of simulation techniques, and parallel tools such as dynamic partitioning to ease implementation on today's computer architectures. However, managing the complexity of interactions between these services on massively parallel distributed memory computers is becoming increasingly difficult, leaving developers little time to focus on the science of their applications. The Interoperable Tools for Advanced Petascale Simulations (ITAPS) center focuses on providing tools to fill specific technology gaps, along with underlying interfaces providing interoperability between these tools. In this talk, I will describe the foundational concepts behind the ITAPS interoperability goals and show several examples of applications leveraging ITAPS technologies to increase simulation accuracy, operate more effectively on complex computational domains, or reduce the total time to solution.

Lori A. Diachin
Lawrence Livermore National Laboratory
diachin2@llnl.gov

IP8

Complexity Reduction in Partial Differential Equations

Mathematical models of complex physical problems can be based on heterogeneous partial differential equations (PDEs), i.e. on boundary-value problems of different kind in different subregions of the computational domain. In different circumstances, especially in control and optimization problems for parametrized PDEs, reduced order models such as the reduced basis method can be used to alleviate the computational complexity. After introducing some illustrative examples, in this presentation several solution algorithms will be proposed and a few representative applications to blood flow modeling, sports design, and the environment will be addressed.

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

CP1

The Discontinuous Galerkin Method for Hyperbolic Problems on Tetrahedral Meshes: A Posteriori Error Estimation

We construct simple, efficient and asymptotically correct a posteriori error estimates for discontinuous finite element solutions of three-dimensional scalar first-order hyperbolic partial differential problems on tetrahedral meshes. We explicitly write the basis functions for the error spaces corresponding to several finite element spaces ($\mathcal{P}_{\surd}, \mathcal{U}_{\surd}, \mathcal{V}_{\surd},$).

The leading term of the discretization error on each tetrahedron is estimated by solving a local problem. The a posteriori error estimates are tested on several linear problems to show their efficiency and accuracy under mesh refinement for smooth solutions.

Idir Mechai
Department of Mathematics, Virginia Tech
mechaii@vt.edu

Sliman Adjerid,
Department of Mathematics
Virginia Tech
adjerids@math.vt.edu

CP1

Assessment of Collocation and Galerkin Approaches to Linear Diffusion Equations with Random Data

We compare the performance of stochastic Galerkin and collocation methods for solving PDEs with random data. The Galerkin method requires solution of a single algebraic system that is dramatically larger than deterministic systems. The collocation method requires many deterministic solves, which facilitates use of existing software. However, the number of unknowns for collocation methods can be considerably larger than for Galerkin methods. Experimental results indicate that for stochastically linear problems, Galerkin methods are highly competitive.

Howard C. Elman, Christopher Miller
University of Maryland, College Park
elman@cs.umd.edu, cmiller@math.umd.edu

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

Raymond Tuminaro
Sandia National Laboratories
rstumin@sandia.gov

CP1

New Hermite Multiwavelet Based Finite Elements for Numerical Solution of Biharmonic Equation

We present a new family of Hermite multiwavelets with significantly improved quantitative properties and an incrementally solved, adaptive finite element method based on these wavelets. The local generalized semiorthogonal wavelets lead to asymptotically optimal conditioning, where the resulting multilevel finite element system is sparse and block diagonal, the system is solved in an incremental and adaptive manner. We finally demonstrate adaptive, incremental solution of a biharmonic equation in two dimensions.

Sarosh M. Quraishi
Department of Mechanical Engg., Institute of Technology
Banaras Hindu University
sarosh.quraishi@gmail.com

CP1

Solving a Nist Suite of PDE Benchmark Problems with Adaptive Low-Order Fem and Hp-Fem

This presentation is based on a recently published NIST Report NISTIR 7668 (2010) by William Mitchell that contains a collection of PDE benchmark problems. These problems feature a spectrum of phenomena that pose challenges to adaptive finite element algorithms. We begin with a brief review of the main differences in automatic adaptivity for low-order and higher-order FEM including hanging

nodes and spatially and polynomially anisotropic refinements. Then we use the benchmarks in the suite to discuss and compare various aspects of automatic adaptivity in low-order FEM and hp-FEM.

Erick A. Santiago, Frank Proa
University of Nevada, Reno
laviticus@sbcglobal.net, proa.fc@gmail.com

CP1

A Lagrangian Vortex Method for the Barotropic Vorticity Equation on a Rotating Sphere

We present a Lagrangian vortex method for the barotropic vorticity equation (BVE) on a rotating sphere. The solution of BVE involves solving a conservative transport equation for the vorticity fields and a Poisson equation for the stream function. The vortex method tracks the flow map and absolute vorticity using Lagrangian particles and panels. The velocity is computed from the Biot-Savart integral on the sphere. An adaptive refinement strategy is implemented to resolve small-scale features and a treecode is used for efficient computation. A fourth-order Runge-Kutta scheme is used for time integration. We start our investigation with point vortex method and the first test case is the Rossby-Haurwitz wave, which is the exact solution for BVE. Convergence study shows that the method is fourth order in time and first order in space for a uniform panel discretization of the sphere. Then we switch to vortex blob method for stability consideration. We also tested the evolution of vortex patch(s), which means the vorticity field is highly nonuniform on the surface of the sphere. Adaptive refinement strategy improves the computational efficiency.

Lei Wang
University of Michigan
olivewl@umich.edu

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

John P. Boyd
University of Michigan
Ann Arbor, MI 48109-2143 USA
jpboyd@umich.edu

CP1

Immersed Finite Element Methods for Moving Interface Problems

We will consider the parabolic equations with moving interface problems with discontinuous coefficients. The interface jump conditions are employed in the IFE functions, and the mesh need not to be aligned with the location of interface. So we can solve the problem on a fixed mesh even though the interface is moving. Several fully discretized algorithms based on the Crank-Nicolson idea will be presented. These algorithms will be compared both analytically and computationally. Numerical examples will show optimal convergence for these schemes.

Xu Zhang
Virginia Tech
xuz@vt.edu

Xiaoming He
Department of Mathematics and Statistics
Missouri University of Science and Technology
hex@mst.edu

Tao Lin
Department of Mathematics, Virginia Tech
tlin@math.vt.edu

CP2

A Discontinuous Immersed Finite Element Method for Interface Problems

We present piecewise quadratic immersed finite element (IFE) spaces that are used with an interior penalty (IP) method for solving two dimensional second order elliptic boundary value problems with discontinuous shape functions without requiring the mesh to be aligned with the material interfaces. Two different IFE spaces are developed, a discontinuous finite element formulation with interior penalty is discussed and numerical experiments are presented to show the optimality of our method.

Mohamed Ben Romdhane
Virginia Tech
mbenromd@vt.edu

Slimane Adjerid
Department of Mathematics
Virginia Tech
adjerids@math.vt.edu

Tao Lin
Department of Mathematics, Virginia Tech
tlin@math.vt.edu

CP2

A Conservative Meshless Framework with Error Bound Minimization

We present a finite-volume-like conservative meshless framework for solving conservation laws. The framework allows the incorporation of existing flux schemes and contains a central-like scheme as a special case. We discuss the necessary conditions for the existence of meshless coefficients and algorithms for generating coefficients that minimize an upper bound of the global approximation error. Numerical examples in fluid mechanics will also be presented to demonstrate the practicality of the framework.

Edmond Kwan-Yu Chiu
Ph.D. Candidate, Department of Aeronautics & Astronatutics
Stanford University
edmondc@stanford.edu

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

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

CP2

An Error Analysis for the Corner Singularity Ex-

pansion of a Compressible Stokes System on a Non-Convex Polygon

In this talk we consider a finite element scheme approximating the regular part and the stress intensity factor, show its unique existence of the discrete solution and derive an (nearly) optimal error estimate. Some numerical experiments confirming these results are given.

Young Pyo Kim
POSTECH
Dept. Mathematics
finance@postech.ac.kr

CP2

Performance of Hp-Adaptive Finite Element Methods

The hp version of the finite element method has been shown to have exponential rates of convergence with respect to the number of degrees of freedom, and consequently can be of optimal efficiency for many problems. Several strategies for the choice between h- and p-refinement have been proposed over the years. In this talk we present some numerical results to demonstrate the performance of these strategies on some test problems.

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

Marjorie McClain
Mathematical and Computational Sciences Division
National Institute of Standards and Technology
marjorie.mcclain@nist.gov

CP2

An Integral Equation Method for the Modified Biharmonic Equation

An integral equation formulation for the modified biharmonic equation in two dimensions is presented. This PDE arises from temporal discretizations of the stream function for the two dimensional incompressible Navier-Stokes equations. Numerical treatment, fast algorithms and preconditioning of the integral equation will be discussed followed by some simple examples.

Bryan D. Quaife
Simon Fraser University
Burnaby, British Columbia, Canada
bquaife@math.sfu.ca

CP2

Multigrid Solution of the Distributed Optimal Control of Semilinear Elliptic Pde.

In this work we design efficient multilevel methods for the distributed optimal control of semilinear elliptic PDEs. Specifically, the control problem is solved using Newton's method, and specially designed multigrid methods are applied to precondition the Hessian of the cost functional. It is shown that the quality of the resulting multigrid preconditioner increases at an optimal rate with respect to the mesh-size, as in the case of optimal control of linear elliptic PDEs.

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

CP3

The Reduced Basis Method for Incompressible Fluid Flow in Parametrized Domains

We present a reduced basis method for rapid and reliable simulation of incompressible fluid flow problems in parametrized domains. We compare different strategies for 1) ensuring the stability of the reduced basis approximation, and 2) the computation of rigorous a posteriori error bounds for both the velocity and the pressure approximations. We apply the method to the Stokes equations, and present results for a model problem relevant in the design of microfluidic devices.

Anna-Lena Gerner, Karen Veroy-Grepl
Graduate School AICES
RWTH Aachen University
gerner@aices.rwth-aachen.de, veroy@aices.rwth-aachen.de

David J. Knezevic
MIT
dknez@mit.edu

CP3

Modeling of Oscillatory Stimulation of Fluid Flows

Oscillatory stimulation is a promising method for increasing drainage of non-Newtonian fluids through porous structures in various applications. We developed a mathematical model for unsteady flow of a bi-viscous incompressible fluid in a circular straight channel. Longitudinal vibrations are superimposed on the flow driven by pressure gradient along the channel. Calculations have been carried out to observe the effect of oscillations of longitudinal pressure gradient on the mean flow rate of non-Newtonian fluids. It is found that oscillations enhance flow rate of a bi-viscous shear-thinning fluid while the effect of oscillations on the shear-thickening fluid is the opposite. The Bingham fluid exhibits even more significant augmentation of the mean flow rate when subjected to vibrations.

Sergey Lapin
Washington State University
Department of Mathematics
slapin@math.wsu.edu

Mohsen Asle Zaeem
Mississippi State University
mohsen@cavs.msstate.edu

Konstantin Matveev
Washington State University
School of Mechanical and Materials Engineering
matveev@wsu.edu

CP3

Recent Advances in Numerical Study of Maxwell's Equations in Metamaterials

Since 2000, there has been a growing interest in the study

of metamaterials across many disciplinaries. In this talk, I'll first introduce the metamaterials, and some of its interesting applications. Then I'll present the metamaterial modeling equations, and some time-domain finite element methods recently developed for solving the time-dependent metamaterial Maxwell's equations. Finally, some numerical results and open issues will be discussed.

Jichun Li
Department of Mathematical Sciences
University of Nevada, Las Vegas, NV 89154-4020
jichun@unlv.nevada.edu

CP3

Uncertainties Quantification and Data Compression in Numerical Aerodynamic

We research how uncertainties in the input parameters (the angle of attack and the Mach number) and in the airfoil geometry propagate in the solution (lift, drag, etc). We show that uncertainties in the Mach number and in the angle of attack weakly affect the lift coefficient ($1\% - 3\%$) and strongly affect the drag coefficient (around 14%). Uncertainties in the geometry influence both the lift and drag coefficients weakly. The RANS solver is TAU code with k-w turbulence model. For quantification of uncertainties we used KLE, PCE and low-rank matrices. The results obtained via collocation method on a sparse Gauss-Hermite grid match to Monte Carlo results, but require much less deterministic evaluations and as a sequence - smaller computing time.

Alexander Litvinenko
TU Braunschweig, Germany
litvinen@tu-bs.de

CP3

Experiences Extending the Cfd Solver of the Pde Framework Peano

The C++ PDE framework Peano [H.-J. Bungartz and M. Mehl and T. Neckel and T. Weinzierl, The PDE framework Peano applied to fluid dynamics: an efficient implementation of a parallel multiscale fluid dynamics solver on octree-like adaptive Cartesian grids, Computational Mechanics 46, pp. 103–114, 2010] has been designed for realising efficient applications on regular or adaptive Cartesian grids. In this contribution, we evaluate Peano's expandability by incorporating thermal energy transfer into the existing incompressible flow solver. A series of numerical benchmarks has been computed to validate the approach. These experiences are currently used to, first, further improve Peano's integration and reuse capabilities by automatic code generation etc. and, second, to achieve fully adaptive simulations of 3D thermal-hydraulic phenomena in the nuclear reactor coolant system.

Tobias Neckel, Michael Lieb, Ralf Sangl
Technische Universität München
neckel@in.tum.de, liebm@in.tum.de, ralf.sangl@mytum.de

Philipp J. Schoeffel, Fabian Weyermann
Gesellschaft fuer Anlagen- und Reaktorsicherheit (GRS) mbH
philipp.schoeffel@grs.de, fabian.weyermann@grs.de

CP4

On the Numerical Solution of Fuzzy Elliptic PDEs by Means of Polynomial Response Surfaces

We consider the solution of elliptic partial differential equations (PDE) with a fuzzy diffusion coefficient. Inspired by popular solution techniques for stochastic elliptic PDEs, we make use of orthogonal sets of polynomials and a Galerkin projection to construct a response surface. We derive an upper bound for the error of this Galerkin approximation in the maximum norm and use it to prove (near-)exponential convergence of this approximation to the exact solution in the fuzzy sense.

Samuel Corveleyn
Katholieke Universiteit Leuven
samuel.corveleyn@cs.kuleuven.be

Stefan Vandewalle
Scientific Computing Research Group
KULeuven
stefan.vandewalle@cs.kuleuven.be

CP4

An Etd Crank-Nicolson Method for Reaction-Diffusion Systems

A novel Exponential Time Differencing (ETD) Crank-Nicolson(CN) method is developed which is stable, second order convergent, and highly efficient as compared to standard CN method and BDF2. Stability and convergence are discussed for semi linear parabolic problems. Numerical experiments are presented for a wide variety of examples, including multi-dimensional chemotaxis and bio-chemical stiff reactions.

Abdul M. Khaliq
Middle Tennessee State University
Department of Mathematical Sciences
akhaliq@mtsu.edu

CP4

Explicit Local Time-Step Method for 3D Maxwell's Equations on Tetrahedral Mesh

An explicit local time-step (LTS) method is developed to solve the time-dependent Maxwell's equations on three-dimensional unstructured tetrahedral meshes. The method is based on the finite volume discretization in space and LTS second order total variation diminishing (TVD) Runge-Kutta scheme in time. The algorithm is an extension of the scheme proposed by Tang and Warnecke to solve conservation laws in one and two-dimensional space. It solves partial differential equations on a finite number of subdomains with LTSs using a simple projection of the solution at each step. In the present work, we remove the restriction for the time-steps to be an integer multiples of the minimum time-step which is present in most electromagnetics literature. We optimize local time steps distribution to further minimize computational time. The application to reference problems for three-dimensional Maxwell equations shows better performance of the proposed method.

Marina Kotovshchikova
University of Manitoba
umkotovs@cc.umanitoba.ca

Dmitry K. Firsov
Geomodeling Corp.
d.k.firsov@gmail.com

CP4

Accuracy Enhancement of Discontinuous Galerkin Solutions over Structured Triangular Meshes

Theoretically we can effectively increase the order of accuracy of a discontinuous Galerkin (DG) solution from order k+1 to 2k+1. However, this is a computationally complex task to execute in an efficient manner. This becomes an even greater issue when we consider non-quadrilateral mesh structures. We present post-processing of DG solutions for accuracy enhancement over structured triangular meshes. We demonstrate that we can obtain 2k+1 order accuracy in an efficient manner for linear hyperbolic equations.

Hanieh Mirzaee
Scientific Computing and Imaging Institute
School of Computing, University of Utah
mirzaee@cs.utah.edu

Jennifer K. Ryan
Delft University of Technology
j.k.ryan@tudelft.nl

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

CP4

Development and Analysis of Computer Simulations of Neuron-Astrocyte Networks to Better Understand the Role of Astrocytes

In the brain, astrocytes had been sidelined to a supporting role in neural networks. The past decade revealed that they influence and propagate neural signals, suggesting an important role for astrocytes in neural networks. By running multiple simulations of neural networks, with and without astrocytes, we observed that they have the ability to synchronize neural activity. Such findings resemble brain recordings of epileptic patients whose neuronal activity is synchronized during seizures correlated with hypertrophied astrocytes.

Tahirali H. Motiwala
Wofford College
motiwalath@email.wofford.edu

Anne Catlla
Wofford College
Dept. of Mathematics
catllaaj@wofford.edu

CP4

Investigations into the Boris Approach for Dealing with Fast Magnetosonic Wave Speeds in An Ale Mhd Modeling Context

The classical magnetohydrodynamic (MHD) approximation admits fast magnetosonic waves speeds which become unbounded as the material density vanishes. This can create wave speeds that are detrimental to an explicit time stepping algorithm in regions which have little relevance to the critical magneto-mechanical issues. Boris proposed a semi-relativistic approach in which the displacement current terms from Ampere's law are kept in the formulation of the momentum equation. This leaves open the possibility that for purposes of modeling efficiency one may choose to limit the maximum computational wave speed by adjusting the value of the speed of light. We describe algorithms of the Boris style in the context of an arbitrary Lagrangian-Eulerian MHD modeling approach and illustrate the effectiveness of these approaches with relevant examples. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Allen C. Robinson
Sandia National Laboratories
acrobin@sandia.gov

CP5

Accelerated Numerical Algorithms for Polyenergetic Digital Breast Tomosynthesis Reconstruction

Tomosynthesis imaging involves acquiring projection images over a limited angular range, which, after reconstruction, results in a pseudo-3D representation of the object. In breast cancer imaging, tomosynthesis is a viable alternative to standard mammography. In this talk, we discuss the mathematical methods for the reconstruction based on a polyenergetic model, focusing on algorithm implementation accelerators using the OpenCL and CUDA frameworks. We will show how vectorization and minimizing communication result in faster reconstruction time.

Veronica M. Bustamante
Emory University
vmejia@emory.edu

Julianne Chung
University of Maryland
jmchung@cs.umd.edu

Ioannis Sechopoulos
Department of Radiology and Winship Cancer Institute
Emory University
isechop@emory.edu

James G. Nagy
Emory University
Department of Math and Computer Science
nagy@mathcs.emory.edu

CP5

Effect Of Flow Topology In Biophysical Interactions

Lagrangian flow topology may affect dynamics of biological tracers, such as chlorophyll. We quantify such effects by simulating carrying capacity-phytoplankton-zooplankton interaction in an archetypical 2D flow characterized by eddies and filaments, representing the ocean surface. We find that the confined geometry of eddies initially help the phytoplankton population grow and filaments enhance carrying-capacity transport, yet do not enhance bioproduction. In both cases, the long time behavior of the phytoplankton population asymptotes.

Juan Durazo
Arizona State University
Juan.Durazo@asu.edu

CP5
Biophysical Models of Brain Cancer Invasion

Glioblastomas are brain cancers characterized by diffusive invasion of brain tissues and a high incidence of recurrence after surgery. The growth and spread of glioblastoma cell populations is frequently modeled as a reaction-diffusion process with spatially varying diffusion. In the present work, the anisotropy of tumor cell diffusion is modeled with consideration of biophysical interactions in the tumor microenvironment and brain fiber tract maps acquired from medical imaging. Results are compared with real patient histories.

John B. Ingraham
Arizona State University
School of Mathematical and Statistical Sciences
john.ingraham@asu.edu

CP5
Quantitative Analysis of Brain Tumor Images

Medical Resonance Images are an essential part of diagnosis and follow-up treatment for patients with brain tumors. Usually two types of MRIs are processed for quantification of tumor and edema volumes. For clinical assessment, T1+Contrast images are used to asses tumor size and T2 for edema. I will describe an approach to separating the edema from the tumor core by processing both sets of images simultaneously, improving the accuracy of edema volume.

Helene A. Rhodes
Arizona State University
School of Mathematical and Statistical Sciences
Helene.Rhodes@asu.edu

CP5
Radiotherapy in Modeling Brain Tumors

Due to how common radiotherapy regiments are in the treatment of glioblastoma and other brain tumors, it is necessary to have an accurate model to be able to simulate potential treatments. The key points of this presentation will be: 1. The different ways that radiotherapy can damage cells 2. The attenuation of a radiation dose across tissue 3. The potential effects of radiotherapy aside from cell death.

Keith A. Voytek
Student
kvoytek@asu.edu

CP6
Analysis of Quasicontinuum Methods for Unconstrained and Circular Chains

We present the results of a theoretical analysis of stability of unconstrained straight and circular chains of atoms held together by nearest and next-nearest neighbor interactions. We study the stability from the point of view of eigenvalues of the Hessian of the quasicontinuum (QC) energy functional, obtained by combining atomistic and continuum models. We present sharp stability estimates for the chain both under compression and expansion, as well as error estimates for the *quasi-nonlocal* QC approximation.

Pavel Belik
Mathematics Department
Augsburg College
belik@augsburg.edu

Mitchell Luskin
School of Mathematics
University of Minnesota
luskin@math.umn.edu

CP6
The Edge Theory

Does there show "assymptotic-uniqueness" along a network using PMA and control and can it be precluded from external sharing. A supposition for application and testing of a controlled system program for use in Polyelectrolyte Multilayer Assemblies. The method of differential analysis is interpolated into a system of equations and regressed onto an original IC schematic of a type serial-parallel. the numericism predicts and enhances material specific results that would be observed of the use of any PMA so far. the essay further describes possible theories that allow the control to have a different aspect of control on the input than does the rating on the load.

Cory J. Cheetham
diamondback Ltd
dch.cory@gmail.com

Mike Norton
colleague
xsideofparadise@yahoo.com

CP6
A Fracture Curve Generated by the Nearest Flaws

In this work, we study a two dimensional model for fragmentation in which the fracture curve is generated recursively by visiting the nearest defect. The following assumptions were also considered: the fracture forces are proportional to the size of the common boundary between two fragments; the material defects are represented by random point flaws; the total mass is conserved. Our main result establish that the visualizations present complex fracture patterns that resemble real systems.

Gonzalo J. Hernandez
University of Valparaiso, School of Industrial Enginnering
gjho@vtr.net

Roberto Leon
USM Department of Informatics
roberto.leon@usm.cl

CP6
An Open Curve Level-Set Based Algorithm for the Motion of Multiple Junctions

We propose an efficient open level-set method based on energy minimization, which captures the behavior of junctions. These junctions arise in the study of material boundaries, where the energy is made up of surface tension and bulk modulus. We solve the PDE using Sobolev gradients, avoiding the need for regularization or re-initialization while also accelerating convergence to the steady state configuration. The algorithm is fairly simple and easy to generalize to more complicated energies.

Rami Mohieddine

UCLA Math department
ramim@math.ucla.edu

Luminita A. Vese
University of California, Los Angeles
Department of Mathematics
lvese@math.ucla.edu

CP6

Two-Dimensional Lagrangian Solver for Elasto-Plastic Deformation of Solids

We present a two-dimensional, cell centered, Lagrangian approach for capturing the response of elasto-plastic deformation of solids under intense loading conditions. In this model, the primitive variables are stored and evolved at the cell centers. The nodal velocities and forces are required to update the mesh are computed in a consitent manner with the numerical fluxes. The hydrodynamic response of the material is modeled using the Mie-Gruneisen equation of state. The deviotaric stress components are evololved in accordance with hypo-elastic stress-strain relations and J2 Von-Mises yield conditions. We demonstrate the capability of the our approach by presenting preliminary one- and two-dimensional problems with and without material models.

Shiv Kumar Sambasivan, Shengnian Luo, Mikhail Shashkov
Los Alamos National Laboratory
shiv@lanl.gov, sluo@lanl.gov, shashkov@lanl.gov

Pierre-Henri Maire
UMR CELIA, Université Bordeaux I
maire@celia.u-bordeaux1.fr

CP6

Augmented Lagrangian-Based Preconditioners for Steady Incompressible Navier-Stokes Equations

We discuss augmented Lagrangian-based preconditioners for linear systems arising from stable finite elements discretization of the steady incompressible Navier-Stokes equations. Spectral properties of the preconditioned matrices are established, and the choice of the augmentation parameter using Fourier analysis is presented. Moreover, we use field-of-values analysis to prove that the rate of convergence of GMRES with this type of preconditioners is independent of the mesh size for suitable choices of the augmentation parameter. The dependence on viscosity will also be discussed.

Zhen Wang
Emory University
zwang26@emory.edu

Michele Benzi
DepartmentõfMathematics and Computer Science
Emory University
benzi@mathcs.emory.edu

Maxim Olshanskii
Department of Mechanics and Mathematics
Moscow State M.V.Lomonosov University, Moscow, Russia
Maxim.Olshanskii@mtu-net.ru

CP7

Modeling Brain Tumor Growth Using Matlab

I will discuss the implementation and visualization of a mathematical model of glioblastoma multiforme (GBM) brain tumor growth. Using Matlab, I have implemented a model developed by Steffen Eikenberry and co-workers to simulate treatment by surgery and radiation. With John Ingraham (another student at ASU working on brain tumor analysis), I have produced code to convert the Matlab simulation results into MRI-looking images that can be visually compared with the actual patient scans.

Eric Adams
Arizona State University, CSUMS
School of Mathematical and Statistical Sciences
e_to_x@yahoo.com

CP7

Implicit Solution of Free-Surface Flows in Glaciology

Ice sheets and glaciers are usually modeled as non-Newtonian Stokes flows with free surface and nonlinear slip on bumpy surfaces. Local conservation is critical for many applications, but most conservative discretizations generate spurious tangent forces which produces non-physical recirculation. We combine a well-balanced conservative discretization of slip with fully-implicit time integration to enable analysis of stability and long-term behavior.

Jed Brown
ETH Zurich
jed@59A2.org

CP7

Robust Renewable Resource Exploitation Strategy with Sustainability

In this paper, we introduce a dynamic Nash game among harvesting firms modeled as robust continuous time variational inequality for renewable resource allocation problems. For formulation, a continuous time fixed-point method is presented. We discuss in detail the application of the proposed framework to the fishery game among firms as well as among nations. We prove that our framework is effective in particular for computationally challenging renewable resource management problems. A numerical example is presented.

Sung Hoon Chung, Terry Friesz
The Pennsylvania State University
shchung@psu.edu, tlf13@psu.edu

CP7

Bayesian Markov Chain Monte Carlo Optimization of a Physiologically Based Pharmacokinetics Model of Nicotine

In an effort to reduce harm caused by tobacco products, the Family Smoking Prevention and Tobacco Act was enacted by Congress and the Food and Drug Administration (FDA) was tasked with regulating tobacco. The FDA is now charged to use the best available science to guide the development and implementation of effective public health strategies to reduce the burden of illness and death caused by tobacco products. Therefore, one significant challenge is to reduce toxicant exposure in tobacco products such as cigarettes. Nicotine is the active drug within tobacco that

is delivered during cigarette smoking. Smoking behavior is driven by the need for nicotine and will influence delivery of other tobacco toxicants into the lungs. To better understand nicotine dosimetry a physiologically-based pharmacokinetics (PBPK) model which captures the essential nicotine metabolism and takes into account the variability of kinetic parameters among individuals was developed.

Rudy Gunawan, Chuck Timchalk, Justin Teeguarden
Pacific Northwest National Laboratory
rudy.gunawan@pnl.gov, charles.timchalk@pnl.gov, jt@pnl.gov

CP7

A Multi-Scale Model of Tumor Growth

Mathematical and computational modeling holds great promise for medicine to predict outcomes such as tumor metastasis or drug response. Fibroblasts and myofibroblasts near the tumor microenvironment are important players in tumor growth and metastasis because of their unique ability to coordinate events which increase cell proliferation especially in breast cancer. It has been experimentally shown that fibroblasts play an important role in promoting tumor growth in vitro. A multi-scale model of this interaction between stroma and transformed epithelial cells near breast duct will be presented. EGF-TGFbeta pathway controls these interactions and our multi-scale model describes these phenomena at different time and spatial time scale, i.e., intracellular dynamics, cell dynamics at cellular level, and mechanical interaction between tumor cells and surrounding stromal tissue (continuum).

Yangjin Kim
University of Michigan, USA
yangjink@umd.umich.edu

Hans G. Othmer
University of Minnesota
Department of Mathematics
othmer@math.umn.edu

Magdalena Stolarska
University of St. Thomas
Department of Mathematics
mastolarska@stthomas.edu

CP7

Finite Difference Time Domain (FDTD) Analysis of Near-Infrared Optical Pulse Propagation in An Adult Head Model for Functional Imaging of Brain Activities

Finite difference time domain (FDTD) analysis is proposed for calculating near infrared optical pulse propagation in an adult head model composed of scattering and clear cerebrospinal fluid (CSF) layer. Integral form of the diffusion equations with additional sources due to radiated photons from an adjoining scattering layer are solved by the FDTD analysis utilizing a new boundary condition at the CSF interface. The numerical results are in excellent agreement with experiments and Monte Carlo simulations.

Tadatoshi Tanifuji
Department of Electrical and Electronics Engineering
Kitami Institute of Technology
tanifuji@mail.kitami-it.ac.jp

CP8

Numerical Solution of Control-State Constrained Optimal Control with An Inexact Smoothing Newton Method

This talk is concerned with a globalized inexact smoothing Newton method for the numerical solution of optimal control problems subject to mixed control-state constraints. The method uses the smoothed Fischer-Burmeister function to reformulate first order necessary conditions and aims at minimizing the squared residual norm using Newton steps and gradient-like steps. Numerical experiments are provided to illustrate the convergence results.

Jinhai Chen
University of Colorado Denver
jinhai.chen@ucdenver.edu

CP8

A Multilevel Monte Carlo for Simulating Extreme Quantiles and Probabilities

Let X denote a generic random vector with probability distribution μ on R^d, and let Φ be a mapping from R^d to R. That mapping can be a black box, e.g., the result from some computer experiments for which no analytical expression is available. Our goal is to estimate the probability $p = P[\Phi(X) > q]$ for any arbitrary real number q, or reversely, when p is fixed, to estimate the quantile q such that $P[\Phi(X) > q] = p$. Naive Monte Carlo estimation of that probability with a prescribed signal to noise ratio becomes computationally intractable for quantiles q lying far out in the right-hand tail of the distribution of the random variable $\Phi(X)$. In this talk we present and analyze a novel simulation algorithm for this problem. It proceeds by successive elementary steps, each one being based on Metropolis-Hastings algorithm. We demonstrate the practical usefulness of our method by applying it to a problem in watermarking.

Nicolas Hengartner
Information Sciences Group
Los Alamos National Laboratory
nick@lanl.gov

Arnaud Guyader
INRIA Rennes
University of Haute Bretagne
arnaud.guyader@uhb.fr

Eric Matzner-Lober
Department of Statistics
University of Haute Bretagne
eml@uhb.fr

CP8

Calculating Derivatives in Software for Designing Optical Fibers

Maxwells equation in one dimension may be used to model an optical fiber. When solving an inverse problem to determine the refractive index profile of the fiber, the computation of the gradient of the dispersion, involving derivatives of the eigenvalues and effective area, an integral of a function of the eigenvector, is more expensive than solving the underlying eigen problem. By interchanging the order of differentiation, taking advantage of the local support of the design parameters, and using separation of variables

we greatly decrease the computation time.

Linda Kaufman, Nick Lworonekin, William Sierchio, William Landon, Seokmin Bang, Ryan Petho, Daniel Savacool
William Paterson University
kaufmanl@wpunj.edu, woronekinn@student.wpunj.edu, sierchiow@student.wpunj.edu, landonw@student.wpunj.edu, bangs@student.wpunj.edu, rspetho@gmail.com, savacoold@student.wpunj.edu

CP8

Time-Optimal Bang-Bang Control of Burger's Equation in 1D

A numerical method for the approximate solution of a time-optimal control problem constrained by viscous burger's equation in one spatial dimension is presented. The control formulation requires a desired, but not fixed, final condition on the state solution. The novel algorithm converges based on iteratively solving the transversality condition for the final time. Numerical results for illustrative cases are given. The method can be adapted to hyperbolic systems of PDEs.

Nathan D. Moshman
Naval Postgraduate School
ndmoshma@nps.edu

CP8

Shape Optimization of Functions of Dirichlet-Laplacian Eigenvalues

We consider the shape optimization of functions of Dirichlet-Laplacian eigenvalues over the set of star-shaped, symmetric, bounded planar regions with smooth boundary. The regions are represented using Fourier-cosine coefficients and the optimization problem is solved numerically using a quasi-Newton method. The method is applied to generalizations of the Payne-Pólya-Weinberger ratio. Optimal values and attaining regions are presented and interpreted as a study of the range of Dirichlet-Laplacian eigenvalues.

Braxton Osting
Columbia University
bro2103@columbia.edu

CP8

Variational Multiscale Proper Orthogonal Decomposition: Convection-Dominated Convection-Diffusion Equations

We introduce a variational multiscale closure modeling strategy for the numerical stabilization of proper orthogonal decomposition reduced-order models of convection-dominated equations. As a first step, the new model is analyzed and tested for convection dominated convection-diffusion equations. The numerical analysis of the finite element discretization of the model is presented. Numerical tests show the increased numerical accuracy over a standard reduced-order model and illustrate the theoretical convergence rates.

Zhu Wang
Department of Mathematics
Virginia Tech
wangzhu@vt.edu

Traian Iliescu
Interdisciplinary Center for Applied Mathematics
iliescu@vt.edu

CP9

Edu @ Jsc

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

Johannes Grotendorst
Research Centre Juelich
j.grotendorst@fz-juelich.de

CP9

Sculpture, Geometry and Computer Science

We present a project for middle and high school students that blends mathematics and computer science with art appreciation. The goal is to model a challenging geometric sculpture, Indiana Arc. After some basic instruction in the Python scripting language, students use their math skills to generate 3-D geometric data and then visualize and analyze it using open source software. The project also pays homage to 'Geometry and the Imagination' (Hilbert and Cohn-Vossen).

Randy Heiland
Open Systems Lab
Indiana University
heiland@indiana.edu

Charles Perry
www.charlesperry.com
coperry@aol.com

Barbara Ream
International School of Columbus
barb.ream@gmail.com

Andrew Lumsdaine
Open Systems Laboratory
Indiana University
lums@cs.indiana.edu

CP9

Assessment Tools for the Computational Science Curriculum

We report on the ongoing development of a diagnostic assessment and a project evaluation rubric intended for use across the computational science curriculum. The diagnostic assessment has predicted success in our introductory computational science course with reasonable reliability. The project evaluation rubric has been used in both introductory and upper-level undergraduate courses and in graduate courses. Emphasis is given to the process of refining the project evaluation rubric through several semesters

of use.

Robert Olsen, Russell Manson
Richard Stockton College of New Jersey
Robert.Olsen@stockton.edu, russell.manson@stockton.edu

CP10

Performance of a Boundary Element Method Solver on General Purpose Graphics Processing Unit

The 3D Laplace boundary element code available from www.intetec.org was adapted to run on an Nvidia Tesla general purpose graphics processing unit (GPU). Global matrix assembly and LU factorization of the resulting dense matrix were performed on the GPU. Out-of-core techniques were used to solve problems larger than available GPU memory. The code achieved over eight times speedup in matrix assembly and nearly 60Gflops/sec in the LU factorization using only 512Mbytes of GPU memory. This research is sponsored by the Office of Advanced Scientific Computing Research; U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725.

Eduardo F. D'Azevedo
Oak Ridge National Laboratory
Mathematical Sciences Section
e6d@ornl.gov

Sylvain Nintcheu Fata
Oak Ridge National Laboratory
nintcheufats@ornl.gov

CP10

Gradient Free Design of Microfluidic Structures on a GPU Cluster

For certain fluid optimization problems such as those of a multi-physics or multi-phase nature, traditional gradient based design methods may break down at sharp interfaces. In these scenarios gradient free methods are necessary. Unfortunately, gradient free methods require more function evaluations than gradient based design methods. In this work we present a design optimization method for microfluidic T-junctions used in lab-on-a-chip devices. We utilize the highly parallel multidirectional search (MDS) optimization procedure along with a multi-core/GPU accelerated flow solver, making the problem ideal for GPU clusters.

Austen C. Duffy
Department of Mathematics
Florida State University
aduffy@math.fsu.edu

CP10

Efficient Numerical Algorithms for Uncertainty Quantification in Computational Mechanics Using Gpus

Graphic processing units (GPUs) are rapidly emerging as much economical and highly competitive alternatives to CPU-based parallel computing. This study implements several existing algorithms, such as, vector convolution, explicit Rünge-Kutta method, multi-dimensional numerical quadrature, and Monte Carlo simulation on a GPU, and discusses their utilization for uncertainty quantification problems arising in computational mechanics and dynamics. Numerical examples are presented to demonstrate the computational efficiency of the GPU-based algorithms compared to corresponding multi-core CPU-based algorithms.

Gaurav Gaurav, Steven Wojtkiewicz
Department of Civil Engineering
University of Minnesota
gaura003@umn.edu, ykvich@umn.edu

CP10

Investigation of Soft Errors As Integrated Components of a Simulation

As architectures become more complex, failures come increasingly into play and simulation behavior becomes less predictable. Especially worrisome are soft errors, i.e. bit ?ips, due to systems operating at such low voltages. Soft errors are especially insidious because they may not even be detected. Linear solvers must be considered as integrated components of the algorithmic stack in a simulation. We investigate the effects of soft errors in the linear solver on the overall simulation.

Victoria Howle
Texas Tech
victoria.howle@ttu.edu

Michael A. Heroux, Patricia D. Hough
Sandia National Laboratories
maherou@sandia.gov, pdhough@sandia.gov

CP10

Parallel Implementation of Adi As Preconditioner

The alternating directions implicit (ADI) method is a classical iterative method for numerically solving linear systems arising from discretizations of partial differential equations. We use ADI as a preconditioner for Krylov subspace methods for problems in two and three space dimensions, because it allows for a highly efficient, completely matrix-free implementation. This talk will demonstrate that effective parallel implementations of the method are possible.

Kyle Stern
University of Maryland, Baltimore County
kstern1@umbc.edu

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

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

CP10

Paladins: A Parallel Algebraic Adaptive Navier-Stokes Solver

In some applications of incompressible fluid dynamics, fast transients, which requires to use small time steps, are

present only in some periods of the overall time interval of interest. A class of second and third order time-accurate splitting schemes has been introduced by Gervasio, Saleri and Veneziani, that features a hierarchical structure prone to time adaptivity. We will present some technical details and parallelization issues of this time adaptive scheme based on algebraic splitting.

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

CP11

OSPRI: A New Communication Runtime System for Global Arrays and Other One-sided Programming Models

We have developed a new communication runtime system named OSPRI (One-Sided PRImitives) for PGAS programming libraries and languages. Significantly improved performance versus ARMCI is observed on the Blue Gene/P architecture for NWChem and Global Arrays benchmarks. The design and implementation strategy of OSPRI targets modern networks in Cray and IBM systems as well as multithreaded and hybrid systems. Finally, interaction with MPI-3 and MPICH-based implementations thereof will be discussed.

Jeff R. Hammond
Argonne National Laboratory
Leadership Computing Facility
jhammond@alcf.anl.gov

CP11

Coupling Multi-Body Fluid-Solid Mechanical Systems

The coupling of multi-physics systems is essential for a variety of problems in engineering and computational physiology. Increasingly, these problems demand the coupling of multiple fluid and solid bodies which may exhibit varied physical responses and interactions through time. We focus on the integration and solution of these systems using a Lagrange multiplier finite element coupling approach and examine – both theoretically and numerically – the accuracy, solvability and uniqueness of these general coupled problems.

David Nordsletten
Massachusetts Institute of Technology
nordslet@mit.edu

CP11

Preconditioning Surface and Subsurface Flow Coupling for Arbitrary Geometries on a Structured Grid

Due to complex dynamics inherent in the physical models, numerical formulation of subsurface and overland flow coupling can be challenging to solve. ParFlow is a subsurface flow code that couples with overland flow via an overland boundary condition prescribed at the top surface. This talk will present a preconditioning approach to discrete systems arising from implicit coupling of these flow regimes in ParFlow. Numerical results will explore the effectiveness of the preconditioner and its cost.

Carol Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

Daniel Osei-Kuffuor
University of Minnesota
dan.osei@gmail.com

Reed M. Maxwell
Department of Geology and Geologic Engineering
Colorado School of Mines
rmaxwell@mines.edu

Steven Smith
Lawrence Livermore National Laboratory
sgsmith@llnl.gov

CP11

Simulations of a Model for Calcium Waves in a Heart Cell Using Comsol Multiphysics

Experiments have shown that calcium waves can occur from spontaneously generated calcium sparks in a heart cell. A model for this process is given by a system of reaction-diffusion equations. We have in the past developed a special-purpose C code with MPI for this model that successfully handles the high-resolution meshes of the desired three-dimensional domains. We will evaluate whether a general-purpose package such as COMSOL Multiphysics can be used at least for lower-resolution studies.

David Trott
University of Maryland, Baltimore County
dtrott1@umbc.edu

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

CP11

Algorithms for Interface Treatment and Load Computation in Embedded Boundary Methods for Fluid-Structure Interaction Problems

This talk focuses on the treatment of fluid-structure interfaces and load computation in embedded boundary methods. First, it presents a numerical method for treating simultaneously the fluid pressure and velocity conditions on static and dynamic embedded interfaces based on the exact solution of local fluid-structure Riemann problems. Next, it describes two consistent and conservative approaches for computing the flow-induced loads on embedded structures. Finally, it discusses applications in aeronautics and underwater implosion.

Kevin Wang
Institute for Compuational and Mathematical Engineering
Stanford University, Stanford, CA
icmewang@stanford.edu

Charbel Farhat
Stanford University

CFarhat@stanford.edu

CP12

Efficient Model Reduction of Large-Scale Nonlinear Systems in Fluid Dynamics

Time critical applications in fluids, such as flow control, demand numerical simulations that are computationally inexpensive, yet extremely accurate. To this end, the authors have developed a model reduction methodology for nonlinear systems that reduces the complexity of evaluating large-scale computational models while demonstrating accuracy and stability. In this talk, results for this method applied to difficult real-world fluids problems are presented for the first time. Orders of magnitude speedups are observed.

Kevin T. Carlberg, David Amsallem
Stanford University
carlberg@stanford.edu, amsallem@stanford.edu

Charbel Bou-Mosleh
Notre Dame University
cboumosleh@ndu.edu.lb

Charbel Farhat
Stanford University
CFarhat@stanford.edu

CP12

Interpolation-Based Model Reduction of Bilinear Systems: New Results and An Approach to Optimal Approximation

Bilinear systems capture the nonlinear features of many problems, ranging from blood circulation to nuclear fission, while retaining much of the simplicity and structure of linear models. In this talk, we will present a method for computing reduced-order models of MIMO bilinear systems satisfying interpolation conditions, and show how to use the systems d-term to recover stability in the reduced-order model. Finally, we present H_2 optimality conditions, and consider their implications in choosing interpolation points

Garret M. Flagg
Virginia Tech
garretf@vt.edu

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

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

CP12

Subspace Tracking for Dimension Reduction in Streaming Data

The real-time analysis of streaming data from sensors can be challenging when the number of sensors is large, the sampling rate is high, and the statistics of the data vary over time. We discuss how we can identify the important data streams using subspace tracking methods such as incremental SVD, PAST, and SPIRIT. These methods track the reduced dimension space as new data come in and the old data become obsolete.

Chandrika Kamath
Lawrence Livermore National Laboratory
kamath2@llnl.gov

CP12

.Direct Method for Calculus of Variations Problems Via Combined Block-Pulse and Orthogonal Functions

In this work, we present a new direct computational method to solve calculus of variations problems. The approach is based of reducing calculus of variations problems into a set of algebraic equations by first expanding the candidate function as a hybrid function with an unknown coefficient. The hybrid function, which consists of combined block-pulse and orthogonal functions, are first introduced. Some properties together with illustrative examples are given.

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

CP12

Approximation of Ml Estimates of Parameters of Hierarchical Logistic Models Via Monte Carlo Simulations

Maximum Likelihood (ML) estimation of parameters of hierarchical logistics models requires the approximation of integrals that do not have a closed form. Several approximation techniques have been proposed. This project explores the performance of four approximations techniques via Monte Carlo simulations when using a hierarchical logistic model to analyze cluster-randomized interventions in Epidemiological studies. Recommendations are made, based on estimation properties, on which approximation technique to use under scenarios where these techniques produce different results.

Rafael E. Diaz
California State University Sacramento
rdiaz@csus.edu

Coskun Cetin
California State University, Sacramento
cetin@csus.edu

Matthew Steinwachs, Lydia Palma, Shawnee Anderson
California State University
Sacramento
mks259@saclink.csus.edu, lp782@saclink.csus.edu, sma274@saclink.csus.edu

CP12

Low-Rank Tensor Methods for Uncertainty Quantification

Uncertainty quantification using spectral stochastic methods like the stochastic Galerkin method involve solving systems with a huge number of unknowns. However, the underlying space has a tensor product structure that can be used for more efficient representation of the quantities involved. We show how to modify standard iterative methods in order to use this structure during the whole solution

process. Error estimates and convergence results for those methods will be given. Further, numerical results will be presented that show how the range of problems feasible to compute can be extended using these methods.

Elmar Zander
TU Braunschweig
Inst. of Scientific Computing
e.zander@tu-bs.de

Hermann G. Matthies
Institute of Scientific Computing
Technische Universität Braunschweig
wire@tu-bs.de

CP13

Recycling Bi-Lanczos Algorithms

Science and engineering problems frequently require solving a long sequence of linear systems. Since CGS and BiCGSTAB are popular methods for solving large nonsymmetric linear systems, we generalize and extend the framework of our Recycling BiCG to CGS and BiCGSTAB. We modify these algorithms to use a recycle space, which is built from left and right approximate eigenvectors. We also propose a Recycling QMR. Initial experiments on various applications give promising results.

Kapil Ahuja, Eric De Sturler
Virginia Tech
kahuja@vt.edu, sturler@vt.edu

CP13

Uintah a Scalable Computational Science Framework for Hazard Analysis

The Uintah Software system was developed to provide an environment for solving a fluid-structure interaction problems on structured adaptive grids on large-scale, long-running, data-intensive problems. Uintah uses a novel asynchronous task-based approach with fully automated load balancing. The application of Uintah to a petascale problem in hazard analysis arising from "sympathetic' explosions in which the collective interactions of a large ensemble of explosives results in dramatically increased explosion violence, is considered. The advances in scalability and combustion modeling needed to begin to solve this problem are discussed and illustrated by prototypical computational results.

Martin Berzins
SCI Institute
University of Utah
mb@sci.utah.edu

Justin Luitjens
SCI Institute and School of Computing
University of Utah
luitjens@cs.utah.edu

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

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

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

Joseph Peterson
Dept of Chemistry
University of Utah
joseph.r.peterson@utah.edu

CP13

Trilinos-Based Solvers for Large Scale Inverse Problems

Solving large scale inverse problems efficiently is integral to computational mathematics. A suite of iterative solvers for linear inverse problems is presented, which has been written utilizing the Trilinos framework. To illustrate the use of the solvers, an application for reducing movement degradation of PET brain scans is described. Removing these blurs through computational post-processing requires solving a large, sparse linear inverse problem.

Sarah Knepper
Emory University
Department of Mathematics and Computer Science
smknepp@emory.edu

James G. Nagy
Emory University
Department of Math and Computer Science
nagy@mathcs.emory.edu

CP13

Real Mathematics

Real Mathematics abstract Mathematics stays today more and more frequently in front of a task for what it is not prepared. It has to its disposal some time series describing objective reality and the task sound: Are these time series a manifestation of a real object Solution The real mathematics (RM), it means mathematics solving this problem, can proceed by the following manner. First of all, it needs to have to its disposal a model of interactions. RM uses the knowledge of the exact philosophy about objects. This says that all real objects have a common structure. The procedure is as follows. If an existence of an object has to be proved, and its properties have to be described, it is necessary to prove that given time series to some extent describe parts of that common structure. RM takes the model of interactions, makes the forecast and measures its confidence using special criterion. If the incorporation of properties of time series been parts of an object improve the forecast, the task is fulfilled. The optimal set of properties for individual time series is found by this manner and RM can determine with what confidence it is an object behind

Martin Vlcek
Ministry of Finance of the Czech Republic
martin.vlcek@mfcr.cz

CP14

For the Stationary Compressible Viscous Navier-Stokes Equations with the No-Slip Condition on a Convex Polygon

Our concern is with existence and regularity of the sta-

tionary compressible viscous Navier-Stokes equations with no-slip condition on convex polygonal domains. Note that $[\breve{,}p] = [0, c]$, c a constant, is the eigenpair for the singular value $l = 1$ of the Stokes problem on the convex sector. It is shown that, except the pair $[0, c]$, the leading order of the corner singularities for the nonlinear equations is the same as that of the Stokes problem. We split the leading corner singularity from the solution and show an increased regularity for the remainder. As a consequence the pressure solution changes the sign at the convex corner and its derivatives blow up. In the numerical point of view, we discretize the stress intensity factor and the regular part of solution for the linearized problem, and show the error estimates of them.

Hyung Jun Choi
Department of Mathematics
POSTECH
choihjs@postech.ac.kr

CP14

A Perfectly Matched Layer for the Boltzmann-BGK Equation and its Application to the Lattice Boltzmann Method

Perfectly Matched Layer (PML) absorbing boundary conditions are proposed for the discreet velocity Boltzmann-BGK equation (DVBE). Following a study of the linear waves supported by DVBE, nonreflecting absorbing boundary conditions are derived using a space-time transformation. Linear analysis shows that the proposed equations are stable for practical numerical calculations. To validate the accuracy of the boundary condition, the DVBE is solved by a finite difference scheme and by a Lattice Boltzmann method (LBM).

Elena Craig
Old Dominion University
ecrai004@odu.edu

Fang Hu
Department of Mathematics
ODU, USA
fhu@odu.edu

CP14

For the Stationary Navier-Stokes Flows with Non-standard Boundary Conditions on Polygonal Domains

In this talk we show existence and regularity for the vorticity-velocity-pressure variables of the stationary Navier-Stokes equations on polygonal domains. Near the non-convex vertex the solution has the same corner singularity order as ths solutions of the Stokes operator with the vorticity boundary condition on the non-convex infinite sector and the regular part that is obtained by splitting the corner singularity from the solution has futher regularities.

Oh Sung Kwon
Postech
jamjjari@postech.ac.kr

CP14

RBF Methods for 2D Fluid Flow on the Sphere

We use a pseudo-spectral method to solve advection problems on the unit sphere. Rather than Chebychev polynomials or Fourier series, we use Radial Basis Functions (RBFs) as an interpolation basis. We investigate the stability of this method with the aim of producing a stable Navier-Stokes solver. Specifically, we are interested in defining the subset of velocity fields for which the advection operator is stable.

Jordan M. Martel
Arizona State University
School of Mathematical and Statistical Sciences
jordan.martel@asu.edu

CP14

Numerical Simulation of Cell/cell and Cell/particle Interaction in Microchannels

A spring model is applied to simulate the skeleton structure of the red blood cell membrane and to study the red blood cell rheology in Poiseuille flow with an immersed boundary method. The lateral migration properties of many cells in Poiseuille flow have been investigated. We also have combined the above methodology with a distributed Lagrange multiplier/fictitious domain method to simulate the interaction of the red blood cells and neutrally buoyant particles in a microchannel for studying the margination of particles.

Tsorng-Whay Pan, Lingling Shi
Department of Mathematics
University of Houston
pan@math.uh.edu, lingling2math.uh.edu

Roland Glowinski
University of Houston
roland@math.uh.edu

CP14

Numerical Treatment of a Slider Bearing Mechanism: Performanance and Evaluation

Numerical solutions are obtained for the thermo hydrodynamic lubrication of a tilted pad slider bearing with heat conduction to the pad as well as to the slider. The fluid film momentum, continuity and energy equations are coupled to the heat conduction eq1uations for the pad and slider, density and viscosity are assumed to be temperature dependent.The results depicted through plots reveal that thermal boundary layer and recirculation zones are noticed in the temperature field.

Pentyala Srinivasa Rao
Indian School of Mines
pentyalasrinivasa@gmail.com

CP15

Framework for Loosely Coupled Fusion Plasma Simulations

We present the Integrated Plasma Simulator (IPS), a highly portable framework for file-based, loose coupling of high performance, parallel multi-physics codes. The framework adopts a component-based approach, where each code is adapted to a standard interface using a thin wrapping layer. Data is exchanged among participating components using a common *Plasma State* layer. Framework services are provided for the management of multi-level concurrent task execution, resource allocation, data movement, fault

tolerance, and asynchronous events propagation.

Wael R. Elwasif
Oak Ridge National Laboratory
elwasifwr@ornl.gov

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

Samantha Foley, Aniruddha Shet
Oak Ridge National laboratory
foleyss@ornl.gov, shetag@ornl.gov

Randall B. Bramley
Indiana University
bramley@cs.indiana.edu

CP15

Lead-Acid Battery Model Under Discharge With Fast Splitting Method

A mathematical model of a valve regulated lead acid battery under discharge is presented as simplified from a standard electrodynamics model. This nonlinear reaction-diffusion model of a battery cell is solved using an operator splitting method to quickly and accurately simulate sulfuric acid concentration. This splitting method preserves continuity over material interfaces encompassing discontinuous parameters. Numerical results are compared with measured data by calculating battery voltage from the acid concentration using the Nernst equation.

R. Corban Harwood, V. S. Manoranjan
Washington State University
rharwood@math.wsu.edu, mano@math.wsu.edu

Dean B. Edwards
University of Idaho
dedwards@uidaho.edu

CP15

Numerical Application Using Invariant Distribution of the Elasto-Plastic Oscillator to Compute Frequency of Deformations

Mean frequency of threshold crossing of the response of an elasto-plastic oscillator under standard white noise are studied in this paper by means of its invariant measure and Rice's formula. Due to the lack of regularity of this measure, an approximation of the Rice's formula is introduced as a solution of a partial differential equation. Finally, a useful criterion is given for engineering problems to evaluate the mean frequencies of deformations.

Laurent Mertz
University Pierre et Marie Curie,Paris 6
mertz@ann.jussieu.fr

Cyril Feau
Commissariat à l'energie atomique
cyril.feau@cea.fr

CP15

3D Shape Optimization in Viscous Incompressible Fluid

A new integral equation for the axially symmetric Oseen equations is obtained based on the Cauchy integral formula for generalized analytic functions. The integral equation has computational advantage over the existing integral equation based on fundamental solutions (Oseenlets). The integral equation constrained optimization approach to finding three-dimensional minimum-drag shapes for bodies translating in viscous incompressible fluid under the Oseen approximation of the Navier-Stokes equations is presented. The approach formulates the Oseen flow problem as a boundary integral equation and finds solutions to this equation and its adjoint in the form of function series. Minimum-drag shapes, being also represented by function series, are then found by the adjoint equation-based method with a gradient-based algorithm, in which the gradient for shape series coefficients is determined analytically. Compared to PDE constrained optimization coupled with the finite element method (FEM), the approach reduces dimensionality of the flow problem, solves the issue with region truncation in exterior problems, finds minimum-drag shapes in semi-analytical form, and has fast convergence. As an illustration, the approach solves three drag minimization problems for different Reynolds numbers: (i) for a body of constant volume, (ii) for a torpedo with only fore and aft noses being optimized, and (iii) for a body of constant volume following another body of fixed shape (torpedo chasing a vessel). The minimum-drag shapes in problem (i) are in good agreement with the existing optimality conditions and conform to those obtained by PDE constrained optimization.

Michael Zabarankin
Department of Mathematical Sciences
Stevens Institute of Technology
michael.zabarankin@stevens.edu

CP16

A Note on the Use of Optimal Control on a Discrete Time Model of Influenza Dynamics

A discrete time Susceptible - Asymptomatic - Infectious - Treated - Recovered (SAITR) model is introduced in the context of influenza transmission. We evaluate the potential effect of control measures such as social distancing and antiviral treatment on the dynamics of a single outbreak. Optimal control theory is applied to identify the best way of reducing morbidity and mortality at a minimal cost. The problem is solved by using a discrete version of Pontryagin's maximum principle. Numerical results show that dual strategies have stronger impact in the reduction of the final epidemic size.

Paula A. Gonzalez-Parra
Program in Computationa Science
The University of Texas at El Paso
pagonzalezparra@miners.utep.edu

Sunmi Lee
Mathematical, Computational and Modeling Sciences Center,
Arizona State University,
mathever@gmail.com

Leticia Velazquez
Computational Science Program
University of Texas, El Paso
leti@utep.edu

Carlos Castillo-Chavez
Arizona State University and
Department of Mathematics
ccchavez@asu.edu

CP16

Shear Flow Instabilities of an Upper Convected Maxwell Model

This work is concerned with the linear stability of viscoelastic shear flows of an upper convected Maxwell fluid under the effect of elasticity. We are focused on the stability problem of a few classes of simple parallel flows in the limit of infinite Weissenberg and Reynolds numbers. We will discuss the numerical stability results. We shall consider plane Couette and Poiseuille flow, the hyperbolic tangent shearlayer and the Bickley jet flows. For all these flows, we shall consider free surface boundary conditions as well as wall boundary conditions. In the inviscid case, all the flows are unstable for free surfaces. For wall bounded flows, the Couette and Poiseuille flows are stable, while stability of the shear layer and Bickley jet depends on the ratio of the channel width to the characteristic length scale of the profile. In all cases, we find that elasticity stabilizes and ultimately suppresses the instability. Our numerical approach is based on the spectral Chebyshev collocation method. We shall also show that some flows, such as plane Poiseuille flow between two parallel free surfaces, also have short wave instabilities. This is in marked contrast to the wall bounded case. In this case, no smooth velocity profiles unstables to short waves are known, and for certain classes of flows there are even results ruling out short wave instability.

Ahmed Kaffel
virginia tech
kaffel07@gmail.com

CP16

Block Preconditioners for Fully-Implicit Integration of the Shallow Water Equations

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-HOMME. We construct a fully implicit framework based on the shallow water equations and view the subsidiary linear system as a block matrix. Formal LU decomposition is performed and block preconditioners are derived based on approximations to the upper triangular block.

P. Aaron Lott
Lawrence Livermore National Laboratory
Aaron.Lott@llnl.gov

Kate Evans
Oak Ridge National Laboratory
evanskj@ornl.gov

Howard C. Elman
University of Maryland, College Park
elman@cs.umd.edu

Carol Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

CP16

A Girsanov Monte Carlo Approach to Particle Filtering for Multi-Target Tracking

We present a novel approach for improving particle filters for multi-target tracking. The suggested approach is based on Girsanov's change of measure theorem for stochastic differential equations. Girsanov's theorem is used to design a Markov Chain Monte Carlo step which is appended to the particle filter and aims to bring the particle filter samples closer to the observations. The numerical results show that the suggested approach can improve significantly the performance of a particle filter.

Panagiotis Stinis
University of Minnesota, Twin Cities
stinis@math.umn.edu

Vasileios Maroulas
University of Tennessee, Knoxville
maroulas@math.utk.edu

CP16

Impice Method for the Simulation of 2D and 3D Compressible Flow Problems in Uintah.

The Implicit Continuous-fluid Eulerian(ICE) method, a semi-implicit finite-volume solver, is a successful and widely used method that applies to flows that range from supersonic to subsonic regimes. The classical ICE method has been expanded to problems in multiphase flow which span a wide area of science and engineering. The ICE method is utilized by the C-SAFE code Uintah written at the University of Utah to simulate explosions, fires and other fluid and fluid-structure interaction phenomena. The implementation of ICE method used in Uintah is described in many papers by Kashiwa at Los Alamos and extended to solve multifield cases by Harman at Utah. We improve the method by using slope limiters and an approximate Riemann Solver to suppress the nonphysical oscillations. The results of the improved ICE method(IMPICE) for compressible flow problems governed by the 2D and 3D Euler equations are shown along with the spatial and temporal error analysis.

Lethuy T. Tran
University of Utah
ltran@cs.utah.edu

Martin Berzins
SCI Institute
University of Utah
mb@sci.utah.edu

CP17

Distributional Properties of Stochastic Shortest Paths for Smuggled Nuclear Material

There are many well studied optimizaton problems on transportation networks, including optimal routing between selected nodes and interdiction. Existing algorithms handle easily weights on both nodes and links. Examples include, Dijkstra's and Floyd Marshall's shortest paths algorithm. We are particularly interested in the scenario where a nuclear material smuggler tries to successfully reach his/her target. The Pathway Analysis, Threat Response and Interdiction Options Tool (PATRIOT), a tool developed at the Los Alamos National Laboratory readily handles this problem using a multi-modal world trans-

portation network by identifying the most likely threat pathway to the target. The identification of this path relies on a set of fixed reliabilities associated with each arc and terminal in the network that represent the adversary's perceived probability of successfully traversing the respective links or nodes. This path thus represents the adversary's most reliable path. In order to account for the adversary's uncertainty and to perform sensitivity analysis we consider random reliabilities. The main quantities of interest are the resulting stochastic most reliable paths or stochastic shortest paths for short, and their properties. These properties may depend on the network, and on the specific origen and destination, nevertheless, to gain a better understanding, we perform some preliminary controlled experiments by performing Monte Carlo simulations on a grid. This paper presents the resulting findings.

Leticia Cuellar, Feng Pan
Los Alamos National Laboratory
leticia@lanl.gov, fpan@lanl.gov

Kevin Saeger, Fred Roach
LANL
saeger@lanl.gov, froach@lanl.gov

CP17

Selection Strategies for Genetic Algorithm to Solve Travelling Salesman Problem

Genetic algorithm is mainly composed of three important operations which are selection, crossover, and mutation. This study explores the performance of various selection strategies in solving travelling salesman problem. Results show that tournament selection strategy outperformed proportional and rank-based selections, achieving the highest solution qualities with low computing times. However, results reveal that tournament and proportional can be superior to the rank-based for smaller problems and become susceptible to premature convergence as problem size increases.

Noraini Mohd Razali
Ms.
norainimbr@ump.edu.my

John Geraghty
Dr.
john.geraghty@dcu.ie

CP17

Scalable Distributed Coalition Formation for Large-Scale Collaborative Multi-Agent Systems

We study an important problem in Distributed AI and collaborative multi-agent systems, that of decentralized coalition formation. We have designed, implemented and extensively analyzed a scalable, fully distributed algorithm for dynamic coalition formation that ensembles of collaborative autonomous agents can use as a basic coordination subroutine. We summarize our algorithm and share recent simulation results that validate our earlier claims about high scalability and efficiency of our approach to distributed coalition formation. We also outline our ongoing and future work on (i) scaling the implementation from hundreds to thousands of autonomous agents, and (ii) analyzing the potential benefits of multi-tiered reinforcement learning on how to coordinate and form coalitions more effectively.

Predrag T. Tosic
University of Houston
ptosic@uh.edu

Naveen Ginne
Department of Computer Science,
University of Houston
ginne.naveen@gmail.com

CP17

Strategies for Iterated Travelers' Dilemma: The Good, The Bad and The Ugly

We study iterated travelers' dilemma (ITD) game from an algorithmic and experimental standpoint. ITD is a two-person game that is not zero-sum, i.e., that in general incorporates both cooperation and competition. We find ITD very interesting due to (i) a unique Nash equilibrium that however corresponds to very low payoffs to both players, (ii) there is a unique pair of strategies that result in maximal social welfare, yet that strategy pair is highly unstable, and (iii) what is an optimal or even *good* play from the individual welfare standpoint critically depends on the adopted notion of individual utility. We propose and then analyze a broad range of IDT strategies via a round-robin, everyone-against-everyone tournament made of a large number of rounds. We share some interesting and, in several instances, surprising findings on the relative performance of various strategies (such as variants of *tit-for-tat*, well-known from the iterated prisoners' dilemma context). We also outline our ongoing and future work focusing on (i) the value of learning and modeling the opponent's behavior for maximizing one's own long-term utility and (ii) co-learning and meta-learning approach to evolving agent pairs whose long-term behavior maximizes the overall, joint utility (that is, social welfare).

Predrag T. Tosic
University of Houston
ptosic@uh.edu

Philip Dasler
Department of Computer Science
University of Houston
philip.dasler@gmail.com

CP17

The Effect of Preconditioning in Interpolatory Model Reduction

We expand upon our results in " Inexact Solves in Interpolatory Model Reduction" by Beattie, Gugercin, Wyatt, 2010, which reported benefits of using a Petrov-Galerkin framework for the linear systems in interpolatory model reduction. We discuss the role of preconditioning in the Petrov-Galerkin framework and the resulting backward error formulation from a systems theory perspective. We also investigate the backward error formulation for the complex case.

Sarah Wyatt
Virginia Tech
sawyatt@vt.edu

Serkan Gugercin
Virginia Tech.
Department of Mathematics

gugercin@math.vt.edu

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

CP18
Explicit Stable Schemes for Advection Equations

Explicit time differencing of advection terms based on augmented spatial stencil is considered both for cases of pure advection and for more complex systems including advection as a part of general phenomena. It is shown that the proposed approximation possesses extended stability and simplifies numerical algorithm without loss of accuracy of numerical solution. Numerical experiments show that the proposed scheme halves required computational time in comparison with the standard leap-frog approximation.

Andrei Bourchtein, Ludmila Bourchtein
Pelotas State University
Mathematics Department
andburstein@gmail.com, bourchtein@gmail.com

CP18
P-Adaptive Hermite Methods for Initial Value Problems

A straightforward order-adaptive (p-adaptive) implementation of Hermite methods for hyperbolic and singularly perturbed parabolic initial value problems is presented, exploiting the fact that Hermite methods allow the degree of the local polynomial representation to vary arbitrarily from cell to cell. Examples illustrating its application to linear and nonlinear wave propagations are included. Stability and convergence analysis of the p-adaptive Hermite methods will be addressed.

XI Chen
University of New Mexico
Department of Mathematics and Statistics
xqroy@unm.edu

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

CP18
Spectral Methods for Systems of Coupled Elliptic Equations

In the presentation I will talk about how our new developed spectral method solvers can be applied to highly nonlinear and high-order evolution equations such as strongly anisotropic Cahn-Hilliard equations from materials science. Both theoretical and empirical points of view will be given on this topic.

Feng Chen
Purdue University, West Lafayette
chen221@math.purdue.edu

Jie Shen
Department of Mathematics
Purdue University
shen@math.purdue.edu

CP18
An Error-Controlled Fast Multipole Method

We present our two-stage error estimation scheme for the fast multipole method (FMM). This scheme can be applied to any particle system with open or periodic boundary conditions. The scaling of the scheme with respect to the number of particles N is $\mathcal{O}(\mathcal{N})$. We show results for homogeneous and clustered distributions consisting of up to one billion particles. The mathematical basics are described briefly.

Holger Dachsel
Juelich Supercomputing Centre
Research Centre Juelich
h.dachsel@fz-juelich.de

Ivo Kabadshow
Research Centre Juelich
i.kabadshow@fz-juelich.de

CP18
A Positivity-Preserving High-Order Semi-Lagrangian Discontinuous Galerkin Scheme for the Vlasov-Poisson Equations

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. We focus on an alternative approach which evolves a grid-based representation through Lagrangian dynamics, followed by a projection onto the original mesh. In particular, we develop a high-order discontinuous Galerkin (DG) method in phase space, and an operator split, semi-Lagrangian method in time. Our novel application of a 4th order split approach maintains mass conservation exactly and is positivity preserving.

David Seal, James A. Rossmanith
University of Wisconsin
Department of Mathematics
seal@math.wisc.edu, rossmani@math.wisc.edu

CP19
Chordal Graph Preconditioners for Solving Large Sparse Linear Systems

We describe a chordal subgraph based preconditioner that can be used to accelerate the convergence of solvers for sparse linear systems. The resulting solvers combine features of both direct and iterative solvers; they provide fast convergence to solution without requiring any additional space for fill-in. We present a parallel algorithm for computing a maximal chordal subgraph and provide experimental results of linear system solutions using the adjacency matrix of the subgraph as preconditioner.

Sanjukta Bhowmick
Department of Computer Science
University of Nebraska, Omaha
sbhowmick@unomaha.edu

Tzu-Yi Chen
Pomona University
tzuyi.chen@pomona.edu

Kanimathi Duraisamy
University of Nebraska at Omaha

kduraisamy@unomaha.edu

Lisa Peairs
Pomona University
lisa.peairs@pomona.edu

CP19

An Integer Programming Formulation for Vertex Elimination Problem

We propose a 0-1 Integer Program formulation for a variant of the optimal Jacobian accumulation problem that employs vertex elimination strategy. The vertex elimination problem is first modeled as a graph problem, and then is casted as a 0-1 IP. The number of variables and constraints of this IP blows up as the number of nodes and edges in the graph increases. Solving such IP for moderate size graph is already a computational challenge. We investigate two different strategies to help solving such large IPs: first, we add valid constraints which help tightening the bounds, as well as symmetry-breaking constraints that lead to less nodes in doing branch-and-bound; second, we use properties derived from the graph to significantly reduce the size of the IP. Computational results and considerations will be discussed.

Jieqiu Chen
Argonne National Laboratory
MCS Division
jieqchen@mcs.anl.gov

Paul D. Hovland
Argonne National Laboratory
MCS Division, 221/C236
hovland@mcs.anl.gov

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

CP19

Modeling Large Nonlinear Anisotropies with Jfnk Methods

Jacobian-Free Newton-Krylov (JFNK) methods have been shown to allow for the accurate computation of nonlinear terms and large time-steps for advancing magnetohydrodynamic (MHD) simulations. We present results on the efficacy of a JFNK method applied to MHD systems with large nonlinear anisotropy due to the dependence of the (thermal) conductivity tensor on the magnetic field and temperature. This massively parallel algorithm is implemented in NIMROD, an extended MHD code supported by DOE Fusion Energy Sciences.

Ben Jamroz
Tech-X
jamroz@txcorp.com

Scott Kruger, Travis M. Austin
Tech-X Corporation
kruger@txcorp.com, austin@txcorp.com

CP19

Multilevel Ilu Preconditioning for Indefinite Linear Systems

Many applications in science and engineering require the solution of ill-conditioned and indefinite linear systems, which can be challenging to solve by iterative methods. This talk will present a multilevel preconditioning approach that combines ideas from multigrid and strategies like modified ILU and diagonal perturbation techniques, to efficiently handle indefinite linear systems. Numerical results will explore the effectiveness of the preconditioner and its cost on problems from electromagnetics and crystal simulation.

Daniel Osei-Kuffuor
University of Minnesota
dan.osei@gmail.com

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

CP19

On the Memory and I/O Requirements of Multifrontal Sparse Factorization

The multifrontal method organizes the factorization of sparse matrices as a tree of tasks. The tasks should be executed in a topological order. Among all topological orders, which one minimizes the memory requirements? If the main memory is not large enough, what is the minimum size of the input-output operations? We investigate the known alternatives and propose a new solution for the first problem. We show that the second problem is NP-hard and propose heuristics.

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

Mathias Jacquelin, Loris Marchal
ENS Lyon
France
mathias.jacquelin@ens-lyon.fr, loris.marchal@ens-lyon.fr

Yves Robert
ENS LYon, France
& Institut Universitaire de France
yves.robert@ens-lyon.fr

CP19

Convergence Analysis of SART by Bregman Iteration and Dual Gradient Descent

We provide two approaches to prove the convergence of simultaneous algebraic reconstruction technique (SART): linearized Bregman iteration and dual gradient descent. These proofs can also be applied to other Landweber-like schemes such as Cimmino's algorithm and component averaging (CAV). Furthermore the noisy case is considered and error estimate is given. The numerical experiments are provided to demonstrate the convergence results.

Ming Yan
University of California, Los Angeles
Department of Mathematics
yanm@math.ucla.edu

CP20

LU Factorization of Finite Difference Matrices in O(N) Operations

The talk describes recently developed techniques for computing a highly accurate LU-factorization of a the stiffness matrix arising from the finite difference discretization of an elliptic PDE. While inspired by nested dissection methods, the new techniques have linear complexity. The new approach is also more stable, versatile, and far faster for problems involving multiple right-hand sides than established O(N) techniques which are based on iterative methods (multigrid, pre-conditioned GMRES, etc).

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

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

CP20

Progress on the Development of a Elliptic Operator Matrix Assembly Extension to the Chombo Amr Library

Block-structured Adaptive Mesh Refinement (AMR) offers the prospect of solving elliptic equations with near linear complexity by locally refining the mesh where needed, applying Gauss-Seidel smoothing on the error recursively from the finest to the coarsest mesh, and prolonging the error from coarse to fine grids. Although this matrix-free procedure is known to be extremely efficient for Laplace type problems, scaling to 10,000s or more processes, the classical multigrid algorithm can encounter difficulties for highly anisotropic elliptic operators and/or when the conductivity tensor is misaligned to the coordinate system. To enable difficult physics problems to leverage multigrid, we are developing an extension to the Chombo AMR library that will build a multi-level representation of the sparse matrix. The sparse linear system can then be solved using algebraic multigrid, among other methods provided by the PETSc library. In addition to extending the multigrid applicability domain, assembling the matrix will allow users to study the discretized properties of the system (determinant, eigenvalues, etc), which are not readily accessible when using a matrix-free approach. Results of highly anisotropic heat equations in two and three dimensions are presented.

Alexander Pletzer
Tech-X Corporation
pletzer@txcorp.com

Ben Jamroz
Tech-X
jamroz@txcorp.com

Travis M. Austin
Tech-X Corporation
austin@txcorp.com

Yongjun Choi
Tech-X
yjchoi@txcorp.com

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

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

Brian Van Straalen
Lawrence Berkeley National Laboratory
Department N E R S C
bvstraalen@lbl.gov

Lois Curfman McInnes
Argonne National Laboratory
curfman@mcs.anl.gov

Louis Howell
Lawrence Livermore National Laboratory
nazgul@llnl.gov

CP20

A Schur Complement Approach for Solving a Nearly Hermitian System

We discuss an approach for solving an $n \times n$ linear system with rank s skew-Hermitian part ($s \ll n$). Such systems arise in discretizations of certain integral equations, e.g., wave scattering applications. Our approach is based on the observation that such a linear system can be interpreted as the Schur complement of a larger $(n+s) \times (n+s)$ system. We can thus solve the original system by solving $s+1$ Hermitian systems of order n, e.g., using MINRES, which leads to significant storage savings. We then directly solve one $s \times s$ system. We present numerical results demonstrating the Schur algorithm's competitiveness with both a standard Krylov method (GMRES) and a short-term recurrence method (IDR). This method can also serve as a preconditioner for systems whose skew-Hermitian part is well-approximated by a rank s matrix.

Mark Embree
Department of Computational and Applied Mathematics
Rice University
embree@rice.edu

Josef Sifuentes
Rice University
josefs@rice.edu

Kirk M. Soodhalter, Daniel B. Szyld
Temple University
Department of Mathematics
ksoodha@temple.edu, szyld@temple.edu

CP20

Solving Symmetric Indefinite Systems with Symmetric Positive Definite Preconditioners

We consider an iterative solution of a symmetric indefinite linear system with a symmetric positive definite preconditioner. We describe a novel preconditioning strategy, which is based on the idea of approximating the inverse of the absolute value of the coefficient matrix (absolute value preconditioners). A simple example of the (geometric) multigrid absolute value preconditioner is constructed for a model problem of the discrete real Helmholtz equation

with a relatively low wavenumber.

Eugene Vecharynski
University of Colorado Denver
yaugen.vecharynski@ucdenver.edu

Andrew Knyazev
CU-Denver
andrew.knyazev@cudenver.edu

CP20

A Stable Runge-Free, Gibbs-Free Rational Interpolation Scheme

We construct a rational interpolation scheme by blending polynomial interpolants on subsets of the interpolation points. For a large class of grids, including uniform points, we prove that our interpolant converges uniformly to analytic functions exponentially fast, and does not have Runge phenomenon. We show that the exponential sensitivity proven for this kind of schemes by Platte and Trefethen does not lead to numerical instability for this method. We also demonstrate that the interpolation does not have Gibbs phenomenon for discontinuous, piecewise analytic functions.

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

CP20

Sensitivity Analysis of Limit Cycle Oscillations

Many unsteady problems equilibrate to periodic or quasi-periodic behavior. For these problems the sensitivity of periodic outputs to system parameters are often desired, and must be estimated from a finite time calculation. We show that sensitivities computed over a finite time can take excessive time to converge, or fail altogether to reach an equilibrium value. Further, we demonstrate that output windowing enables the accurate computation of periodic output sensitivities.

Joshua A. Krakos, Qiqi Wang
Massachusetts Institute of Technology
joshua.krakos@gmail.com, qiqi@mit.edu

David l. Darmofal
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
darmofal@mit.edu

Steven Hall
Massachusetts Institute of Technology
srhall@mit.edu

MS1

Molecular Dynamics Simulations of Biomolecules on GPUs using the Multilevel Summation Method

Molecular dynamics simulations of biomolecules require at each time step the costly calculation of electrostatic forces between all pairs of N atoms. Of the fast methods for solving the N-body problem, the multilevel summation method (MSM) offers particular advantages to molecular dynamics. MSM also turns out to be well-suited to GPU computation, with a short-range part calculated entirely by square root, multiply, and add instructions and a long-range part arranged as localized grid calculations.

David J. Hardy
Theoretical Biophysics Group, Beckman Institute
University of Illinois at Urbana-Champaign
dhardy@ks.uiuc.edu

MS1

High-Order Discontinuous Galerkin Methods by GPU Metaprogramming

I will describe techniques and tools to tap the enormous performance potential of GPUs for discontinuous Galerkin finite element solvers. Particular emphasis will be on the advantages that high-order discretizations offer on modern SIMD-like architectures. I will explain design considerations and tricks that enabled sustained single-chip performance of above 200 GFlops/s across a wide range of discretization parameters and equation types. I will also briefly touch upon methods and tools for run-time code generation and empirical optimization from a high-level language which were crucial to the present effort's success.

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

Timothy Warburton
Department of Computational and Applied Mathematics
Rice University
timwar@caam.rice.edu

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

MS1

A Case Study of GPUs in Scientific Computing: Low-Order FEM

We present a model for scientific applications based upon a Python framework, in which linear algebra and solver work is handled by libraries, and integration, or physics, is handled by massively parallel accelerators, in this case a GPU, through the PyCUDA framework. We illustrate this paradigm by discussing the step-by-step development of a high performance, portable engine for low-order FEM integration.

Matthew G. Knepley
University of Chicago
knepley@mcs.anl.gov

MS1

Large Scale Multi-GPU FMM for Bioelectrostatics

Abstract not available at time of publication.

Rio Yokota
University of Bristol
rio.yokota@bristol.ac.uk

MS2

On the Construction of Preconditioners for Sys-

tems of PDEs

The purpose of this talk is to discuss a general approach to the construction of preconditioners for the linear systems of algebraic equations arising from discretizations of systems of partial differential equations. We construct the preconditioners based on the mapping properties of the differential operator in Sobolev spaces. In particular, the exact preconditioners can be seen as the Riesz isomorphisms between properly chosen Sobolev spaces, while computationally efficient preconditioners are constructed as spectrally equivalent isomorphisms by using multigrid or domain decomposition methods. Finally, we discuss how to employ this technique in FEniCS.

Kent-Andre Mardal
Simula Research Laboratory and Department of Informatics
University of Oslo
kent-and@simula.no

MS2

An Efficient Implementation of Nitsche's Method on Overlapping Meshes in 3D

Frequent remeshing is required when large deformations occur in ALE-based solution algorithms for Fluid-Structure-Interaction problems. Therefore, fixed fluid mesh methods provide a promising approach when dealing with large deformations. Nitsche's method can be used to derive a systematic finite element formulation for problems with overlapping meshes that is stable and has optimal order. Furthermore, it avoids the introduction of Lagrange multipliers used in other domain decomposition approaches to enforce the interface conditions. Despite its promising properties, Nitsche's method on overlapping meshes has so far only been implemented in 2D with several restrictions on how meshes overlap. This is mainly because of the complexity of arbitrary mesh intersections in 3D. In this contribution, we discuss in detail the general aspects of an efficient 3D implementation of Nitsche's method for overlapping meshes, ranging from computing mesh intersections to integration on intersected elements. Finally, a realization within the FEniCS framework is discussed and we present numerical examples that include the Poisson equation and linear elasticity.

Andre Massing, Anders Logg
Simula Research Laboratory
massing@simula.no, logg@simula.no

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

MS2

Automated Goal-Oriented Error Control

In many areas of computer simulation, the assessment of the quality of a computed numerical approximation can be a vital task. It can however also be a challenging task, both mathematically, computationally, and from the programmer's view point, implementationally. In this talk, I will present an automated framework for goal-oriented adaptivity and error control for finite element methods. Since the framework is automated, a user can take advantage of state-of-the-art adaptivity techniques with minimal implementational effort. The framework is implemented within the FEniCS project, a project for the development of concepts and software for automated solution of differential equations.

Marie E. Rognes, Anders Logg
Simula Research Laboratory
meg@simula.no, logg@simula.no

MS2

Automated Solution of Optimal Control Problems for Coupled PDEs

We will present various approaches to solving optimal control problems with multiple coupled equations. Different solution methods can be implemented and examined rapidly using automated code generation, and in particular using automatic differentiation. For 'one-shot methods', when changing from one equation to another it is possible to simply change the functional, in a high-level domain-specific language, whose stationary points correspond to the solution of the control problem. The necessary derivatives will be computed automatically, and the remainder of the program will be generated and compiled automatically. Other methods for finding stationary points of a functional can be constructed easily using a collection of basic building blocks that mirror mathematical operations. A number of examples using different algorithms will be presented for Stokes flow and flow through porous media.

Andi Merxhani
Department of Engineering, University of Cambridge
am920@cam.ac.uk

Garth Wells
University of Cambridge
gwn20@cam.ac.uk

MS3

Algorithms for Linear Elasticity on Overlapping Grids

In this talk we describe, evaluate, and compare two numerical approaches for solving the equations of linear elasticity on composite overlapping grids. In the first approach we solve the elastodynamic equations posed as a second-order system using a conservative finite difference approximation. In the second approach we solve the equations written as a first-order system using a high-order characteristic-based (Godunov) finite-volume method.

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

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

Daniel Appelo
Division of Engineering and Applied Science
California Institute of Technology
appelo@caltech.edu

MS3

Hybrid and Dynamically Adaptive Higher-Order Shock-Capturing Methods for Compressible Gas Dynamics

We present an advanced high-resolution approach for simulating turbulence with shock interaction in multi-component gases that aims at approximating interfacial mixing with centered schemes and capturing shock waves with higher-order upwind methods. Numerical stability is achieved by the utilization of deep hierarchies of block-structured adaptive mesh refinement combined with a reliable scheme switching criterion. The design of the software will be discussed and large-scale parallel direct numerical and large eddy simulation examples will be presented.

Ralf Deiterding
Oak Ridge National Laboratory
Computer Science and Mathematics Division
deiterdingr@ornl.gov

Jack L. Ziegler, Manuel Lombardini, Dale I. Pullin
Graduate Aeronautical Laboratories
California Institute of Technology
jackalak@caltech.edu, manuel@caltech.edu, dale@caltech.edu

MS3

Deforming Composite Grids for Fluid Structure Interactions

We describe the use of deforming composite overlapping grids for the solution of problems coupling fluid flow and deforming solids. The method is based on a mixed Eulerian Lagrangian technique. Local moving boundary-fitted grids are used near the deforming interface and these overlap non-moving grids which cover the majority of the domain. The approach is described and validated for some fluid structure problems involving high speed compressible flow and linear elastic solids.

William D. Henshaw
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
henshaw@llnl.gov

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banks20@llnl.gov

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

MS3

Stable Grid Refinement for Seismic Wave Simulation

We will present an energy conserving extension of our finite difference method for the elastic wave equation to a composite grid, consisting of a set of structured rectangular component grids with hanging nodes on the grid refinement interface. We prove that the resulting, non-dissipative, method with variable coefficients is stable.

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

Anders Petersson
Center for Applied Scientific Computing
Lawrence Livermore National Lab
andersp@llnl.gov

MS4

Full Waveform Inversion using Diffraction Waves

Diffracted and reflected waves are fundamentally different physical phenomena. Most of seismic imaging, as practiced in the industry, is tuned to processing reflected waves, which carry most of the information about subsurface. However, diffracted waves are also important, both because they are a direct response of small geologically-important subsurface features and because they behave differently from reflections in the process of imaging. We propose to separate diffractions from reflections in the recorded reflection seismic data and to use focusing of diffractions in the imaging process as an objective function for full waveform inversion for seismic velocity parameters. Preliminary tests indicate that the focusing objective function may allow for an efficient gradient-based optimization.

Sergey Fomel
University of Texas at Austin
sergey.fomel@beg.utexas.edu

MS4

Full-waveform Inversion with Compressive Updates

Full-waveform inversion relies on large multi-experiment data volumes. While improvements in acquisition and inversion have been extremely successful, the current push for higher quality models reveals fundamental shortcomings handling increasing problem sizes numerically. To address this fundamental issue, we propose a randomized dimensionality-reduction strategy motivated by recent developments in stochastic optimization and compressive sensing. In this formulation conventional Gauss-Newton iterations are replaced by dimensionality-reduced sparse recovery problems with source encodings.

Felix J. Herrmann
The University of British Columbia
Department of Earth and Ocean Sciences
fherrmann@eos.ubc.ca

Aleksandr Aravkin, Tristan van Leeuwen, Xiang Li
the University of British Columbia
saravkin@eos.ubc.ca,
tristan van leeuwen ¡tleeuwen@eos.ubc.ca¿a, xiang li ¡lxiang@eos.ubc.ca¿

MS4

A Discontinuous Galerkin Method for Coupled Elastic-acoustic Inverse Wave Propagation

Optimization-based methods for full-waveform inversion to recover locally varying seismic wave speeds will be presented. The forward equation is discretized by a higher-order discontinuous Galerkin (dG) method, which uses an upwind numerical flux. The computation of discretely consistent gradients leads to a dG discretization for the adjoint equation based on a downwind flux. To quantify the uncertainty in the reconstruction, an estimate of the variance

at the maximum a posteriori point is computed.

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

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, Lucas Wilcox
University of Texas at Austin
georgst@ices.utexas.edu, lucasw@ices.utexas.edu

MS4

From Simulation to Inversion: Algorithm and Software Organization for Imaging and Inversion

We describe a few mild design constraints which permit rapid adaptation of simulation code for linear wave problems to inversion or design optimization applications, retaining the parallel and other performance enhancements of the underlying simulator. We also describe a framework taking advantage of these concepts which we have used to build a variety of inversion applications. Wave inverse problems tend to be afflicted by a variety of features, including extreme ill-conditioning and nonlinearity, which degrade the performance of straightforward optimization methods. Variants of data-fitting inversion, motivated by standard methods in exploration geophysics, may relieve some of these difficulties. The framework approach also accommodates these extensions to standard inversion.

William Symes, Marco Enriquez, Dong Sun, Rami Nammour, Xin Wang
Rice University
symes@caam.rice.edu, marco.enriquez@caam.rice.edu, dong.sun@rice.edu, rami.nammour@rice.edu, wangxin.tom@caam.rice.edu

Igor Terentyev
Dept. of Computational and Applied Mathematics
Rice University
igtrnt@caam.rice.edu

MS5

Treed Gaussian Processes for Mixed-Integer Surrogate Modeling

The ability to understand and analyze computer simulation models can rely heavily on the ability to approximate the model with a good statistical surrogate. While traditional approaches have focused on the case of only continuous input variables, we present new developments with treed Gaussian processes making them applicable for both discrete and continuous inputs. Thus they are an ideal surrogate for mixed-integer problems, as they allow modeling of nonstationarity and quantification of uncertainty.

Yuning He
UC Santa Cruz
yuning@soe.ucsc.edu

Herbie Lee
University of California, Santa Cruz
Dept. of Applied Math and Statistics
herbie@ams.ucsc.edu

MS5

Evaluation of Mixed Continuous-discrete Surrogate Approaches

Evaluating the performance of surrogate modeling approaches is essential for determining their viability in optimization or uncertainty analysis. To this end, we evaluated categorical regression, ACOSSO splines, and treed GPs on a set of test functions. We describe the principles and metrics we used for this evaluation, the characteristics of the test functions we considered, and our software testbed. Additionally, we present our numerical results and discuss our observations regarding the merits of each approach.

Herbie Lee
University of California, Santa Cruz
Dept. of Applied Math and Statistics
herbie@ams.ucsc.edu

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

Patricia D. Hough
Sandia National Laboratories
pdhough@sandia.gov

MS5

Design and Analysis of Computer Experiments with Qualitative and Quantitative Factors

We propose a simple yet efficient approach for building Gaussian process models for computer experiments with both qualitative and quantitative factors. This approach uses the hypersphere parameterization to model the correlations of the qualitative factors, thus avoiding the need of directly solving optimization problems with positive definite constraints. The effectiveness of this method is successfully illustrated by several examples. Also will be discussed are new classes of space-filling designs for building mixed-integer surrogate models.

Peter Qian, Qiang Zhou, Shiyu Zhou
University of Wisconsin - Madison
peterq@stat.wisc.edu, qzhou3@wisc.edu, szhou@engr.wisc.edu

MS5

Functional ANOVA Decomposition and Discrete Inputs in Computer Model Emulation

Abstract not available at time of publication.

Curtis Storlie
Los Alamos National Laboratory
storlie@lanl.gov

Brian Reich
North Carolina State University
brian_reich@ncsu.edu

MS6

Parameter Estimation for Photosynthesis

We have developed an extensive model of (leaf) photosynthesis based on ordinary differential equations. Using this model, we have shown how carbon fixation might be doubled (more biofuel), and we can analyze questions from evolutionary biology and the influence of climate change. An important problem is estimating the many parameters that govern the photosynthesis process, but that cannot be measured directly inside a cell, and that may vary considerably among different plants. This is a collaboration with Xinguang Zhu (Plant Systems Biology Group, Chinese Academy of Sciences, Shanghai), and Stephen Long (Plant Biology and Crop Science, University of Illinois at Urbana-Champaign)

Eric De Sturler, Eun Chang
Virginia Tech
sturler@vt.edu, changer@math.vt.edu

MS6

Numerical Solution for Time Dependent Optimal Transport

In this talk we present a new computationally efficient numerical scheme for the computation of the optimal L2 mass transport mapping. Our starting point is the Benamou and Brenier approach to the problem which computes the complete transport path. We review the approach and discuss its numerical shortcomings. We then derive an efficient discretization and a solution technique for the problem. We demonstrate the effectiveness of our approach using a number of numerical experiments.

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

MS6

Rank-Deficient Nonlinear Least Squares Problems and Subset Selection

We examine the local convergence of the Levenberg-Marquardt method for the solution of nonlinear least squares problems that are rank deficient and have nonzero residual. We show that replacing the Jacobian by a truncated singular value decomposition can be numerically unstable. We recommend instead the use of subset selection. We corroborate our recommendations by perturbation analyses and numerical experiments.

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

Carl T. Kelley
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu

Scott Pope
SAS Institute
scott.pope@gmail.com

MS6

Regularized Gauss-Newton for Parameter Estimation with Applications in Tomographic Imaging

We present a new algorithm for the solution of nonlinear least squares problems arising from low-order, parametric tomographic imaging models. The ill-conditioning of the Jacobian, together with the presence of noise in the data, motivates us to devise a regularized, trust-region-based Gauss-Newton approach for determining search directions. Examples show the success of our approach relative to well-known alternatives.

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

Eric De Sturler
Virginia Tech
sturler@vt.edu

MS7

Optimization based Modeling. Part I. Additive Decomposition of Multiphysics Problems

We formulate an approach, based on additive operator decomposition and reconnection via optimization, which allows to synthesize robust and efficient solvers for a coupled multiphysics problems from simpler solvers for their constituent components. To illustrate the scope of the approach we show how a robust and efficient solver for nearly hyperbolic PDEs can be derived from standard, off-the-shelf algebraic multigrid solvers for the Poisson equation that do not work for the original equations.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

MS7

Optimization-based Modeling, Part II: Monotone, Bound Preserving Transport and Remap

Remap, broadly defined as the transfer of simulation data between different computational meshes subject to physical constraints is a critical component in a number of numerical algorithms, e.g. ALE transport schemes. We present a new mathematical framework for remap, based on ideas from constrained optimization. Optimization-based remap (OBR) is formulated as the solution of an optimization problem, in which accuracy considerations, handled by an objective functional, are separated from monotonicity considerations, handled by a carefully defined set of inequality constraints. As such, OBR naturally applies to unstructured meshes and meshes comprised of arbitrary polyhedral cells. In addition, we demonstrate that OBR can generate significantly more accurate solutions at or below the cost of the best currently used remapping techniques.

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

Pavel Bochev

Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

MS7

Advanced Variational Methods on Tetrahedral Meshes for Hyperbolic Systems: Stabilization, Compatibility, Accuracy

This talk describes a new variational multiscale formulation for Lagrangian-ALE shock hydrodynamics for tetrahedral finite elements. To the author knowledge this is the only robust and accurate tetrahedral formulation developed to date for shock hydrodynamics computations. The formulation preserves conservation of mass/momentum/total energy, which can be interpreted as compatibility relationships. A specific ALE remap algorithm based on flux-corrected transport (FCT) of systems of equations invokes another compatibility relationship, the Geometric Conservation Law, which is shown to be necessary in ensuring monotonicity of the remapped fields.

Guglielmo Scovazzi
Sandia National Laboratories
gscovaz@sandia.gov

MS7

Optimization-Based Synchronized Flux-Corrected Conservative Interpolation (remapping) of Mass and Momentum for Arbitrary Lagrangian-Eulerian Methods

A new optimization-based synchronized flux-corrected conservative interpolation (remapping) of mass and momentum for arbitrary Lagrangian-Eulerian hydro methods is described. Fluxes of conserved variables - mass and momentum - are limited in synchronized way to preserve local bounds of primitive variables - density and velocity.

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

Richard Liska
Czech Technical University
Fac Nuclear Science & Phys Eng
liska@siduri.fjfi.cvut.cz

Pavel Vachal
Czech Technical University in Prague
vachal@galileo.fjfi.cvut.cz

Burton Wendroff
Los Alamos National Lab
bbw@lanl.gov

MS8

Finite Element Method (FEM) for the Modified Poisson Nernst Planck Equations (MPNPE)

The dynamics of ions in implicit solvent models are described by the coupled Poisson-Nernst-Planck equations (PNPE). The PNPE are an alternative to the computationally expensive fully explicit particle based methods. However, PNPE may exhibit unbounded concentration of ions. The MPNPE removes this drawback by taking steric effects into account. This talk addresses FEM for the MPNPE for ion flow through a protein pore. We show MPNPE leads to more accurate calculation of ionic current than the classical PNPE.

Jehanzeb H. Chaudhry
University of Illinois
Urbana-Champaign
jhameed2@illinois.edu

MS8

Goal-Oriented Error Estimation for the Poisson-Boltzmann Equation

Studying solvation effects in biomolecules is important to accurate representation of their physical environment. In this regard, a meaningful quantity is the solvation free energy, which can be written as a linear functional of the solution to the nonlinear PDE known as the Poisson-Bolzmann equation (PBE). In this talk we present an adaptive mesh refinement algorithm that drives refinement in an effort to control the error in the solvation free energy. Key to this algorithm is the development of goal-oriented error indicators. We show how these indicators are developed and give the steps needed for their computation and use in adaptive mesh refinement. An additional aspect of this refinement algorithm, is the development of a split-domain marking strategy that proved critical to its sucess. To show the efficacy of the adaptive refinement algorithm, we present results based on computing the solvation free energy of the protein Fasciculin-1.

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

Stephen D. Bond
Sandia National Laboratories
sdbond@sandia.gov

Burak Aksoylu
TOBB University of Economics and Technology, Turkey
tba@siam.org

Michael Holst
University of California, San Diego, CA, USA
tba@siam.org

MS8

Recovery-type Error Estimators for the Adaptive Finite Element Approximation of the Poisson-Boltzmann Equation

In this talk, we discuss recovery-type a posteriori error estimators for use in the adaptive finite element method applied to the Poisson-Boltzmann equation. Specifically, we derive and study both gradient and flux recovery methods. These quantities are important in molecular dynamics simulations and the computation of reaction rates. Finally, we will present numerical results illustrating convergence of the adaptive method using each indicator.

Long Chen
University of California at Irvine
chenlong@math.uci.edu

Ryan Szypowski, Michael Holst
UC San Diego
Department of Mathematics
rszypows@math.ucsd.edu, mholst@ucsd.edu

Yunrong Zhu
Department of Mathematics
University of California at San Diego
zhu@math.ucsd.edu

MS8

Adaptive Multiscale Methods for the Poisson-Boltzmann and Nerst-Planck Equations

In this lecture, we give an overview of several related projects based at UCSD involving the design and analysis of high-resolution, high-fidelity mathematical and numerical modeling techniques for solvation and diffusion phenomena in biophysics. We first outline some new theoretical results on the robust numerical discretization of the Poisson-Boltzmann equation using adaptive finite element techniques. We then give an analysis of a coupled solvation-large-deformation nonlinear elasticity model, describe an iterative method for its simulation, and give some convergence results. We then consider reaction diffusion models such as the Poisson-Nernst-Planck (PNP) and Smoluchowski-Poisson-Boltzmann (SPB) models, and describe our recent work on the development of robust simulation tools using modern geometric modeling techniques and adaptive finite element methods.

Michael J. Holst
Univ. of California, San Diego
Department of Mathematics
mholst@math.ucsd.edu

Yongcheng Zhou
Department of Mathematics
Colorado State University
yzhou@math.colostate.edu

MS9

STK-Mesh Example Computation Setup and Parallel Execution

HPC computational kernels applied to parallel heterogeneous unstructured meshes require setup of and access to discretization knowledge, mesh data structures, and computational fields. A STK-Mesh example of problem setup and hybrid parallel execution of a non-trivial computational kernel is presented.

Todd Coffey, H. Carter Edwards
Sandia National Laboratories
tscoffey@sandia.gov, hcedwar@sandia.gov

MS9

STK-Mesh Domain-driven Design

An unstructured mesh providing hybrid parallel (distributed + threaded), heterogeneous discretization, and dynamic modifications for HPC scientific and engineering computations has significant inherent complexity. Managing this complexity through domain-driven design has been critical to the success of the STK-Mesh project. The STK-Mesh domain model (conceptual design) developed to manage this complexity is presented.

H. Carter Edwards
Sandia National Laboratories
hcedwar@sandia.gov

MS9

STK-Mesh Example Dynamic Mesh Modification

Evolving solutions to problems may require advanced modeling and simulation strategies to dynamically modify the discretization of the problems unstructured mesh. Dynamic load balancing, domain cracking or erosion, and solution feature resolution through mesh refinement exemplify advanced problems that require mesh modification. A STK-Mesh example of robust and efficient parallel dynamic mesh modification is presented.

Daniel Sunderland, H. Carter Edwards
Sandia National Laboratories
dsunder@sandia.gov, hcedwar@sandia.gov

MS9

Sierra Toolkit Capabilities and Future

The STK-Mesh is the first component deployed within a larger SIERRA Toolkit project. This project will be releasing advanced HPC CS&E capabilities in the public domain through Sandia National Laboratory's Trilinos (http://trilinos.sandia.gov) repository. This presentation summarizes the current set of STK-Mesh and other SIERRA Toolkit component capabilities and future plans for the SIERRA Toolkit.

Alan B. Williams
Sandia National Laboratory
Distributed Systems Research Department
william@sandia.gov

H. Carter Edwards
Sandia National Laboratories
hcedwar@sandia.gov

MS10

Latency? Not a Problem When You Use FG

If your computation works with large amounts of data, then it may very well suffer from high latency due to disk I/O and/or interprocessor communication. You can mitigate the effect of high latency by overlapping high-latency operations with other, useful work—but implementing programs that do so can be a painstaking, time-consuming, and error-prone venture. In this talk, I will describe the FG system, which makes it easy to overlap high-latency operations with other work.

Thomas Cormen
Department of Computer Science
Dartmouth College
thc@cs.dartmouth.edu

MS10

A Parallel Fast Algorithm for Applying Fourier Integral Operators

Fourier Integral Operators (FIOs) encompass a wide range of transforms including several variants of the Fourier Transform and the Generalized Radon Transform. A new scheme for the distributed-memory parallelization of the butterfly algorithm for FIOs is presented and results are shown on thousands of cores. This is a joint work with L. Demanet and N. Maxwell.

Jack Poulson
Institute for Computational Engineering and Sciences

The University of Texas at Austin
jack.poulson@gmail.com

Lexing Ying
University of Texas
Department of Mathematics
lexing@math.utexas.edu

Nicholas Maxwell
The University of Houston
nicholas.maxwell@gmail.com

Laurent Demanet
Mathematics, MIT
demanet@gmail.com

MS10

Evaluation of OpenMP Task-Based Parallelism for the Adaptive Fast Multipole Algorithm

Parallel execution of the Adaptive Fast Multipole Algorithm on a multicore shared-memory computer presents complex load balancing and memory locality issues. We report on the use of OpenMP 3.0 task parallel directives to express the nested parallelism in the algorithm in a clear and concise fashion, and report on the execution times obtained using a novel OpenMP runtime system to measure and tune overall performance as well as parallel scaling.

Jan F. Prins
Dept. of Computer Science
University of North Carolina at Chapel Hill
prins@cs.unc.edu

Robert Overman, Stephen Olivier
UNC Chapel Hill
reoverma@email.unc.edu, olivier@cs.unc.edu

Allan Porterfield
RENCI
UNC Chapel Hill
akp@renci.org

Bo Zhang
Duke University
zhangb@cs.duke.edu

Jinfang Huang
UNC Chapel Hill
Applied Mathematics
huang@email.unc.edu

Michael Minion
University of North Carolina-Chapel Hill
Mathematics Department
minion@email.unc.edu

MS10

Parallelization of New-Version Adaptive Fast Multipole Method on Multicore Machines

We present a parallelization scheme for the adaptive new version fast multipole method on the multicore machines. We organized the algorithm into upward, interaction, and downward stages, where all involved operators are implemented to have intrinsic mutual exclusion property. A spatio-temporal partition strategy is developed. The scheme has been implemented for Laplace and Yukawa kernels and is shown to be highly efficient in several numerical results.

Bo Zhang
Duke University
zhangb@cs.duke.edu

Jingfang Huang
Department of Mathematics
University of North Carolina, Chapel Hill
huang@amath.unc.edu

Nikos Pitsianis
Department of Electrical & Computer Engineering
Duke University
nikos.p.pitsianis@duke.edu

Xiaobai Sun
Department of Computer Science
Duke University
xiaobai@cs.duke.edu

MS11

Next- Generation Sequencing

Abstract not available at time of publication.

Srinivas Aluru
Iowa State University
aluru@iastate.edu

MS11

Sensor Placement

Abstract not available at time of publication.

William E. Hart
Sandia National Labs, USA
wehart@sandia.gov

MS11

Towards Hypothesis Testing for Communities in Networks

Community detection methods for social network analysis must generally account for uncertainty in relationship observation. Community detection algorithms approximately optimize an objective that leads to higher internal connectivity than external. Many believe that communities should not have statistically significant sub-communities and communities should not be readily explained as statistical fluctuations in larger communities. We describe hypothesis testing methods to help determine if a set of communities has been computed to the correct level of resolution.

Jonathan Berry, Cynthia Phillips
Sandia National Laboratories
jberry@sandia.gov, caphill@sandia.gov

MS12

Velvetrope: a Parallel, Bitwise Algorithm for Finding Homologous Regions within Multiple Se-

quences

Abstract not available at time of publication.

Scott Clark
Cornell University
sc932@cornell.edu

MS12

Parallel, Precision-Limited PDE Evolution using Variable-Order Series Expansion

Abstract not available at time of publication.

Hal Finkel
Yale University
hal.finkel@yale.edu

MS12

Subgraph Isomorphism Algorithms for Multithreaded Shared Memory Architectures

Abstract not available at time of publication.

Claire Ralph
Cornell University
ccr53@cornell.edu

MS12

Bringing OpenCL to Supercomputing: Finite Volume Solvers on GPU Clusters

Abstract not available at time of publication.

Noah F. Reddell
University of Washington
noah.reddell@gmail.com

MS13

Hierarchical Matrix Preconditioners for Saddle Point Problems

Hierarchical (H-) matrices provide a powerful technique to compute and store approximations to dense matrices in a data-sparse format. The basic idea is the approximation of matrix data in hierarchically structured subblocks by low rank representations. The usual matrix operations can be computed approximately with almost linear complexity. We use such an H-arithmetic to set up preconditioners for the iterative solution of sparse linear systems of saddle point type as they arise in the finite element discretization of systems of PDEs.

Sabine Le Borne
Tennessee Technological University
Department of Mathematics
SLeborne@tntech.edu

MS13

Fast Direct Solvers for Elliptic PDEs

That the linear systems arising upon the discretization of elliptic PDEs can be solved very rapidly is well-known, and many successful iterative solvers with linear complexity have been constructed (multigrid, Krylov methods, etc). Interestingly, it has recently been demonstrated that it is possible to directly compute an approximate inverse to the coefficient matrix in linear (or close to linear) time. The talk will survey some recent work in the field and demonstrate the advantages of direct solvers in terms of stability, and, sometimes, speed.

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

MS13

Exploiting the Natural Block-Structure of Hierarchically-Decomposed Variational Discretizations of Elliptic Boundary Value Problems

Hierarchical bases for higher-order finite element discretizations provide a natural block structure with two key properties which suggest that a block Gauss-Seidel (or Jacobi) approach will be effective: only the linear-linear diagonal block is ill-conditioned, so it is the only part which requires sophisticated techniques; and the off-diagonal coupling is weak enough that (at least) a fixed error reduction is guaranteed in each iteration. The performance of the approach is demonstrated using H-matrix techniques.

Jeffrey S. Ovall
University of Kentucky, Mathematics
jovall@ms.uky.edu

MS14

Hybrid Hermite-Discontinuous-Galerkin Methods for Hyperbolic Systems

A class of hybrid methods for hyperbolic problems based on the combination of Hermite-Taylor-Runge-Kutta methods and nodal Discontinous Galerkin methods is presented. Examples illustrating its application to the wave equation and Maxwell's equation will be given.

Daniel Appelo
Division of Engineering and Applied Science
California Institute of Technology
appelo@caltech.edu

Ronald Chen
Department of Mathematics
The University of New Mexico
xqroy@math.unm.edu

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

MS14

Implicit High-order Compact Schemes for Incompressible Flow

This talk describes an implicit, high-order accurate, method for incompressible flow combining compact spatial discretizations with approximate factorization schemes on overlapping grids. The approach introduces techniques that minimize the number of factors used in the implicit scheme while maintaining up to fourth order spatial accuracy at physical boundaries. Implicit time discretization is achieved via a second order accurate approximately factored Crank-Nicolson method that incorporates the compact spatial approximations into the banded solves.

Kyle Chand

Lawrence Livermore National Laboratory
chand1@llnl.gov

MS14
Accurate Methods for Time-Domain Scattering

Volume-based methods for solving scattering problems in the time domain require an accurate near-field radiation boundary condition. Complete radiation boundary conditions represent an essentially optimal solution to this problem for isotropic systems. They are local and inexpensive, they can be stably implemented on polygonal boundaries, and their accuracy is guaranteed by a priori error estimates. We will review their construction and analysis and describe our efforts to produce generally usable implementations.

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

MS14
Using Adaptive Overset Grid Methods for Shape Optimization of Medical Devices

Adaptive overset grid methods are well-suited for shape optimization problems. Rapid and local grid generation techniques facilitate the low-cost construction of computational grids required during the optimization procedure. This work couples the overset grid method with a derivative free optimization algorithm to determine the optimal shape and configuration of blood clots that are trapped by medical devices. The methods are fully automated and demonstrate an adaptive design optimization framework that is broadly applicable.

Michael Singer
Lawrence Livermore National Laboratory
msinger@llnl.gov

Stephen Wang
Kaiser Permanente
stephen.wang@alumni.duke.edu

Darin Diachin
Kanoga Technologies
dpdinca@comcast.net

MS15
On the Application of Parallel Scalable Solvers for the Integrated Hydrologic Model ParFlow

Integrated models with coupled nonlinear physics have particular solver requirements. Here, the ParFlow model, a parallel hydrology model with integrated land-surface processes, and its solver and model framework will be discussed. The different solver requirements for different scale of application, the disparate timescales of physical processes and the different degree of model coupling between components will be highlighted. Finally, physical applications and parallel scaling will be addressed.

Reed M. Maxwell
Department of Geology and Geologic Engineering
Colorado School of Mines
rmaxwell@mines.edu

MS15
A Subdomain-based Parallel PCG Solver for the Cell Centered Finite Difference Groundwater Flow Equations using Incomplete Cholesky Preconditioning

An algorithm based in non-overlapping subdomains for solving sparse symmetric matrix systems that might result from the cell-centered finite difference (CCFD) algorithm has been developed for a multi-processor environment. This scheme is based in the PCG algorithm using incomplete Cholesky preconditioning with zero fill (IC(0)); parallelization is based in non-overlapping subdomains, whereby application over each subdomain corresponds to a computational process. As the IC(0) preconditioner does not lend itself easily to parallelization, certain compromises in parallelizations must be accommodated. The parallelization itself is instituted by means of standard Message-Passing Interface (MPI) functions. As a result, this parallel solver can be run on any Beowulf cluster or any multi-processor machine on which an implementation of MPI has been installed. The parallel PCG scheme with IC(0) preconditioning is compared with similar parallel PCG schemes where preconditioning is instituted by either block Jacobi or block Gauss-Seidel methods.

Richard L. Naff
U S Geological Survey
rlnaff@usgs.gov

MS15
Newton-GMRES Solvers for the Integrated Water Flow Model (IWFM) with Adaptive Stopping Criteria

We will discuss the use of inexact Newton method in conjunction with GMRES to solve systems of nonlinear equations in IWFM, a water resources management and planning model developed by California State Department of Water Resources. A strategy to adaptively control GMRES accuracy based on information of inexact Newton iterates is proposed. Numerical results show that the new strategy could save a significant number of GRMES iterations. Efficiency of several preconditioners will be also presented.

Hieu Nguyen
University of California, Davis
htrnguyen@ucdavis.edu

Zhaojun Bai
University of California
bai@cs.ucdavis.edu

Matthew F. Dixon
Department of Mathematics, UC Davis
mfdixon@ucdavis.edu

Charles Brush, Emin Dogrul, Tariq Kadir, Francis Chung
California State Department of Water Resources
cbrush@water.ca.gov, dogrul@water.ca.gov, kadir@water.ca.gov, chung@water.ca.gov

MS15
Improvement of Performance and Applicability of MODFLOW-2005: New NWT Solver and xMD Matrix Solver Package

The efficiency of Picard's method used by MODFLOW de-

grades when a model system is strongly nonlinear. MODFLOW has difficulties solving problems that involve larger hydraulic conductivity contrasts among its geological units. These difficulties are caused by a hard-to-solve matrix. The Newton method is combined with a new preconditioned conjugate-gradient type matrix solver to improve MODFLOW's performance. This talk will present results, including the models ability to provide a solution for a difficult unconfined groundwater-flow problem.

Rich Niswonger
US Geological Survey
tba@siam.org

Sorab Panday
AMEC Geomatrix Inc
sorab.panday@amec.com

Motom Ibaraki
The Ohio State University
ibaraki.1@osu.edu

MS16

A Fast Algorithm for the Time Harmonic Elastic Inverse Medium with Multiple Events

We consider the inversion of the 3-D time-harmonic elastodynamic equation in a lossy medium. In particular, we focus on broadband multi-point illumination problems for low frequency regimes. Such a problem finds many applications in geosciences. We use an integral-equation formulation for the forward problem and consider only small perturbations of the background medium (Born approximation). To solve this inverse problem we use a least squares formulation. If N_ω is the number of excitation frequencies, N_s the number of incoming waves, N_d the number of detectors, and N the discretization size for the scatterer, a dense SVD for the overall input-output map will have $[\min(N_s N_\omega N_d, N)]^2 \times \max(N_s N_\omega N_d, N)$ cost. We have developed a fast algorithm that brings the cost down to $O(NN_\omega N_s + NN_\omega N_d)$ thus, providing orders of magnitude improvements. Also, we propose an adaptive method in space to optimize the ratio accuracy versus number of degrees of freedom in space (the number of spatial discretization points).

Stephanie Chaillat
POems-ENSTA
stephanie.chaillat@ensta-paristech.fr

George Biros
Georgia Institute of Technology
biros@gatech.edu

MS16

Large Scale Inversion for Electromagnetic Wave Propagation Problems

In this talk we discuss the inversion of electromagnetic signals. In particular we discuss time domain EM inversion. We address the main computational bottlenecks and show how to reduce them.

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

MS16

Solving Time-Harmonic Inverse Medium Optimization Problems on the Cray Xe6 Supercomputer

We formulate the inverse medium problem as a PDE-constrained optimization problem, where the underlying wave field solves the time-harmonic Helmholtz equation. Ill-posedness is tackled through regularization while the inclusion of inequality constraints is used to encode prior knowledge. The resulting nonconvex optimization problem is solved by a primal-dual interior-point algorithm with inexact step computation. Numerical results both in two and three space dimensions on the Cray XE6 illustrate the usefulness of the approach.

Olaf Schenk
Department of Mathematics and Computer Science
University of Basel, Switzerland
olaf.schenk@unibas.ch

Grote Marcus, Johannes Huber
University of Basel
marcus.grote@unibas.ch, johannes.huber@uniabs.ch

MS16

A Collection of Computational Experiments in Parameter Estimation

Full Wavefield Seismic Inversion (FWI) represents a challenging parameter estimation problem. In this problem, we require the solution of an optimization problem with one PDE constraint for every seismic source. As a result, applying fast converging algorithms that require us to store second order information becomes challenging. In the following presentation, we present several computational experiments where we compare and contrast a variety of first and second order low memory algorithms applied to FWI.

Joseph Young, Scott Collis, Bart G. Van Bloemen Waanders
Sandia National Laboratories
josyoun@sandia.gov, sscoll@sandia.gov, bartv@sandia.gov

MS17

Basis Design for Polynomial Regression with Derivatives

We discuss polynomial regression with derivative (PRD) information; a method for approximating the stochastic response of a complex system. We demonstrate that the method needs less sampling information compared to its derivative-free version. Nevertheless, the method also poses other challenges in that the polynomial basis of choice for derivative-free methods is no longer suitable in this case. We propose a novel basis based on orthogonalization with respect to an inner product that involves gradient information. We demonstrate that the basis results in more accurate approximation at the same information level.

Yiou Li
Illinois Institute of Technology
yli88@stuart.iit.edu

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

Oleg Roderick
Argonne National Laboratory
roderick@mcs.anl.gov

MS17
Ensemble Emulators

Constructing Gaussian process/Kriging models requires the repeated inversion of an M by M matrix where M is the number of design points. For $M = \mathcal{O}(\infty\prime)$ points, the cost is $\mathcal{O}(\infty\prime^{\infty\forall})$ operations and the matrix is sure to be numerically singular. Using an ensemble of $\mathcal{O}(\infty\prime)$ small emulators made from $\mathcal{O}(\infty\prime\prime)$ points each reduces the cost to $\mathcal{O}(\infty\prime^{\infty\in})$ operations, avoids singularity issues, allows the ensemble emulator to represent non-stationary processes, and enables its concurrent construction and evaluation.

Keith Dalbey
Sandia National Laboratories
kdalbey@sandia.gov

Matthew Jones
Center for Comp. Research,
Univ. at Buffalo, Buffalo, NY 14260
jonesm@ccr.buffalo.edu

Abani Patra
Dept. of Mechanical and Aerospace Engineering
University at Buffalo
abani@buffalo.edu

Eliza Calder
University of Buffalo
ecalder@buffalo.edu

MS17
Countering the Curse of Dimensionality Using Higher-order Derivatives

Surrogate model approaches can reduce the computational cost for uncertainty analysis dramatically since their estimated function values can be used for an inexpensive Monte Carlo simulation. The surrogate model construction can be enhanced by using higher-order derivatives, whereby the information gain with higher dimensionality at reduced additional cost through adjoint methods can be exploited. Some uncertainty analysis examples applying these ideas will be given involving analytic test functions as well as computational fluid dynamics simulations.

Markus P. Rumpfkeil
University of Dayton
Markus.Rumpfkeil@udatyon.edu

Wataru Yamazaki
Nagaoka University of Technology
wyamazak@uwyo.edu

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

MS17
Improving the Scalability of Simplex Stochastic Collocation using Gradient-enhanced Response Surface Approximation

Simplex Stochastic Collocation (SSC) is an efficient method for the propagation of multiple aleatoric uncertainties based on an adaptive Delaunay triangulation of probability space and higher degree polynomial interpolation. SSC is here extended to include gradient-enhanced response surface approximation for improving the scalability to higher dimensional probability spaces using adjoint information. The computational complexity is further reduced by a pointwise construction of the response surface instead of the Delaunay triangulation.

Gianluca Iccarino
CTR, Stanford University, Stanford,
CA 94305-3035, USA
jops@ctr.stanford.edu

Jeroen Witteveen
Stanford University
Center for Turbulence Research
jasw@stanford.edu

MS18
Proper Generalized Decomposition based Dynamic Data-Driven Application

Dynamic Data-Driven Application Systems (DDDAS) appear as a new paradigm in the field of applied sciences and engineering, and in particular in simulation-based engineering sciences. By DDDAS we mean a set of techniques that allow the linkage of simulation tools with measurement devices for real-time control of systems and processes. DDDAS entails the ability to dynamically incorporate additional data into an executing application, and in reverse, the ability of an application to dynamically steer the measurement process. DDDAS needs for accurate and fast simulation tools making use if possible of off-line computations for limiting as much as possible the on-line computations. We could define efficient solvers by introducing all the sources of variability as extra-coordinates in order to solve off-line only once the model to obtain its most general solution to be then considered in on-line purposes. However, such models result defined in highly multidimensional spaces suffering the so-called curse of dimensionality. We proposed recently a technique, the Proper Generalized Decomposition (PGD), able to circumvent the redoubtable curse of dimensionality. The marriage of DDDAS concepts and tools and PGD off-line computations could open unimaginable possibilities in the field of dynamics data-driven application systems. In this work we explore some possibilities in the context of parameter estimation.

Francoise Masson
GEM - Centrale Nantes, France
f.masson@crossdata.it

Elias Cueto
I3A, University of Zaragoza, Spain
ecueto@unizar.es

Francisco Chinesta
Ecole Centrale de Nantes
EADS Corporate Foundation Internantional Chair
francisco.chinesta@ec-nantes.fr

David Gonzalez
I3A, University of Zaragoza, Spain
gonzal@unizar.es

MS18

Towards Optimal Interpolatory Model Reduction for Parameterized Systems

Optimal interpolatory model reduction has received great attention recently due to numerically effective methods for proving optimal point selection strategies. In this talk, after briefly reviewing the interpolation framework for parameterized systems, we show how to extend the concept of optimality to the parametric setting.

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

MS18

Applications of DEIM in Nonlinear Model Reduction

A dimension reduction method called Discrete Empirical Interpolation (DEIM) is described and shown to dramatically reduce the computational complexity of the popular Proper Orthogonal Decomposition (POD) method for constructing reduced-order models for parametrized nonlinear partial differential equations (PDEs). POD reduces dimension in the sense that far fewer variables are present, but the complexity of evaluating the nonlinear term remains that of the original problem. DEIM is a modification of POD that reduces complexity of the nonlinear term of the reduced model to a cost proportional to the number of reduced variables obtained by POD. The method applies to arbitrary systems of nonlinear ODEs, not just those arising from discretization of PDEs. Applications from Neural Modeling, Porous Media Flow, and Shape Optimization will be presented to illustrate the wide applicability of the DEIM approach.

Danny C. Sorensen
Rice University
sorensen@rice.edu

MS18

A Comparison of Reduced Basis Methods and POD for an Optimal Control Problem

In this contribution, a linear-quadratic optimal control problem governed by the Helmholtz equation is considered. For the computation of suboptimal solutions, two different model reduction techniques are compared: the reduced basis method (RBM) and proper orthogonal decomposition (POD). By an a-posteriori error estimator for the optimal control problem the accuracy of the suboptimal solutions is ensured. The efficiency of both model reduction approaches is illustrated by a numerical example for the stationary Helmholtz equation.

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

MS19

Isogeometric Analysis for Solids and Structures

A Reissner-Mindlin shell formulation based on a degenerated solid is implemented for NURBS-based isogeometric analysis. Its performance is examined on a set of elastic and nonlinear elasto-plastic benchmark examples. The analyses were performed with LS-DYNA, a general-purpose finite element code, for which a user-defined shell element capability was implemented. This new feature, to be reported on in subsequent work, allows for the use of NURBS and other non-standard discretizations in a sophisticated nonlinear analysis framework.

David J. Benson
Department of Structural Engineering
University of California, San Diego
dbenson@ucsd.edu

MS19

Finite Element Methods for Nonlocal Models of Diffusion and Mechanics

Abstract not available at time of publication.

Max Gunzburger
Florida State University
School for Computational Sciences
gunzburg@csit.fsu.edu

MS19

Nonlocal Continuum Balances in Continuum Mechanics

I review the nonlocal balances of linear momentum, angular momentum and energy. The nonlocal balances represent extensions to the classical balances and this relationship is given. I also highlight the current mathematical and phenomenological understanding of the nonlocal balances.

Richard B. Lehoucq
Sandia National Laboratories
rblehou@sandia.gov

MS19

Computational Peridynamics

Peridynamics is a nonlocal extension of classical continuum mechanics. Whereas classical continuum mechanics is governed by familiar partial differential equations, peridynamics is governed by an integro-differential equation. We discuss the impact of this nonlocal formulation upon the computational structure of the problem, reviewing discretization techniques, conditioning results, and solution methods. We also survey the state-of-the-art in computational peridynamics, discussing available codes and showing demonstration problems.

Michael L. Parks
Sandia National Laboratories
mlparks@sandia.gov

MS20

Sparse Bayesian Kernel Techniques for the Solution of SPDEs

The sparse grid method has been widely utilized in uncertainty propagation problems of SPDEs with well known limitations. As the need of alternative approaches is evident, we explore the performance of sparse Bayesian kernel techniques to the problem. We employ sophisticated model selection methods to adaptively identify the functional form of the kernels and estimate the scale parameters. The use of Bayesian variance enables us to select highly informative data points which minimizes the calls

to the FE solver.

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

Xinzeng Feng, Nicholas Zabaras
Cornell University
xf39@cornell.edu, zabaras@cornell.edu

MS20

Bayesian Inversion without Markov Chains

Bayesian inference provides a natural framework for quantifying uncertainty in PDE-constrained inverse problems, for fusing heterogeneous sources of information, and for conditioning successive predictions on data. In this setting, simulating from the posterior via Markov chain Monte Carlo (MCMC) constitutes a fundamental computational bottleneck. We present a new technique that entirely avoids MCMC by constructing functional approximations that enable rapid sampling and characterization of the posterior distribution. The approximations are implemented efficiently using optimization methods.

Tarek El Moselhy, Youssef M. Marzouk
Massachusetts Institute of Technology
tmoselhy@mit.edu, ymarz@mit.edu

MS20

Uncertainty Quantification given Discontinuities, Long-tailed Distributions, and Computationally Intensive Models

Conventional global spectral methods for uncertainty quantification are challenged when computer model predictions are discontinuous. On the other hand, local methods can be inefficient given computationally intensive models. The presence of fat-tail distributions in model predictions is also challenging, as excessive sampling can be prohibitive when forward models are computationally intensive. To circumvent these challenges we demonstrate a methodology that employs Bayesian inference to locate discontinuities in the model output, followed by efficient spectral propagation of uncertainty using domain mapping. We also illustrate how the use of tailored basis functions improves the convergence of spectral expansions in the tail regions.

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

MS20

An Application of Rare Event Tools to a Bimodal Ocean Current Model

I will review some recent results on importance sampling methods for rare event simulation as well as discuss their application to the filtering problem for a simple model of the Kuroshio current. Assuming a setting in which both observation noise and stochastic forcing are present but small, I will demonstrate numerically and theoretically that sophisticated importance sampling techniques can achieve very small statistical error. Standard filtering methods deteriorate in this small noise setting.

Jonathan Weare
NYU
weare@cims.nyu.edu

MS21

Thread Level Parallelism of Modern Architectures: New Degrees of Freedom for Mixed Parallel and Serial Processing Optimization of Algorithm Computation

The massively parallel SIMD architecture of the GPU combined with its ability to barrier synchronize computation has enabled a new area of loop optimization. In particular, nested loops with dependencies can be decomposed into different levels of granularity of parallelism and serial processing. Applying unimodular linear transformations to the dependence matrix that represents the order of computation can optimize certain classes of algorithms. Restating the problem in a dual forward path dynamic programming optimization can be used to discover algorithm optimization.

Dennis Braunreiter
SAIC
dennis.c.braunreiter@saic.com

MS21

Optimization of Cartesian Treecodes

Many choices go into designing a fast method for computing long-range particle interactions. Here we discuss a relatively simple particle-cluster treecode that uses Cartesian Taylor series to implement the far-field expansions. An update is provided on our efforts to optimize the serial and parallel performance of this algorithm. One application of special interest is the Poisson-Boltzmann model for solvated biomolecules.

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

Weihua Geng
Department of Mathematics
University of Michigan
geng@umich.edu

MS21

Parallel Rank-1 Updates of QR Factorizations

The QR factorizations of matrices under successive rank-1 updates are used in adaptive signal processing. The traditional approach is highly sequential. We review a recently introduced algorithm that enables efficient parallel computation, has lower complexity in total arithmetic operations, and bahaves better numerically. We extend the methodology to a few other array operations frequently used in

signal processing.

Xiaobai Sun
Department of Computer Science
Duke University
xiaobai@cs.duke.edu

MS21
Parallelizing the Fast Gauss Transform

Fast Gauss transform allows for calculation of the sum of N Gaussians at M points in O(N +M) time. We present new algorithms for its efficient parallelization. Computing the transform to six-digit accuracy at 120 billion points took approximately 140s using 4096 cores on the Jaguar supercomputer. Our algorithms can also be used for other "Gaussian-type" kernels and thereby form a new class of core computational machinery for solving parabolic PDEs on massively parallel architectures.

Shravan Veerapaneni
Courant Institute
New York University
shravan@cims.nyu.edu

Rahul Sampath
ORNL
sampathrs@ornl.gov

Hari Sundar
Siemens Corp.
hari.sundar@siemens.com

MS22
The General Curvilinear Environmental Model

General Curvilinear Environmental Model (GCEM) is a very high-resolution model composed of two sub models (1) General Curvilinear Atmosphere Model (GCAM) and (2) General Curvilinear Coastal Ocean Model (GCCOM). GCEM is written in fully 3D curvilinear coordinate and can work with both non-orthogonal and orthogonal grid in all three dimensions. It uses non-hydrostatic approach to solve for pressure and Large Eddy Simulation (LES) to solve primitive Navier-Stokes equations with Boussinesq approximations.

Mohammad Abouali
CSRC
San Diego State University
mabouali@sciences.sdsu.edu

MS22
Using the Distributed Coupling Toolkit (DCT) to Couple Model Components of Generalized Curvilinear Environmental Model

The General Curvilinear Environmental Model (GCEM) is an on-going modeling project to simulate a high resolution earth model system. Current GCEM components are a General Curvilinear Coastal Ocean Model (GCCOM) and a General Curvilinear Atmospheric Model (GCAM). Both models can run in parallel or Sequential computational environments. The Distributed Coupling Toolkit (DCT) is a library to couple multi-physics and multi-resolution models in a truly distributed manner. The DCT has a user-friendly interface to formulate the coupling of variables and fields within pairs of model components. DCT distributed approach guarantees scalability both at the model complexity and parallel processing levels. Here we use the DCT to weakly couple different components of GCEM. Also, we present some new capabilities implemented in DCT. Lastly, we show some preliminary performance results of DCT in a GCEM based application.

Dany De Cecchis
Computational Science Research Center
San Diego State University
dececchi@sciences.sdsu.edu or dcecchis@uc.edu.ve

Leroy A. Drummond
Lawrence Berkeley National Laboratory
Computational Research Division
ladrummond@lbl.gov

Jose E. Castillo
Computational Science Research Center
SDSU
castillo@myth.sdsu.edu

MS22
Data Asimilation Techniques for the GCEM

Abstract not available at time of publication.

Mariangel Garcia
CSRC
San Diego State University
mgarcia@sciences.sdsu.edu

MS22
A Cybeinfrastructure-based Computational Environment for the GCEM

The GCEM computational environment (CE) allows clients (human or application) to perform tasks including: running simulations across heterogeneous computing environments; hosting sub-models as services for applications; nesting sub-models within other models. Utilizing the SDSU Cyberinfrastructure Web Application Framework (CyberWeb), the CE provides middleware and backend services including: dynamic job creation, submission, history, tracking, and management; data migration and management; quick visualization of results; and resources and services management. In this talk, we describe the design, architecture and status of the project.

Mary Thomas
San Diego State University
mthomas@sciences.sdsu.edu

MS23
Accelerated Simulation of Analog Systems Using Linear and Nonlinear Robust Compact Models

One common approach for accelerating the simulation of complex analog systems is to replace subsystems of the circuit with simpler compact models. However, such procedures often result in unstable and non-physical behavior of the resulting system, particularly when employing nonlinear models. Using recently developed optimization-based model reduction techniques, it is possible to generate robust models, for both linear and nonlinear subsystems, that are guaranteed to be stable and passive, ensuring simula-

tion of the overall system remains well-behaved.

Brad Bond
Sandia National Labs
bbond@sandia.gov

MS23

A Nonlinear Timing/Phase Macromodeling Technique for General Circuits

We extend the concept of timing/phase macromodels, previously established rigorously only for oscillators, to apply to general circuits, both non-oscillatory and oscillatory. We derive a timing/phase macromodel via nonlinear perturbation analysis, and provide numerical methods to compute the macromodel. The macromodel that emerges is a scalar, nonlinear time-varying equation that accurately characterizes the system's phase/timing responses. We validate the technique on non-oscillatory circuits and show that our macromodel accurately captures the timing response.

Chenjie Gu, Jaijeet Roychowdhury
University of California, Berkeley
gcj@eecs.berkeley.edu, jr@eecs.berkeley.edu

MS23

Parallel Circuit Simulation Using Multi-Level Newton as a Graph Mitigation Strategy

Parallel SPICE-style circuit simulation is challenging for a number of reasons. Traditional circuit simulation involves implicitly solving a potentially large set of differential-algebraic equations (DAE's), which requires constructing and solving a linear system at each Newton step. Unfortunately, modern integrated circuit topologies frequently lead to matrix structures that are problematic for preconditioned iterative solvers. In a previous work, a combination of techniques, including singleton removal, Dulmage-Mendelsohn decomposition, and hypergraph partitioning has been shown to be an effective preconditioning strategy for some integrated circuits. However, there are circumstances where this strategy is ineffective, as it depends upon the Jacobian matrices having a particular structure. Circuit features that can break the expected structure include highly connected nodes such as non-ideal power supply and clock nodes (common in many integrated circuits), as well as feedback loops typical in phase-locked loops (PLL's). Through strategic application of multi-level Newton methods it is possible to remove problematic structures from the matrix and/or circuit graph, and enable the preconditioner to be effective.

Eric Keiter
Sandia National Laboratories
erkeite@sandia.gov

MS23

Efficient Preconditioners for Large-Scale Parallel Circuit Simulation

While direct linear solvers have long been regarded as a requirement for successful circuit simulation, the parallel transistor-level simulation of large-scale integrated circuits necessitates the use of iterative linear solvers. However, the linear systems generated through circuit simulation are challenging for conventional matrix ordering, load balancing, and preconditioning techniques. We will discuss the challenges presented by these linear systems, the current techniques used for preconditioning circuit matrices, and the ongoing work in developing efficient preconditioners for scalable circuit simulation using Xyce.

Heidi K. Thornquist
Sandia National Laboratories
hkthorn@sandia.gov

MS24

A Computational Quest for Quantum Subsystem Codes

Quantum error correction is important because it allows us to build quantum computers that work despite having noise in their components. Previous approaches for investigating quantum error correcting codes have focused on applying theoretical analysis to look for interesting codes and to investigate their properties. In this talk we present an alternative approach that uses computational analysis to accomplish the same goals. Specifically, we present an algorithm that computes the optimal quantum subsystem code that can be implemented given an arbitrary set of measurements that are tensor products of Pauli operators. We then demonstrate the utility of this algorithm by performing a systematic investigation of the quantum subsystem codes that exist in the setting where the measurements are limited to 2-body measurements between neighbors on lattices derived from the convex uniform tilings of the plane.

Gregory M. Crosswhite
University of Washington
gcross@phys.washington.edu

MS24

A Large Time Step and Low Communication Finite Volume Method for Atmospheric Simulation

Abstract not available at time of publication.

Matthew Norman
North Carolina State University
Raleigh, NC
matthew.ross.norman@gmail.com

MS24

Strategies for In-situ Analysis and Visualization in Large-scale Cosmological Simulation

Abstract not available at time of publication.

Paul Matthew Sutter
University of Illinois at Urbana-Champaign
psutter2@illinois.edu

MS25

Gradient and Hessian Consistency in Discontinuous Galerkin Solution of Inverse Wave Propagation Problems

Abstract not available at time of publication.

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

MS25

The Inverse Medium Problem in Site Characteri-

zation

I discuss recent progress in the full-waveform imaging of probed solids/soils, with site characterization applications in mind that typically involve arbitrarily heterogeneous domains, and mandate domain truncation via Pefectly-Matched-Layers (PMLs). I discuss a variational, symmetric, hybrid, non-convolutional, unsplit-field approach for wave simulations in the combined truncated domain and PML layer, and report on numerical experiments, including inversion attempts for the Marmousi profile.

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

MS25

On the Effect of Boundary Conditions in Seismic Full Waveform Inversion

In seismic full waveform inversion, the enormous size of the computational earth models that are of interest to the oil and gas industry makes solving this problem challenging, even with today's supercomputers. In order to reduce the computational domain of earth models to a manageable size, the forward simulation of wave propagation on the unbounded domain of the problem is truncated to a finite computational domain and some form of absorbing boundary conditions (ABCs) are applied. The truncated computational domain is typically a minimal volume spanned by the receiver spread at the earth surface. By doing so, however, one implicitly discards reflections that would come from near the outside of the truncated domain, caused by the inherent inhomogeneities of the earth's material properties. We present results from numerical experiments on the inversion of material properties from acoustic wave propagation data as obtained through typical geometries used in standard reflection seismology. The forward problem is discretized with discontinuous Galerkin finite elements, and we compare results obtained from a very large computational domain to these obtained on a trunctated computational domain with absorbing boundaries modeled either with the simple characteristic ABC, or with the state-of-the-art perfectly matched layer (PML). We conclude that while the PML conditions are clearly superior for modeling a purely absorbing boundary, its effect on the inversion does not always justify the increased computational cost. Our work suggests that the error introduced in the inversion by the use of a truncated domain that cannot account for the incoming reflections coming from outside the domain of interest and present in the signal, could often be larger than the residual reflections due to imperfectly absorbing boundaries.

Dimitar Trenev
IMA Univ of Minnesota
trenev@ima.umn.edu

Volkan Akcelik, Huseyin Denli, Alex Kanevsky, Kines K Patel, Laurent White, Martin-Daniel Lacasse
ExxonMobil
volkan.akcelik@exxonmobil.com,
huseyin.denli@exxonmobil.com,
alex.kanevsky@exxonmobil.com,
kinesh.k.patel@exxonmobil.com,
laurent.white@exxonmobil.com,
martin.d.lacasse@exxonmobil.com

MS25

Full Wave Form Inversion using Discontinuous Galerkin

We demonstrate the use of Discontinuous Galerkin (DG) to solve a linear least squares problem constrained by full wave form propagation. An adjoint provides sensitivity information which is combined with Nonlinear Conjugate Gradient and phase encoding to solve a large multi right-hand side problem. DG offers flexibility to provide added resolution in areas where the subsurface is complex. We demonstrate an initial implementation using a synthetic two dimensional dataset.

Bart G. Van Bloemen Waanders, Scott Collis, Curt Ober
Sandia National Laboratories
bartv@sandia.gov, sscoll@sandia.gov, ccober@sandia.gov

MS26

Real-Time Parametric Adaptation of Reduced-Order Models by Consistent Interpolation on a Manifold

The concept of parametric adaptation of reduced-order bases using a database and interpolation algorithms on manifolds is extended to linearized reduced-order models. Two steps are involved: (1) the transformation of the reduced-operators in consistent bases, and (2) the interpolation of the transformed operators on a suitable manifold, resulting in a fully on-line method. It is illustrated with applications to parametric studies of structural and aeroelastic systems in subsonic to supersonic regimes. Its ability to detect mode crossing and veering is highlighted.

David Amsallem, Charbel Farhat
Stanford University
amsallem@stanford.edu, CFarhat@stanford.edu

MS26

Parametric Adaptation of Reduced-Order Bases by Interpolation on a Manifold

Physics-based Reduced-Order Models can enable real-time computations but remain expensive to generate. They also generally lack robustness with respect to variations of underlying model parameters. To address these variations, a computational approach based on an offline database of reduced-order bases equipped with online algorithms for interpolation on appropriate manifolds is presented. The proposed approach incorporates a machine learning-based training algorithm. It is demonstrated here by application to the flutter analysis of an F-16 aircraft configuration.

Charbel Farhat, David Amsallem
Stanford University
CFarhat@stanford.edu, amsallem@stanford.edu

Julien Cortial
Stanford University
Institute for Computational and Mathematical Engineering
jcortial@stanford.edu

MS26

The Certified Reduced Basis Method for High-Fidelity Simulations in Real-Time Deployed Appli-

cations

We employ the certified reduced basis method to achieve high-fidelity numerical simulation of parametrized partial differential equations on "lightweight' deployed devices. The computational approach is divided into a computationally intensive Offline stage (performed on a supercomputer, for example) in which the reduced order model is generated, and a very inexpensive Online stage in which quantities of interest and rigorous error bounds are evaluated in real-time. This methodology allows us to consider "in the field' inverse and design problems and we present a number of examples from heat transfer, fluid mechanics and acoustics.

David J. Knezevic
MIT
dknez@mit.edu

MS26

Empirical Operator Interpolation for Reduced Basis Approximations of Nonlinear Evolution Equations

In this contribution we present a new approach to treat nonlinear operators in reduced basis approximations of parametrized evolution equations. The approach is based on empirical interpolation of nonlinear differential operators and their Frechet derivatives. Efficient online/offline decomposition is obtained for discrete operators that satisfy an H-independent DOF dependence for a certain set of interpolation functionals, where H denotes the dimension of the underlying high dimensional discretization space. The resulting reduced basis method is applied to nonlinear parabolic and hyperbolic equations based on explicit or implicit finite volume discretizations. We show that the resulting reduced scheme is able to capture the evolution of both smooth and discontinuous solutions. In case of symmetries of the problem, the approach realizes an automatic and intuitive space-compression or even space-dimensionality reduction. We perform empirical investigations of the error convergence and runtimes. In all cases we obtain a runtime acceleration of at least one order of magnitude.

Martin Drohmann
WWU Muenster
mdrohmann@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

MS27

Edge Functions – A Sound Basis

The reconstruction map $\mathcal{I}$ – also known as the Whitney map – plays an important role in mimetic/compatible discretizations, as described in [1]. $\mathcal{I}$ maps cochains onto differential forms and needs to satisfy the following criteria:

1. $\mathcal{I}$ is linear;
2. $\mathcal{R}\mathcal{I} \equiv Id$, the reconstruction map is the right inverse of the reduction map $\mathcal{R}$;
3. $\mathcal{I}\mathcal{R} = Id + O(h^p)$, the reconstruction map is an approximate left inverse of the reduction map;
4. $\mathcal{I}$ satisfies $\mathrm{d}\mathcal{I} = \mathcal{I}\delta$, where d is the exterior derivative acting on differential forms and δ is the coboundary operator acting on cochains.

This talk will address *edge functions* which are high order polynomials which satisfy the criteria for reconstruction maps listed above. All properties of the reconstruction map are retained when highly deformed curvilinear spectral elements are concerned, [2].

[1] Bochev and Hyman, *Principles of Mimetic Discretizations of Differential Operators*, Vol. 142 of the IMA Volumes in Mathematics and its Applications, Springer, Berlin, pp. 89-119, 2006 [2] Gerritsma, *Edge functions for spectral element methods*, In: Spectral and High Order Methods for Partial Differential Equations. Eds: Jan S. Hesthaven and Einar M. Rønquist, Lecture notes in Computational Science and Engineering, 76, Springer, pp. 199-208, 2010.

Mark Gerritsma
Delft University of Technology
m.i.gerritsma@tudelft.nl

MS27

Fast Finite Element Algorithms using Bernstein Polynomials

Bernstein polynomials on the d-simplex possess special structure that allows spectral-complexity matrix-free application of finite element operators, both constant and variable coefficients. I shall describe this structure and discuss application of these ideas for the de Rham complex of H(grad), H(curl) elements, and H(div).

Robert C. Kirby
Texas Tech University
robert.c.kirby@ttu.edu

MS27

A Locally Conservative, Discontinuous Least-Squares Finite Element Method for the Stokes Equations

Conventional least-squares finite element methods (LSFEMs) for incompressible flows do not lead to exact conservation of mass in the resulting approximation. In this talk we formulate a new, locally conservative LSFEM for the Stokes equations wherein a discontinuous velocity field is computed that is point-wise divergence free on each element. The effect of the new LSFEM approach on improved local and global mass conservation is compared with a conventional LSFEM for the Stokes equations employing standard C^0 elements.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

James Lai, Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
jhlai2@illinois.edu, lukeo@illinois.edu

MS27

Application of a Discontinuous Petrov-Galerkin Method to the Stokes Equations

The discontinuous Petrov-Galerkin finite element method proposed by L. Demkowicz and J. Gopalakrishnan guarantees the optimality of the solution in what they call the *energy norm*. An important choice that must be made in the application of the method is the definition of the inner product on the test space. In this work, we apply the DPG method to the Stokes problem in two dimensions, analyzing it to determine appropriate inner products, and perform a series of numerical experiments.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Nathan Roberts
University of Texas at Austin
nroberts@ices.utexas.edu

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

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

MS28

A Posteriori Error Estimates for Polynomial Chaos Expansions of Response Surfaces for Differential Equations

We develop computable a posteriori error estimates for linear functionals of a solution to a general nonlinear stochastic differential equation with random model/source parameters. These error estimates are based on a variational analysis applied to stochastic Galerkin methods for forward and adjoint problems. The result is a representation for the error estimate as a polynomial in the random model/source parameter. The advantage of this method is that we use polynomial chaos representations for the forward and adjoint systems to cheaply produce error estimates by simple evaluation of a polynomial. By comparison, the typical method of producing such estimates requires repeated forward/adjoint solves for each new choice of random parameter. We present numerical examples showing that there is excellent agreement between these methods.

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

Tim M. Wildey
The University of Texas at Austin
Austin, USA
twildey@ices.utexas.edu

MS28

Numerical Solutions of SPDES with Multiplicative Noise Forcing Terms

This talk focuses on the numerical solution of stochastic parabolic SPDES with multiplicative noise forcing terms. We first convert the SPDE to forward-backward doubly stochastic differential equations based on the theory of forward-backward stochastic differential equations (FBSDEs) and forward-backward doubly stochastic differential equations (FBDSDEs). We propose a new numerical algorithm, called binomial tree method, for solving FBDSDEs. Error analysis as well as numerical experiments will be presented. We shall demonstrate that our method is superior in comparison with standard finite difference method.

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

Feng Bao
Auburn University
fzb0005@auburn.edu

Weidong Zhao
shandong university
wzhao@yahoo.com

MS28

On the use of ANOVA Expansions in UQ

We consider the use of ANOVA expansions in the context of uncertainty quantification and parameter compression in high-dimensional problems. The discussion includes attention to both Lebesgue and Dirac ANOVA expansions and in the latter case we highlight a close connection between sparse grid integration and the optimal choice of the anchor point. The accuracy and efficiency of the proposed techniques are illustrated through examples of dynamical systems with high-dimensional parametric uncertainty.

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

Zhen Gao
Ocean University of China, Qiandao, China
gzwd9272@sina.com

MS28

Stochastic Elliptic Modeling based on the Wick Product

Based on the study of two commonly used stochastic elliptic models: I: $-\nabla\cdot(a(x,\omega)\cdot\nabla u(x,\omega)) = f(x)$ and II: $-\nabla\cdot(a(x,\omega)\diamond\nabla u(,\omega)) = f(x)$, we constructed a new Wick-type stochastic elliptic model III: $-\nabla\cdot\left(\left(a^{-1}\right)^{\diamond(-1)}\diamond\nabla u(x,\omega)\right) = f(x)$. The difference between models I and II is twofold: a scaling factor induced by the way of applying the Wick product and the regularization induced by the Wick product itself. We show that model III has the same scaling factor as model I, and present a detailed discussion about the difference between models I and III with respect to the two characteristic parameters of the random coefficient, i.e., the standard deviation σ and the correlation length l_c. Numerical results are presented

for both one- and two-dimensional cases.

Xiaoliang Wan
Louisiana State University
Department of Mathematics
xlwan@math.lsu.edu

MS29

Extending the Time Scales of Biomolecular Simulations on Emerging Computing Architectures

One of the significant challenges faced with molecular dynamics (MD) simulation is to extend its attainable time scale. There are two ways to do this: to make each sequential step run faster by using fast algorithms, and to improve the parallel efficiency so that more processors can be used. In this talk, I will discuss more scalable treatment of long-range electrostatic interactions in simulation using Ewald-mesh based explicit models and PB/GB based implicit models.

Xiaolin Cheng
Oak Ridge National Laboratory
chengx@ornl.gov

MS29

Accelerating Separable Tensor Operations with Throughput-Oriented Processors

Streamlining memory performance is a critical issue for utilizing the full computational bandwidth of throughput-oriented processors such as graphics cards, especially for memory-bound and high-dimensional applications. There are significant related scientific computing challenges, both within the context of the underlying architecture and at the application level. In this work we present a technique for improving the memory efficiency of separable tensor computations that matches application requirements against architectural constraints.

Christopher McKinlay
University of California Los Angeles
mckinlay@math.ucla.edu

MS29

The Parallel Full Approximation Scheme in Space and Time

I will discuss a strategy for the parallelization of numerical methods for partial differential equations in both the spatial and temporal directions. The method is based on an iterative multilevel approach whereby spectral deferred correction sweeps are applied to a hierarchy of discretizations at different spatial and temporal resolutions. Connections to the parareal algorithm and space-time multigrid methods will be discussed, and the parallel efficiency and speedup for three dimensional problems will be presented.

Michael Minion
University of North Carolina-Chapel Hill
Mathematics Department
minion@email.unc.edu

MS29

On the Parallelization of Non-Uniform Convolutions

The non-uniform convolution (NUCONV) refers to the discrete convolution with non-equally spaced data at input, or output, or both. This irregularity in sample locations complicates the data structure and the concurrent relation in parallel computation. We introduce a systematic methodology for parallel NUCONV in one to three dimensions, and present an auto-tuning library for NUCONV on multicore processors using Posix threads.

Nikos Pitsianis
Department of Electrical & Computer Engineering
Duke University
nikos.p.pitsianis@duke.edu

Xiaobai Sun
Department of Computer Science
Duke University
xiaobai@cs.duke.edu

MS30

Flipping Edges and Vertices in Graphs

We study a certain random process on a graph G which is a variation of a classical voter model and is also a special case of the so-called Tsetlin library random walk. Initially each vertex of G is colored either in blue or red. At each step an edge is chosen at random and both endpoints change their colors to blue with probability p and to red otherwise. This edge-flipping process corresponds with a random walk on the associated sta te graph in which each coloring configuration is a node. We show that the eigenvalues for the random walk on the state graph can be indexed by subsets of the vertex set of G. For example, for the uniform case of $p = 1/2$, for each subset T of the vertex set V of G, the eigenvalue λ_T (with multiplicity 1) is the ratio of the number of edges in the induced subgraph of T over the total number of edges in G. We analyze the stationary distribution of the state graph of colorings of G for several special families of graphs, such as paths, cycles and trees. We also mention related problems in connection with memoryless games.

Fan Chung-Graham
Department of Mathematics
University of California at San Diego
fan@ucsd.edu

Ron Graham
UCSD
graham@ucsd.edu

MS30

The Spectre of the Spectrum: Spectral Insights Across Graphs

In this talk, we present insights from analyzing a suite of common discrete graph metrics (connected components, clustering coefficients, cores, for example) across a range of real-world graphs and synthetic graph models. We also analyze spectral properties of the adjacency, Laplacian, and modularity matrices of these graphs in normalized and unnormalized versions. This includes finding complete spectra for graphs with hundreds of thousands of nodes.

David F. Gleich
Sandia National Labs
dgleich@stanford.edu

MS30

Eigenspace Analysis for Subgraph Detection

We describe statistical tests for anomalous subgraph detection and localization using spectral properties of a graph's modularity matrix. In addition to a Chi- squared test in the matrix's principal two-dimensional subspace, we use a test based on L1 properties of less significant eigenvectors to detect weaker anomalies. The techniques are extended to dynamic graphs and demonstrated on real and simulated data.

Nadya Bliss, Benjamin Miller
MIT Lincoln Laboratory
nt@ll.mit.edu, bamiller@ll.mit

Patrick J. Wolfe
Statistics and Information Sciences Laboratory
Harvard University
patrick@seas.harvard.edu

MS30

The Laplacian Paradigm: Emerging Algorithms for Massive Graphs

This presentation describes an emerging paradigm for the design of efficient algorithms for massive graphs. This paradigm, which we will refer to as the Laplacian Paradigm, is built on a recent suite of nearly-linear time primitives in spectral graph theory developed by Spielman and Teng, especially their solver for linear systems $Ax = b$, where A is the Laplacian matrix of a weighted, undirected n-vertex graph and b is an n-place vector. In the Laplacian Paradigm for solving a problem (on a massive graph), we reduce the optimization or computational problem to one or multiple linear algebraic problems that can be solved efficiently by applying the nearly-linear time Laplacian solver. So far, the Laplacian paradigm already has some successes. It has been applied to obtain nearly-linear-time algorithms for applications in semi-supervised learning, image process, web-spam detection, eigenvalue approximation, and for solving elliptic finite element systems. It has also been used to design faster algorithms for generalized lossy flow computation and for random sampling of spanning trees. The goal of this presentation is to encourage more researchers to consider the use of the Laplacian Paradigm to develop faster algorithms for solving fundamental problems in combinatorial optimization (e.g., the computation of matchings, flows and cuts), in scientific computing (e.g., spectral approximation), in machine learning and data analysis (such as for web-spam detection and social network analysis), and in other applications that involve massive graphs.

Shanghua Teng
University of Southern California
shanghua@usc.edu

Daniel Speilman
Yale University
spielman@cs.yale.edu

MS31

Integrating Efficiency Into Designing Performance Models

Characterizing the performance of scientific applications is essential for effective code optimization. Limited tools exist to support the model building process that involves extracting detailed information about the application, platforms, and their interactions. We have designed a suite of tools to automate the model building process, providing tools to design models, validate models, and measure model parameters. We will present an example in designing a performance model for selecting messaging schemes for an electromagnetic simulation.

Van Bui
University of Illinois at Urbana-Champaign
vanbui1@illinois.edu

Boyana Norris
Argonne National Laboratory
norris@mcs.anl.gov

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

MS31

Predictive Models of Memory Subsystem Performance for Sparse Scientific Computing

A large number of computational modeling and simulation applications rely on sparse data structures and algorithms. Such sparse computations are inherently scalable yet difficult to tune and adapt for energy-aware high-performance on modern multicore architectures, since the irregular and complex access patterns that classify these codes create many bottlenecks. Consequently, we develop memory-based interaction models between such applications and hardware to enable performance optimizations at the software and hardware layers of advanced computing systems.

Michael Frasca
The Pennsylvania State University
mrf218@cse.psu.edu

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

MS31

Performance Modeling for Systematic Performance Tuning

The performance of parallel scientific applications depends on many factors. Especially on large systems, it is too expensive to explore the solution space with a series of benchmarks. Analytical models allow estimating and extrapolating their execution performance, bottlenecks, and the potential impact of optimization options. We propose to use such performance modeling techniques from the beginning of the applications design. We will motivate the use of performance modeling with several examples.

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

Torsten Hoefler
University of Illinois
htor@illinois.edu

Marc Snir
Uniersity of Illinois
snir@illinois.edu

MS31

A Dynamic Runtime Optimization Framework for OpenUH's OpenMP Runtime

We present the design and implementation of an OpenMP dynamic optimization framework based on the OpenUH Collector Interface. This framework allows user-defined dynamic libraries to register callback functions for predefined events with the purpose of gathering runtime information and affecting runtime system behavior. Users are able to focus on implementing various optimizations without the needs to modify the program. Interesting applications of this framework related to performance modeling are discussed.

Besar Wicaksono, Ramachandra Nanjegowda, Brett Estrade, Barbara Chapman
University of Houston
besarwicaksono@gmail.com, cn.ramachandra@gmail.com, estrabd@cs.uh.edu, bchapman@uh.edu

MS32

Fast Higher Order FD like Schemes with Sobolev Type Norm Minimization

Fast and stable numerical discretization schemes based on a Sobolev like norm minimization are presented. These are used to construct Finite Difference like weights for differentiation and integration and used to construct higher-order PDE Solvers for Elliptic Boundary value problems in planar domains. Specific examples derived from engineering areas are considered over complex domains; the related stability and convergence issues are discussed. Performance of the schemes are compared to the FEM techniques in terms of convergence and computational speed.

Shivkumar Chandrasekaran
University of California, Santa Barbara
Department of Electrical and Computer Engieering
shiv@ece.ucsb.edu

Karthik Jayaraman Raghuram
Dept. of Electrical Engineering
University of California, Santa Barbara
jrk@ece.ucsb.edu

Joseph Moffitt
Dept. of Electrical and Computer Engineering
U C Santa Barbara
jmoffitt@umail.ucsb.edu

Hrushikesh Mhaskar
California State University, Los Angeles
Dept. of Mathematics
hmhaska@gmail.com

Ming Gu
University of California, Berkeley
Mathematics Department
mgu@math.berkeley.edu

MS32

Structured Low-Rank Matrix Recovery via Convex Optimization

In this talk, we discuss new algorithms for computing structured low-rank matrix approximations to a given matrix using convex optimization techniques. The structures we consider include Hankel matrices and banded matrices. We show how our techniques can be used to solve practical problems in signal processing applications.

Ming Gu
University of California, Berkeley
Mathematics Department
mgu@math.berkeley.edu

MS32

Existence of $\mathcal{H}$-matrix Approximants to the Inverse Finite-Element Matrix of Electrodynamic Problems and $\mathcal{H}$-Based Fast Direct Finite-Element Solvers

We prove that the inverse of the sparse matrix resulting from a finite-element-based analysis of electrodynamic problems has a data-sparse H-matrix approximation. We thus develop an $\mathcal{H}$-matrix-based direct finite-element-solver having $O(kN \log N)$ storage units and $O(k^2 N \log^2 N)$ operation counts for electromagnetic-analysis, where k is a parameter that is adaptively determined based on accuracy requirements. The complexity is further reduced to $O(M \log M)$ in storage and $O(Nlog^2 M)$ in time for layered-structures, where M is the number of single-layer unknowns.

Dan Jiao
Electrical and Computer Engineering
Purdue University
djiao@purdue.edu

Haixin Liu
Purdue University
Electrical and Computer Engineering
haixin@purdue.edu

MS32

Parallel Algorithms for Hierarchically Semi-separable Structures

Much progress has been made in fast algorithms for structured linear systems, such as those involving hierarchically semi-separable (HSS) matrices. Nearly linear time factorization algorithms have been developed to solve these systems. A key idea behind these algorithms is to fully exploit numerical low rankness in these structured matrices. In this talk, we present new parallel algorithms for HSS matrix operations and their use in the context of factorization-based sparse solvers and preconditioners.

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

Shen Wang
Purdue University
wang273@purdue.edu

Jianlin Xia
Department of Mathematics
Purdue University
xiaj@math.purdue.edu

MS32

Efficient Structured Solution of Large Sparse Linear Systems

We present some new techniques for the structured solution of sparse linear systems. They include new semiseparable structure generation under different circumstances, new structured factorizations, and their applications to sparse problems. The efficiency, robustness and scalability of this sparse solution are analyzed. Numerical examples in terms of PDEs and some real-world applications are shown.

Jianlin Xia
Department of Mathematics
Purdue University
xiaj@math.purdue.edu

Ming Gu
University of California, Berkeley
Mathematics Department
mgu@math.berkeley.edu

MS33

Cool@hpc

A non-standard numerical approach called COOL (Constraint Oriented Library) is presented. It enables to solve P.D.E. without generating non-physical solutions. External constraints such as $\nabla.\mathbf{u}= 0$, can handled by an elimination process, thus reducing the number of variables to the number of the physical problem. Operators including singularities in the spectrum as in Maxwells are well represented. It is shown how this method will be implemented to run on HPC machines.

Medji Azaiez
Universite de Bordeaux
azaiez@enscbp.fr

MS33

An Eigen-Based High Order Expansion Basis for Spectral Elements

We present an efficient high-order expansion basis for the spectral element approach. This belongs to the category of modal basis, but it is not hierarchical. The interior modes are constructed by solving a small generalized eigenvalue problem, while the boundary modes are constructed based on such eigen functions in lower dimensions. We compare this expansion basis with the commonly-used Jacobi polynomial-based expansion basis, and demonstrate the significantly superior numerical efficiency of the new basis in terms of conditioning and the number of iterations to convergence for iterative solvers.

Xiaoning Zheng, Suchuan Dong
Purdue University
zheng3@math.purdue.edu, sdong@math.purdue.edu

MS33

Radial Basis Functions for Planetary Scale Flows: Recent Developments

The talk will concentrate on the development of radial basis function methods for fluid modeling. Applications will be geared towards the geosciences. Recent advances in algorithm development to make the methodology more computationally effective will be discussed.

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

MS33

The Lower Bounds for Eigenvalues of Elliptic Operators - by Nonconforming Finite Element Methods

Abstract. We propose a condition and prove that it is sufficient to guarantee the nonconforming finite element methods to produce the lower bounds for eigenvalues of the symmetric elliptic operators. We show that this condition holds for the most used nonconforming elements, e.g., the Wilson element, the nonconforming linear element by Crouzeix and Raviart, the nonconforming rotated Q1 element by Rannacher and Turek, and the enriched nonconforming rotated Q1 element by Lin, Tobiska and Zhou for the second order elliptic operators, the Morley element, the Adini element and the enriched Adini element by Hu and Shi for the fourth order elliptic operators, and the Morley-Wang-Xu element for the 2m-th order elliptic operator. Whence they will give lower bounds for eigenvalues of these operators. Moreover, we follow the sufficient condition to propose two new classes of nonconforming elements for the second order elliptic operators and prove that they will yield the lower bounds for eigenvalues.

Yunqing Huang
Xiangtan University, China
Department of Mathematics
huangyq@xtu.edu.cn

MS33

New Efficient Spectral Methods for High-Dimensional PDEs

Many scientific, engineering and financial applications require solving high-dimensional PDEs. However, traditional tensor product based algorithms suffer from the so called "curse of dimensionality". We shall construct a new sparse spectral method for high-dimensional problems, and present, in particular, rigorous error estimates as well as efficient numerical algorithms for some typical PDEs.

Jie Shen
Department of Mathematics
Purdue University
shen@math.purdue.edu

MS34

Comparison of Search Strategies in Empirical Performance Tuning of Linear Algebra Kernels

We have developed an empirical tuning system, Orio, which is aimed at improving both application performance and developer productivity by generating and evaluating multiple optimizations based on developer annotations of key computations. The size of the search space that Orio explores empirically depends exponentially on the number of performance parameters, making exhaustive search impractical in most cases. We compare different search strategies in the context of tuning several linear algebra kernels on multiple platforms.

Thomas Nelson
University of Colorado at Boulder

Argonne National Laboratory
thomas.nelson@colorado.edu

Qian Zhu, Boyana Norris, Prasanna Balaprakash
Argonne National Laboratory
qzhu@mcs.anl.gov, norris@mcs.anl.gov, tba@anl.gov

MS34
Automatic Performance Tuning with BTO BLAS

Current hardware trends require that basic linear algebra routines take advantage of multiple cores. The BTO tool fuses multiple linear algebra operations and automatically recognizes opportunities for parallelism as it generates basic linear algebra routines. This talk describes how the BTO compiler extracts parallelism and shows that such parallelism can achieve further speedups for fused routines.

Jeremy Siek
Department of Electrical and Computer Engineering
University of Colorado at Boulder
jeremy.siek@colorado.edu

Geoffrey Belter
Dept. of Electrical, Computer, and Energy Engineering
University of Colorado at Boulder
geoffrey.belter@colorado.edu

MS34
Diagnosis, Tuning, and Redesign for Multicore

We describe how we improved the within-node scalability of the fast multipole method (FMM) through a systematic sequence of modeling, analysis, and tuning steps. On a quad-socket Intel Nehalem-EX system, we show speedups of of 1.7 over the previous best multithreaded implementation, 19.3 over a sequential but highly tuned (e.g., SIMD-vectorized) code, and match or outperform a GPGPU implementation. Our study suggests a more general tuning process that practitioners and tools could themselves apply.

Aparna Chandramowlishwaran, Richard Vuduc
Georgia Institute of Technology
aparna@cc.gatech.edu, richie@cc. gatech. edu

Kamesh Madduri
Lawrence Berkeley National Lab
kmadduri@lbl.gov

MS34
Modeling Loop Fusion on Shared Memory Parallel Architectures

The performance of many scientific applications is most accurately expressed in terms of data movement. Loop fusion is a technique that can decrease or increase data movement. In this talk we describe a model that accurately captures large performance effects of fusing loops on shared memory parallel architectures. We show how the model economically compares different implementations of the same routine with varying amount of loop fusion within a compiler framework without sacrificing kernel efficiency.

Ian Karlin
Department of Computer Science
University of Colorado at Boulder
ian.karlin@colorado.edu

Elizabeth Jessup
University of Colorado at Boulder
Department of Computer Science
jessup@cs.colorado.edu

MS35
Presenters To Be Announced

Abstract not available at time of publication.

TBD A1
TBD
neckel@in.tum.de

MS36
Sparse Interpolatory Models for Molecular Potential Energy Surfaces

We describe a parallel algorithm for generating interpolatory approximations to molecular potential energy surfaces. We show how that algorithm can be applied to efficiently model a transition from a stable ground state, to an excited state, and finally to a different stable ground state.

Carl T. Kelley
North Carolina State Univ
Department of Mathematics
tim_kelley@ncsu.edu

Alexei Bykhovski
Dept of Electrical and Computer Engineering
NC State University
abykhov@ncsu.edu

David Mokrauer
Dept of Mathematics
NC State University
dsmokrau@ncsu.edu

MS36
Optimal Placement of Gamma-monitoring Beacons using the Mesh Adaptive Direct Search Algorithm

The deployment of gamma-monitoring (GMON) beacons is considered by the Research Institute of Hydro-Quebec in order to improve the snowpack estimate accuracy in order to manage the hydrological forecast throughout the year, especially at critical times such as spring snowmelt. The placement of these beacons is critical and it may be seen as an optimization problem in which the GMON locations are the variables and where the objective is to minimize the approximation error. Map constraints as well as the error computation categorize this problem as a difficult blackbox problem for which the Mesh Adaptive Direct Search (MADS) is designed. In addition, the fact that the variables correspond to two-dimensional locations of objects suggests grouping variables in order to decompose the problem. Different grouping and regrouping strategies are developed and supported by a convergence analysis.

Sebastien Le Digabel
GERAD - Polytechnique Montreal
Sebastien.Le.Digabel@gerad.ca

Stephane Alarie

IREQ
alarie.stephane@ireq.ca

Charles Audet
École Polytechnique de Montréal - GERAD
Charles.Audet@gerad.ca

Vincent Garnier
Polytechnique Montreal
vincent.garnier@polymtl.ca

Louis-Alexandre Leclaire
IREQ
leclaire.louis-alexandre@ireq.ca

MS36

Robust Design for Cardiovascular Surgery Applications using the Surrogate Management Framework

Recent work has demonstrated substantial progress in patient-specific cardiovascular flow simulations, and a need for customization of treatment plans for individual patients. We present a unified framework for derivative-free optimization of cardiovascular geometries with pulsatile flow that incorporates uncertainties. Optimization is performed using the surrogate management framework with mesh adaptive direct search. We incorporate uncertainties using adaptive stochastic collocation to perform robust design. We apply these tools to several pediatric and adult cardiovascular surgery applications.

Alison Marsden
Department of Mechanical and Aerospace Engineering
University of California, San Diego
amarsden@ucsd.edu

MS36

Applied Optimization Without Derivatives

Through a set of diverse applications, we provide an overview of challenges faced when optimizing functions of complex simulations. These difficulties are compounded by the computational expense of many simulations and include computational noise, constraints without relaxations, mixtures of continuous and discrete variables, and availability of some derivatives. These examples, and the applications presented throughout the minisymposium, illustrate the potential for methods in this area to make significant contributions to computational science and engineering.

Stefan Wild, Jorge More'
Argonne National Laboratory
wild@mcs.anl.gov, more@mcs.anl.gov

MS37

Stability of Methods for Quasiseparable Matrices

In this talk we discuss computation of the matrix–vector product for a quasiseparable matrix, a computation at the heart of many fast algorithms involving the class. The usual parametrization of a quasiseparable matrix can introduce errors, even when algorithms are structured backward stable. We derive a nested factorization involving Householders to give a fast, stable algorithm. This talk is based on joint work with Vadim Olshevsky and Michael Stewart.

Tom Bella
University of Rhode Island
Department of Mathematics
tombella@math.uri.edu

Vadim Olshevsky
University of Connecticut
Department of Mathematics
olshevsky@uconn.edu

Michael Stewart
Georgia State University
Department of Mathematics
mastewart@gsu.edu

MS37

Polynomial Representations for Matrices: How they Work and What One can do with Them

An input vector presented to a matrix can be viewed as a sequence of data that is processed in order of appearance, producing a similarly linearly ordered output vector. Block-diagonal matrices can then be viewed as instantaneous operators producing an output at the same index number as the input. In a further step, a matrix can be viewed as a collection of shifted (block-)diagonals. A banded matrix is then polynomial in the shift operator. The question arises whether there exist (minimal) rational representations for matrices in a sense that parallels rational representations for functions in a variable, whether general matrices can be so approximated, whether there exist banded matrices with banded inverses and whether interpolation theories can be derived to obtain low complexity approximations of general matrices. The results are in many ways surprising as most classical theories have matricial counterparts. The theory produces new representations for the celebrated semi-separable matrices and a new vista on classical interpolation problems s.a. Lwner and Schur-Takagi interpolation.

Patrick Dewilde
Institute for Advanced Study, Technische Universität München
and CAS Section, EEMCS, TU Delft
p.dewilde@ewi.tudelft.nl

MS37

A Structured Linear Algebra Problem in Space Imaging

Resolution quality of image deblurring algorithms depends on accurate knowledge of the blurring operator, or point spread function (PSF). When imaging objects in space using telescopes, a wavefront sensor (WFS) is used to estimate the PSF. However the estimation quality is limited because the WFS collects low resolution measurements. We describe a new approach to obtain high resolution PSFs from multiple low resolution WFS measurements. Efficiency is obtained by exploiting structured and sparse matrix computations.

James G. Nagy
Emory University
Department of Math and Computer Science
nagy@mathcs.emory.edu

MS37

Orthogonal Methods for the Solution of Quasisep-

arable Systems

This talk describes a nested product decomposition of a quasiseparable matrix A. The decomposition represents A as a nested product of Householders and a sequence of very sparse matrices matrices that each have only a few nonzeros. The cost of computing the decomposition from an $n \times n$ matrix with off-diagonal ranks bounded by r is $O(r^2n^2)$. The cost of using the decomposition to solve a system of equations or multiply a vector by A is $O(rn)$. The decomposition and the system solver are both backward stable and a proof of this fact is briefly summarized in the talk.

Michael Stewart
Georgia State University
Department of Mathematics
mastewart@gsu.edu

MS38

Modelling, Algorithm and Simulation of Wave Motion in Quantum and Plasma Physics

In this talk, I begin with a review of several mathematical models for describing wave motion in quantum and plasma physics. Computational difficulties for simulating wave propagation and interaction in quantum and plasma physics are discussed. Efficient and accurate numerical algorithms for computing ground and excited states as well as the dynamics of the nonlinear Schroedinger equation are presented. Extensive simulation results of wave propagation and interaction in quantum and plasma physics are reported. Finally, some conclusions are drawn.

Weizhu Bao
National University of Singapore
Department of Mathematics
bao@math.nus.edu.sg

MS38

Optimal Error Estimates of Finite Difference Methods for the Gross-Pitaevskii Equation with Angular Momentum Rotation

We analyze finite difference methods for the Gross-Pitaevskii equation with an angular momentum rotation term in two and three dimensions and obtain the optimal convergence rate, for the conservative Crank-Nicolson finite difference (CNFD) method and semi-implicit finite difference (SIFD) method, at the order of $O(h^2 + \tau^2)$ in the l^2-norm and discrete H^1-norm with time step τ and mesh size h. Besides the standard techniques of the energy method, the key technique in the analysis for the SIFD method is to use the mathematical induction, and resp., for the CNFD method is to obtain *a priori* bound of the numerical solution in the l^∞-norm by using the inverse inequality and the l^2-norm error estimate.

Yongyong Cai
National University of Singapore
caiyongyong@nus.edu.sg

MS38

Title Not Available at Time of Publication

Abstract not available at time of publication.

Laurent Di Menza
University Paris 11
laurent.dimenza@math.u-psud.fr

MS38

Critical Angular Velocities of Vortex Nucleation in Rotating Bose-Einstein Condensation

Recently, nucleation of vortices in rotating Bose-Einstein condensates (BEC) has been the subject of intensive experimental and theoretical research. For a two-dimensional rotating BEC, we numerically study the critical angular velocities of vortex nucleation. The Gross-Pitaevskii equation in a rotating frame is used as the mathematical model, which is solved by using the time-splitting Fourier spectral method. Our numerical results agree with the theoretical predictions in the literature.

Yanzhi Zhang
Missuri University of Science and Technology
zhangyanz@mst.edu

MS39

Structure Preserving Interpolatory Model Reduction

Dynamical systems are the basic framework for modeling and control of an enormous variety of complex systems. Rational Krylov methods are often capable of providing nearly optimal approximating subspaces for efficient reduction of large-scale, complex dynamical systems. An interpolation framework for model reduction is presented that includes rational Krylov methods as a special case. This broader framework allows retention of special structure in the reduced order models that is often encoded in the system parameterization such as symmetry, internal delays, and port-Hamiltonian structure.

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

MS39

Structure-preserving Model Reduction for MEMS

Simulations of damped vibration in micro-electro-mechanical systems (MEMS) typically lead to non-Hermitian problems which depend nonlinearly on a frequency parameter. In this talk, we discuss model reduction methods which use three types of structure present in the full discrete model: algebraic structure, such as complex symmetry of the system matrices, that is inherited from the underlying PDEs; geometric structures such as symmetry groups or the presence of beam-like or plate-like components; and perturbative structure that arises from disparate time scales in the physics. We illustrate our methods with models of a several resonant microstructures.

David Bindel
Cornell University
bindel@cs.cornell.edu

MS39

Model Order Reduction for RC Networks with Many Terminals

A novel approach for reducing multi-terminal RC circuits is proposed. Using graph-based circuit partitionings in combination with state-of-the-art model reduction methods, structure and sparsity in the reduced model are enhanced.

The reduced models are easily converted to a circuit representation. These contain many fewer nodes and circuit elements than the original circuit, allowing faster simulations at little accuracy loss.

Roxana Ionutiu
TU Eindhoven
r.ionutiu@gmail.com

Joost Rommes
NXP Semiconductors, The Netherlands
joost.rommes@nxp.com

Wil Schilders
Technical University of Eindhoven
w.h.a.schilders@TUE.nl

MS39

Model Reduction and Vertex Cuts

We report on an algorithm for model reduction that relies heavily on the computation of small cardinality vertex cuts in a graph representing resistances between the external nodes of a circuit. The approach we propose is unified and conceptually simple. Furthermore our results are at least as good as those previously reported in the literature.

Odile Marcotte
CRM and UQAM
odile.marcotte@gerad.ca

Suzanne M. Shontz
The Pennsylvania State University
shontz@cse.psu.edu

Wil Schilders
Technical University of Eindhoven
w.h.a.schilders@TUE.nl

Petko Kitanov
University of Guelph
pkitanov@uoguelph.ca

MS40

Reduced Basis Method for Radar Cross Section Computation of a Pacman Scattering Problem

We consider the scattering of TM-polarized electromagnetic waves by a perfectly conducting 2D cylinder with a cut-out wedge. The parameter of this problem is the angle of the wedge. An interesting phenomenon of this problem is that the fields, and thus the bistatic radar cross section (RCS), change dramatically with only a small change in the wedge angle. Reduced basis method (RBM) is applied to this nonlinear problem. An extensive test of the algorithm shows exponential convergence of the RB solutions with roughly constant effectivity index. A study of the monostatic scattering as a function of the wedge angle reveals the capability of RBM to capture the critical wedge angle for the optimal reduction of the backscatter.

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

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

Yvon Maday
Université Pierre et Marie Curie
and Brown university
maday@ann.jussieu.fr

MS40

Reduced Basis A Posteriori Error Bounds for Linear Quadratic Optimal Control Problems

We present a posteriori error bounds for the solution of linear quadratic optimal control problems using reduced basis methods. We consider error bounds for the optimal control input and the corresponding output of interest as well as for the error in the optimal cost function. Our bounds are online efficient, i.e., the computation cost is independent of the underlying truth approximation, and are provable upper bounds for the quantities of interest.

Martin Grepl, Daniel Amat, Mark Kaercher
RWTH Aachen University
grepl@igpm.rwth-aachen.de, amat@igpm.rwth-aachen.de, mark.kaercher@web.de

MS40

Reduced Basis and Freeform Deformation for Shape Optimization

We present a new approach for shape optimization that combines two different types of model reduction: a suitable low-dimensional parametrization of the geometry (yielding a geometrical reduction) combined with reduced basis methods (yielding a reduction of computational complexity). More precisely, free-form deformation techniques are introduced for the geometry description and its parametrization, while reduced basis methods are used upon a finite element discretization to solve systems of parametrized partial differential equations. This allows an efficient flow field computation and cost functional evaluation during the iterative optimization procedure, resulting in effective computational savings with respect to usual shape optimization strategies. This approach is very general and can be applied for a broad variety of problems. To prove its effectivity, we apply it to find the optimal shape of aorto-coronaric bypass anastomoses based on vorticity minimization in the down-field region.

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

Gianluigi Rozza
Mathematics Institute of Computational Science and Eng, EPFL
Switzerland
gianluigi.rozza@epfl.ch

Andrea Manzoni
Ecole Pol. Fed. de Lausanne
MATHICSE-CMCS
andrea.manzoni@epfl.ch

MS40

An "Uncertainty Region" Reduced Basis Approach

to Parameter Estimation

We present a novel reduced basis approach for the solution of parameter estimation problems. The method allows efficient calculation of the "uncertainty region' in parameter space, where the uncertainty region is the set of all parameter values consistent with measurements. The method exploits the reduced basis method to rapidly solve the underlying parametrized PDEs. The computed uncertainty region incorporates (reduced-order) modelling and experimental error, needs no regularization, and can handle nonconvex and non-simply connected domains.

Karen Veroy
AICES
RWTH Aachen University
veroy@aices.rwth-aachen.de

Martin Grepl
RWTH Aachen University
grepl@igpm.rwth-aachen.de

MS41

On the Application of Estimation Theory to Complex System Design Under Uncertainty

Complex system design is a multifidelity multidisciplinary task that often begins with large initial uncertainties in all aspects of the system. We present a method for systematically reducing these uncertainties by casting the design process as a discovery procedure built on the tools of estimation theory. Further, we incorporate methods of global sensitivity analysis to manage the multifidelity multidisciplinary nature of the system throughout the process.

Doug Allaire, Karen Willcox
Massachusetts Institute of Technology
dallaire@mit.edu, kwillcox@mit.edu

MS41

Advanced Dynamically Adaptive Algorithms for Stochastic Simulations on Extreme Scales

In spite of the recent and continuing growth in computational resources available to the scientific community, full-scale resolution of most complex stochastic models is completely intractable. The situation becomes even more serious because of the need to conduct uncertainty quantification. The traditional deterministic models of complex systems are now stochastic models in higher dimensional spaces, whose dimensionality is determined by a potentially very large set of random variables. We discuss the set of challenges needed to overcome in order to perform advanced numerical algorithms for modeling uncertainty in complex stochastic systems at extreme scales.

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

Ralf Deiterding
Oak Ridge National Laboratory
Computer Science and Mathematics Division
deiterdingr@ornl.gov

Dongbin Xiu
Purdue University
dxiu@purdue.edu

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

MS41

Fast Generation of Nested Space-filling Latin Hypercube Sample Designs

When "binning optimal" sample designs are recursively divided into uniform grids of cube bins, every same-sized bin contains the same number of points until the bins are small enough to all contain either 1 or 0 points. Binning optimal designs are space-filling. We construct sequences of Latin Hypercube Sample designs in which subsequent designs inherit all points from previous designs. These designs are binning optimal in the highest, lowest and certain intermediate subsets of dimensions.

Keith Dalbey
Sandia National Laboratories
kdalbey@sandia.gov

George Karystinos
Department of Electronic and Computer Engineering
Technical University of Crete
karystinos@telecom.tuc.gr

MS41

Optimal Uncertainty Quantification

We propose a rigorous framework for Uncertainty Quantification (UQ) in which the UQ objectives and the assumptions/information set are brought to the forefront. This framework, which we call *Optimal Uncertainty Quantification* (OUQ), is based on the observation that, given a set of assumptions and information about the problem, there exist optimal bounds on uncertainties: these are obtained as extreme values of well-defined optimization problems corresponding to extremizing probabilities of failure, or of deviations, subject to the constraints imposed by the scenarios compatible with the assumptions and information. In particular, this framework does not implicitly impose inappropriate assumptions, nor does it repudiate relevant information. Although OUQ optimization problems are extremely large, we show that under general conditions, they have finite-dimensional reductions. As an application, we develop *Optimal Concentration Inequalities* of Hoeffding and McDiarmid type. Surprisingly, contrary to the classical sensitivity analysis paradigm, these results show that uncertainties in input parameters do not necessarily propagate to output uncertainties. In addition, a general algorithmic framework is developed for OUQ and is tested on the Caltech surrogate model for hypervelocity impact, suggesting the feasibility of the framework for important complex systems. This is a joint work with C. Scovel, T. Sullivan, M. McKerns and M. Ortiz. A preprint is available at http://arxiv.org/abs/1009.0679v1

Houman Owhadi
Applied Mathematics
Caltech
owhadi@caltech.edu

Clint Scovel
Los Alamos National Laboratory
jcs@lanl.gov

Timothy J. Sullivan, Michael McKerns, Michael Ortiz

California Institute of Technology
tjs@caltech.edu, mmckerns@caltech.edu, ortiz@aero.caltech.edu

MS42

The Method of Regularized Stokeslets with Applications to Biological Flows

Biological flows, such as those surrounding swimming microorganisms or beating cilia, are often modeled using the Stokes equations due to the small length scales. The organism surfaces can be viewed as flexible interfaces imparting force on the fluid. I will present the latest results on the Method of Regularized Stokeslets and other elements that are used to compute Stokes flows interacting with immersed flexible bodies or moving through obstacles. The method treats the flexible bodies as sources of force or torque in the equations and the resulting velocity is the superposition of flows due to all the elements. A set of images is used to compute flows bounded by a plane. Exact flows are derived for forces that are smooth but supported in small spheres, rather than point forces. I will present the idea of the method, some of the known results and several examples from biological applications.

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

MS42

Front Tracking Method on Fluid Structure Interaction

Front tracking is a Lagrangian tool for the propagation of material interface and is known for its excellent preservation of interface geometry. In this presentation, we will discuss how to use this software tool to track the motion of rigid body structure and airfoil. We will present the simulation and comparison with experiments of two interesting problems, the windmill power generator and the inflation of 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@ams.sunysb.edu, yli@ams.sunysb.edu

MS42

Immersed Finite Element Methods for Interface Problems

This presentation will give a quick survey on the research of immersed finite element methods for solving interface problems. In engineering and sciences, many simulations need to be carried out over domains consisting multiple homogeneous materials separated from each other by curves or surfaces. This often leads to the so called interface problems of partial differential equations whose coefficients are piecewise constants. Traditional finite element methods can be used to solve interface problems satisfactorily with meshes constructed according to the material interfaces; otherwise, convergence cannot be guaranteed. This requires each element in a mesh to be occupied essentially by one of the materials. In other words, each element needs to be on one side of a material interface. Therefore, the mesh in a traditional finite element method for solving an interface problem has to be unstructured to handle non-trivial interface configurations. This restriction usually causes a substantial negative impact on the simulation cost for applications whose material interfaces are not fixed because the involved simulation domain has to be remeshed repeatedly whenever the interfaces are adjusted. In this talk, we will discuss a class of recently developed immersed finite element (IFE) methods intended to alleviate this limitation of traditional finite element methods. The basic idea of the IFE methods will be described through particular applications and numerical examples will also be presented to illustrate features of IFE methods.

Tao Lin
Department of Mathematics, Virginia Tech
tlin@math.vt.edu

MS42

Numerical Methods for Elliptic Systems with Moving Interfaces

Linear constant-coefficient elliptic systems of partial differential equations occur frequently in computational science, and challenge standard solution methodologies. We present new "locally-corrected spectral boundary integral methods' for their accurate discretization and fast solution. Arbitrary elliptic systems such as Poisson, Yukawa, Helmholtz, Maxwell, Stokes and elasticity equations are transformed to an overdetermined first-order form amenable to unified solution. A simple well-conditioned boundary integral equation, for solutions satisfying arbitrary boundary conditions posed on complex interfaces, is derived from first principles. A fast Ewald summation formula for the periodic fundamental solution is derived by Fourier analysis, linear algebra, and local asymptotic expansion. Ewald summation evaluates box, volume and layer potentials efficiently, and separates the boundary integral equation into a low-rank system with regular spectral structure, followed by a simple local correction formula. A new "geometric nonuniform fast Fourier transform' (GNUFFT) produces accurate Fourier coefficients of discontinuous piecewise-polynomial data on a d-dimensional simplicial tessellation in R^D, for arbitrary dimensions d and D, and costs as little as a few standard FFTs.

John A. Strain
Department of Mathematics
University of California at Berkeley
strain@math.berkeley.edu

MS43

Multilevel Aggregation Methods for Small-World Graphs with Application to Random-Walk Ranking

We describe multilevel aggregation in the context of using Markov chains to rank the nodes of graphs. Aggregation successfully generates efficient multilevel methods for various eigenproblems from discretized PDEs, which involve mesh-like graphs. Our goal is to extend the applicability to similar problems on small-world graphs. For a class of small-world model problems, we show how multilevel hierarchies formed with non-overlapping aggregation are efficiently employed to accelerate convergence of methods that

calculate the ranking vector.

Hans De Sterck
University of Waterloo
Applied Mathematics
hdesterck@uwaterloo.ca

Van Henson
Lawrence Livermore National Laboratory
henson5@llnl.gov

Geoffrey D. Sanders
Center for Applied Scientific Computing
Lawrence Livermore National Lab
sanders29@llnl.gov

MS43

Directed Graph Embedding: Continous Limit of Laplacian-based Operators

We consider the problem of embedding directed graphs in Euclidean space while retaining directional information. We view a directed graph as a finite sample from a diffusion process on a manifold endowed with a vector field. The algorithms we design separate and recover the features of this process: the geometry of the manifold, the data density and the vector field. The application of our method to both artificially constructed and real data highlights its strengths.

Marina Meila
Statistics
University of Washington
mmp@stat.washington.edu

Dominique Perrault-Joncas
Department of Statistics
University of Washington
dcpj@stat.washinton.edu

MS43

Algebraic Multigrid for Spectral Calculations on Large Complex Networks

Approximations to the expected commute time between any two vertices within a graph are quickly obtained using an accurate set of eigenpairs corresponding to the lower eigenvalues of the graph Laplacian matrix. This and other examples motivate the development of scalable eigensolvers for matrices associated with scale-free graphs. We investigate the use of Lanczos iteration on such matrices and develop a multilevel restart method, which is based on aggregation multigrid, to accelerate the Lanczos convergence.

Geoffrey D. Sanders
Center for Applied Scientific Computing
Lawrence Livermore National Lab
sanders29@llnl.gov

Panayot Vassilevski, Van Henson
Lawrence Livermore National Laboratory
vassilevski1@llnl.gov, henson5@llnl.gov

MS43

Counting Triangles in Real-World Networks

The number of triangles is a computationally expensive graph statistic which is frequently used in complex network analysis (e.g., transitivity ratio), in various random graph models (e.g., exponential random graph model) and in important real world applications such as spam detection, uncovering of the hidden thematic structure of the Web and link recommendation. In this talk we will present a family of spectral algorithms for triangle counting which perform efficiently on 'real-world' networks, due to the special spectral properties which they typically possess. Finally we will present Triangle Sparsifiers, i.e., sparse graphs which approximate the original graph G with respect to the count of triangles within a factor of ϵ with high probability as long as G contains $\tilde{\Omega}(n)$ triangles. Triangle Sparsifiers allow us to justify significant expected speedups (e.g., if G has $O(n^{1.5+\epsilon})$ triangles, we obtain in expectation a speedup of $O(n)$) and obtain excellent running times for large networks in practice. Finally, we will provide potential research directions.

Charalampos Tsourakakis
CARNEGIE MELLON UNIVERSITY
tsourolampis@gmail.com

Mihalis Kolountzakis
Department of Mathematics
University of Crete
kolount@gmail.com

Gary Miller
Carnegie Mellon University
glmiller@cs.cmu.edu

MS44

Energy and Performance Tuning Framework for Accelerator Environment

Accelerator computing with GPUs for HPC has been wide spread. Energy consumption of applications became as important as execution performance. We propose a framework to perform tuning of energy consumption and execution performance for accelerated computer environment. The framework contains user specified policy to indicate execution policy including energy consumption and execution performance. The framework tunes applications according to the user specified policy.

Shoichi Hirasawa
The University of Electro-communications
hirasawa@is.uec.ac.jp

Satoshi Ohshima
Univ. of Tokyo
ohshima@cc.u-tokyo.ac.jp

Hiroki Honda
The University of Electro-Communications
honda@is.uec.ac.jp

MS44

Development of C Language Version of ABCLibScript - The Impact To Auto-tuning Software

The much eagerly-awaited techniques on current computer environment are auto-mated system for arbitrary user programs. Compiler support of tuning is limited due to complex architectures: multicore, heterogeneous configuration of CPU and GPU, and very deep hierarchy of cashes. One of solutions is to supply general computer language and preprocessor for auto-tuning. The ABCLibScript was de-

signed to match the requirement of auto-tuning. However, the previous released version was only supported for source codes written by Fortarn90. We hence now are developing a C language version. In this presentation, the author will explain the impact of the C version to HPC software on heterogeneous environments, like CPU with GPU acceleration.

Takahiro Katagiri
Information Technology Center, The University of Tokyo
katagiri@cc.u-tokyo.ac.jp

MS44

Development of a Multiple Precision Version of LAPACK and Performance Measurement

We have been developing a multiple precision version of BLAS and LAPACK (are widely used linear algebra packages) called MPACK at http://mplapack.sourceforge.net/ . This project is intended to (i) provide reference implementation of multiple precision version of BLAS and LAPACK (ii) rewritten by C++ (iii) can use many multiple precision arithmetic libraries like mpfr, gmp and qd. (iv) currently 76 BLAS routines and 100 LAPACK routines are implemented and tested, (v) distributed under 2-clause BSD license. In this presentation we present current status of MPACK and provide some benchmark results.

Maho Nakata
RIKEN
maho@riken.jp

MS44

Computer Assisted Proof of Non-singularity of a Floating-point Matrix based on High Performance Functions

This talk is concerned with verified numerical computations, especially, verification of non-singularity of a floating-point matrix A is focused on. Let R be an approximate inverse matrix of A, then the problem is to prove that a maximum norm of $RA - I$ is smaller than 1, where I is the identity matrix. Our algorithm does as much works as necessary to prove it by using high performance functions supported by BLAS.

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

Takeshi Ogita
Tokyo Woman's Christian University
ogita@lab.twcu.ac.jp

Shin'ichi Oishi
Waseda University
oishi@waseda.jp

MS45

Presenters To Be Announced

Abstract not available at time of publication.

TBD A6
TBD
neckel@in.tum.de

MS46

Designing Derivative-Free Hybrid Optimization Methods for Hydrological Applications

Derivative-free methods have emerged as invaluable for finding solutions to black-box optimization problems, and each has distinct advantages and disadvantages. Often, these strengths and weaknesses are problem dependent. We consider hybrid approaches which combine beneficial elements of multiple methods in order to more efficiently search the design space. In this talk, we will examine some common derivative-free optimization approaches and describe some appropriate hybrids which address the specific characteristics of hydrology applications. Using numerical examples, we will illustrate how these hybrids can find solutions under conditions which the traditional optimal search methods fail.

Genetha Gray
Sandia National Laboratories
gagray@sandia.gov

Katie Fowler
Clarkson University
kfowler@clarkson.edu

MS46

Exploiting Uncertainty Quantification in Derivative-free Optimization

We describe a model-based trust-region algorithm for derivative free optimization using weighted regression. The approach provides a simple mechanism for exploiting uncertainty quantification, allowing greater weight to be given to function evaluations that are known with greater certainty, resulting in better local models. We describe a strategy for choosing weights based on this idea, but which takes other factors into account, including proximity to the trust-region center and the geometry of the sample set.

Stephen C. Billups, Jeffrey Larson
University of Colorado Denver
stephen.billups@cudenver.edu,
jeffrey.larson@ucdenver.edu

MS46

Parameter Estimation of Complex Simulations

Model-based methods evaluate the objective function at trial points and construct a model (or surrogate) of the function. We discuss algorithmic and performance issues in a new model-based trust-region algorithm (Pounder) that constructs a quadratic model of least change by interpolating the function at a selected set of previous trial points. We discuss performance issues on a set of model problems and on parameter estimation problems in nuclear energy density functional theory.

Jorge J. More
Argone National Laboratory
more@mcs.anl.gov

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS46

Using Derivative-free Algorithms to Identify Surrogate Models of Energy Systems

We propose a method of simulation-based optimization with accurate, low-complexity surrogate models. Derivative-free and derivative-based algorithms are used to identify, improve, and validate a set of surrogate models of simulations. The methodology is applied to a power plant simulation, where a set of surrogate models is identified by separating the plant into less complex components, modeling the components, then combining the surrogate models to formulate an algebraic nonlinear program for optimization.

Alison Cozad
Carnegie Mellon University
alison.cozad@gmail.com

Nick Sahinidis
sahinidis@cmu.edu
carnegie mellon university

MS47

A Well-Conditioned Hierarchical Basis for Triangular H(Curl) -Conforming Elements for Maxwell's Equations

We construct a well-conditioned hierarchical basis for triangular H(curl) conforming elements with selected orthogonality. The basis functions are grouped into edge and interior functions, and the later is further grouped into normal and bubble functions. In our construction, the trace of the edge shape functions are orthonormal on the associated edge. The interior normal functions, which are perpendicular to an edge, and the bubble functions are both orthonormal among themselves over the reference element. The construction is made possible with classic orthogonal polynomials, viz., Legendre and Jacobi polynomials. For both the mass matrix and the quasi-stiffness matrix, better conditioning of the new basis is shown by a comparison with the existing high order basis.

Wei Cai
University of North Carolina, Charlotte
wcai@uncc.edu

Steve Xin
UNC Charlotte
jxin@uncc.edu

MS47

A Parareal in Time Algorithm

The parareal algorithms are efficients tools to solve evolution problems in real time. The principle is the following: one first approximates the solution on a coarse time grid using a Coarse Solver, and then locally solves the problem on fine time subgrids using a Fine Solver on parallel computers. The associated iterative procedure ensures an accuracy which is of same order as the accuracy of the fine scheme when applied on the fine time grid. We present new applications of the method.

Sidi M. Kaber
University Pierre & Marie Curie
kaber@ann.jussieu.fr

MS47

Maximum-principle-satisfying and Positivity-preserving High Order Discontinuous Galerkin and Finite Volume Schemes for Conservation Laws

We construct uniformly high order accurate discontinuous Galerkin (DG) and weighted essentially non-oscillatory (WENO) finite volume (FV) schemes satisfying a strict maximum principle for scalar conservation laws and passive convection in incompressible flows, and positivity preserving for density and pressure for compressible Euler equations. A general framework (for arbitrary order of accuracy) is established to construct a limiter for the DG or FV method with first order Euler forward time discretization solving one dimensional scalar conservation laws. Strong stability preserving (SSP) high order time discretizations will keep the maximum principle and make the scheme uniformly high order in space and time. One remarkable property of this approach is that it is straightforward to extend the method to two and higher dimensions. The same limiter can be shown to preserve the maximum principle for the DG or FV scheme solving two-dimensional incompressible Euler equations in the vorticity stream-function formulation, or any passive convection equation with an incompressible velocity field. A suitable generalization results in a high order DG or FV scheme satisfying positivity preserving property for density and pressure for compressible Euler equations. Numerical tests demonstrating the good performance of the scheme will be reported.

Xiangxiong Zhang
Brown University
zhangxx@dam.brown.edu

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

MS47

Numerical Methods of the Space-Time Fractional Differential Equations

We consider initial boundary value problems of fractional differential equations and its numerical solutions. Two definitions, i.e. Riemann–Liouville definition and Caputo one, of the fractional derivative are considered in parallel. A theoretical framework for the weak solutions of the space-time fractional diffusion equations is developed, and the well-posedness of the weak problems is proved. Moveover, based on the proposed weak formulation, we construct an efficient spectral method for numerical approximations of the weak solution.

Chuanju Xu
School of Mathematical Sciences
Xiamen University
cjxu@xmu.edu.cn

MS47

Spectral Collocation/p-Version Finite Element Methods for Hamiltonian Dynamical Systems

We make a systematical study of spectral Galerkin methods (or p-version finite element methods) and spectral collocation methods in numerically solving the Hamiltonian dynamical systems. Different strategies including Legendre-Lobatto collocation, Chebyshev-Lobatto colloca-

tion, spectral Galerkin/p-version, are discussed and compared, especially with symplectic methods. Numerical tests on some benchmark nonlinear problems are provided.

Zhimin Zhang, Nairat Kenyamee
Wayne State University
zzhang@math.wayne.edu, ax7812@wayne.edu

MS48

Artificial Boundary Conditions for Schrödinger Equations with General Potentials and Nonlinearities

This talk addresses the construction of absorbing boundary conditions for the two-dimensional Schrödinger equation with a general variable repulsive potential or with a cubic nonlinearity. Semi-discrete time schemes, based on Crank-Nicolson approximations or relaxation scheme, are built for the associated initial boundary value problems. Finally, some numerical simulations give a comparison of the various absorbing boundary conditions to analyse their accuracy and efficiency.

Christophe Besse
Universite des Sciences et Technologies de Lille
christophe.besse@math.univ-lille1.fr

MS48

Propagators for the Time-dependent Schrodinger Equation in an Adaptive, Discontinuous, Multiwavelet Basis

Abstract not available at time of publication.

Jun Jia
Oak Ridge National Laboratory
jiaj@ornl.gov

MS48

Fourth Order Time-stepping for Kadomtsev-Petviashvili and Davey-Stewartson Equations

Purely dispersive partial differential equations as the Korteweg-de Vries equation, the nonlinear Schrödinger equation and higher dimensional generalizations there of can have solutions which develop a zone of rapid modulated oscillations in the region where the corresponding dispersionless equations have shocks or blow-up. To numerically study such phenomena, fourth order time-stepping in combination with spectral methods is beneficial to resolve the steep gradients in the oscillatory region. We compare the performance of several fourth order methods for the Kadomtsev-Petviashvili and the Davey-Stewartson equations, two integrable equations in 2+1 dimensions: exponential time-differencing, integrating factors, time-splitting, and Driscoll's IMEX method. The accuracy in the numerical conservation of integrals of motion is discussed.

Christian Klein
Institut de Mathématiques de Bourgogne
9 avenue Alain Savary, BP 47870 21078 DIJON Cedex
christian.klein@u-bourgogne.fr

Kristelle Roidot
IMB
Université de Bourgogne
kr0083vi@u-bourgogne.fr

MS48

Fast Gaussian Wavepacket Transforms and Gaussian Beams for the Schrodinger Equation

Abstract not available at time of publication.

Jianliang Qian
Department of Mathematics
Michigan State University
qian@math.msu.edu

MS49

Toward Extreme Scale: Performance Analysis on the IBM BG/P for High-Order Electromagnetic Modeling

We demonstrate a massively parallel, open-source electromagnetic code, based on advanced high-order numerical algorithms targeting production runs on the petascale computer architecutres. Computational results include cutting-edge simulations in accelerator physics and nanotechnology-based science applications,involving research and development in applied mathematics, computer science, and software. We present convergence and stability analysis, efficient parallel communication kernel, I/O performance, and realistic simulations and validation of the computational results in comparison with other methods. A detailed analysis on the parallel performance of the spectral-element discontinous Galerkin (SEDG) method for the time-domain electromagnetic simulations based on explicit timestepping scheme will be discussed, provided with strong and weak scalings up to 131K cores on the Argonne BG/P.

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Jing Fu
Rensselaer Polytechnic Institute
fuj@cs.rpi.edu

MS49

Large Scale Computational Dynamics on Next Generation Heterogeneous Clusters using DDEP

The talk discusses a Dynamic Data Exchange Protocol (DDEP), a programming framework developed to assist in large-scale computational dynamics simulations on hetergeneous GPU-accelerated clusters. The problem domain is subdivided spatially with each subdivision assigned to a specific GPU; the developed framework facilitates the movement of body data between GPUs when bodies cross these subdivisions. It is shown that using DDEP, a GPU-accelerated cluster may decrease simulation time by an order of magnitude.

Andrew Seidl
University of Wisconsin-Madison
aaseidl@wisc.edu

Hammad Mazhar, Toby Heyn, Dan Negrut
University of Wisconsin Madison
Department of Mechanical Engineering
hmazhar@wisc.edu, heyn@wisc.edu, negrut@cae.wisc.edu

MS49

Finite-Element Electromagnetic Modeling of Accelerator at Extreme Scale

The past decade has seen tremendous advances in electromagnetic modeling for accelerator applications with the use of high performance computing on state-of-the-art supercomputers. Under the support of the DOE SciDAC computing initiative, a comprehensive set of parallel electromagnetic codes based on the finite-element method has been developed aimed at tackling some of the most computationally challenging problems in accelerator RD. Aided by collaborative efforts in computational science, these powerful tools have enabled large-scale simulations of complex systems to be performed with unprecedented details and accuracy. Significant progress has been made in scalable eigensolvers for accelerator cavity prototyping and optimization, calculations of wakefields for large accelerator systems and for ultra-short charged particle bunches, and self-consistent particle-in-cell simulations for space-charge dominated devices. The impact of these capabilities on current and future accelerator projects worldwide will be described.

Cho-Kuen Ng
Stanford Linear Accelerator Center
cho@slac.stanford.edu

MS49

Large Scale Ab Initio Calculations for Photovoltaic Materials

Large scale ab initio calculation can play an important role in understanding the underlying mechanism in solar cell. Not only the ab initio calculation can be used to study the atomic structures of many systems (e.g., the interface and defects), it can also be used to study the electronic structures, optical properties and carrier transport. In order to study these properties, large scale computations are needed. In this talk, the computational requirements for different types of calculations will be discussed, and the main challenges will be presented. I will also discuss how such calculations can be scaled to hundreds of thousands of processors.

Lin-Wang Wang
Lawrence Berkeley National Lab
lwwang@lbl.gov

MS50

Multi-GPU Calculations in Lattice Quantum Chromodynamics

Multi-GPU Calculations in Lattice Quantum Chromodynamics Quantum chromodynamics (QCD) is the fundamental theory describing the interactions of quarks and gluons. In the lattice formulation of QCD, the equations governing these interactions are solved numerically on a four-dimensional spacetime grid. In this talk, I will discuss the implementation of "QUDA," a library of linear solvers tailored for QCD that is capable of leveraging many graphics processing units (GPUs) in parallel. MPI is used for communication, while kernels executing on the GPUs are written in CUDA C. I will discuss the strategies we employed both to obtain high performance on a single GPU and maintain efficiency on large clusters.

Ronald Babich
Center for Computational Science
Boston University
rbabich@bu.edu

Kipton Barros
Center for Nonlinear Studies
Los Alamos National Laboratory
kbarros@gmail.com

Richard Brower
Department of Physics
Boston University
brower@bu.edu

Michael Clark
Harvard-Smithsonian Center for Astrophysics
mikec@seas.harvard.edu

Balint Joo
Jefferson Laboratory
bjoo@jlab.org

Claudio Rebbi
Department of Physics
Boston University
rebbi@bu.edu

MS50

Overlapping Communication and Calculation with a PGAS Language

The high complexity of climate applications poses difficulties for porting and validation on new platforms. Efforts for improving programmability and performance of such applications cannot be disruptive and must lead to clear benefits. Partitioned Global Address Space (PGAS) languages have emerged as alternative programming models for highly parallel systems and climate applications offer opportunities for assessing the efficiency of these models. Here we report on overlapping communication and calculation in an ocean circulation model.

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

MS50

Multi-GPU MapReduce

MapReduce greatly simplifies many processing tasks. We modify portions of MapReduce such as chunking and implement extensions such as Partial Reduction to better fit the GPU and show how to efficiently implement MapReduce on GPUs. We demonstrate the ability of a GPU MapReduce implementation through several benchmarks including Word Occurrence and Matrix Multiplication.

Jeff Stuart
Department of Computer Science
University of California, Davis
stuart@cs.ucdavis.edu

MS50

Using UPC for Hybrid Programming

Unified Parallel C (UPC) is a Partitioned Global Address Space (PGAS) language that provides programming con-

venience similar to shared-memory programming models and can deliver scalable performance on diverse computer architectures ranging from commodity multi-core systems to customized supercomputers. Two aspects of using UPC for hybrid programming will be presented: 1) combining UPC with MPI or OpenMP for programming multi-core clusters and 2) a UPC extension with CUDA/OpenCL for programming GPU clusters.

Yili Zheng
Lawrence Berkeley National Laboratory
yzheng@lbl.gov

MS51

Uncertainty Quantification and Numerical Error Control in Simulation of Coupled Porous Flow and Mechanical Deformation

We present recent work on interactions between numerical error and uncertainty quantification (UQ), especially in the context of multiphysics applications. Our approach involves error correction using goal-oriented error estimation, especially over the large parameter space induced by random variables and random spatial fields. We also investigate approaches for generating surrogate models of numerical error that can be used to correct model realizations in UQ calculations.

Brian Carnes
Sandia National Laboratories
bcarnes@sandia.gov

MS51

A Jacobian-Free Newton Krylov Method for Mortar-Discretized Thermomechanical Contact Problems

The thermomechanical contact approach studied here is based on the mortar finite element method, where Lagrange mulipliers are used to enforce weak continuity constraints at participating interfaces. In this formulation, the heat equation couples to linear mechanics (in weak form) through a thermal expansion term. Lagrange multipliers are used to formulate the continuity constraints for both heat flux and the normal component of interface traction at contact interfaces. The resulting system of nonlinear algebraic equations is normalized and cast in residual form for solution of the transient problem. A preconditioned Jacobian-free Newton Krylov method is used to obtain the solution of the coupled constraint, thermal contact, and heat equations.

Glen A. Hansen
Idaho National Laboratory
Glen.Hansen@inl.gov

MS51

Conservative and Noise Resistant Data Remapping for Coupled Regional Climate Modeling

The coupling of WRF (the Weather Research and Forecast Model) and CAM (Community Atmospheric Model) for regional climate modeling requires transferring data (such as wind velocities, temperature, moisture, etc.) between the grids of WRF and CAM. These grids are in general nonmatching and may use different map projects. We investigate a common-refinement based data transfer method that satisfies physical conservation, numerical accuracy, and noise resistance.

Ying Chen
Stony Brook University
yingchen@ams.sunysb.edu

Xiangmin Jiao
Stony Brook University
Stony Brook, NY
xjaio@sunysb.edu

Wuyin Lin
Brookhaven National Laboratory
wlin@bnl.gov

Minghua Zhang, Juanxiong He
Stony Brook University
mzhang@notes.cc.sunysb.edu, hejuanxiong@gmail.com

MS51

Multiscale Simulation of the Effect of Irradiation-Induced Microstructure Evolution on Reactor Fuel Performance

Fuel performance in nuclear reactors is highly dependent on irradiation-induced microstructure evolution. Therefore, a fuel performance code must consider atomistic and mesoscale effects in order to provide a predictive capability. In this work, we present the multiscale fuel performance modeling approach currently being employed at Idaho National Laboratory (INL). Atomistically-informed mesoscale phase field simulations are used to determine the effect of irradiation-induced microstructure evolution on bulk properties, such as thermal conductivity and density. Continuum expressions describing the effect of irradiation on the bulk properties as a function of temperature are then determined from the mesoscale simulations. Finally, these expressions are used in INLs BISON fuel performance code to model fuel pellet behavior.

Michael Tonks, Paul Millet, Derek Gaston, Bulent Biner
Idaho National Laboratory
michael.tonks@inl.gov, paul.millet@inl.gov,
derek.gaston@inl.gov, bulent.biner@inl.gov

Anter El-Azab
Florida State University
aelazab@fsu.edu

MS52

Lessons Learned and Open Issues from the Development of the Proteus Toolkit for Coastal and Hydraulics Modeling

The Proteus toolkit evolved to support research on new models for coastal and hydraulic processes and improvements in numerical methods. The models considered include multiphase flow in porous media, shallow water flow, turbulent free surface flow, and flow-driven processes such as sediment and species transport. Python was used for implementing high-level class hierarchies and prototyping new algorithms, while performance critical sections were optimized using compiled languages. We discuss the toolkit design, performance, and open issues.

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 Centerq
matthew.w.farthing@usace.army.mil

MS52

Python, Clawpack, and PyClaw

Clawpack (Conservation Laws Package, www.clawpack.org) is an open source softwarepackage for solving nonlinear conservation laws and other hyperbolic systems of PDEs. The core computational routines are in Fortran, but Python is now used for the user interface, both on the input end and for graphics and visualization of results. Recent developments will be summarized with discussion of some lessons learned that may be useful to developers of other packages.

Randall J. LeVeque
Applied Mathematics
University of Washington (Seattle)
rjl@uw.edu

Kyle T. Mandli
University of Washington
Dept. of Applied Mathematics
mandli@amath.washington.edu

MS52

FEniCS: An Attempt to Combine Simplicity, Generality, Efficiency and Reliability

FEniCS allows PDE problems expressed in terms of variational forms to be stated in Python, with a syntax that closely resembles the mathematics. The Python definition of the problem is compiled to highly efficient, problem-specific C++ code and linked with general libraries for finite elements and linear algebra. Recently, FEniCS can also automatically generate tailored error estimators and provide error control. Many examples on using FEniCS to solve PDEs will be shown.

Kent-Andre Mardal
Simula Research Laboratory and Department of Informatics
University of Oslo
kent-and@simula.no

Hans Petter Langtangen
Center for Biomedical Computing
Simula Research Laboratory and University of Oslo
hpl@simula.no

Anders Logg
Simula Research Laboratory
logg@simula.no

Garth N. Wells
Dept. of Engineering,
Univ. of Cambridge
gnw20@cam.ac.uk

MS52

FEMhub Online Numerical Methods Laboratory

The FEMhub Online Numerical Methods Laboratory (http://lab.femhub.org) is an Ext-JS-based web interface to FEMhub (http://femhub.org), an open source distribution of PDE solvers with a unified Python interface. The objective of the Online Lab is to make the codes in FEMhub widely accessible without installation. We will explain the structure of FEMhub (Python wrappers, Python build system etc.), describe the Online Lab, and perform sample finite element computations over the web. We will also present a new FEMhub Mesh Editor that makes it possible to generate 2D geometries and finite element meshes in the web browser.

Pavel Solin, Ondrej Certik
Department of Mathematics and Statistics
University of Nevada, Reno
solin@unr.edu, ondrej@certik.cz

Mateusz Paprocki
Technical University of Wroclaw
Wroclaw, Poland
mattpap@gmail.com

Aayush Poudel
University of Nevada, Reno
aayushpoudel@gmail.com

MS53

ReALE: A Reconnection-based Arbitrary-Lagrangian-Eulerian Method

We present a new reconnection-based arbitrary Lagrangian-Eulerian (ALE) method. The main elements in as standard ALE simulation are an explicit Lagrangian phase in which the solution and grid is updated, a rezoning phase in which a new grid is defined, and remapping phase in which the Lagrangian solution is conservatively interpolated onto new mesh. In our new method we allow connectivity of the mesh to change in rezone phase, which leads to general polygonal mesh. Rezone strategy is based on using Voronoi mesh. In our talk we will discuss details of rezoning phase and show some numerical results.

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

Raphael Loubere
University of Toulouse, France
raphael.loubere@math.univ-toulouse.fr

Pierre-Henri Maire
UMR CELIA, Université Bordeaux I
maire@celia.u-bordeaux1.fr

Jerome Breil
UMR CELIA
breil@celia.u-bordeaux.fr

Stephan Galera
INRIA Bordeaux - Team BACCHUS
stephane.galera@inria.fr

MS53

Vertex Reordering for Local Mesh Quality Improvement of Tetrahedral Meshes

Vertex reordering can be performed within the context of local mesh quality improvement in order to decrease the amount of time spent on the mesh optimization. We propose several static vertex reordering schemes, based on the performance of the optimization algorithm and the mesh

quality, and investigate the trade-offs between ordering and overall performance of the optimization algorithm. We employ Laplace smoothing, within the Mesquite package, to optimize tetrahedral mesh quality.

Jeonghyung Park, Suzanne M. Shontz
The Pennsylvania State University
jxp975@cse.psu.edu, shontz@cse.psu.edu

Patrick M. Knupp
Sandia National Laboratories
pknupp@sandia.gov

MS53

New Developments on Image-based Mesh Generation

We have developed a high-fidelity image segmentation approach using a variant of the graph-cut method. From the segmented boundaries, high-quality and smooth surface meshes are generated with the center-to-center connection scheme. Tetrahedral meshes are then generated using an octree-based approach. The experiments have shown promising results in mesh generation from 3D imaging data. A user-friendly GUI has also been developed to encapsulate all these algorithms, which will be made available at www.bimos.org.

Zeyun Yu
University of Wisconsin-Milwaukee
yuz@uwm.edu

MS53

Guaranteed-Quality Quadrilateral Mesh Generation for Planar Domains

Provably good-quality triangular mesh generation were well developed. However, fewer algorithms exist for all-quadrilateral mesh generation, and most of these algorithms are heuristic. Here, two novel approaches are presented to generate guaranteed-quality all-quadrilateral meshes. It is proved that for any planar domain, all the angles in the constructed mesh are within $[45^\circ \pm \varepsilon, 135^\circ \pm \varepsilon]$ ($\varepsilon \leq 5^\circ$) for the quadtree-based method and $[60^\circ \pm \varepsilon, 120^\circ \pm \varepsilon]$ for the hexagon-based method.

Yongjie Zhang
Department of Mechanical Engineering
Carnegie Mellon University
jessicaz@andrew.cmu.edu

MS54

Adaptive High Dimensional Smoothing using Iterative bias Reduction Algorithms

Smoothing high dimensional noisy data is challenging because of the curse of dimensionality. This talk presents a simple numerical procedure that iteratively fits the residuals to correct for the bias of the current smoother. While conceptually simple, this method is a practical algorithm for achieving adaptativity to the underlying smoothness of the true regression function. The iterative algorithm works for thin plate splines, and for kernel smoother for positive defined kernels. In practical examples, our smoother has smaller out of sample prediction error than competing state-of-the-art multivariate smoothers. Finally, the algorithm is distributed as the ibr R package available at CRAN.

Nicolas Hengartner
Information Sciences Group
Los Alamos National Laboratory
nick@lanl.gov

MS54

Inverse Methods for Event Reconstruction Problems

When a contaminant has been released into the atmosphere, event reconstruction algorithms are used to determine where, when and how much has been released. A limited number of measurements are available from environmental sensors deployed in urban areas, and inverse algorithm are used to find atmospheric release parameters, based on these data. These parameters can then be used in flow simulations to predict current and future locations of the material. Many inverse algorithms use Bayesian inference to solve the inverse problem. This can be viewed as a stochastic process and typically Markov chain Monte Carlo sampling is used to generate an ensemble of predictive simulations. We will give an overview of these methods, and also discuss adjoint methods from the Bayesian perspective and their potential impact on these problems. Relevant issues include non-Gaussian distribution of parameters and measurements, nonlinear models, uncertainty is data, parameters, models and solution, and efficient algorithms for real-time emergency response.

Jodi Mead
Boise State University
Department of Mathematics
jmead@boisestate.edu

MS54

Projected Conjugate Gradient Methods on GPUs for Improving Efficiency in Large Scale Image Reconstruction

Graphical processing units improve the capability for large scale computations using low-end computing platforms. For the solution of linear systems of equations, much effort has been devoted to their use for efficient Krylov subspace-based solvers. Here we discuss implementation of a projected CG algorithm, taking advantage of BLAS3 operations, and Krylov subspace recycling. It is extended for solution of regularized least squares problems arising in large scale image restoration situations. Numerical results are presented.

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

Tao Lin
Department of Mathematical Sciences
New Jersey Institute of Technology
tl48@njit.edu

MS54

Rapid-response Simulations of Forward and Inverse Problems in Atmosperic Transport and Dispersion

Event reconstruction of an atmospheric contaminant plume, detected by a sensor network, is a challenging in-

verse problem in threat reduction field. Dispersion modeling in urban areas and over complex terrain constitutes a forward model, and benefits from computationally intensive prognostic models. To this end, we present the development and performance of an incompressible wind solver with a geometric multigrid algorithm on GPU clusters and discuss its efficient use within an inverse methodology.

Inanc Senocak
Mechanical & Biomedical Engineering
Boise State University
senocak@boisestate.edu

Dana Jacobsen
Boise State University
danajacobsen@u.boisestate.edu

MS55

Auto-tuning for BLAS-based Matrix Computations

We discuss performance tuning of matrix computations based on BLAS routines. In such computations, blocked algorithms are generally used and how to partition the target matrices into submatrices determines the performance of each BLAS routine. To achieve high performance, blocking should be done depending on the characteristics of the BLAS routines. In this talk, we clarify issues toward auto-tuning of BLAS-based computations and explain our approach, called generalized recursive blocking, that tries to optimize the partitioning using dynamic programming.

Takeshi Fukaya
Department of Computational Science and Engineering
Nagoya University
t-fukaya@na.cse.nagoya-u.ac.jp

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

Shao-Liang Zhang
Nagoya University, Japan
zhang@na.cse.nagoya-u.ac.jp

MS55

Performance Evaluation for a Dense Eigenvalue Solver for the Next-generation Petascale System

We developed a high performance and high scalable eigenvalue solver for the next-generation petascale system, 'K-computer', which is supposed to have almost million cores. We adopt MPI-OpenMP hybrid parallelization, two-dimensional topological communication, and auto-tuning technology in order to achieve higher performance. Currently, our solver performs on T2K supercomputer housed at the University of Tokyo with 256 AMD Barcelona processors (4096 cores). At the workshop, we will report the performance on other systems and estimate the performance by using full system of 'K-computer'.

Toshiyuki Imamura, Huu Phuong Pham
The University of Electro-Communications
imamura@im.uec.ac.jp, phuong@i.im.uec.ac.jp

Susumu Yamada, Masahiko Machida
Japan Atomic Energy Agency
yamada.susumu@jaea.go.jp, machida.masahiko@jaea.go.jp

MS55

Autotuning of Sparse Matrix-Vector Multiplication by Selecting Storage Schemes on GPU

Sparse matrix vector multiplication is one of the most often used functions in computer science. The storage schemes for sparse matrices have been proposed, however, each sparse matrices have an optimal storage scheme. In this talk, we propose an auto-tuning algorithm of sparse matrix vector multiplication by selecting storage schemes automatically on GPU. We evaluated our algorithm using Conjugate Gradient solver. As a result, we found that our algorithm was effective in many sparse matrices.

Yuji Kubota, Daisuke Takahashi
Graduate School of Systems and Information Engineering
University of Tsukuba
kubota@hpcs.cs.tsukuba.ac.jp, daisuke@cs.tsukuba.ac.jp

MS55

Development of Xabclib: A Sparse Iterative Solver with Numerical Computation Policy Interface

Matrix libraries have many parameters as inputs by the user. They include problem parameters what are difficult to set values and the approach of automatically setting them is needed. In this talk, we propose matrix libraries named "Xabclib" with numerical computation policy interface. By using Xabclib, users can specify their numerical policy, such as minimizing computation time, memory saving, and accuracy requirement for the residual of solution without difficult parameter setting. A result of performance evaluation with one node of the T2K Open super computer shows the proposed approach can reduce over 60% of parameters and satisfy user needs.

Takao Sakurai
Central Research Laboratory
Hitachi Ltd.
takao.sakurai.ju@hitachi.com

Ken Naono
HITACHI Ltd.
ken.naono.aw@hitachi.com

Takahiro Katagi
University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

Hisayasu Kuroda
Univ. of Tokyo / Ehime U.
kuroda@cs.ehime-u.ac.jp

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

Satoshi Ohshima, Shoji Itoh
Univ. of Tokyo
ohshima@cc.u-tokyo.ac.jp, itosho@cc.u-tokyo.ac.jp

Mitsuyoshi Igai
Hitachi ULSI Systems Corporation
mitsuyoshi.igai.bf@hitachi.com

MS56

Fluid-Structure Interaction of Wind Turbines in

3D: Methods and Applications

In this talk I will present a collection of numerical methods combined into a single framework for wind turbine rotor modeling and simulation. I will cover structural discretization for wind turbine blades and details of the fluid?structure interaction computational procedures focusing on the challenges of FSI coupling and mesh motion in the presence of large rotation. I will present simulations of the NREL 5MW offshore baseline wind turbine rotor, including validation against published data.

Yuri Bazilevs
aDepartment of Structural Engineering
University of California, San Diego
yuri@ucsd.edu

MS56

Time Adaptive SDIRK Methods for Thermal Fluid Structure Interaction

To model cooling processes in steel manufacturing, we consider the unsteady thermal coupling of the compressible Navier-Stokes equations with the heat equation. In time, stiffly stable SDIRK methods are employed. The coupled problem on each stage is solved using partitioned FSI methods. The SDIRK schemes allow for error estimation, saving considerable computing time and guaranteeing a maximal time integration error. The method is tested on parallel flow along a heated plate, allowing for verification by experiments.

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

Stefan Hartmann
Clausthal University of Technology
Institute of Applied Mechanics
stefan.hartmann@tu-clausthal.de

Karsten J. Quint
University of Kassel
Department of Mechanics
karsten.quint@gmx.de

Detlef Kuhl
University of Kassel
Institute of Statics and Dynamics
kuhl@uni-kassel.de

MS56

Multiphysics Coupling Via Lime: Lightweight Integrated Multiphysics Environment

LIME is an enabling solution technology built as extensions to the Trilinos Solver Framework. It accommodates codes ranging from mature legacy applications to novel applications undergoing active development. LIME's coupling algorithms range from loose fixed-point iterations to various implementations of Newton's method. It is lightweight by only requiring an application's response to inputs but can incoroprate additional information, eg residuals, Jacobians, etc., to effect commensurate improvements in robustness and efficieny of coupled solutions.

Russell W. Hooper
Sandia National Laboratories
rhoope@sandia.gov

Roger Pawlowski
Sandia National Labs
rppawlo@sandia.gov

Kenneth Belcourt, Rodney Schmidt
Sandia National Laboratories
kbelco@sandia.gov, rcschmi@sandia.gov

MS56

Stable and Efficient Coupling for Fluid-Structure Interactions

Partitioned fluid-structure interaction simulations exhibit big challenges in particular for incompressible fluids as instabilities caused by the added mass effect cannot be reduced by smaller time steps. Different approaches ranging from Gauss-Seidel coupling iterations with constant underrelaxation over Aitken underrelaxation to quasi-Newton methods are known as possible remedies. We propose a new hierarchical coupling method inducing very low overhead and at the same time yielding fast convergence of the coupling iterations.

Miriam Mehl
Department of Computer Science
Technische Universitaet Muenchen
mehl@in.tum.de

MS56

Partitioned Fluid-Structure Interaction Using Multilevel Acceleration

In partitioned fluid-structure interaction sub-iteration are required for strongly coupled problems, increasing computational time as flow and structure have to be resolved multiple times every time step. In this paper we apply a multilevel acceleration technique to the sub-iterations, which resolves a defect-correction in the flow domain on a coarse mesh level to reduce computing time. For different sub-iteration techniques computational efficiency is determined in combination with multilevel acceleration.

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

MS57

A Highly Scalable Method for Coulomb Interactions Based on Multigrid

The simulation of particle interactions due to electrostatics can be reformulated as the solution of a Poisson equation in a consistent way plus near-field correction terms. As multigrid methods are well-known to be optimal Poisson solvers, a method incorporating multigrid was developed. The resulting method scales linearly in the number of particles. The parallelization is based on a domain-splitting approach obeying the underlying computer architecture,

thus a very good scaling behavior is achieved.

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

MS57

Maxwell Equations Molecular Dynamics - A Local Algorithm for a Global Interaction

The presented electrostatics algorithm is designed to propagate electric field through the simulation system via a local update scheme rather than solving the Poisson equation globally. This directly yields algorithmic advantages such as linear scaling with system size and trivial parallelization. In the talk, the algorithm is presented and the scaling behaviour of a simple implementation is shown. Finally, it will be discussed, for which systems it might advantageous to use.

Florian W. Fahrenberger
University of Stuttgart
florian.fahrenberger@icp.uni-stuttgart.de

MS57

A Parallel Fast Multipole Method for Particle Simulations

In this talk we present the parallel version of our error-controlled FMM implementation for long-range interactions in particle simulations. The algorithm can handle open as well as 1D, 2D and 3D periodic boundary conditions efficiently. We will show results and timings of low-precision and high-precision computations for up to several billion particles. The cross-over point with the direct summation as well as the scaling behavior on different supercomputers (e.g. BlueGene) of the code will be discussed.

Ivo Kabadshow
Research Centre Juelich
i.kabadshow@fz-juelich.de

Holger Dachsel
Juelich Supercomputing Centre
Research Centre Juelich
h.dachsel@fz-juelich.de

MS57

Parallel Fast Ewald Summation Based on Nonequispaced Fourier Transforms

The fast calculation of long-range interactions is a demanding problem in particle simulation. For periodic boundary conditions, Ewald proposed to split interactions into two fast converging parts. We show that the reciprocal space part can be efficiently evaluated by nonequispaced fast Fourier transforms. The resulting algorithms are very similar to particle-particle particle-mesh methods. Furthermore, we present a massively parallel implementation of the nonequispaced fast Fourier transform, which yields a straightforward parallelization of our Ewald method.

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

MS58

Recent Progress in Radial Basis Functions

RBFs are a meshless spectral method that have been applied successfully in complicated geometry. We review recent developments in analytic approximations to RBF cardinal functions, asymptotics of RBF coefficients, RBFs and group theory, boundary layer effects, RBFs and the Runge Phenomenon, solving the vorticity equation for flow on a rotating sphere, and damping and dealiasing.

John P. Boyd
University of Michigan
Ann Arbor, MI 48109-2143 USA
jpboyd@umich.edu

MS58

Computing at the Speed of Thought: Solving Differential Equations in Chebfun

Chebfun is a free MATLAB-based software system for computing with functions as readily as numbers by using fast and powerful Chebyshev series approximations. Using operators and boundary conditions expressed in a compact, natural syntax, many boundary-value, eigenvalue, and continuation problems can be solved by spectral collocation methods to high accuracy in moments, without requiring knowledge of the numerical discretizations. The ability to interact with such models in realtime accelerates the process of understanding their solutions.

Tobin Driscoll
University of Delaware
Mathematical Sciences
driscoll@udel.edu

MS58

Petascale Applications of the Spectral Element Method

We describe recent advances in the spectral element code Nek5000 for simulation of incompressible and low-Mach number flows. The SEM can significantly reduce computational complexity, particularly for large-scale problems enabled by peta- and exascale architectures where cumulative effects dictate tight control of local discretzation errors. We describe our underlying discretization and solution strategies, including a scalable multigrid solver and present several recent simulation and performance results including sustained performance exceeding 20% of peak and 70% efficiency for P=262000.

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

MS58

High-Order Integral Equations Methods for High-Frequency Scattering by Diffraction Gratings

In this talk we present a new high-order Integral Equation algorithm for the simulation of high-frequency scattering returns by diffraction gratings. For shallow gratings (those for which Geometric Optics indicates that there will be no multiple reflections) the method amounts to a phase-extraction technique resulting in a slowly-varying amplitude as unknown which requires only a small number of degrees of freedom to resolve. For deeper gratings we follow the work of Bruno, Reitich, and collaborators (e.g.,

Phil. Trans. Roy. Soc. London A 362, 2004) who utilize Geometric Optics corrections to iteratively update the rapidly varying amplitude which consists of many slowly-varying components. Our current contribution shows that the iterative update scheme can be eliminated and replaced with a simultaneous solution procedure. While our central ideas can be extended to the full vector electromagnetic time-harmonic Maxwell equations, we focus upon the case of two-dimensional linear acoustics for simplicity.

David P. Nicholls
University of Illinois at Chicago
nicholls@math.uic.edu

Fernando Reitich
School of Mathematics
University of Minnesota
reitich@math.umn.edu

MS59

WaLBerla: HPC Software Design for Computational Engineering Simulations

WaLBerla is a massively parallel software framework supporting a wide range of physical phenomena. We describe the software designs realizin the major goal of the framework, a good balance between expandability and scalable, highly optimized, hardware-dependent, special purpose kernels. In this talk, we discuss the coupling of our Lattice-Boltzmann fluid flow solver and a method for fluid structure interaction. Additionally, we show a software design for heterogeneous computations on GPU and CPU utilizing optimized kernels.

Christian Feichtinger
Chiar for System Simulation
University of Erlangen-Nuremberg, Germany
christian.feichtinger@informatik.uni-erlangen.de

Jan Götz
Universität Erlangen-Nürnberg
jan.goetz@informatik.uni-erlangen.de

Stefan Donath
Computer Science 10 - System Simulation
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
stefan.donath@informatik.uni-erlangen.de

Harald Köstler
Universität Erlangen-Nürnberg
harald.koestler@informatik.uni-erlangen.de

Ulrich Rüde
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
ulrich.ruede@informatik.uni-erlangen.de

MS59

Parallel Performance of Spectral-Element Discontinuous Galerkin Lattice Boltzmann Method (SEDG-LBM) for Flow Simulations on BG/P

We present the spectral-element discontinuous Galerkin lattice Boltzmann method for 3D flow simulations based on 19-velocity model. Our SEDG-LBM employs body-fitted unstructured meshes that enable us to deal with complex geometry with high-order accuracy.Explicit time marching, diagonal mass matrix, and minimum communication cost for numerical flux make our SEDG-LBM highly efficient in parallel. We demonstrate scalable parallel algorithms and parallel performance on the IBM BG/P at large scale.

Taehun Lee
Department of Mechanical Engineering
The City College of New York
thlee@ccny.cuny.edu

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

MS59

Million Body Simulations of Granular Dynamics on the GPU

This paper describes the formulation and implementation of an algorithm capable of simulating dynamic systems with over one million bodies on the GPU. The problem is formulated as a cone complementarity problem, and the corresponding quadratic optimization problem is solved via an iterative method. The GPU is leveraged via NVIDIAs CUDA architecture, solving the dynamics and the collision detection tasks in parallel. An example simulation of a tracked vehicle moving on granular terrain is shown.

Hammad Mazhar, Toby Heyn
University of Wisconsin Madison
Department of Mechanical Engineering
hmazhar@wisc.edu, heyn@wisc.edu

Alessandro Tasora
Università di Parma
tasora@ied.unpr.it

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

Dan Negrut
University of Wisconsin Madison
Department of Mechanical Engineering
negrut@cae.wisc.edu

MS59

In-Situ Visualization for Large-Scale Turbulent Combustion Simulations

As scientific supercomputing moves toward petascale and exascale, in-situ visualization clearly becomes one of the most feasible ways to enable scientists to see the full extent of the data generated by their simulations. This is particularly crucial for capturing and understanding highly intermittent transient phenomena, e.g. ignition/extinction events in turbulent combustion. However, integrating the visualization process into a simulation pipeline poses several challenges which have not been fully addressed. In this talk, we present our study of in-situ visualization for large-scale turbulent combustion simulations. We describe our design decisions and optimization strategies on domain decomposition, rendering, and image compositing. We present an in-depth evaluation of our implementation on the Cray XT5 at the NCCS/ORNL, and discuss the lessons

learned through the study. We demonstrate that in-situ visualization is a feasible solution and our work provides a valuable reference to those who will pursue a similar solution.

Hongfeng Yu
Sandia National Laboratory
hyu@sandia.gov

MS60

An MPI/OpenMP Implementation of 3D FFTs for Plane Wave First Principles Materials Science Codes

Plane wave first principles codes based on 3D FFTs are one of the largest users of supercomputer cycles in the world. I will present results for an MPI/OpenMP version of our specialized 3D FFT that gives improved scaling over a pure MPI version on mulitcore machines. I will also present some preliminary results for the electronic structure code PARATEC using the new 3D FFT and threaded libraries to gain improved scaling to large core counts.

Andrew M. Canning
Lawrence Berkeley National Laboratory
acanning@lbl.gov

MS60

Hybrid Programming Models for Computational Chemistry

We describe and compare two programming models and their performance on the Cray XT-5 up to 120K cores on a number of simulations from computational chemistry and density functional theory. The first is a more traditional model based on a NUMA programming paradigm Global Array/ARMCI for NWChem and the second is a message-passing/multi-threaded task based model for MADNESS, an adaptive pseudo-spectral code.

George Fann
Oak Ridge National Laboratory
fanngi@ornl.gov

MS60

Evaluation of MPI/OpenMP Efficiency in Particle-in-cell Beam Dynamics Simulation

We present performance analysis of the particle-field decomposition method and the domain decomposition method in a parallel particle-in-cell beam dynamics simulation. Both parallelization approaches can be effective depending on the ratio of macroparticles to grid points. In both cases, the performance of the flat MPI implementation deteriorates rapidly with increasing number of cores per node. Using mixed MPI/OpenMP model reduces both communication and intra-node memory contention, leading to large performance gain.

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

Ji Qiang
Lawrence Berkeley National Laboratory
jqiang@lbl.gov

Robert D. Ryne
LBNL
RDRyne@lbl.gov

MS60

Multicore Parallelization of Determinant Quantum Monte Carlo Simulations

QUEST is a software package that implements the determinant Quantum Monte Carlo method for simulations of interacting electrons. Its main computational kernel is the evaluation of Green's functions. The existing approach is based on computing a graded decomposition of a long product of small matrices using the QR decomposition with column pivoting (QRP). Current implementations of QRP are limited by BLAS level 2 performance and do not scale well in modern multicore processors. In this talk, we will first discuss a number of alternative approaches and then focus a structured orthogonal factorization by exploiting the problem structure and avoiding the communication cost of pivoting. This new approach show promising parallel performance on multicore architectures.

Che-Rung Lee
National Tsing Hua University, Taiwan
cherung@gmail.com

Andres Tomas, Zhaojun Bai
University of California
andres@cs.ucdavis.edu, bai@cs.ucdavis.edu

Richard Scalettar
Department of Physics,
University of California, Davis, USA
scalettar@physics.ucdavis.edu

MS61

A Posteriori Analysis of a Multirate Numerical Method for Evolution Problems

We consider multirate integration methods for evolution models that present significantly different scales within the components of the model. We describe both an a priori error analysis and a hybrid a priori-a posteriori error analysis. Both analyses distinguish the effects of the discretization of each component from the effects of multirate solution while the effects on stability are reflected in perturbations to certain associated adjoint operators.

Donald Estep
Colorado State University
estep@math.colostate.edu or don.estep@gmail.com

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

S Tavener
Dept. of Mathematics
Colorado State university
tavener@colostate.edu

MS61

On the use of Newton-Krylov Methods in a General, Moment-based, Scale-bridging Algorithm

We are developing a consistent, scale-bridging algorithm

which accelerates an accurate solution of the fine scale equations to the point where coarse time and length scales are achievable. Our prototype fine scale problem has time, configuration space, and phase space as independent variables. This fine scale model is often referred to as the "kinetic" problem, and the discretized version of this problem will be referred to as the High Order (HO) problem. The kinetic problem could be solved by either a deterministic or Monte-Carlo approach. We will accelerate the solution to this fine problem using a Low Order (LO) problem as a coarse space preconditioner. The LO problem is derived from a small number of phase space moments of the HO problem, and can also be solved on a coarser configuration space mesh. We self-consistently determine the higher-order moments required by the LO problem with the local (in time and space) phase-space solution of the HO problem. The potential impact for these algorithmic advancements is best summarized by a selected list of applications. These could include: neutron transport; photon transport; plasma kinetic simulation; rarefied gas dynamics; transport in condensed matter e.g., semiconductors to name a few.

Dana Knoll
Los Alamos National Laboratory
nol@lanl.gov

MS61

Simulating Multi-Scale Multi-Physics Particle Accelerator Problems with the Adaptive Discontinuous Galerkin Method

Simulating the dynamics of charged particles in large accelerator facilities poses a variety of problems. It requires the simultaneous and self-consistent solution of Maxwell's equations and the relativistic equations of motion of the charge carriers. Moreover, the problem exhibits a pronounced multi-scale character. Short particle bunches (μm to a few mm) excite electromagnetic fields up to the THz regime. This is contrasted by a length of some meters of the accelerator sections considered. We present a time-domain discontinuous Galerkin (DG) method with *hp*-adaptation for solving such problems. Details of the both refinement types, the local regularity estimation and the dispersive properties are presented.

Sascha Schnepp
Graduate School of Computational Engineering
Technische Universitaet Darmstadt
schnepp@gsc.tu-darmstadt.de

Thomas Weiland
Institut fuer Theorie Elektromagnetischer Felder, TEMF
Technische Universitaet Darmstadt
weiland@temf.tu-darmstadt.de

MS61

Adaptive Solution of Multiphysics Coupled Problems Using the HERMES Library

Hermes2D is a C++ library for rapid development of space- and space-time adaptive hp-FEM/hp-DG solvers. Novel PDE-independent hp-adaptivity and multimesh assembling algorithms allow the user to solve a large variety of PDE problems ranging from stationary linear equations to complex time-dependent nonlinear multiphysics PDE systems. The library comes with a free interactive online lab powered by UNR computing facilities. Detailed tutorial enhanced with many benchmarks and examples allows the user to employ Hermes2D without being expert in object-oriented programming, finite element methods, or in the theory of partial differential equations. There is a very active user community where help can be obtained quickly. The code is distributed under the GNU General Public License.

Pavel Solin
Department of Mathematics and Statistics
University of Nevada, Reno
solin@unr.edu

MS62

Mpi4py and Petsc4py: Using Python to develop Scalable PDE Solvers

PETSc is a suite of routines and data structures for the scalable (parallel, MPI-based) solution of scientific applications modeled by PDE's. PETSc for Python (aka petsc4py) is a Cython-based wrapper to PETSc components like distributed vectors and matrices, Krylov-based linear solvers, Newton-based nonlinear solvers, and basic timesteppers. petsc4py facilitates the Cython/SWIG/f2py wrapping of C/C++/Fortran codes using PETSc. New mechanisms are being worked on to cross language boundaries and easily employ Python as an extension language.

Lisandro Dalcin
Centro Int. de Métodos Computacionales en Ingeniera
dalcinl@gmail.com

MS62

Paper to GPU: Optimizing and Executing Discontinuous Galerkin Operators in Python

I will present a domain-specific language that represents discontinuous Galerkin discretizations of linear and nonlinear hyperbolic partial differential equations. While trying to remain faithful to the human-readable, paper-based representation of such operators uncluttered by implementation detail, an optimizing processing pipeline brings this human-suited representation into a form that is suited to high-performance machine execution. The core of the talk is devoted to the design and components of this pipeline, which emits optimized C/C++ code that can be targeted at CPUs and GPUs in single-core and distributed memory arrangements.

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

MS62

A Distributed Architecture for Very Large-scale Geometric Computing

Field modelling and simulation dominate computational science and engineering, but entered biology only recently, to help understanding the mechanisms of life on a hierarchy of scales, from proteins to cells to tissues, organs and systems. Computer biomodels require to integrate petascale geometric data, multiscale and multiphysics simulations, concurrency-based adaptive behaviour and functional specialization. We introduce a computational format capable of combining geometric and physical representations and algorithms on distributed environments of any size.

Alberto Paoluzzi

Roma Tre University
paoluzzi@dia.uniroma3.it

Valerio Pascucci
University of Utah
pascucci.valerio@gmail.com

Giorgio Scorzelli
University of Utah and
Università Roma Tre
scrgiorgio@cogesic.it

MS62

Making a Python based PDE Solver Work Efficiently in Parallel with the Available Open Source Interfaces to MPI

This talk will describe our experiences in porting a serial python-based PDE solver (FiPy, http://www.ctcms.nist.gov/fipy) to run efficiently in parallel. These experiences include migrating a serial test suite to parallel, effective partitioning and interfacing with the available open-source parallel tools, while making minimal alterations to the original code base. Results from materials science models simulated on large clusters will be used to demonstrate the efficiency benefits and overhead associated with using these tools.

Daniel Wheeler
National Institute of Standards and Technology
daniel.wheeler@nist.gov

Jonathan Guyer
NIST
Metallurgy Division
jonathan.guyer@nist.gov

James O'Beirne
George Mason University
jobeirn1@gmu.edu

MS63

Mesh Quality and Monge-Kantorovich Optimization

Equidistribution of a monitor function has been a guiding principle in mesh generation for a long time. While in 1D equidistribution determines the mesh uniquely, an infinite number of meshes can satisfy a given equidistribution principle in 2D and 3D, implying that there is room for optimization. In this context, Monge-Kantorovich (MK) optmization choses one equidistributed mesh by minimizing the L2 norm of the grid point displacement of such mesh relative to an initial mesh. This procedure gives rise to a very robust and efficient multidimensional mesh generation/adaptation method, based on the solution of the nonlinear Monge-Ampere equation. The latter is solved with the multigrid preconditioned Jacobian Free Newton-Krylov method, delivering a scalable algorithm under grid refinement [G.L. Delzanno, et al., JCP 227, 9841 (2008)]. We will present mesh quality comparisons against alternative equidistribution methods for 2D static problems, showing that MK optimization produces good quality grids with nearly minimal grid cell distortion. It is also very robust against mesh tangling, maintaining good quality grids when other competitor methods fail. We will also present results on the application of MK optimization to 3D time-dependent mesh adaptation for the advection of a passive scalar with various levels of flow shear. Our results confirm the effectiveness and robustness (algorithmic and against mesh tangling) of the method, in addition to good mesh quality.

Gian Luca Delzanno
LANL
delzanno@lanl.gov

Luis Chacon
Oak Ridge National Laboratory
chaconl@ornl.gov

John M. Finn
Los Alamos National Laboratory
finn@lanl.gov

MS63

Mesh Generation for Modeling and Simulation of Carbon Sequestration Processes

To support modeling and simulation of carbon sequestration processes at specific geologic sites, we introduce an algorithm to generate Voronoi meshes conforming to non-convex domains that might include internal fractures. The Voronoi mesh is required to be non-regular to mimic properties found in geologic fracture systems. Thus the algorithm first generates a random point cloud using maximal Poisson disk sampling. Our target is to generate one million Voronoi cells in less than thirty seconds using a single processor.

Mohamed S. Ebeida, Patrick M. Knupp, Vitus Leung
Sandia National Laboratories
msebeid@sandia.gov, pknupp@sandia.gov, vjleung@sandia.gov

MS63

Generation and Optimization of Prismatic-tetrahedral Hybrid Meshes for Complex Biomedical Geometries

We present a method to generate and optimize hybrid meshes with prisms and tetrahedra. Our method first generates a layered tetrahedral mesh, and then generates prismatic boundary layers by advancing the surface and shrinking the tetrahedra using a variational procedure that optimizes the prisms while preserving the quality of tetrahedra. The resulting mesh is further improved using local edge flipping and smoothing. Experimental results demonstrates its effectiveness for complex biological geometries.

Xiangmin Jiao
Stony Brook University
Stony Brook, NY
xjaio@sunysb.edu

Daniel R. Einstein
Pacific Northwest National Laboratory
daniel.einstein@pnl.gov

Vladimir Dyedov
Department of Applied Mathematics & Statistics
Stony Brook University
vladimir@ams.sunysb.edu

Andrew Kuprat

Pacific Northwest National Laboratory
andrew.kuprat@pnl.gov

Navamita Ray
Stony Brook University
nray@ams.sunysb.edu

MS63

A New Strategy for Untangling Meshes via Node-Movement

A new mesh optimization strategy for untangling quadrilateral meshes, based on node-movement, is investigated. The strategy relies on a set of Propositions which show that, for certain non-negative quality metrics within the Target-matrix paradigm, if the value of the metric is less than 1, then the local area is positive. The Propositions are exploited to devise a new strategy for simultaneous mesh untangling and quality improvement. Numerical results confirm the expected behavior.

Patrick M. Knupp
Sandia National Laboratories
pknupp@sandia.gov

Jason Franks
University of Washington
Seattle, WA
jasonwfranks@gmail.com

MS64

AC-SAMMM - The Aachen Platform for Structured Automatic Manipulation of Mathematical Models - A Case Study

We apply second-order adjoint sensitivity analysis to an industrial polymerization process given by about 2000 parametric differential-algebraic equations (DAEs). A Modelica model of the process is automatically transformed into a C-implementation of the DAE residuals. The derivative code compiler dcc is applied to this C-implementation to yield first-order tangent-linear and adjoint derivatives. Second-order adjoint derivatives are obtained by reapplication of dcc to its own output. The obtained derivatives models are embedded into an extension of the ESO class (equation set object class) of the CAPE-OPEN standard.

Ralf Hannemann
Aachener Verfahrenstechnik
RWTH Aachen University
ralf.hannemann@avt.rwth-aachen.de

Boris Gendler, Michael Förster
STCE
RWTH Aachen University
bgendler@stce.rwth-aachen.de,
foerster@stce.rwth-aachen.de

Jutta Wyes, Moritz Schmitz, Wolfgang Marquardt
Aachener Verfahrenstechnik
RWTH Aachen University
jutta.wyes@avt.rwth-aachen.de,
moritz.schmitz@avt.rwth-aachen.de,
wolfgang.marquardt@avt.rwth-aachen.de

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

MS64

Three Dimensional Tomography in Atmospheric Remote Sensing

The upper troposphere / lower stratosphere plays a key role in the climate system and is one of the least understood atmospheric regions. Next generation remote sensing instruments allow for tomographic measurement of this part of the atmosphere and will improve our understanding of dynamical, radiative, and chemical processes in this region significantly. The amount of data obtained by these new instruments is enormous and new retrieval schemes have to be developed. In this talk we present a new algorithm optimized for massive 3-D retrievals of several hundred thousands of measurements and atmospheric constituents. Different regularization approaches and minimizers to solve the non-linear problem are discussed. A thematic focus will be on automatic differentiation by operator overloading and source code transformation.

Martin Kaufmann
Institute for Energy and Climate Research - Stratosphere
Juelich Research Centre
m.kaufmann@fz-juelich.de

Joern Ungermann, Lars Hoffmann, Joerg Blank
Institute of Chemistry and Dynamics of the Geosphere
Juelich Research Centre
j.ungermann@fz-juelich.de, l.hoffmann@fz-juelich.de,
j.blank@f-juelich.de

Martin Riese
nstitute of Chemistry and Dynamics of the Geosphere
Juelich Research Centre
m.riese@fz-juelich.de

Johannes Lotz
LuFG Software and Tools for Computational Engineering
RWTH Aachen University
johannes.lotz@rwth-aachen.de

Ebadollah Varnik
Software and Tools for Computational Engineering
RWTH Aachen University
varnik@stce.rwth-aachen.de

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

MS64

A Gentle Introduction to Algorithmic Differentiation and Discrete Adjoints

Potentially very large gradients and Hessian of multivariate functions $F : R^n \to R$ play an often crucial role in modern CSE algorithms. Traditional numerical differentiation by finite differences as well as exact (up to machine accuracy) differentiation using tangent-linear models exhibit a computational complexity of $O(n) \cdot Cost(F)$ ($O(n^2) \cdot Cost(F)$) for the evaluation of the gradient (Hessian) which is likely to become infeasible for large-scale numerical simulations. Discrete adjoints compute the same gradients (Hessians) with machine accuracy at $O(1) \cdot Cost(F)$ ($O(n) \cdot Cost(F)$), thus saving an order of maginitude in the computational

cost. Their (semi-)automatic generation by tools for Algorithmic Differentiation (AD) yield a number of challenges - both for AD-tool developers and users. This talk is meant to set the stage for the following talks by giving a gentle introduction to AD, discrete adjoints, and their use within various numerical algorithms.

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

MS64

Error Estimation in Geophysical Fluid Dynamics with Algorithmic Differentiation

Uncertainty quantification is an essential part of model development in the geosciences. We present in this talk variations of dual weight error estimation for goals calculated from the solution of a shallow water model. The necessary adjoint solutions are obtained with an Algorithmic Differentiation tool. The local error estimates are obtained with deterministic and stochastic functionals of the flow. We show that the stochastic interpretation of dual weight error estimation allows ensemble-type uncertainty quantification for physical quantities derived from a single model run.

Florian Rauser
Max Planck Institute for Meteorology
Ocean in the Earth System
florian.rauser@zmaw.de

Peter Korn
Max Planck Institute for Meteorology
peter.korn@zmaw.de

MS65

Programming Motifs 2: Adaptive Mesh Refinement (High Level)

Efforts to write adaptive mesh refinement code in Chapel have been strikingly successful, demonstrating a reduction in code size by factors of 10 to 20 compared to existing implementations. This talk will present a high-level look at block-structured AMR, detailing its various geometric challenges and the Chapel features used to address them.

Jonathan Claridge
University of Washington
Applied Mathematics
claridge@gmail.com

MS65

Programming Motifs 2: Adaptive Mesh Refinement (Details)

This talk will explore in greater detail several core elements of the code. This includes Chapel's capacity for dimension-independent programming, our method for cleanly handling unions of rectangular index sets, and a remarkably simple implementation of the Berger-Rigoutsos partitioning algorithm.

Jonathan Claridge
University of Washington
Applied Mathematics
claridge@gmail.com

MS65

Programming Motifs 1: Linear Algebra

Our evaluation of Chapel is built around implementations of the Berkeley programming motifs, which we discuss briefly. We illustrate Chapel's power with some specific examples from blocked dense linear algebra algorithms and from graph algorithms relevant to sparse matrix computations.

John G. Lewis
Cray, Inc.
jglewis@cray.com

MS65

An Overview of Chapel

Chapel is a new programming language, under development by Cray Inc., which is designed to simplify parallel programming on multicore workstations, commodity clusters, and supercomputers alike. Chapel strives to increase productivity for users of all levels by supporting greater abstraction than current parallel programming models, while also providing performance that meets or exceeds current technologies. This talk will provide an overview of the Chapel language, focusing on features utilized throughout the remainder of the minisymposium.

Jonathan Turner
University of Colorado
Computer Engineering
jonathan.turner@colorado.edu

MS66

Accelerating Linear Algebra Calculations using Statistical Techniques

We illustrate with two applications how linear algebra calculations can be enhanced by statistical techniques. The first example is related to condition numbers in linear systems or linear least-squares for which computation might be as expensive as the solution itself. In this case the computational cost can be reduced from one order of magnitude using statistical estimates. Another application concerns linear systems solution in which preconditioning by random matrices enables us to avoid pivoting.

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

MS66

Solving Linear Systems of Equations on Parallel Computers with Multiple Levels of Parallelism

In this talk we present an algorithm for solving dense systems of equations on hierarchical parallel machines, that exploits multiple levels of parallelism. The algorithm is based on CALU, a communication optimal LU factorization algorithm initially developed for distributed memory mchines. We evaluate its performance on clusters of multicore processors. We show that an implementation based on MPI and Pthreads leads to a better performance compared to routines implementing LU factorization in well-known numerical libraries.

Laura Grigori
INRIA
France

Laura.Grigori@inria.fr

Simplice Donfack
INRIA, France
sidonfack@gmail.com

MS66
Towards High Performance Algorithms for the Symmetric Eigenvalue Problems

The goal of this talk is to solve the symmetric eigenvalue problem on multicore/GPU architecture more quickly than existing implementations. The bulk of the existing algorithms is the reduction to tridiagonal form. It becomes judicious to efficiently develop new algorithms. In order to achieve this goal, we have developed an approach based on tiles algorithm that reduces a symmetric matrix to banded form and that find the eigenvalues using a band divide and conquer algorithm.

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

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

MS66
Speeding Householder Bidiagonalization

Iterations to find singular values of bidiagonal matrices are quick, due to the small number of nonzero entries. Householder reduction to bidiagonal form is relatively slower than QR or LU decomposition, which can be accomplished by BLAS-3 (matrix matrix multiplications) at near peak computational speeds. In LAPACK Householder bidiagonalization, half the operations are matrix vector multiplications (BLAS-2), which are relatively slow because only two floating operations are performed for each matrix element fetched to a CPU or GPU. One way to speed the computation is to combine two matrix vector multiplications (BLAS-2.5) Householder reduction from full to banded triangular from is BLAS-3, thus fast. This talk addresses the "bottleneck" reduction from banded to bidiagonal form.

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

MS67
Recovery-Based A Posteriori Error Estimation

Estimators of the recovery type possess a number of attractive features: their ease of implementation, generality, and asymptotical exactness, that have led to their popularity in the engineering community. However, it is well-known that for interface problems with large jumps, existing estimators of the recovery type over-refine regions where there are no errors and, hence, fail to reduce the global error. In this talk, I will first explain why they fail and how to fix this structural failure. I will then introduce new recovery-based estimators for conforming, nonconforming, mixed, and discontinuous Galerkin elements. It is shown theoretically and numerically that these estimators are robust with respect to the jumps of diffusion coefficients. Moreover, these estimators do not require triangulation being aligned with physical interfaces, which is essential for their applications to nonlinear interface problems of practical interest.

Zhiqiang Cai
Purdue University
Department of Mathematics
zcai@math.purdue.edu

MS67
Adjoint Error Estimation Formulations for Nonlinear Algorithms Applied to Linear Advection and Advection-Diffusion Equations

A theoretical framework is developed for adjoint-based methods of *a posteriori* error estimation with nonlinear algorithms. Application is restricted to linear advection and advection-diffusion equations to delineate the technical difficulties introduced specifically by the algorithmic nonlinearity. Development of the correct adjoint operator and efficiency in implementation are discussed, with computational examples. This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Jeffrey M. Connors
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
connors4@llnl.gov

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banks20@llnl.gov

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

Carol Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

MS67
A Posteriori Error Estimates and Adaptive Error Control for Probabilistic Sensitivity Analysis

We consider the nonparametric density estimation problem both the forward and inverse sensitivity analysis of differential equations. We describe an a posteriori error analysis for computed probability distributions that accounts for both stochastic and deterministic sources of error. The error estimate is precise and can be used for a general adaptive error control that balances all sources of error. We describe applications to both ordinary differential equations and elliptic problems.

Troy Butler
Institute for Computational Engineering and Sciences
University of Texas at Austin
tbutler@ices.utexas.edu

Donald Estep
Colorado State University
estep@math.colostate.edu or don.estep@gmail.com

Axel Malqvist
Department of Information Technology
Uppsala University
axel.malqvist@it.uu.se

Simon Tavener
Colorado State University
tavener@math.colostate.edu

MS67

A Posteriori Error Measures and Adaptive Local Refinement with First-Order System Least-Squares (FOSLS) Finite-Element Methods

Standard FOSLS formulation of a system of PDEs enjoys a simple a posteriori error measure that is both locally sharp and globally reliable. One would prefer an error measure that is also a local upper bound on the error. FOSLS local error measure provides an upper bound in a special semi-norm of the error, similar to an H^1 semi-norm. This talk will include a discussion of implications of this distinction. A hybrid FOSLS/FOSLL* formulation will be introduced that reduces local L^2 error in the approximation, while retaining the FOSLS semi-norm bounds.

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

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

Lei Tang, Kuo Liu
University of Colorado
lei.tang@colorado.edu, kuo.liu@colorado.edu

MS68

Immersed Elastic Structure Dynamics in Viscoelastic Fluids

Many biological fluids are non-Newtonian and exhibit viscoelastic responses. Here we discuss recent results obtained using an immersed boundary framework to study the interaction between immersed elastic structures and surrounding viscoelastic fluids.

John Chrispell
Mathematics Department
Tulane University
jchrispe@tulane.edu

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

MS68

A Monolithic, Conservative Finite Element Method for Darcy-Stokes Coupling: Towards an Optimal Error Analysis

Recently, finite element methods for Darcy-Stokes coupling based on Raviart-Thomas elements have been introduced. These elements have the right continuity properties at the interface to implement the Beavers-Joseph-Saffman conditions. Indeed, mass conservation holds strongly independent of permeability, and the possibility of tangential slip avoids boundary layers in the porous medium. We show optimal convergence of the method in L^2 on the whole domain. Furthermore, we address the question of preconditioning the resulting discrete problems.

Guido Kanschat
Department of Mathematics
Texas A&M University
kanschat@tamu.edu

MS68

Error Estimate for Discontinuous Finite Element Approximation to Interface Problems under Minimum Regularity

An error analysis of discontinuous Galerkin methods is derived for the interface problems under minimum regularity assumption. This is achieved by using the technique involving a posteriori error estimate introduced by Gudi [1]. [1] T. Gudi, A new error analysis for discontinuous finite element methods for linear elliptic problems, Math. Comp., electronically published on April 12, 2010.

Xiu Ye
University of Arkansas, Little Rock
xxye@ualr.edu

MS68

A High Order Cartesian Grid Method for the Helmholtz Equation with Geometrically Complicated Material Interfaces

Across the dielectric interfaces, the electromagnetic wave solutions are usually non-smooth or even discontinuous, so that our effort in designing efficient algorithms is easily foiled, unless the complex interfaces are properly treated. We have recently introduced a novel higher order finite difference method – the matched interface and boundary (MIB) method, for solving the Helmholtz equation with arbitrarily curved dielectric interfaces based on a simple Cartesian grid. Like other Cartesian grid methods, the MIB method in some sense fits the numerical differentiation operators to the complicated geometries. Nevertheless, the MIB method distinguishes itself from the existing interface methods by avoiding the use of the Taylor series expansion and by introducing the concept of the iterative use of low order jump conditions. The difficulty associated with other interface approaches in extending to ultra high order is thus bypassed in the MIB method. In solving waveguides with straight interfaces, the MIB interface treatment can be carried out systematically so that the proposed approach is of arbitrarily high order, in principle. Orders up to 12 are confirmed numerically for both transverse magnetic (TM) and transverse electric (TE) modes. In solving waveguides with curved interfaces, the enforcement of jump conditions couples two transverse magnetic field components, so that the resulting MIB method becomes a full vectorial approach. The full vectorial MIB method has been shown to be able to deliver a fourth order of accuracy in treating arbitrarily curved interfaces.

Shan Zhao
Department of Mathematics
University of Alabama
szhao@bama.ua.edu

MS69

Comparison of Some Aspects of Nodal and hp-FEM Methods for Nuclear Reactor Simulations

The highly efficient nodal methods have originated in the nuclear engineering community, where they quickly overshadowed the more demanding FVM or low-order FEM. To our best knowledge, though, nodal methods received little attention in other engineering fields, where the FEM were more dominant. In an effort to clarify their relationship, we will solve a neutron diffusion problem with both an hp-adaptive FEM and a transverse integrated nodal method and discuss their similarities and differences.

Milan Hanus
Department of Mathematics
University of West Bohemia, Czech Republic
mhanus@kma.zcu.cz

MS69

Combined Adaptive Multimesh Hp-Fem/hp-Dg for Multiphysics Coupled Problems Involving Compressible Inviscid Flows

During the last decade, Discontinuous Galerkin (DG) methods have been studied by numerous researchers in the context of different problems ranging from linear elliptic equations to Euler equations of compressible inviscid flow. It is well known that DG methods yield larger discrete problems than standard continuous finite element methods (FEM). On the other hand, their implicit stabilization through embedded numerical fluxes makes them particularly well suited for hyperbolic flow problems. Thus for multiphysics problems involving compressible flow we propose to combine the best of both worlds: DG is used for the flow part only while standard FEM is employed for second-order equations where it works very well.

Lukas Korous
Charles University in Prague
Faculty of Mathematics and Physics
lukas.korous@gmail.com

Milan Hanus
Department of Mathematics
University of West Bohemia, Czech Republic
mhanus@kma.zcu.cz

MS69

Fully-coupled AMG Preconditioned Newton-Krylov Solution Methods for Strongly-coupled Multi-physics Applications

This study considers the performance of a fully-coupled algebraic multilevel preconditioner for Newton-Krylov solution methods. The performance of the preconditioner is demonstrated on a set of challenging multiphysics PDE applications: drift-diffusion approximations for semiconductor devices; a low Mach number formulation for flow, transport and non-equilibrium chemical reactions; and a low flow Mach number resistive magnetohydrodynamics (MHD) system. The AMG preconditioner is based on an aggressive-coarsening graph-partitioning of the nonzero block structure of the Jacobian matrix.

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

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

MS69

Rheagen: The Rheology Application Engine

The FEniCS project with Automated Scientific Computing aims to present a higher level of abstraction to user of scientific software. Equations are entered directly as weak forms in Python, from which efficient back-end finite element code is generated. We have enriched the language in order to accommodate many complex fluid models. We will demonstrate this for several difficult benchmark problems in Oldroyd fluids using numerous discretization and stabilization schemes.

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

Matthew G. Knepley
University of Chicago
knepley@mcs.anl.gov

Ridgway Scott
Department of Computer Science
University of Chicago
ridg@cs.uchicago.edu

MS70

Analyzing Optimization Models Developed with the Pyomo Modeling Software

We describe Pyomo, an open source tool for modeling optimization applications in Python. Pyomo can be used to define symbolic problems, create concrete problem instances, and solve these instances with standard solvers. Pyomo provides a capability that is commonly associated with algebraic modeling languages such as AMPL, AIMMS, and GAMS, but Pyomo's modeling objects are embedded within a full-featured high-level programming language with a rich set of supporting libraries. Pyomo leverages the capabilities of the Coopr software library, which provides interfaces to optimizers that can analyze a rich array of optimization models, including linear and integer programs, stochastic programs and generalized disjunctive programs.

William E. Hart
Sandia National Labs, USA
wehart@sandia.gov

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

David Woodruff
University of California, Davis
dlwoodruff@ucdavis.edu

MS70

A Flexible Python Environment for PDE-Constrained Optimization

This talk gives an overview of NLPy, a toolkit for optimization in Python, and decribes how NLPy interacts with the FEniCS framework for the description of variational problems to yield a convenient and powerful environment for PDE-constrained optimization. A new solver and its implementation are covered. Several applications are used to illustrate the concepts and demonstrate usage of the software.

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

Nick Gould
Numerical Analysis Group
Rutherford Appleton Laboratory
nick.gould@stfc.ac.uk

Sue Thorne
Rutherford Appleton Laboratory
sue.thorne@stfc.ac.uk

MS70

Algorithmic Differentiation in Python

We discuss what kind of derivatives can be evaluated efficiently using Algorithmic Differentiation and give a brief overview of univariate Taylor polynomial arithmetic. We also describe how numerical linear algebra functions like the QR decomposition fit into the framework, describe how sparsity in derivative tensors can be exploited and show how arbitrary mixed partial derivatives can be evaluated using interpolation methods. We show examples and give pointers to available software in Python.

Sebastian F. Walter
Humboldt-Universität zu Berlin
sebastian.walter@gmail.com

MS70

Stochastic Nonlinear Programming with Pyomo

We describe PySP, an open-source extension of Pyomo - a Python-based modeling language for mathematical programming - that enables modeling and solution of stochastic mixed-integer programs. PySP contains a number of generic decomposition-based solution strategies made possible through Python language features such as introspection. We discuss the design and implementation of these generic strategies, in addition to computational results on standard stochastic benchmarks.

Carl Laird
Assistant Professor
Texas A&M University
carl.laird@tamu.edu

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

MS71

System Requirements for Enabling Graph Algorithms on Massively Parallel Computers

Many scientific, national security, and business applications need to process irregular, unstructured data. Unlike scientific applications based on linear algebra routines, these applications comprise large graph computations with irregular memory patterns, low computation to memory ratios, and abundant fine grain parallelism that synchronizes frequently. Traditional architectures optimized to run floating point intensive simulations are inadequate. In this talk I discuss the requirements for graph based applications motivated by a clustering algorithm for generating Delaunay meshes.

John Feo, Jace Mogill
Pacific Northwest Laboratory
john.feo@pnl.gov, jace.mogill@pnl.gov

MS71

Approximate Graph Operations on Parallel Platforms

Abstract not available at time of publication.

Ananth Grama
Purdue University
Department of Computer Science
ayg@cs.purdue.edu

MS71

Multithreaded Algorithms for Graph Coloring

Graph coloring is an abstraction for partitioning a set of binary related objects into few independent subsets. It is used, among others, to discover concurrency in parallel computing. There exist linear-time heuristic algorithms with good solutions for this problem, but they are challenging to parallelize due to limited concurrency and a high ratio of data access to computation. We present two kinds of multithreaded coloring algorithms: the first is targeted for the Cray-XMT and the second for general shared-memory architectures. We present results on carefully chosen input graphs, covering various classes of graphs, with up to billions of edges.

Mahantesh Halappanavar
Pacific Northwest National Lab
mahantesh.halappanavar@pnl.gov

Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
umit@bmi.osu.edu

John Feo
Pacific Northwest Laboratory
john.feo@pnl.gov

Assefaw H. Gebremedhin
Purdue University
agebreme@purdue.edu

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

MS71

Optimizing Short-read Genome Assembly Algorithms for Emerging Multicore Platforms

We present a new parallel implementation for de novo assembly problem of large-scale genomic data on multicore clusters. This new approach belongs to the family of Eulerian path-based methods to de novo genome assembly, and involves construction, traversal, and simplification of a large string graph. We will discuss parallelization of various steps of the assembly process, focusing on techniques for scalable graph construction and simplification.

Kamesh Madduri
Lawrence Berkeley National Lab
kmadduri@lbl.gov

MS72

Emulating the Nonlinear Matter Power Spectrum for the Universe

Many of the most exciting questions in astrophysics and cosmology, including most observational probes of dark energy, rely on an understanding of the nonlinear regime of structure formation. In order to fully exploit the information available from this regime and to extract cosmological constraints accurate theoretical predictions are needed. Currently such predictions can only be obtained from costly, precision numerical simulations. This work is aimed at constructing an accurate predictor of the nonlinear mass power spectrum on Mpc scales for a wide range of currently viable cosmological models, including dark energy. We use the Coyote Universe simulation suite which comprises nearly 1,000 N-body simulations at different force and mass resolutions, spanning 38 wCDM cosmologies. This large simulation suite enables us to construct a prediction scheme for the nonlinear matter power spectrum accurate at the percent level for large wave numbers. We present this scheme and discuss the tests we have done to ensure its accuracy, and discuss how it can be used to estimate cosmological parameters.

Dave Higdon
Los Alamos National Laboratory
Statistical Sciences
dhigdon@lanl.gov

Earl Lawrence, Katrin Heitmann, Salman Habib
Los Alamos National Laboratory
earl@lanl.gov, heitmann@lanl.gov,
salmanhabib@earthlink.net

Martin White
Lawrence Berkeley National Lab
mwhite@berkeley.edu

Christian Wagner
Astrophysikalisches Institut Potsdam (AIP), Potsdam, Germany
cwagner@aip.de

MS72

Gradient Enhanced Universal Kriging Model for Inexpensive Uncertainty Quantification in Reactor Safety Simulations

In this work we discuss an approach for uncertainty propagation through computationally expensive physics simulation codes. Our approach incorporates gradient information to provide a higher quality surrogate with fewer simulation results compared with derivative-free approaches. In turn, we create an Gaussian probabilistic model for the system response which, coupled with input uncertainty information provides a complete uncertainty approach when the physics simulation code can be run at only a small number of times. We demonstrate that explicitly modeling the mean process substantially increases the efficiency of the approach when compared with ordinary kriging and that using a Gaussian process substantially reduces the error when compared to regression approaches using the same information. We demonstrate our findings on synthetic functions as well as nuclear reactor models.

Brian Lockwood
University of Wyoming
blockwoo@uwyo.edu

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

MS72

Parallelization Schemes for Constructing Effective Surrogate Models of Large Scale Computer Simulations

We will present in this talk some effective computational strategies for constructing surrogate models (Bayes Linear Models) from an ensemble of expensive computer models. Such surrogates are crucial to uncertainty quantification of many classes of models of physical systems modeled by partial differential equations. Careful strategies for parallel construction of the inverse of the covariance matrix are used to render this computation tractable. While similarities do exist between these and more classical iterative solution methods there are also interesting differences.

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

Keith Dalbey
Sandia National Laboratories
kdalbey@sandia.gov

E. Bruce Pitman
Dept. of Mathematics
Univ at Buffalo, Buffalo, NY 14260
pitman@buffalo.edu

Elena Stefanescu
SUNY at Buffalo Dept of Mechanical Engineering
ers32@buffalo.edu

Matthew Jones
Center for Comp. Research,
Univ. at Buffalo, Buffalo, NY 14260
jonesm@ccr.buffalo.edu

MS72

Bayesian Model Analysis Using Parallel Adaptive Multi-Level Sampling Algorithms

We present a parallel adaptive algorithm for sampling multimodal distributions and computing Bayesian model evi-

dences. The algorithm samples a sequence of distributions (levels) that converge to the final multimodal distribution, communicating sampling information among adaptively selected levels in order to correctly capture the volume proportions among final modes. The parallel version of the algorithm considers load balancing as well. We show parallel computational examples on the analysis of some turbulence models.

Sai Hung Cheung
Institute for Computational Engineering and Sciences
saihung@ices.utexas.edu

Todd Oliver
PECOS/ICES, The University of Texas at Austin
oliver@ices.utexas.edu

Ernesto E. Prudencio
Institute for Computational Engineering and Sciences
prudenci@ices.utexas.edu

Karl W. Schulz
Institute for Computational Engineering and Sciences
The University of Texas at Austin
karl@ices.utexas.edu

MS73

Computational Islet Simulations of Pancreatic Beta Cells

One of the causes of diabetes is the failure of beta cells to secrete insulin in response to blood glucose levels. Beta cell is the most prevalent cell type in the islets of the endocrine system. By using robust numerical methods and efficient programming techniques, we have created an extensible, efficient, and functional computation islet simulator in Matlab that can simulate islets with 1000 or more beta cells, sufficient for physiologically representative simulations. Advisor: Dr. Matthias Gobbert

Christopher Raastad
University of Washington
raastad@u.washington.edu

MS73

Solving a NIST Suite of PDE Benchmark Problems with Adaptive Low-Order FEM and hp-FEM

This paper is based on a recently NIST Report by Mitchell containing a collection of PDE benchmark problems. These feature various phenomena that pose challenges to adaptive finite element algorithms. After a brief review of main differences in automatic adaptivity for low-order and higher-order FEM including hanging nodes and spatially and polynomially anisotropic refinements, we use the benchmarks in the suite to discuss and compare various aspects of automatic adaptivity in low-order FEM and hp-FEM. Advisor: Dr. Pavel Solin; University of Nevada, Reno

Erick A. Santiago
University of Nevada, Reno
laviticus@sbcglobal.net

MS73

The Physics and Engineering of Skateboardings MegaRamp

The MegaRamp was designed by a professional skateboarder based on his experience. There has been no formal research performed to analyze the ramp. This research uses lumped-parameter modeling techniques supplemented with motion tracking data to develop equations of motion that accurately model the dynamics of a skateboarder on the MegaRamp. This model will be used as a quantitative tool to determine if the ramp can be modified to improve a skateboarders performance. Advisors: Kevin Fite (Mechanical Engineering) and Aaron Luttman (Mathematics)

Emily A. Stefano
Clarkson University
stefanea@clarkson.edu

MS73

Developing Computational Tools to Predict Material Behavior

By solving the Schrdinger Equation for many different atoms, we can predict the behavior of materials without having to make higher level assumptions. However, it is currently difficult or even impossible to solve the equation for millions of atoms. We present a method that uses a Discrete Wavelet Transform to downsize the matrix arising from the discretization of the Schrdinger Equation in order to compute the ground state energy of the system more efficiently. Advisors: Dr. Malena Espanol and Prof. Michael Ortiz

Stephanie Tsuei
California Institute of Technology
stsuei@caltech.edu

MS73

A Generalized Monte Carlo Loop Algorithm for Frustrated Ising Models

Monte Carlo (MC) simulation of some frustrated 2D lattice spin models (Ising models) at low temperatures can be prohibitively slow due to an extensive number of ground-state or near-ground-state spin configurations separated by large energy barriers to single spin flips. In this paper, we introduce a Generalized Loop Move (GLM) that uses the dual graph of a 2D Ising model to overcome this slowness and demonstrate its effectiveness in several cases where standard MC is ineffective. Advisors: Hans De Sterck, University of Waterloo; Roger Melko, University of Waterloo

Yuan Wang
University of Waterloo
tba@siam.org

MS74

Adjoint Methods for Error Control and Adaptivity in Ocean Models

We present an error control framework, based on the solutions of adjoint problems on potentially coarser scales, that allows for reliable error estimation and grid adaption. The framework supports explicit, multirate time integration. In addition, resource saving methods such as block adaptivity and compensated domain decomposition can be used to solve the adjoint problem. We illustrate the method on

a variety of single and multi-stack shallow water models.

Varis Carey, Donald Estep
Colorado State University
carey@math.colostate.edu, estep@math.colostate.edu or don.estep@gmail.com

MS74

A Multirate Time Integration Scheme and A Posteriori Error Estimates for Ocean Circulation Modeling

n ocean circulation modeling, the system of modeling equations is often split into a 2D, depth-independent barotropic subsystem and a 3D, depth-dependent baroclinic system. In this presentation we describe a multirate discontinuous Galerkin method for integrating the subsystems in time. The scheme allows us to use different time steps for different subsystems. We also present an a posteriori estimate of the error at the final time f computation. This estimator may be used for adaptive meshing

Brandon Chabaud
Pennsylvania State University
chabaud@math.psu.edu

MS74

A Scale-invariant Formulation of the Anticipated Potential Vorticity Method

This is the first of a series of efforts to develop scale-aware subgrid parametrization schemes on variable-resolution grids. We focus on the anticipated potential vorticity method (APVM) on quasi-uniform grids with varying resolutions. By a scale analysis technique and the phenomenological theories for two-dimensional turbulence, we derive a scale-invariant formulation of the APVM, which depends on one single parameter. Then utilizing a basic optimization technique, we determine the optimal value of the parameter for the APVM through numerical experiments.

Qingshan Chen
Florida State University
qchen3@fsu.edu

Max Gunzburger
Florida State University
School for Computational Sciences
gunzburg@csit.fsu.edu

Todd Ringler
Los Alamos National Laboratory
ringler@lanl.gov

MS74

Title Not Available at Time of Publication

Abstract not available at time of publication.

Todd Ringler
Los Alamos National Laboratory
ringler@lanl.gov

MS75

Automatic Differentiation and Geometry Adaptation Techniques for Shape Optimization of Fluid-Based Systems

In the context of shape optimization of fluid-based systems involving complex fluids, the aspect of sensitivity of optimal shapes with respect to the fluid model parameters is examined. The parametric sensitivities are obtained by applying automatic differentiation tools to the entire optimization tool chain. A second aspect of practical geometry modification for complex shapes is also addressed. Applications include blood flow devices and plastics extrusion dies.

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

Stefanie Elgeti, Mike Nicolai
RWTH Aachen University
elgeti@cats.rwth-aachen.de, nicolai@cats.rwth-aachen.de

Markus Probst
RWTH Aachen University
Chair for Computational Analysis of Technical Systems
probst@cats.rwth-aachen.de

MS75

Development and Application of a Discrete Adjoint Approach for Unsteady Flow Control

We discuss the development and application of an efficient discrete adjoint solver for unsteady, viscous, incompressible flow control problems. The discrete adjoint flow solver is developed by the use of automatic differentiation in reverse mode. Here, the binomial checkpointing algorithm 'revolve' is combined with the adjoint flow solver for optimal reduction in memory requirements.

Nicolas R. Gauger, Anil Nemili, Emre Oezkaya
Department of Mathematics and CCES
RWTH Aachen University
gauger@mathcces.rwth-aachen.de, nemili@mathcces.rwth-aachen.de, oezkaya@mathcces.rwth-aachen.de

MS75

Heliostat Layout Optimization via Automatic Differentiation: Adjoints and Convex Relaxations

Concentrated solar thermal power generation with a central receiver is considered. Its major costs are the land area required and the heliostats used to concentrate the insolation. Heliostats are optimally placed via a model developed. Automatic differentiation techniques are used for the calculation of derivatives, convex relaxations and their subgradients. The theory of convex relaxation of algorithms is summarized and numerical results from the heliostat placement and other case studies are presented.

Corey Noone, Alexander Mitsos
Department of Mechanical Engineering
Massachusetts Institute of Technology
noone@mit.edu, mitsos@mit.edu

MS75

Memory-Efficient Newton/Pantoja Method for Optimal Control Problems

We investigate a memory-efficient New-

ton/Pantoja method for the nu- merical solution of optimal control problems. Using a local minimum princi- ple we derive the first order necessary optimality conditions, that are equiv- alent to a nonlinear equation in appropriate Banach spaces. This equation is solved by a combination between the Newton and Pantoja methods. Each it- eration of the Newton method, i.e. each application of the Pantoja method to evaluate a search direction, contains three alternative sweeps through a time horizon, with a specific information dependence between different sweeps, so that the straightforward implementation of the algorithm would require a huge amount of memory to store all intermediate variables. To reduce this memory requirement we develop nested checkpointing techniques and prove some theoretical results concerning them. Finally, we discuss numerical ex- periences considering an optimal control problem for laser surface hardening of steel.

Julia Sternberg-Kaletta
Department of Mathematics
University of Hamburg
julia.sternberg@uni-hamburg.de

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

MS76

Shape Optimization of Chiral Structures for Low Reynolds Number Propulsion

Recent advances in micron-scale fabrication techniques allow for the construction of helically shaped and magnetically controlled artificial micro-swimmers. To facilitate the design of these novel transport devices, we conduct shape optimization simulations based on a boundary integral representation of the Stokes equations and a variational approach for the optimization problem. We determine those swimmer shapes that maximize speed in a particular direction and identify key improvements to the shapes fabricated in experimental studies.

Eric E. Keaveny
Courant Institute
New York University
ekeaveny@cims.nyu.edu

MS76

A Two-dimensional Numerical Study of a Permeable Capsule under Stokes Flow Condition

In the present work, a healthy red blood cell and a malaria-infected red blood cell in the presence of membrane permeability is modeled using two-dimensional immersed interface method. The fluid properties of the enclosed fluid and the surrounding fluid are assumed as the same. The results show that the healthy red blood cell gradually moves away from the vessel wall while the malaria-infected red blood cell rolls on the vessel wall due to adhesion. It is found that the resistance on the blood flow given by the malaria-infected red blood cell is higher than the corresponding resistance given by the healthy red blood cell. Moreover, the mass transfer characteristics of the healthy and malaria-infected red blood cells are investigated for various flow fields and the membrane properties.

Pahala Jayathilake, B.C. Khoo
National University of Singapore
Singapore
g0501295@nus.edu.sg, mpekbc@nus.edu.sg

MS76

Wetting Dynamics and Particle Deposition for An Evaporating Colloidal Drop: A Lattice Boltzmann Study

A three-dimensional (3D) lattice Boltzmann method (LBM) has been developed for multiphase (liquid and vapor) flows with solid particles suspended within the liquid phases. The method generalizes our recent 2D model to 3D, extends the implicit scheme to include inter-particle forces and introduces an evaporation model to simulate drying of the colloidal drop. The LBM is used to examine the dynamical wetting behavior of drops containing suspended solid particles on homogeneous and patterned substrates. The influence of the particle volume fraction and particle size on the drop spreading dynamics is studied as is the final deposition of suspended particles on the substrate after the carrier liquid evaporates. The final particle deposition can be controlled by substrate patterning, adjusting the substrate surface energies and by the rate of evaporation.

Ying Sun, Abhijit Joshi
Drexel University
ysun@coe.drexel.edu, joshi1974@gmail.com

MS76

Comparing Slender-body Formulation with the Method of Regularized Stokeslets

In this work we compare the slender-body formulation for an elastic slender body immersed in a viscous fluid to the method of regularized Stokeslets. Quantitative comparison is conducted, and both formulations will be applied to model and simulate the dynamics of primary cilium to elucidate both formulations. In particular, the experimental data on primary cilium are utilized and comparative results will be presented. This work is supported by NSF/CBET-0853673.

Yuan-Nan Young
Department of Mathematical Sciences
NJIT
yyoung@oak.njit.edu

MS77

Scalable Algorithms for Large-scale Inverse Wave Propagation

We consider algorithms for geophysical inverse problems with the goal of achieving large-scale parallel scalability. To this end, all parts of the computational pipeline need to be examined, and redeveloped where necessary. We present inexact Newton-Krylov iterative methods, where the Hessian is applied via the solution of forward and adjoint problems. These are solved in parallel using continuous and discontinuous Galerkin methods, where dynamic mesh adaptivity is applied to both the state and parameter fields.

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

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, Lucas Wilcox
University of Texas at Austin
georgst@ices.utexas.edu, lucasw@ices.utexas.edu

MS77

Parallel Multiscale Optimization Techniques in Nonlinear Mechanics

The parallel numerical solution of realistic mechanical problems, such as large-deformation contact between an elastic body and a rigid obstacle, often gives rise to nonlinear and possibly non-convex optimization problems. Thus, in order to succeed in computing a local minimizer of such optimization problems their solution is most oftenly carried out employing globalization strategies, i.e., Trust-Region and Linesearch methods. As is well-known, the paradigm of globalization strategies is to compute and to damp a search direction in order to achieve a descent in the value of a given objective function. Usually, search directions are computed as the solution of constrained quadratic programming problems. But, even if these quadratic programming problems are solved exactly, the damping of the search directions might yield a slow convergence of the overall scheme. Unfortunately, this effect often increases with the size of the optimization problem. Therefore, we present a class of nonlinear preconditioning strategies where a possible slow convergence is bypassed by computing search directions in parallel. In particular, the paradigm of these strategies is to locally solve certain nonlinear programming problems employing either Trust-Region or Linesearch methods. But, since these globalization strategies *asynchronously* compute local corrections, we must furthermore take care of the overall convergence of the method. As it turns out, this can be done by employing global control strategies yielding globally convergent, inherently parallel Linesearch (APLS) and Trust-Region (APTS) strategies. In this talk, we will therefore review the concept of nonlinear additively preconditioned globalization strategies and focus on the application of such strategies to the parallel solution of large-scale nonlinear optimization problems arising from the discretization of large deformation contact problems.

Christian Gross
Universita' della Svizzera Italiana
Institute of Computational Science
christian.gross@usi.ch

Rolf Krause
Universita' della Svizzera Italiana
rolf.krause@usi.ch

MS77

On the Treatment of Uncertainties in Aerodynamic Shape Optimization

The unavoidable presence of uncertainties poses several difficulties to the numerical treatment of optimization tasks. In this talk, we discuss a novel approach towards aleatory uncertainties for the specific application of optimal aerodynamic design under uncertainties. An appropriate robust formulation of the underlying deterministic problem and efficient approximation techniques of the probability space are investigated. Finally, algorithmic approaches based on multiple-setpoint ideas in combination with one-shot methods are presented as well as numerical results.

Claudia Schillings
University of Trier, Germany
Claudia.Schillings@uni-trier.de

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

MS77

HPC Issues in Shape and Topology Optimization

Shape and Topology optimization problems are a special sub-class of PDE constrained optimization. Concerning shape optimization problems, special care must be taken to treat the resulting mesh sensitivities properly, as otherwise the sensitivity information of the objective with respect to a change of the shape of the underlying domain can become prohibitively expensive to compute for a fine parameterization of the domain. One possible remedy here is shape calculus which can be used to arrive at a surface formulation of the shape gradient. The resulting numerical procedure is very efficient allowing one to use every surface mesh node position as the unknown for the shape. Concerning topology optimization, the special structure of the SIMP- method creates a very regular memory access pattern, which also leads to potentially very efficient schemes for novel hardware. As such, the implementation of this topology optimization approach is studied on modern GPUs.

Stephan Schmidt
University of Trier
Stephan.Schmidt@uni-trier.de

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

MS78

Adjoint-Based Numerical Error Estimation for the Unsteady Compressible Navier-Stokes Equations

We present practical adjoint-based strategies for estimating numerical discretization errors in scalar outputs obtained from unsteady simulations of the compressible Navier-Stokes equations. The discretization is discontinuous Galerkin in space and time, and the adjoint is obtained in a discrete fashion. We investigate various approximations of the fine-space adjoint in terms of error estimate cost and accuracy, and we present a method for separating effects of the spatial and temporal discretizations.

Krzysztof Fidkowski
University of Michigan
kfid@umich.edu

Isaac Asher
Iasher@umich.edu
university of michigan

MS78

Space-Time Error Estimation Via Dual Problems for CFD Problems

Space-time *a-posteriori* error estimates for derived outputs $M(u)$ are expressed in terms of residuals of the primal numerical problem u_h and weights involving the solution of a locally linearized dual problem ϕ

$$|M(u) - M(u_h)| < \sum_{time} \sum_{space} \mathbf{W}(\phi) \cdot \text{Residual}(u_h) \ .$$

This error representation can also be used as the basis for a mesh adaptivity and error control. Primal physical systems with genuine linearity may exhibit local instability (e.g. hydrodynamic instability, MHD instability, etc) while remaining globally stable. This often leads to a rapid growth in the dual problem and the weights $\mathbf{W}(\phi)$ appearing in *a-posteriori* error estimates. This rapid growth can make the control of errors in certain computed outputs problematic. In this presentation, we investigate the growth of dual problems for nonlinear systems such as the compressible Navier-Stokes equations. Particular attention is given to

- the choice of computed output $M(u)$,
- approximate local linearization,
- long time integration.

Numerical calculations include compressible flow simulation at low to moderate Reynolds number and preliminary calculations arising in 2-fluid plasma simulation.

Timothy J. Barth
NASA Ames Research Center
Timothy.J.Barth@nasa.gov

MS78

Error Estimation for Aleatoric Uncertainty Propagation using Stochastic Adjoints

Stochastic adjoint equations are introduced within a non-intrusive discrete sampling framework to enable efficient propagation of aleatoric uncertainties in systems governed by algebraic or differential equations involving random parameters. Adjoints are used to estimate the error (due to inexact reconstruction of the solution in stochastic space) in statistical moments of interest and the procedure is shown to exhibit super-convergence, in accordance with the underlying theoretical rates. Goal-oriented error indicators are then built using the adjoint solution and used to identify regions for adaptive sampling.

Karthik Duraisamy
Stanford
dkarthik@stanford.edu

MS78

Adjoint-based Discretization Error Estimation for Time Dependent Problems

Adjoint-based error estimation techniques for functional outputs in time-dependent computational fluid dynamics problems are discussed. Unsteady fluid flow problems on dynamically deforming unstructured meshes in two dimensions are considered, where the governing flow equations are discretized in arbitrary Lagrangian Eulerian (ALE) form. The discrete adjoint for the coupled fluid flow/mesh deformation equations is derived and solved using a backwards integration in time. The adjoint solution is then used to derive error estimates for spatial, temporal and algebraic error in the simulation, and for driving adaptive refinement strategies for reducing these error sources. Prospects for extending these techniques to more complex multiphysics simulations will also be considered.

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

MS79

Numerical Methods for Subsurface Flow in Karst Aquifer

In a karst aquifer, free flow and porous media flow are tightly coupled together, for which the Stokes-Darcy model has higher fidelity than either the Darcy or Stokes systems on their own. The Stokes-Darcy model has attracted significant attention in the past ten years since it also arises in many other applications such as surface water flows, petroleum extraction and industrial filtration. However, coupling the two constituent models leads to a very complex system. This presentation discusses numerical methods for solving several types of Stokes-Darcy system. Computational results are presented to illustrate their features and some convergence analysis is demonstrated.

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

Max Gunzburger
Florida State University
School for Computational Sciences
gunzburg@csit.fsu.edu

Xiaoming He
Department of Mathematics and Statistics
Missouri University of Science and Technology
hex@mst.edu

Xiaoming Wang
FSU
wxm@math.fsu.edu

MS79

Numerical Method for Tracking Interfaces in Fluids

We describe the sliding interface method (SIM) for tracking a sharp interface separating two materials as it moves through a structured logically rectangular grid in two spatial dimensions. The interfaces between the different solution regions are represented by a piecewise linear curve along grid cell edges or diagonals. Accurate numerical approximations to partial differential equations (PDEs) with moving discontinuous interfaces require methods that explicitly account for the movement of the interface and the interaction between the interface and the solution of the underlying PDE. The best existing numerical algorithm to account for these discontinuities is problem dependent. We will review the strengths and weakness of other interface tracking methods to identify which problems where our new sliding interface methods is most appropriate. The SIM provides a simple approach to incorporate a sharp approximation of moving interfaces that separate regions with vastly different properties within existing methods based

on finite volume discretizations of the underlying system of PDEs. We describe the advantages of the SIM on a series of problems where a moving interface immersed in a fluid can be accurately approximated.

Michael Nichola
Mathematics Department
Tulane University
mnichol@tulane.edu

MS79

Ciliary Dynamics and Chamlydomonas

Chlamydomonas is a tiny unicellular green alga with two hair-like flagella for propelling the cell body through its fluid environment. The flagella with both sensory and motility functions are very similar to animal cilia, which makes Chlamydomonas an outstanding model organism to study. In recent years, Chlamydomonas has been intensively used for the study of numerous fundamental biological processes in cell and molecular biology, including cilia-pathology in the human health. However the mechanism governing the bi-flagellar motilities is still poorly understood. We present a fluid-dynamical model, using computational approaches, that examine swimming and a variety of the bi-flagellar motilities of Chlamydomonas. This model couples the time-dependent fluid dynamics and the internal force generation mechanism by ATP-induced molecular motor proteins.

Xingzhou Yang
Department of Mathematics and Statistics
Mississippi State University
xyang@math.msstate.edu

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

MS79

Matched Interface and Boundary (MIB) Method for Solving Multi-Flow Navier-Stokes Equations with Applications to Geophysics

Based on Matched-Interface and Boundary (MIB) method, we propose a high-order numerical approach for solving incompressible Navier-Stokes equations with discontinuous viscosity and singular forces at internal interfaces. The continuity of the viscous stress across the interface is enforced by coupling all velocity components in modifying the local finite difference scheme of computing spatial derivatives of any velocity component. This allows the application of projection method to ensure a velocity field free of divergence, and hence is suitable for applications to real geophysical problems with strong non-linearity.

Yongcheng Zhou, James Liu
Department of Mathematics
Colorado State University
yzhou@math.colostate.edu, liu@math.colostate.edu

Dennis Harry
Department of Geosciences
Colorado State University
dharry@cnr.colostate.edu

MS80

Cython: Compiled Code meets Dynamic Python

Cython is an extension to the Python language that allows explicit type declarations and is compiled directly to C. This addresses Pythons large overhead for numerical loops and makes it very easy to make effiicient use of existing C, C++ and Fortran code, which Cython code can interact with natively. Cython combines the speed of C with the power and simplicity of Python, and is one of the core tools for Python-based scientific computation.

Lisandro Dalcin
Centro Int. de Métodos Computacionales en Ingeniera
dalcinl@gmail.com

Robert Bradshaw
Google, Seattle
robertwb@math.washington.edu

MS80

Matplotlib - from Interactive Exploration to Publication Graphics

matplotlib is a python library for generating scientific graphs and visualizations. While often used for ease of use in interactive and scripted environments, the library has a full featured API for publication quality graphics and production applications. In this talk we explore some of the new features of matplotlib such as the HTML5 canvas, with an emphasis on publication quality features such as sophisticated axes grid layouts, dropped spines, fonts and mathematical text.

John Hunter
Tradelink, Inc.
jdh2358@gmail.com

MS80

Capabilities and Recent Developments of NumPy for Scientific Computing

NumPy is a library that provides a flexible and powerful N-dimensional array system, along with core numerical functionality, and lays the foundation for scientific computing in Python. The array system allows creation of an array of arbitrary but strongly-typed data along with basic math functions that can operate at compiled speeds on the entire array in an element-by-element fashion. This talk will describe NumPy as well as the ports to the Python3 and IronPython platforms.

Travis E. Oliphant
Enthought, Inc.
oliphant@enthought.com

MS80

Why Modern, High-performance Networking Matters for Interactive Computing

IPython is widely used as a terminal-based environment for interactive scientific computing. Recently, we have redesigned its architecture around the ZeroMQ library for high-performance networking, to provide vastly enhanced capabilities. Based on a protocol for frontends communicating with kernels that execute user code, we can build multiple user interfaces, provide facilities for interactive remote collaboration and much more. We will present this

design and showcase some of the possibilities for novel interactive interfaces it enables.

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

Brian Granger
Cal Poly, San Luis Obispo
bgranger@calpoly.edu

Evan Patterson
Department of Physics
California Institute of Technology
epatters@caltech.edu

MS81

Implicit Newton-Krylov Drift-Diffusion Semiconductor Simulations on Multicore Architectures

This talk will discuss the performance of a massively parallel simulation code for semiconductor devices on a few current multicore architectures. The drift-diffusion equations are discretized by a finite element method, and solved with a fully-implicit Newton-Krylov approach. While this approach is robust, it is also heavily dependent on sparse matrix algorithms. We discuss the impact of preconditioners on the scaling and performance, and the efficiency of the MPI-only code on multicore architectures.

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

MS81

Multithreaded Hybrid Solver for Sparse Linear Systems

Abstract not available at time of publication.

Sivasankaran Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

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

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS81

TraceMin: A Scalable Parallel Symmetric Eigensolver

The Trace Minimization algorithm (TraceMIN) obtains a few of the smallest eigenpairs of a symmetric generalized eigenvalue problem. It is ideally suited for implementation on parallel architectures. We illustrate its scalability on a variety of problems including the simulation of car body dynamics at high frequencies. In addition, we show the effectiveness of TraceMIN in computing the Fiedler vector for weighted spectral reordering of large sparse matrices.

Ahmed Sameh
Purdue University
sameh@cs.purdue.ed

MS81

Parallel Scalability Analysis of Sparse Hybrid Linear Solvers on the Cray XE6

Achieving high parallel scalability of sparse linear system solvers implemented on large-scale computing platforms comprised of tens of thousands of multicore processors is a task that offers many challenges. Towards achieving such a solver, we build on the success of the PARDISO-SPIKE family of parallel solvers sparse linear systems. In this talk, I present a generalization of a hybrid family of schemes for handling general sparse linear systems. Application results will be presented from seismic inversion and 3D oil reservoir modeling.

Olaf Schenk
Department of Mathematics and Computer Science
University of Basel, Switzerland
olaf.schenk@unibas.ch

MS82

Parallel Stochastic Newton MCMC for Large-scale Statistical Inverse Problems

We present a new Parallel MCMC method for the solution of the Bayesian statistical inverse problem. Local Hessian and gradient information is used to adaptively construct a radial basis function approximation (RBF) of the posterior, which is used as proposal distribution for the Metropolis-Hastings algorithm across several parallel chains. Parallelism allows for rapid convergence of this approximation, and thus minimal sample correlation in the resulting MCMC chains.

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

Carsten Burstedde
The University of Texas at Austin
carsten@ices.utexas.edu

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

Lucas Wilcox
University of Texas at Austin
lucasw@ices.utexas.edu

MS82

Benchmarking of Discrete Stochastic Galerkin Solver for Computational Fluid Dynamics

In this paper we describe an unintrusive approach to

Galerkin-based generalized polynomial chaos (gPC) using templating and operator overloading. We have developed a CFD solver using this approach which is easy to read and can address new classes of physics automatically, without extensive re-implementation of the gPC equations. We benchmark the performance of the code for canonical CFD problems, and present results for accuracy, timing and scaling on large-scale parallel computing platforms.

Sanjay Mathur
Purdue University
smathur@purdue.edu

Jayathi Murthy
Purdue University, Dept. of Mechanical Engineering
West Lafayette, IN 47907-1288
jmurthy@ecn.purdue.edu

MS82

Uncertainty Quantification and Robust Design of Cardiovascular Bypass Graft Surgeries

We present a framework for performing uncertainty quantification and stochastic optimization of systems governed by partial differential equations. A novel and non-intrusive adaptive stochastic collocation technique is used to account for uncertainties. Optimization is performed using a derivative-free technique called the Surrogate Management Framework (SMF), in which Kriging interpolation functions are used to approximate the objective function. Applications to the method on computationally intensive simulation and design of cardiovascular bypass graft surgeries are presented.

Sethuraman Sankaran
University of California at San Diego
sesankar@ucsd.edu

Alison Marsden
Department of Mechanical and Aerospace Engineering
University of California, San Diego
amarsden@ucsd.edu

MS82

High-Performance Preconditioning for Instrusive Stochastic Projections

Several solution methods for stochastic Galerkin discretization of partial differential equations with random input data are compared. Less intrusive approaches based on Jacobi and Gauss-Seidel iterations are compared with more intrusive Krylov-based approaches. A set of preconditioners for Krylov-based methods are examined, including mean-based, Gauss-Seidel, approximate Gauss-Seidel and approximate Jacobi preconditioners. Krylov-based approach using approximate Gauss-Seidel and Jacobi preconditioners is found to be most effective. Sandia's Trilinos software is used to implement above algorithms.

Ramakrishna Tipireddy
University of Southern California
Aerospace and Mechanical Engineering and Civil Engineering
tipiredd@usc.edu

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

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

MS83

Reduced Models of Multiscale Kinetic Systems under Uncertainty

Chemical kinetic systems contain both uncertain rate parameters *and* dynamics at multiple time scales. The latter feature aids model reduction in the deterministic case, but model reduction under uncertainty raises new challenges. We use computational singular perturbation to calculate 'importance indices' for species-reaction pairs. Distributions of these indices are used to form reduced models that (1) yield predictions *within* probabilistic bounds determined by the full model, or (2) preserve entire output distributions of the full model.

Thomas Coles
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
tcoles@mit.edu

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

MS83

Coupling Algorithms for Stochastic Multiphysics

We present recent developments for the numerical resolution of coupled partial differential equations with stochastic coefficients. These equations arise naturally in the context of UQ for multiphysics problems. The proposed algorithms explore different probabilistic representations of information exchanged between the coupled equations. Issues of accuracy and efficiency, including model and dimension reduction, are stressed. An application to a problem arising in nuclear safety is used to demonstrate the mathematical constructions and associated algorithms.

Maarten Arnst
University of Southern California
arnst@usc.edu

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

John Red-Horse
Validation and Uncertainty Quantification Processes
Sandia National Laboratories
jrredho@sandia.gov

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

MS83

Stochastic Atomistic-to-Continuum Coupling using

Bayesian Inference

We perform a multiscale simulation in a model operating under uncertainty. This latter has the form of both parametric uncertainty and sampling noise intrinsic in atomistic simulations. We present a mathematical formulation that enables the exchange of information and propagation of uncertainty between the discrete and continuum components. We implement a Bayesian inference machinery to build the polynomial chaos expansion (PCE) of the exchanged variables. We consider a simple Couette flow model where the variable of interest is the wall velocity. Results show convergence toward the analytical solution at a reasonable computational cost.

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

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

Reese Jones, Bert Debusschere,
Co-author
rjones@sandia.gov, bjdebus@sandia.gov

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

Helgi Adalsteinsson
Sandia National Laboratories, Livermore CA
hadalst@sandia.gov

MS83

Kernel Principal Component Analysis for Stochastic Input Generation of Multiscale Systems

We apply kernel principal component analysis (KPCA) to construct a reduced-order stochastic input model for the material property variation in heterogeneous media. KPCA can be considered as a nonlinear version of PCA. Through use of kernel functions, KPCA enables the preservation of high-order statistics of the random field. Thus, this method can model non-Gaussian, non-stationary random fields. In addition, polynomial chaos (PC) expansion is used to represent the random coefficients in KPCA which provides a parametric stochastic input model. Thus, realizations, which are consistent statistically with the experimental data, can be generated in an efficient way. We showcase the methodology by constructing a low-dimensional stochastic input model to represent channelized permeability in porous media.

Nicholas Zabaras, Xiang Ma
Cornell University
zabaras@cornell.edu, xm25@cornell.edu

MS84

ToppGene Computational Analysis: Identifying Potential Genes Necessary for Acute Myeloid Leukemia Viability

Acute myeloid Leukemia (AML) is a disease characterized by various cytogenetic abnormalities and poor clinical outcome. The focus of this study was to synthesize the diverse gene expression data stemming from these abnormalities and their subsequent AML subtypes. To accomplish this, we utilized the ToppGene Computational Analysis Software Suite, which provided statistically significant identification of potentially therapeutic targets. As a result, we identified MCL1, KLF44, KLF6, and the miR-29 family as likely targets for therapy. Advisors: Dr. Bruce Aronow, Cincinnati Children's Hospital

Zach Beaver
Wofford College
beaverzw@email.wofford.edu

MS84

Understanding How Errors in Model Parameters Impact the Simulation of Subsurface Temperature Profiles

We consider simulation of heat flow in the shallow subsurface. This work is motivated by trying to match temperature profiles taken at meteorology stations using a simulation-based approach and analytic approaches. Specifically we determine soil parameters using derivative-free optimization to solve the nonlinear-least squares problems. We also study how errors in the initial and boundary conditions propagate over time. Advisor: K.R.Fowler

Xiaojing Fu
Clarkson University
fux@clarkson.edu

MS84

Optimization for Modeling Brain Activity During Threatening Scenarios

The focus of this work is determining the architecture representing how the brain detects threats. A model was developed containing seventeen connections between sensory, response attention, and threat detection nodes for somatosensory and visual tasks. Optimization was used to fit these parameters but results indicated that the model contained undesirable local minima. To account for this occurrence physical constraints were adapted to the model and an analysis of variance determined the sensitivities of the parameters. Advisors: Dr Katie Fowler, Dr Robert Dowman

Brian Leventhal
Clarkson University
leventbc@clarkson.edu

MS84

Development and Analysis of Computer Simulations of Neuron-Astrocyte Networks to Better Understand the Role of Astrocytes

In the brain, astrocytes had been sidelined to a supporting role in neural networks. The past decade revealed that they influence and propagate neural signals, suggesting an important role for astrocytes in neural networks. By running multiple simulations of neural networks, with and without astrocytes, we observed that they have the ability to synchronize neural activity. Such findings resemble brain recordings of epileptic patients whose neuronal activity is synchronized during seizures correlated with hypertrophied astrocytes. Advisors: Dr. Anne J. Catlla

Tahirali H. Motiwala
Wofford College

motiwalath@email.wofford.edu

MS84

Multimaterial Multiphase Deflagration-to-Detonation Model in PetaScale Computational Framework

High-performance science-based computer simulation tools can predict explosion violence in accidents involving high explosives and propellants. Our model combines steady-state reaction models for thermally-activated combustion and pressure-activated detonation in the Uintah computational framework, which computes coupled fluid/structure interactions in multiphase multimaterial domains. The unified reaction model was validated against aluminum-flyer impact, Steven and exploding cylinder tests for PBX9501. Preliminary results show that leveraging the massively parallel capability of Uintah allows prediction in large domains with high fidelity. Advisors: Chuck A. Wight and Martin Berzins

Joseph Peterson
University of Utah
tba@siam.org

MS85

Title Not Available at Time of Publication

Abstract not available at time of publication.

Marino Arroyo
Universitat Politecnica de Catalunya
marino.arroyo@upc.edu

MS85

An Afem for Fluid-Structure Interaction

A parametric FEM for free boundary problems is discussed. Some applications include geometric and fluid-membrane interaction in bio-membranes. With slight modification the method can be successfully applied in shape optimization problems such the design of an obstacle with minimal drag or a by-pass design. The last in the context of inexact sequential quadratic programming. Inexactness is a consequence of using adaptive finite element methods (AFEM) to approximate the state equation, update the boundary, and compute the geometric functional.

Miguel S. Pauletti
Texas A&M University
pauletti@math.tamu.edu

MS85

Electric-field-induced Ordering and Pattern Formation in Colloidal Suspensions

The long-time dynamics and pattern formation in semi-dilute suspensions of colloidal spheres in a viscous electrolyte under a uniform electric field are investigated using numerical simulations. The rapid chain formation that occurs in the field direction as a result of dielectrophoretic interactions is found to be followed by a slow coarsening process by which chains coalesce into hexagonal sheets and eventually rearrange to form mesoscale cellular structures, in agreement with recent experiments. The morphology and characteristic wavelength of the equilibrium phases that emerge are shown to depend on suspension volume fraction, electrode spacing and field strength, suggesting novel ways of controlling effective suspension properties in practical applications.

Jae Sung Park, David Saintillan
Department of Mechanical Science and Engineering
University of Illinois at Urbana-Champaign
park80@illinois.edu, dstn@illinois.edu

MS85

The Microfluidics of Particle/ Vesicle Mixtures in the Microvasculature

Many dispersions of colloidal and noncolloidal particles with application in medicine contain elongated particles. For example, such particles are useful at the nanoscale for the delivery of anti-cancer agents to tumor endothelial cells. In a different context, hemostasis in the small vessels relies on naturally occurring blood particles such as platelets being concentrated near the vessel walls. We will review our computer simulations of these processes with a view toward understanding and engineering these therapies.

Eric S. Shaqfeh
Professor, Chemical Engineering & Mechanical Engineering
Stanford University, Stanford, CA
esgs@stanford.edu

Hong Zhao
Department of Mechanical Engineering
Stanford University, Stanford, CA 94305
hongzhao@stanford.edu

Andrew Spann
Institute of Computational and Mathematical Engineering
Stanford University, Stanford CA 94305
spann@stanford.edu

MS86

On Multilevel Diffusion-based Load Balancing for Parallel Adaptive Numerical Simulations

Load balancing is an important requirement for the efficient execution of parallel numerical simulations. In particular when the simulation domain changes over time, the mapping of computational tasks to processors needs to be modified accordingly. Here we further explore the very promising diffusion-based graph partitioning multilevel algorithm DibaP, which uses two different strategies (algebraic multigrid and matchings) for constructing a multilevel hierarchy. The presented experiments with graph sequences that imitate adaptive numerical simulations demonstrate the applicability and quality of DibaP for load balancing by repartitioning.

Henning Meyerhenke
Georgia Institute of Technology
College of Computing
henningm@gatech.de

Burkhard Monien
University of Paderborn
Department of Computer Science
bm@upb.de

MS86

Multilevel and Graph-based Methods in Data-mining

A number of nonlinear dimensionality reduction methods which exploit *affinity graphs* have been developed in the past for solving various problems in data-mining. This talk will explores a multilevel framework based on graph coarsening whose goal is to reduce the cost of these techniques. Among the applications, we will consider the problem of unsupervised manifold learning and spectral clustering. An application to information retrieval will also be discussed.

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

MS86

Relaxation-based Coarsening and Multiscale Graph Organization with Applications to Linear and Compression-friendly Orderings

In this talk we generalize and improve the multiscale organization of graphs by introducing a measure that quantifies the "closeness' between two nodes. The calculation of the measure is linear in the number of edges in the graph and involves just a small number of relaxation sweeps. A similar notion of distance is then calculated and used at each coarser level. We demonstrate the use of this measure in multiscale methods for several important graph/matrix linear and compression-friendly ordering problems and discuss the multiscale graph organization.

Achi E. Brandt
Weizmann Institute of Science
Dept of Applied Mathematics
achi.brandtweizmann.ac.il

Jie Chen
Department of Computer Science and Engineering
University of Minnesota
jchen@cs.umn.edu

Dorit Ron
The Weizmann Institute of Science
dorit.ron@weizmann.ac.il

Ilya Safro
Argonne National Laboratory
safro@mcs.anl.gov

MS86

Engineering Multilevel Graph Partitioning Algorithms

We describe two different approaches to multi-level graph partitioning (MGP). The first algorithm is based on the extreme idea to contract only a single edge on each level of the hierarchy. The second algorithm is an approach to parallel graph partitioning that scales to hundreds of processors. Quality improvements compared to previous systems are due to better priorization of edges to be contracted and FM local search algorithms that work more locally than previous approaches.

Peter Sanders
Karlsruhe Institue of Technology
Faculty of Informatics
sanders@kit.edu

Christian Schulz
Karlsruhe Institut of Technology
Faculty of Informatics
christian.schulz@kit.edu

MS87

Title Not Available at Time of Publication

Abstract not available at time of publication.

M Adams
Texas A&M University
mladams@tamu.edu

MS87

A Posteriori Error Estimation via Error Transport

Error estimation for time dependent hyperbolic problems is challenging for theoretical and practical reasons. In these systems, error propagates long distances and produces effects far from the point of generation. In addition, nonlinear interaction of error plays an important role. Time marching and non-linear discretizations must also be addressed. In this talk we investigate the use of error equations for a-posteriori error estimation. These auxiliary PDEs are treated numerically to yield field estimates of error.

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banks20@llnl.gov

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

MS87

Quantification of Numerical Uncertainty for Entire Curves or Fields

A new capability is proposed to assess the asymptotic convergence of entire solution fields without having to reduce them to scalars. The technique estimates lower and upper bounds of solution error when the exact solution is unknown. The technique also generalizes the Grid Convergence Index to estimate bounds of numerical uncertainty for entire fields. Applications are presented using test problems analyzed with finite element or hydro-dynamics codes.

Francois M. Hemez
Los Alamos National Laboratory
hemez@lanl.gov

MS87

The Role of Theory in Calculation Verification

Error estimation is a recommended practice in computa-

tional science and engineering. The approach employs a sequence of discretizations, a measure from the solution, and extrapolation to estimate errors. The standard application of this methodology usually does not involve practices beyond the rote application of the methodology, but often produces results that defy explanation. The use of appropriate theories of numerical approximation can be used to clarify results, and suggest improvements to the simulation methodology.

William J. Rider
Sandia National Laboratory
wjrider@sandia.gov

MS88

A Two-dimensional hp-adaptive Discontinuous Galerkin Model for the Shallow Water Equations

Unstructured meshes are becoming more and more popular in geophysical flow models. We present a two-dimensional model solving the shallow water equations on unstructured meshes. The latter is dynamically adapted using the AMR technique to minimize the discretization error. The interpolation order is also adapted during the solution process. Classical test cases on the sphere are used to validate the model, as well as the global simulation of the 2010 tsunami in Chile.

Sébastien Blaise
National Center for Atmospheric Research
sblaise@ucar.edu

Amik St-Cyr
National Center for Atmospheric Research
Institute for Mathematics Applied to the Geosciences
amik@ucar.edu

David Hall
National Center for Atmospheric Research
1850 Table Mesa Drive, Boulder CO 80305
halld@ucar.edu

MS88

A Comparison of Gauss-Legendre and Gauss-Legendre-Lobatto Discontinuous Galerkin Spectral Element Methods

In this talk, different implementations of the collocation discontinuous Galerkin spectral element method, namely the Gauss Legendre and the Gauss-Lobatto Legendre variants, are discussed and compared. Analysis of the different wave propagation properties as well as computational aspects are presented. An application to the simulation of a simple turbulent flow is used to illustrate the differences of both methods.

Gregor Gassner
Institute for Aerodynamics and Gasdynamics
Universitaet Stuttgart
gassner@iag.uni-stuttgart.de

David Kopriva
Department of Mathematics
The Florida State University
kopriva@math.fsu.edu

MS88

Two-level Optimized Schwarz Preconditioning for SEM-based MHD

We present a new two-level preconditioner for a "pseudo-Laplacian' operator that emerges from an explicit spectral-element discretization of the incompressible magnetohydrodynamics equations. Extending work in JCP,133:84(1997), we use new overlap stenciling and corner communication for the fine grid whose preconditioner utilizes a theoretical result, SISC,29(6):2402, enabling trivial conversion to an optimized overlapping Schwarz method. Results are presented for the (pseudo-)Poisson equation, and for magnetohydrodynamics simulations of the Orszag-Tang vortex.

Amik St-Cyr
National Center for Atmospheric Research
Institute for Mathematics Applied to the Geosciences
amik@ucar.edu

Duane Rosenberg
NCAR
Institute for Mathematics Applied to Geosciences
duaner@ucar.edu

James Lottes
Oxford University
lottes@maths.ox.ac.uk

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

Annick Pouquet
NCAR
pouquet@ucar.edu

MS88

GPU Accelerated Discontinuous Galerkin Methods

We will discuss the trend towards many-core architectures in modern computers and in particular how current non-uniform memory hierarchy should be factored in to comparisons of competing formulations for numerically solving partial differential equations. We will focus primarily on the discontinuous Galerkin methods, in particular a customized version that is designed to deliver an accurate treatment of curvilinear domains with relatively low storage overhead. Examples simulations from electromagnetics and gas dynamics will be shown with benchmarks indicating the performance obtained on current graphics processing units.

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

MS89

Real-time Classification of Astronomical Events with Python

Over the next several years, massive sky surveys will record more astronomical transient events than all previously recorded by mankind. *Discovering* those events (and maximizing the scientific value of a subset of them) is rapidly becoming a needle-in-the-haystack problem that cannot be tackled with humans in the loop. Here I describe a Python-based framework for ingesting massive astronomical datastreams, applying machine-learning techniques on

time-series-derived data, and producing probabilistic classifications of transients.

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

Dan Starr, Joseph Richards, Nathaniel Butler
UC Berkeley
dstarr@astro.berkeley.edu, joeyrichar@gmail.com, nat@astro.berkeley.edu

Dovi Poznanski
UC Berkeley
Lawrence Berkeley National Laboratory
dovi@astro.berkeley.edu

MS89

High-Order, Adaptive Finite Element Methods for Atomic Structure Calculations

We present high-order, adaptive finite element solvers for radial Schroedinger and Dirac equations. These solvers are written in C++ and Fortran, and exposed to Python using Cython and Fwrap. We explain the importance and the technical details of this exposition. Moreover, we highlight how the resulting solvers may advance quantum chemistry by providing a robust, variational alternative to conventional shooting (or Gaussian-based) methods, finding all desired states simultaneously, with accuracy and orthogonality approaching machine precision.

Ondrej Certik, Pavel Solin
Department of Mathematics and Statistics
University of Nevada, Reno
ondrej@certik.cz, solin@unr.edu

MS89

FEMhub Online Numerical Methods Laboratory

The FEMhub Online Numerical Methods Laboratory (http://lab.femhub.org) is an Ext-JS application based on Codenode, whose objective is to facilitate remote scientific computing with open source packages included in FEMhub. Currently, these packages include FiPy, Hermes, Phaml and SfePy, and their number will grow in the future. We will mention how the Online Lab can be used to enhance teaching of numerical methods courses and perform sample finite element computations with codes included in FEMhub. We will also present a new FEMhub Mesh Editor that makes it possible to generate 2D finite element meshes inside a web browser window.

Pavel Solin, Ondrej Certik
Department of Mathematics and Statistics
University of Nevada, Reno
solin@unr.edu, ondrej@certik.cz

Mateusz Paprocki
Technical University of Wroclaw
Wroclaw, Poland
mattpap@gmail.com

Aayush Poudel
University of Nevada, Reno
aayushpoudel@gmail.com

MS89

Interactive Parallel Python with ZeroMQ

ZeroMQ is an advanced, lightweight message passing library, written in C/C++. It is based around rich Sockets that provide simple, but powerful, messaging primitives including publish/subscribe, request/reply, point-to-point and broadcast. These building blocks can be used to construct sophisticated and high-performance messaging architectures. IPython is an open-source package that provides high-level parallel computing capabilities in the Python programming language. By moving both the serial and parallel IPython computing models to ZeroMQ, we gain access to higher performance, and move to a fundamentally different programming model with many advantages. We will present some details of ZeroMQ, how it is used in IPython, and the new model for interactive parallel computing.

Min Ragan-Kelley
UC Berkeley AS&T
minrk@berkeley.edu

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

Brian Granger
Cal Poly, San Luis Obispo
bgranger@calpoly.edu

MS90

Stochastic Galerkin and Collocation Methods for PDEs with Random Coefficients

We focus on PDE-based models with random input parameters and consider multivariate polynomial approximations of the solution as a function of the random parameters. We review and compare Galerkin and Collocation-type approximations with particular attention to the choice of the polynomial space depending on the structure of the PDE and the input probability measure. These methods typically require lots of "deterministic" solves and demand for HPC. We also comment on implementation issues and parallelization.

Fabio Nobile
MOX, Dip. di Matematica
Politecnico di Milano
fabio.nobile@polimi.it

Joakim Back
Applied Mathematics and Computational Science, KAUST
joakim.back.09@ucl.ac.uk

Lorenzo Tamellini
MOX, Department of Mathematics
Politecnico di Milano
lorenzo.tamellini@mail.polimi.it

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

MS90

Practical Considerations in Modeling Random Variables, Vectors and Fields Using Expansion Methods

Expansion-based probability methods are at the core of Sandia's HPC random field simulation capabilities development, and are key elements in this regard to both the SIERRA and Trilinos software suites. We address many of the relevant probability-related aspects that are implicit in these approaches, including the existence of certain transformations of random variables and the sense of equality under which these transformations are understood. We also discuss various aspects of equality and conditions on existence for, and the convergence of, approximations for these random variables.

John Red-Horse
Validation and Uncertainty Quantification Processes
Sandia National Laboratories
jrredho@sandia.gov

Philippe Pebay
Sandia National Laboratories
pppebay@sandia.gov

Paul Boggs
Sandia National Lab
ptboggs@sandia.gov

MS91

Application of Control Theory to Spacecraft Attitude Stability

We describe results from a year-long capstone research problem in satellite control posed by Space Systems / Loral of Palo Alto, CA. Our goal is to find criteria that guarantee stability for a given feedback control system, even in the presence of parameter uncertainties, delays in the signals, and disturbances to the system. We have constructed a computer simulation and derived an analytical transfer function for the spacecraft control system. We present a comparison of numerical and theoretical results.

Jacob Bouricius, Max Lee, Andrea Levy, Maggie Rogers
Harvey Mudd College
jbouricius@hmc.edu, mjlee@hmc.edu, andrea_b_levy@hmc.edu, mrogers@hmc.edu

MS91

Spectral Counting Functions of Atomic Measures

For this talk, we will be working with a family of atomic measures supported on the boundary of a generalized Cantor-like string. We will develop and analyze geometric counting functions and spectral zeta functions for a family of fractal strings. We will also discuss regularity, partition zeta functions, and complex dimensions as they pertain to our example.

Kate Ellis
California State University, Stanislaus
kellis1@csustan.edu

MS91

Simulating Stochastic Inertial Manifolds by a Backward-Forward Approach

We construct stochastic inertial manifolds numerically for stochastic differential equations with multiplicative noises. After splitting the stochastic differential equations into a backward part and a forward part, we use the theory of solving backward stochastic differential equations to achieve our goal.

Xingye Kan, Jinqiao Duan
Illinois Institute of Technology
xkan@iit.edu, duan@iit.edu

Yannis Kevrekidis
Dept. of Chemical Engineering
Princeton University
yannis@princeton.edu

Anthony J. Roberts
University of Adelaide
anthony.roberts@adelaide.edu.au

MS91

Aspects on Robust Shape Optimization in CFD

The proper treatment of uncertainties in the context of aerodynamic shape optimization is a very important challenge to ensure a robust performance of the optimized airfoil under real life conditions. This talk will propose a general framework to identify, quantize and include the uncertainties in the overall optimization procedure. Efficient discretization techniques of the probability space, algorithmic approaches based on multiple-setpoint ideas in combination with one-shot methods as well as numerical results will be presented.

Claudia Schillings
University of Trier, Germany
Claudia.Schillings@uni-trier.de

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

MS91

Integral Tau Leap Method for Stochastic Chemical Systems

Tau leaping methods can efficiently simulate models of stiff stochastic chemical systems. However, most existing methods do not naturally preserve some chemical structures, such as integer-valued and nonnegative molecular population states. In this talk, I will present structure preserving tau methods for simulating stochastic chemical systems. We illustrate the new methods through a number of biochemically motivated examples, and provide comparisons with existing implicit tau methods.

Yushu Yang
University of Maryland Baltimore County
yushu1@umbc.edu

Muruhan Rathinam
University of Maryland, Baltimore County
muruhan@umbc.edu

Jinglai Shen
University of Maryland Baltimore County

shenj@umbc.edu

MS91

A Diffuse Interface Model of Multicomponent Vesicle Adhesion and Fusion

Adhesion and fusion of lipid bilayer vesicle membranes are important biological process, which play key roles in exocytosis, endocytosis etc. In this talk, a diffuse interface approach is introduced to model the multicomponent vesicle adhesion and fusion. Various equilibrium configurations are examined for the adhered vesicles, pre-fusion and post-fusion between vesicles. Our model predicts the fact that adhesion can promote phase separation of multicomponent vesicle membrane. The effect of spontaneous curvatures, bending rigidities, and adhesion strength on adhered vesicles and fusion process are discussed.

Yanxiang Zhao
Penn State University
zhao@math.psu.edu

Sovan Das
Indian Institute of Technology Kanpur
sovandas@iitk.ac.in

MS92

StarPU: A Runtime System for Scheduling Tasks on Accelerator-Based Multicore Machines

We introduce StarPU, a runtime system designed to exploit accelerator-based multicore machines. StarPU provides a simple task-based API, allowing programmers to focus on algorithmic issues. The task scheduling engine automatically balances the tasks across available processing units and uses prefetching to hide data transfers latency. We evaluate our system with dense linear algebra kernels, using a scheduling policy based on auto-tuned performance models. We detail how our approach provides portability of performance.

Cedric Augonnet
University of Bordeaux - INRIA Bordeaux
cedric.augonnet@inria.fr

MS92

Scheduling Strategies in StarSs

StarSs programming model enables to write sequential applications which inherent concurrency is exploited at runtime. OpenMP-like pragmas are inserted to annotate tasks, indicating the directionality of the subroutine parameters (input, output or inout). A data-dependence graph of the application is dynamically built and scheduled in the different cores of a multicore platform. The talk will present the scheduling strategies that have been implemented in the different StarSs runtimes (i.e., for the Cell, SMP or GPU).

Rosa M. Badia
Barcelona Supercomputing Center
rosa.m.badia@bsc.es

MS92

CnC for HPC

We discuss our experiences in expressing and tuning computations from high-performance computing (HPC) in Intel Concurrent Collections (CnC), a dataflow programming model. Our case studies include asynchronous-parallel dense linear algebra algorithms, where we can match or exceed Intel MKL for representative computations, as well as more irregular computations like the kernel-independent fast multipole method.

Aparna Chandramowlishwaran
Georgia Institute of Technology
aparna@cc.gatech.edu

Kathleen Knobe
Intel
Kath.knobe@intel.com

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

MS92

Development of a Task Execution Environment Driven by Linear Algebra Applications

We describe the design and development of QUARK (QUeuing And Runtime for Kernels), which is an multicore execution environment for applications composed of tasks with data dependencies between the tasks. The development of QUARK is driven by the needs of the PLASMA linear algebra library, and some of the optimizations specific to linear algebra will be presented. QUARK is intended to be easy-to-use, efficient and scalable, and we discuss our progress on those goals. Performance results on linear algebra applications will be presented.

Asim YarKhan
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
yarkhan@eecs.utk.edu

Jakub Kurzak
University of Tennessee Knoxville
kurzak@eecs.utk.edu

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

MS93

Importance of Parallel Linear Solvers in Battlespace Environments

Characterizing the battlespace environment is crucial to US Department of Defense mission planning because it enhances safety and warfighting effectiveness. ERDC is in a unique position because of its stewardship of surface and groundwater codes solving large real-world problems. Often parallel linear solvers become performance bottleneck of these simulations using implicit time scheme. The performance of these applications may be capable of advancing toward exasacale by understanding the effect of physical attributes and linear operators.

Jing-Ru C. Cheng
Scientific Computing Research Center
U.S. Army Engineer Research and Development Center (ERDC)
Ruth.C.Cheng@usace.army.mil

MS93

Using GPU-Based Accelerators to Prepare for Exascale Computing in Defense Applications

CPU clock speeds are peaking due to thermodynamic limits and other limitations, requiring the use of parallelism to continue gaining improved performance. In order to gain exascale performance, we must take advantage of extreme parallelism requiring millions/billions of threads. One approach towards reaching exascale involves hybrid architectures using many-core CPUs with GPU-based accelerators. A 2-D finite difference time domain simulation using multiple GPUs demonstrates GPU programming techniques and performance.

Paul Eller
U.S. Army Engineer Research and Development Center
paul.r.eller@usace.army.mil

MS93

The ParalleX Execution Model for Exascale Computation

Achieving Exaflops sustained performance by the end of this decade will require dramatic improvements in efficiency, scalability, and power. ParalleX is a new model of computation being explored to expose significant parallelism and provide latency hiding. It provides a framework for co-design of all future system layers from programming models, through system software, to parallel system and core architectures. This talk will present the foundational ideas of ParalleX and early results in its use on conventional parallel systems for Adaptive Mesh Refinement as applied to numerical relativity computations.

Thomas Sterling
LSU Center For Computation and Technology (CCT)
tron@cct.lsu.edu

MS93

Opportunities in the Present and Future of Defense Applications

Abstract not available at time of publication.

John West
U.S. Army Engineer Research and Development (ERDC) Center
Information Technology Laboratory
john.e.west@usace.army.mil

MS94

An Energy Based Stochastic Mapping between Model Parameters for Complementary Scales

In multiscale modeling and multifidelity model representations, one of the recent research interests is proper characterization of fine-to-coarse scale model parameters relationship. The present work focuses on how this relationship can be characterized, particularly when a limited amount of highly accurate fine scale/high fidelity information (computational/experimental) is available. The proposed work treats this relationship as a stochastic mapping that results in *stochastic* coarse model parameters even when the fine scale information is *deterministic*. The uncertainty induced by the stochastic mapping reflects the 'effects of loss of information' in constructing the coarse scale model parameters. The proposed scheme shifts the domain of knowledge from the deterministic (highly accurate) fine scale/high fidelity regime to the stochastic coarse scale/low fidelity regime. The stochastic mapping is developed such that it is constrained by highly accurate fine scale/high fidelity energy observables in a certain sense.

Sonjoy Das, Youssef M. Marzouk
Massachusetts Institute of Technology
sonjoy@mit.edu, ymarz@mit.edu

MS94

Stochastic Multiscale Modeling from Molecular Reactions to Catalytic Reactor

Integrating first-principles kinetic Monte Carlo (KMC) simulation with a stochastic continuum model, a stochastic hybrid model was developed to study the effects of heat and mass transfer on the heterogeneous reaction kinetics. The stochastic hybrid reaction model consists of a surface phase domain where catalytic surface reactions occur and a gas-phase boundary layer domain imposed on the catalyst surface where the temperature and pressure gradients exist. The surface phase is described using the site-explicit first-principles KMC simulations. The heat and mass flux fluctuation in the gas-phase boundary layer domain, which is represented by thermal and molecular diffusion, is characterized using the stochastic grid-based continuum model. At each time step, the heat and mass exchanges between two domains are simulated simultaneously until the steady-state reaction condition is reached. At the steady-state reaction condition, the activity, the surface coverage of each reaction species, as well as the temperature and pressure gradient profiles in the gas-phase boundary domain are statistically constant with very small fluctuations. In order to illustrate that the stochastic hybrid model is more accurately elucidate the experimentally observed reaction kinetics by considering the fluctuation of heat and mass transfer, we investigated the surface kinetics of CO oxidation over the RuO2(110) catalysts with various operating reaction conditions. Simulation results indicate that the stochastic hybrid model captures the coupling effect between the heat and mass flux fluctuation in the gas-phase boundary layer and the reaction kinetics more accurately than the traditional KMC simulation results.

Guang Lin, Donghai Mei
Pacific Northwest National Laboratory
guang.lin@pnl.gov, donghai.mei@pnl.gov

MS94

Bayesian Inference of Atomic Diffusivity in a Binary Ni/Al System based on Molecular Dynamics

Atomic mixing in Ni/Al nanolaminates is characterized using MD computations. The diffusivity is extracted using a Bayesian inference framework, based on contrasting the moments of the cumulative distribution functions of the constituents with corresponding moments of a dimensionless concentration evolving according to a Fickian process. Posterior estimates of the diffusivity are compared with experimental measurements, and conclusions are drawn regarding the potential of the present framework to refine continuum modeling approaches.

Francesco Rizzi
Department of Mechanical Engineering
Johns Hopkins University
frizzi1@jhu.edu

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

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

Rong-Guang Xu
Department of Physics and Astronomy
Johns Hopkins University
xrg117@pha.jhu.edu

Michael Falk, Timothy P. Weihs, Gregory Fritz
Materials Science and Engineering
Johns Hopkins University
mfalk@jhu.edu, weihs@jhu.edu, gfritz2@jhu.edu

Omar M. Knio
Dept. of Mech. Eng.
Johns Hopkins University
knio@jhu.edu

MS94

PC-based Uncertainty Propagation in the Gulf of Mexico using HYCOM

Abstract not available at time of publication.

Carlisle Thacker
National Oceanic and Atmospheric Administration
carlisle.thacker@noaa.gov

MS95

Computation of Spherical Delaunay Triangulations and Centroidal Voronoi Tessellations in Parallel

Spherical centroidal Voronoi tessellations (SCVT) are used in many applications in a variety of fields, one being climate modeling. They are a natural choice for spatial discretizations on the Earth. New modeling techniques have recently been developed which allow the simulation of ocean and atmosphere dynamics on arbitrarily unstructured meshes, which could include SCVTs. Since these communities are beginning to focus on exa-scale computing for large scale climate simulations, a need is brought to light for fast and efficient grid generators. Current high resolution simulations on the earth call for a spatial resolution of about 0.1^o which corresponds to about 11.1km. In terms of a SCVT this corresponds to a quasi-uniform SCVT with roughly 2 million generators. Computing this grid in serial is incredibly expensive, and can take on the order of weeks to converge sufficiently for the needs of climate modelers. Utilizing conformal mapping techniques, as well as planar triangulation algorithms, and basic domain decomposition, this presentation outlines a new algorithm that can be used to compute SCVTs in parallel, thus reducing the overall time to convergence. This reduces the actual time needed to create a grid on the Earth, as well as allows for new techniques to be explored when modeling the ocean and atmosphere.

Doug Jacobson
Florida State University
dwj07@fsu.edu

MS95

Modeling Storm Surges Using the Multilayer Shallow Water Equations

Storm surge prediction is becoming more integral to the safety of the increasing populations on the coastlines. Current efforts have focused on using the single layer shallow water equations to balance the dominant physics of surges with computational efficiency. We are examining the advantages of using the multilayer shallow water equations to model storm surges. We will present idealized simulations using this new model and discuss their effectiveness in capturing additional storm surge physics.

Kyle T. Mandli
University of Washington
Dept. of Applied Mathematics
mandli@amath.washington.edu

Randall J. LeVeque
Applied Mathematics
University of Washington (Seattle)
rjl@uw.edu

MS95

Title Not Available at Time of Publication

Abstract not available at time of publication.

Chris Newman
Los Alamos National Lab
newman@lanl.gov

MS95

A Numerical Study of Boundary Value Problems for the Shallow Water Equations with Topography

Limited Area Models (LAMs) have been used to achieve high resolution over a region of interest in geophysical fluid dynamics. The model equations under consideration are the nonviscous Shallow Water equations with topography in space dimension one. In this work, our goals are twofold: first to find boundary conditions which are physically suitable, namely they let the waves move freely of the domain without spurious reflecting waves at the boundary in a nonphysical way; second to to numerically implement these boundary conditions in a numerically effective way. This is achieved by applying a suitable extension of the central-upwind method for the spatial discretization and the Rung-Kutta method of second order for the time discretization. Several numerical simulations for which we tested the proposed boundary conditions and the numerical schemes are presented.

Ming-Cheng Shiue
Indiana University
mshiue@indiana.edu

Jacques Laminie
Université des Antilles et de la Guyane
jacques.laminie@univ-ag.fr

Roger M. Temam
Inst. f. Scientific Comput. and Appl. Math.
Indiana University
temam@indiana.edu

Joseph J. Tribbia
Nat'l Center for Atmospheric

Research
tribbia@ncar.ucar.edu

MS95

Some Issues Related to the Boundary Conditions for the Inviscid Equations of the Atmosphere and the Oceans

In this lecture we will present some issues related to the boundary conditions for the primitive equations of the atmosphere and the oceans. We will present some nonlocal boundary conditions for these equations which appear to be well suitable, and show the results of three dimensional numerical simulations performed with these boundary conditions based on joint works with Q. Chen, M.-C. Shiue and J. Tribbia).

Roger M. Temam
Inst. f. Scientific Comput. and Appl. Math.
Indiana University
temam@indiana.edu

MS96

A High-Resolution Fast, Spectral Boundary-Integral Method for Multiple Interacting Blood Cells

We discuss a boundary integral method for simulation of flowing red blood cells. A particle-mesh-Ewald approach (PME) achieves an overall $O(N \log N)$ scaling. The cell shapes are represented by spherical harmonics for their perfect resolution and because such global basis functions facilitate control of aliasing errors without explicit filtering or implicit numerical dissipation. Example simulations include a cell flowing through a constriction, blood flow in round tubes which reproduce measurements, white cell transport by many red cells, and flow in a model vessel network.

Jonathan B. Freund
UIUC
jbfreund@uiuc.edu

MS96

Front Tracking Method on Precipitation with Subsurface Flow

We present the simulation results of solute precipitation with subsurface flow using the FronTier software library. The FronTier interface propagation tool is used to model the growing crystal surface in pore scale simulation. The phase transition code is coupled with the incompressible Navier-Stokes solver for the convection due to subsurface flow. We measure the upstream and downstream growth rates on different flow velocities and the effect of precipitation on Darcy's law.

Xiaolin Li, Yijing Hu, Saurabh Joglekar
Department of Applied Math and Stat
SUNY at Stony Brook
linli@ams.sunysb.edu, yhu@ams.sunysb.edu, joglekar@ams.sunysb.edu

MS96

Stabilizing fluid-fluid Interfaces using Colloidal Particles

Bicontinuous interfacially jammed emulsion gels ('bijels') were proposed in 2005 as a hypothetical new class of soft materials in which interpenetrating, continuous domains of two immiscible fluids are maintained in a rigid state, by a jammed layer of colloidal particles at their interface. Such gels offer an important route to materials with unique combinations of properties not available in a single phase material. In 2007, the first bijels were created experimentally. In joint work with Sebastian Aland and Axel Voigt, we develop a continuum model for such systems which combines a Cahn-Hilliard-Navier-Stokes model for the macroscopic two-phase fluid system with a surface Phase-Field-Crystal model for the microscopic colloidal particles along the interface. We demonstrate the feasibility of this approach and present numerical simulations that confirm the ability of the colloids to stabilize interfaces for long times.

John Lowengrub
Department of Mathematics
University of California at Irvine
lowengrb@math.uci.edu

MS96

Large Scale Simulations of Vesicles Suspended in 3D Viscous Flows

Vesicles are locally-inextensible closed membranes that possess tension and bending energies. Vesicle flows model numerous biophysical phenomena that involve deforming particles interacting with a Stokesian fluid. We will present new schemes for simulating the three-dimensional hydrodynamic interactions of large number of vesicles. They incorporate a stable time-stepping scheme, high-order spatio-temporal discretizations, spectral preconditioners, and a new reparameterization scheme capable of resolving extreme mesh distortions in dynamic simulations.

Shravan Veerapaneni
Courant Institute
New York University
shravan@cims.nyu.edu

Abtin Rahimian, George Biros
Georgia Institute of Technology
rahimian@gatech.edu, biros@gatech.edu

Denis Zorin
Computer Science Department
Courant Institute, New York University
dzorin@cs.nyu.edu

MS97

Video Background Subtraction Using Communication-Avoiding QR on a GPU

Reducing communication between the GPU cores and memory can give us substantial speedups, turning a bandwidth-bound problem into a compute-bound problem. Communication-Avoiding QR (CAQR) is a recent algorithm for solving a QR decomposition that is optimal with regard to communication. This talk will describe the implementation of CAQR on the GPU. As an application, we use CAQR to efficiently get the SVD of a tall-skinny matrix, which allows us to perform robust background subtraction on surveillance videos.

Michael Anderson
University of California Berkeley
mjanders@eecs.berkeley.edu

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

MS97

A Critical Path Approach to Excluding Algorithmic Variants

Many algorithms can provide competing variants simply by changing the mathematical approach. We introduce a critical path metric for tiled algo- rithms on multi-core architectures which can help determine the performance of these competing variants. Our metric provides a more intuitive grasp of the performance of a variant within the early development to access whether or not more effort should be applied to tune the algorithmic variant. One of our case study is the two-sided tile reduction to block Hessenberg form which is the main topic of this minsymposium.

Henricus M. Bouwmeester
Department of Mathematics
University of Colorado at Denver
henricus.bouwmeester@ucdenver.edu

Julien Langou
University of Colorado at Denver and Health Sciences Center
julien.langou@cudenver.edu

MS97

Solving the Two-Stage Symmetric Eigenvalue Problem on Multicore Architectures

This talk will describe the two-stage symmetric eigenvalue problem on multicore architectures. The first stage consists in reducing the symmetric matrix to band tridiagonal form using level 3 BLAS operations. The second stage further reduces the band form to the required tridiagonal form using a "Left-Looking" bulge chasing technique to reduce memory traffic and improve data-locality. A dynamic runtime system is used to concurrently schedule both stages on multicore architectures.

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

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

MS97

Hybrid Two-Sided Transformations using GPU Accelerators

We present GPU algorithms for three main two-sided dense matrix factorizations - bidiagonalization for SVD, reductions to upper Hessenberg and tridiagonal forms for general and symmetric eigenvalue problems. These fundamental algorithms are not yet adequately accelerated on homogeneous multicores. Our approach, implemented in the MAGMA library, is based on "hybridization" of the LAPACK algorithms, where the corresponding sequential algorithms are split into tasks and the tasks properly scheduled for execution over the hybrid hardware components.

Stanimire Tomov
Innovative Computing Laboratory
University of Tennessee; Knoxville
tomov@eecs.utk.edu

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

MS98

Direct Numerical Simulation and A Priori Analysis for Large Eddy Simulation of a Two-Phase Multicomponent Mixing Layer

Direct Numerical Simulation (DNS) is performed to study the physics of evaporating multicomponent fuel drops in a transitional, compressible, gaseous mixing layer by using a coupled Eulerian/Lagrangian approach. The Navier-Stokes equations for multiphase flows are extended by four equations to compute the evolution of the statistical vapour composition. The DNS database is used during an a priori study for Large Eddy Simulation to examine two different formulations of the filtered composition equations. The study identifies the modeling requirements of the subgrid scale (SGS) contributions and tests the performance of two different SGS models.

Michael Gloor
Seminar for Applied Mathematics
ETH Zurich
gloor@ifd.mavt.ethz.ch

Josette Bellan
Jet Propulsion Laboratory
California Institute of Technology, Pasadena, CA 91109
josette.bellan@jpl.nasa.gov

MS98

Adaptive Finite Element Methods for Compressible Flows

We present adaptive finite element methods for the compressible Euler equations using continuous piecewise linear approximation in space and time, stabilized by the standard SUPG, and entropy viscosity. To capture and resolve features such as shocks, rarefaction waves, contact discontinuities and boundary layers, the mesh must be sufficiently refined. Hence, construction of adaptive algorithms are necessary for efficient simulations. We present a duality based adaptive finite element method for compressible flow in three space dimensions.

Murtazo Nazarov
PhD student
murtazo@kth.se

MS98

Introductory Talk

The CSE community has grown enormously over the last decade and so have the education programs in this field. Beside individual characteristics of the many and diverse CSE programs, including their strategies, concepts and perspectives, this talk will address new and innovative ideas for MSc and PhD education and training. A par-

ticular focus of this contribution is on factors and parameters for sustainability and a high level of quality in CSE education.

Martin Ruess
Technische Universität München
Computation in Engineering
ruess@tum.de

MS98

Uncertainty Quantification Methods in Data Assimilation for Numerical Weather Prediction

Uncertainty quantification plays a crucial role in data assimilation: overestimating background uncertainty leads to overfit of observations, while underestimating uncertainty restricts the applicability of observations, especially where they are needed most. Using the Weather Research and Forecasting (WRF) model and in collaboration with Argonne National Laboratory, methods to estimate the background error covariance matrix in the context of data assimilation in numerical weather prediction including the "NMC" method, estimation using a Gaussian based on geostrophic balance correlation lengths, an adjoint technique, and a front-based method are presented. The broader impact of these improvements will be briefly discussed as well.

Jeffrey L. Steward
Department of Scientific Computing
Florida State University
jeffsteward@gmail.com

MS99

Introduction to the ITAPS Field Interface

Scientific computing applications focus on physical tensors: scalar temperature, vector velocity, stress tensors, etc. These are formally tensor fields, assigning a tensor quantity to each point in a domain. We are working to develop a simple, flexible applications programming interface (API) for tensor fields that will support the needs of all common families of numerical discretizations with a modest-sized data model and API. Domain support will leverage our previously developed mesh API.

Carl Ollivier-Gooch
University of British Columbia
cfog@mech.ubc.ca

Mark Miller
Lawrence Livermore National Laboratory
miller86@llnl.gov

Fabien Delalondre
Scientific Computation Research Center
Rensselaer Polytechnic Institute
delalf@scorec.rpi.edu

MS99

ITAPS Tools in Computational Evaluation of Alternative Methods of Nuclear Fusion

Computational model and software for the simulation of plasma liner driven magnetoinertail fusion have been developed based on ITAPS front tracking libraries. The formation of plasma liners via the merger of several hundreds of plasma jets, the liner implosion and target compression have been simulated and compared with theoretical predictions and other available numerical studies. Simulations relevant to the future Plasma Liner Experiment at LANL have been obtained as well as simulations at extreme liner energies leading to larger fusion energy gains.

Roman Samulyak
Brookhaven National Laboratory
rosamu@bnl.gov

Lingling Wu
SUNY at Stony Brook
llwu@ams.sunysb.edu

Paul Parks
General Atomics
parks@fusion.gat.com

MS99

Tools to Support Unstructured Meshes on Massively Parallel Computers

A set of tools for managing unstructured meshes on massively parallel computers are being developed. These tools interact with a mesh infrastructure for parallel unstructured meshes based on the DOE ITAPS iMeshP interface. Mesh management tools developed to support parallel adaptive meshing operations include neighborhood aware message packing, predictive load balancing, and incremental partition improvement. The iMeshP interface and mesh management tools are available for download from http://www.tstt-scidac.org/.

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

Kenneth Jansen, Min Zhou, Misbah Mubarak
Rensselaer Polytechnic Institute
jansen@scorec.rpi.edu, zhoum@scorec.rpi.edu,
mubarm@scorec.rpi.edu

Seegyoung Seol
Rensselaer Polytechnic Institute
Scientific Computation Research Center
seol@scorec.rpi.edu

TIng Xie, Aleksandr Ovcharenko
Rensselaer Polytechnic Institute
txie@scorec.rpi.edu, shurik@scorec.rpi.edu

MS99

MeshKit: An Open-Source Toolkit for Mesh Generation

Mesh generation typically requires the use of a collection of algorithms, to address application-specific constraints on the mesh, to apply different meshing approaches in different parts of the domain, or to apply pre- or post-generation algorithms like smoothing. From the algorithm developer point of view, interactions with other meshing algorithms, or support services outside the algorithm being developed, are often necessary. In this talk, we present MeshKit, an open-source library for mesh generation. MeshKit is designed to support both end users and mesh generation algorithm developers. For the former, MeshKit provides both tri/tet and quad/hex algorithms, along with services to co-

ordinate and automate the mesh generation process. For the latter, MeshKit provides important pre-/post-meshing services like support for geometric models, smoothing, and mesh I/O. MeshKit uses the ITAPS mesh and geometry interfaces internally, and therefore can interact with any geometry and mesh databases providing those interfaces.

Timothy J. Tautges
Argonne National Labory
tautges@mcs.anl.gov

Jason Kraftcheck
The University of Wisconsin
kraftche@cae.wisc.edu

James Porter
University of Wisconsin
jvporter@wisc.edu

MS100

Simulations of ICRF Heating using the Delta-f PIC Method

Heating using the ion cyclotron range of frequencies (ICRF) has been explored in a number of tokamak devices, and, in addition, is being considered for ITER. ICRF power can be transferred from the edge to the core without destroying the favorable wave properties of the plasma. We present recent progress on modeling ICRF heating in a tokamak using the δf PIC method. Both 1D and 2D results will be presented from the VORPAL Computational Framework.

Travis M. Austin, David N. Smithe, Chet Nieter, Scott Sides, C. D. Zhou
Tech-X Corporation
austin@txcorp.com, smithe@txcorp.com, nieter@txcorp.com, swsides@txcorp.com, czhou@txcorp.com

MS100

Towards a Particle-in-Cell Solver based on Discontinuous Galerkin Methods

We discuss the ongoing development of a particle-in-cell (PIC) method for the modeling of kinetic plasma phenomena in which the field solver is based on a general high-order accurate discontinuous Galerkin method. The presentation shall motivate the need for such a development, present a detailed discussion of the current state of the effort and an account of open algorithmic challenges. The performance and accuracy will be highlighted through 1D-3D benchmarks to illustrate the potential for such an approach to serve as a modeling tool for complex high-speed kinetic plasma problems in complex geometries.

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

Gustaaf Jacobs
Department of Aerospace Engineering
San Diego State University
gjacobs@mail.sdsu.edu

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

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

MS100

High-Order Vlasov Simulation with Adaptive Mesh Refinement for Laser-Plasma Interaction

We will report on our development of algorithms for Vlasov-Maxwell discretization with adaptive mesh refinement (AMR) for the simulation of laser-plasma interactions. Our approach is based on an explicit, high-order, nonlinear, finite-volume discretization that is discretely conservative, controls oscillations, and can explicitly enforce positivity. AMR algorithms particular to Vlasov simulation, e.g. those required for inter-dimensional reductions and injections, will be discussed. Physically-motivated results in 1+1D and 2+2D will be presented to demonstrate our progress.

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

Jeffrey W. Banks
Lawrence Livermore National Laboratory
banks20@llnl.gov

Bruce Cohen
Fusion Energy Program
Lawrence Livermore National Laboratory
coehn1@llnl.gov

Richard Berger
AX DIVISION
Lawrence Livermore National Laboratory
berger5@llnl.gov

Stephan Brunner
Centre de Recherches en Physique des Plasmas
Ecole Polytechnique Fédérale de Lausanne
stephan.brunner@epfl.ch

MS100

Dynamic Particle Weighting and Velocity Distributions

In this presentation we will discuss our recent work in implementing novel dynamic particle weighting schemes in an electrostatic PIC code. Examples of prior approaches for addressing similar problems include radial particle weighting and the CPK method. Our method is based on selecting pairs of particles to merge based on velocity phase-space matching criteria. A short history of our efforts and our current status, including sheath modeling, will be presented.

Matthew Hopkins
Sandia National Laboratories
mmhopki@sandia.gov

Jeremiah Boerner
Nanoscale and Reactive Processes
Sandia National Laboratories

jjboern@sandia.gov

Paul Crozier, Thomas Hughes
Sandia National Laboratories
pscrozi@sandia.gov, thughe@sandia.gov

Matthew Bettencourt
Sandia National Laboratories
US Air Force Research Laboratory
mbetten@sandia.gov

MS101

3D Nonlinear Electromagnetic Inversion using a Compressed Implicit Jacobian Scheme

We present a compressed implicit Jacobian scheme for the regularized Gauss-Newton inversion algorithm for reconstructing three-dimensional conductivity distribution from electromagnetic data. In this scheme, the Jacobian matrix, whose storage usually requires a large amount of memory, is decomposed in terms of electric fields excited by sources located and oriented identically to the physical sources and receivers. As a result, the memory usage for the Jacobian matrix reduces from O(NFNSNRNP) to O[NF (NS +NR)NP], in which NF is the number of frequencies, NS is the number of sources, NR is the number of receivers and NP is the number of conductivity cells to be inverted. Moreover, we apply the adaptive cross approximation (ACA) to compress these fields in order to further reduce the memory requirement and to improve the efficiency of the method. This implicit Jacobian scheme provides a good balance between the memory usage and the computational time and renders the Gauss-Newton algorithm more efficient. We demonstrate the benefits of this scheme using numerical examples including both synthetic and field data for both cross-well and surface electromagnetic applications.

Aria Abubakar
Schlumberger Doll Research
aabubakar@slb.com

Maokun Li, Guangdong Pan, Jianguo Liu
Schlumberger-Doll Research
mli7@slb.com, gpan@slb.com, jliu17@slb.com

Tarek Habashy
Schlumberger
thabashy@slb.com

MS101

Finite-difference Optimal Gridding Approach for 3D EM Modeling of Geophysical Applications in Time- and Frequency Domain

I present an automatic optimal-gridding tool for finite-difference (FD) 3D modeling of various geophysical electromagnetic (EM) applications: resistivity logging, marine/land EM survey, cross-well EM, ground-penetrating radar. The latest development of the FD technique allows applying a medium-independent automatic grid refinement that allows accurate computation of multi-frequency/multi-time and multi-spacing responses in one run at the approximate cost of a single-frequency/time and single-spacing run. The method allows analyzing 3D models of various geological formations and enables handling exceptionally challenging test cases. Examples of various geophysical applications are presented. The method enabled development of a new separation technique for new-generation triaxial induction/propagation logging data, in order to separate different azimuthal effects to which the older-generation logging tools have weak sensitivity (the tool eccentricity effect, the dipping anisotropy and dipping boundary effects). The technique is based on the symmetrization and rotation of the tensor measurements. This allows extraction of some formation properties directly from the logging data, that makes reduced-dimension inversion schemes applicable. The method also enabled development of a new Focused Source EM (FSEM) survey that allows separation of effects of deep resistive bodies from unwanted shallow effects, based on using a proper combination of measurements. FSEM allows simple visual interpretation of deep reservoir responses, due to higher spatial resolution and higher sensitivity to deep resistors than the standard EM surveys can offer.

Sofia Davydycheva
Rice University and 3DEM Soft
sdavydycheva@smith.com

MS101

Efficient TEM Forward Modelling using Krylov Subspace Approximations

In this talk we describe the solution of the time-dependent quasi-static Maxwell equations to simulate the transient electromagnetic (TEM) method for geophysical exploration. Our approach is based on Ndlec finite element discretization, an exact boundary condition at the air-earth interface and rational Krylov subspace methods for the time-stepping. We present numerical experiments to demonstrate the effectiveness of our method.

Oliver G. Ernst
TU Bergakademie Freiberg
Fakultaet Mathematik und Informatik
ernst@math.tu-freiberg.de

MS101

Large-Scale Joint Imagingof Geophysical Attributes

Large-scale three-dimensional (3D) modeling and imaging is receiving considerable attention in the interpretation of controlled source electromagnetic (CSEM) and magnetotelluric (MT) data in offshore hydrocarbon exploration and geothermal exploration in complex volcanic terrains. The need for 3D modeling and imaging is necessary as the search for energy resources now increasingly occurs in highly complex situations where hydrocarbon effects and geothermal fluids are subtle aspects of their particular geological environment. Further complicating matters is the realization that electrical anisotropy also needs to be incorporated directly into the imaging process. Failure to properly treat anisotropy can produce misleading and sometimes un-interpretable results when broadside/wide azimuth CSEM data is included. Merely excluding broadside data detecting antennas is frequently an issue when 3D coverage is desired. In this workshop we discuss a 3D modeling and imaging approaches that treats both CSEM and MT data jointly as well as transverse anisotropy, which appears to be relevant for many practical exploration scenarios. We discuss effective strategies for large scale modeling and imaging using multiple levels of parallelization on distributed machines, and new developments in porting our 3D modeling and imaging algorithms to GPU compu-

tational platforms.

Gregory Newman
LBL
GANewman@lbl.gov

MS102
Implicit Particle Filter

Particle filters are presented in the setting of nonlinear stochastic differential equations (SDE). The task is to estimate the solution of the SDE conditioned by noisy observations. Traditional particle filters approximate the conditional probability by sequential Monte Carlo; at each step in time, one first guesses a "prior" incorporating the information in the SDE, which is then corrected by sampling weights determined by the observations, yielding a "posterior" density. The catch is that, in common weighting schemes, most of the weights become very small very fast, leading to a catastrophic growth in the number of required particles, especially if the dimension of the state space is large. The implicit filter overcomes this problem by reversing the standard procedure. Rather than find samples and then determine their probabilities, the implicit filter picks probabilities and then generates samples that assume them. The posterior density for each new particle position x^{n+1} is written as $\exp(-F(x^{n+1}))$ (this defines a function F for each particle). We then represent x^{n+1} as a function of a fixed Gaussian variable ξ by solving the equation $F(x^{n+1}) - \phi = \xi^T\xi/2$, where $\phi = \min F$ and T denotes a transpose. The solution of this equation maps high probability samples of ξ onto high probability samples x^{n+1}; x^{n+1} is a sample of the posterior density with weight $\exp(-\frac{1}{2}\phi)J$, where J is the Jacobian of the map $\xi \to x^{n+1}$. There is a great deal of freedom in choosing a map $\xi \to x^{n+1}$ that satisfies the equation, and we use this freedom to produce maps that are easy to perform with a Jacobian that is easy to evaluate. Nothing is assumed in advance about the pdf one is sampling. We demonstrate the power of implicit filters with several challenging examples, including a Kuramoto-Sivashinski equation driven by white noise and a Lorenz system with very sparse data.

Alexander J. Chorin
University of California, Berkeley
Mathematics Department
chorin@math.berkeley.edu

Matthias Morzfe
Lawrence Berkeley National Laboratory
mmo@berkeley.edu

Xuemin Tu
University of Kansas
xuemin@math.berkeley.edu

MS102
A Measure Theoretic Approach to Inverse Sensitivity Analysis

We discuss a numerical method for inverse sensitivity analysis of a deterministic map from a set of parameters to a randomly perturbed quantity of interest. The solution method has two stages: (1) approximate the unique set-valued solution to the inverse problem using derivative information and (2) apply measure theory to compute the approximate probability measure on the parameter space that solves the inverse problem. We discuss convergence and numerical analysis of the method.

Jay Breidt
Department of Statistics
Colorado State University
jbreidt@stat.colostate.edu

Troy Butler
Institute for Computational Engineering and Sciences
University of Texas at Austin
tbutler@ices.utexas.edu

Donald Estep
Colorado State University
estep@math.colostate.edu or don.estep@gmail.com

Jeff Sandelin
Department of Mathematics
Colorado State University
tsandelin@comcast.net

MS102
Stochastic Collocation based on Interpolation with Arbitrary Nodes

We present a generalized algorithm for the 'least interpolant' method of Carl de Boor and Amos Ron for polynomial interpolation on arbitrary data nodes in multiple dimensions. Our variation on the least interpolant produces an interpolant that can be tailored for various probability distributions. We empirically analyze conditioning of the associated Vandermonde-like matrices and present a few examples illustrating utility of the method for generalized Polynomial Chaos collocation methods and for generating response surfaces.

Dongbin Xiu, Akil Narayan
Purdue University
dxiu@purdue.edu, akil@purdue.edu

MS102
Predictability in Stochastic Reaction Networks

Predictability analysis in stochastic reaction networks is typically challenged by intrinsic noise. We utilize non-intrusive spectral expansions to efficiently propagate input parametric uncertainties in the presence of intrinsic stochasticity. To address the curse of dimensionality, orthogonal spectral projections are performed using a sparse quadrature approach that is shown to perform better than High Dimensional Model Representation (HDMR) for the benchmark problem. The methodology is illustrated for the gene regulation network of the *Bacillus Subtilis* bacterium.

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

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

MS103

Distributed Control of Cascading Blackouts

We describe continuing experience with online, adaptive control algorithms for minimizing the impact of a cascading blackout. Assuming that a contingency has taken place, the algorithms compute a locally-optimum control that will be applied as the cascade unfolds. The control will take as inputs observable quantities, such as line overloads, and output control actions such as load shedding. The objective is to stop the cascade in a stable state with minimal loss of demand. We present computational experience using large real-life grids.

Daniel Bienstock
Columbia University IEOR and APAM Departments
IEOR Department
dano@columbia.edu

MS103

Models and Algorithms for N-k Survivable Grid Design

Not available at time of publication.

Richard Li-Yang Chen
Sandia National Laboratories
Livermore, CA
rlchen@sandia.gov

Amy Cohn
University of Michigan, Ann Arbor
amycohn@umich.edu

Ali Pinar
Sandia National Labs
apinar@sandia.gov

MS103

Flexibility in Electricity Markets

Greater flexibility is especially important in the context of a smarter electric grid. In this talk we examine the flexibility of electric assets in electric network optimization models, specifically the unit commitment and dispatch problems. We review how flexibility is represented now, how it might be represented better, and issues that need to be addressed.

Kory Hedman
Arizona State University
kory.hedman@asu.edu

MS103

Protecting Electric Power Grids from Terrorist Attack: Solving Full-scale, Three-stage Stackelberg Games

Abstract not available at time of publication.

Kevin Wood
Naval Postgraduate School
Monterey, CA
kwood@nps.edu

MS104

Preconditioned Krylov Subspace Methods on GPU Applied to Multiphase CFD

By using C for CUDA environment, we have successfully implemented and evaluated various kinds of preconditioning methods for Krylov subspace solvers on GPU, including SOR variants, Multigrid, Line-by-Line and incomplete Cholesky decomposition. These preconditioned Krylov subspace methods has then been applied to multiphase CFD simulation and real-time visualization on GPU.

Hidetoshi Ando, Takuma Shibuya, Koji Toriyama
University of Yamanashi
ando@yamanashi.ac.jp, g09mkg14@yamanashi.ac.jp, toriyama@yamanashi.ac.jp

MS104

Building the Next Generation of Scalable Manycore Applications and Libraries

Multicore nodes have become the standard building block of scalable computers systems, and core counts per node promise to increase dramatically over the next decade. Although single level MPI-only programming approaches can work for some applications, we will need different approaches going forward. In this talk we make observations about existing successful approaches to parallel application development and highlight some key requirements for success in developing the next generation of scalable applications.

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS104

Efficient Use of GPUs and a Large Number of CPUs in a Dense Linear Algebra Library

The efforts of implementing dense linear algebra on multicore and accelerators have been pursued in two different directions, one that emphasizes the efficient use of multicore processors, represented by the PLASMA project, and another that emphasizes the use of accelerators, represented by the MAGMA project. While the former makes great usage of multicores, it is void of support for accelerators. While the latter makes great usage of GPUs, it seriously underutilizes CPU resources. This presentation introduces an approach for efficiently combining GPU accelerators with a large number of classic multicores for dense linear algebra computations.

Jakub Kurzak
University of Tennessee Knoxville
kurzak@eecs.utk.edu

MS104

Shared Memory and GPU Algorithms for Finite-element Linear-system Assembly

Computing element-stiffness matrices in a finite-element analysis application and assembling them into a global sparse linear system is a computationally intensive problem which could benefit from parallelism. In a multicore CPU setting it seems natural to have separate threads compute each element-matrix, but the challenge is that threads could then collide when making contributions to the same global matrix row. Algorithms suitable for multicore CPUs and GPUs will be discussed, along with performance re-

sults.

Alan B. Williams
Sandia National Laboratory
Distributed Systems Research Department
william@sandia.gov

MS105

An h-p adaptive Discontinuous Galerkin Method for the Navier-Stokes Equations

We discuss the implementation of an h-p adaptive approach for high-order accurate discontinuous Galerkin discretizations of the Navier-Stokes equations. The discretization operates on mixed element triangular and quadrilateral meshes in two dimensions and h-adaptivity is incorporated using non-conforming elements. Both the primal (Navier-Stokes) and dual (adjoint) equations are converged using an efficient h-p multigrid solver. The adjoint solution is used to estimate discretization error in selected simulation output functionals which in turn drives the adaptation process. A smoothness indicator is also implemented in order to choose between h and p refinement options. Predicted and observed error reduction at each refinement level are compared and used to gauge convergence of the error estimate. Dual consistency of the discretization and its effect on error estimates and adaptive convergence are also discussed.

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

MS105

High Fidelity Simulations of Flapping Wings Designed for Energetically Optimal Flight

We present our work on a multi-fidelity framework for inverse design of flapping wings. A panel method-wake only energetics solver is used to define the energetically optimal wing shape and flapping kinematics. Candidate designs are then simulated using our high-order accurate discontinuous Galerkin solver, based on fully unstructured curved meshes of tetrahedra, a mapping-based ALE solver, and efficient parallel Newton-Krylov solvers, to gain insight into practical wing designs and the influence of viscous effects.

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

MS105

On the Implementation of Rosenbrock-W Methods

Abstract not available at time of publication.

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

MS105

New Practices on High Order Numerical Methods

High order numerical methods have been applied in practical applications and their advantages have been shown in terms of accuracy, flexibility and efficiency. Multi-grid technique has been used to improve the performance of the high order method. Different high order bases have been used in the simulations of different physical phenomena. Their performances have been analyzed and compared. Practice and challenges on applying high order methods in large scale computing will be shown and discussed.

Jin Xu
Math. & Computer / Physics Division, Argonne National Lab.
jin_xu@anl.gov

MS106

A Fast Optimization Method for the Chan-Vese Model in Image Segmentation

The Chan-Vese model is used to segment images that are approximately piecewise continuous (a large class of images of practical interest, e.g. in medical imaging). The goal of this work is to devise a fast and robust optimization method to compute the boundaries minimizing the Chan-Vese model. To achieve this, we compute the second shape derivative of the Chan-Vese model and incorporate this in the gradient descent process in order to reduce the number of iterations needed for convergence. We discretize the gradient descent using the finite element method. We also employ adaptive discretizations for both the curves and the domains, resulting in further reductions in computations.

Gunay Dogan
Theiss Research
National Institute of Standards and Technology
gunay.dogan@nist.gov

MS106

Multilevel Methods for Image Deblurring

We present multilevel methods for discrete ill-posed problems arising from the discretization of Fredholm integral equations of the first kind. In particular, we present wavelet-based multilevel methods for signal and image restoration problems as well as for blind deconvolution problems. We show results that indicate the promise of these approaches.

Malena I. Espanol
Graduate Aeronautical Laboratories
California Institute of Technology
mespanol@caltech.edu

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

MS106

Adaptive Methods in Total Variation based Image Restoration

Abstract not available at time of publication.

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

MS106

Multiphase Scale Segmentation and a Regularized

K-means

A typical Mumford-Shah-based image segmentation is driven by the intensity of objects in a given image, and we consider image segmentation using additional scale information. Using the scale of objects, one can further classify objects in a given image from using only the intensity value. We develope a fast automatic data clustering method for the data clustering. The model automatically gives a reasonable number of clusters by a choice of a parameter. We explore various properties of this classification model and present different numerical results.

Sung Ha Kang
Georgia Inst. of Technology
Mathematics
kang@math.gatech.edu

Berta Sandberg
Adel Research, Inc.
berta.sandberg@adelresearch.com

Andy Yip
Department of Mathematics
National University of Singapore
andyyip@nus.edu.sg

MS107

High Resolution Simulation of the Intracranial Arterial Network

Simulating the human arterial tree is a grand challenge requirung state-of-the-art mathematical algorithms and computers. In this talk, we will discuss modeling of arterial flow in a patient-specific intracranial arterial tree and present a methodology we have developed. Our focus is on ultrascale parallel algorithms based on a two-level domain decomposition (2DD) method for the solution of Navier-Stokes equations with billions of unknowns on thousands of computer processors.

Leopold Grinberg
Brown University
lgrinb@dam.brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
gk@cfm.brown.edu

MS107

Vascular Modelling Toolkit: An Open-source Framework for Image-based Modeling and Analysis of Blood Vessels

Thanks to the development of the imaging modalities, image-based computational fluid dynamics (CFD) has shown its potentiality in elucidating the role of haemodynamics in cardiovascular diseases. However, the process from images to CFD is still a time-consuming and operator-dependent task, while the variability of real-anatomies challenges CFD application to large patient populations. VMTK represents a concrete step in addressing these concerns: robust and objective computational techniques will be presented for geometric and hemodynamics modeling.

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

Luca Antiga
Mario Negri Institute, Bergamo,Italy
Orobix, Bergamo, Italy
luca.antiga@orobix.com

MS107

Some Recent Challenges in Patient-Specific Cardiovascular Mathematics: from Forward to Inverse Problems

Combination of data coming from medical images and numerical solutions is a fundamental step for patient-specific simulations in computational hemodynamics. We will consider some examples of data assimilation for merging velocity measures and numerical simulations and for estimating physical parameters such as vessel compliance. Emphasis will be focused on controllability problems for the incompressible Navier-Stokes equations and the impact of noise of measures on the reliability of the numerical results.

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

Marta D'Elia
MathCS, Emory University
mdelia2@mathcs.emory.edu

Mauro Perego
MathCS, Emory University, Atlanta, GA
Politecnico di Milano, Italy
mauro@mathcs.emory.edu

Christian Vergara
Universita' degli Studi di Bergamo
Bergamo, Italy
christian.vergara@unibg.it

MS107

Surgical Modeling of the Fontan Procedure: Growing the Knowledge Base One Patient at a Time

Long-term morbidities of single ventricle patients can be often related to the blood dynamics through the surgical construct (the total cavopulmonary connection - TCPC). To understand/improve the outcomes of these patients, we have generated patient-specific anatomical (experimental and computational) models and developed a software to mimic surgical different alternatives for the TCPC design, to optimize energetic efficiency/flow distribution. We review our experience, and demonstrate the value of a patient-specific surgical approach.

Ajit P. Yoganathan
Department of Biomedical Engineering
Georgia Institute of Technology
ajit.yoganathan@bme.gatech.edu

Diane De Zélicourt
Department of Biomedical Engineering
Georgia Institute of Technology, Atlanta, Georgia;
gtg258f@mail.gatech.edu

Christopher Haggerty, Maria Restrepo, Kartik Sundareswaran
Department of Biomedical Engineering

Georgia Institute of Technology, Atlanta, Georgia
chaggerty3@mail.gatech.edu, mrestrepo3@gatech.edu, gte771s@mail.gatech.edu

Jarek Rossignac
Georgia Institute of Technology
jarek@cc.gatech.edu

Kirk Kanter
Division of Cardiothoracic Surgery
Emory University School of Medicine, Atlanta, Georgia
kkanter@emory.edu

William Gaynor, Thomas Spray
Division of Cardiothoracic Surgery
Childrens Hospital of Philadelphia, Philadelphia
gaynor@email.chop.edu, spray@email.chop.edu

MS108

Direct Numerical Simulation of Particulate Flows on 294912 Processor Cores

High performance on current supercomputers is mainly obtained by highly optimized kernels, which handle one specific problem with given restrictions. waLBerla in contrast is a large C++ software framework and aims at high performance. It is our extensible implementation of a parallel coupled fluid-structure solver. In this talk, the suitability of waLBerla for current and upcoming supercomputers is discussed and performance values on up to 294 912 processor cores of the Blue Gene/P are given.

Jan Goetz
System Simulation , Friedrich-Alexander-University Erlangen-Nuernberg, Germany
jan.goetz@informatik.uni-erlangen.de

Klaus Iglberger
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
klaus.iglberger@informatik.uni-erlangen.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ruede@informatik.uni-erlangen.de

MS108

Massively Parallel Solution of 3D Ice Sheet Flow Models

The flow of polar ice sheets is characterized by a wide range of length scales, with localized flow features many orders of magnitude smaller than continental scales. To capture this wide range of scales in 3D ice sheet models, we employ adaptive mesh refinement and higher-order methods. We investigate the application of multigrid preconditioners to the resulting systems, as well as issues arising from capturing the free surface at the ice/air boundary.

Tobin Isaac
University of Texas at Austin
tisaac@ices.utexas.edu

Carsten Burstedde
The University of Texas at Austin
carsten@ices.utexas.edu

Georg Stadler, Omar Ghattas
University of Texas at Austin
georgst@ices.utexas.edu, omar@ices.utexas.edu

MS108

Massively Parallel Stabilized Finite Element Methods for Fluid Dynamics

Abstract not available at time of publication.

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

Kenneth Jansen
Rensselaer Polytechnic Institute
jansen@scorec.rpi.edu

MS108

Multiscale Simulation of Blood Flow in the Coronary Arteries

We present a computational model for blood flow in coronary arteries. The simulation uses the Lattice Boltzmann method coupled with microscopic Molecular Dynamics modeling of the red blood cells, which interact with one another and the surrounding fluid, to provide a multiphysics and multiscale representation of the flow in patient specific geometries. We will present the modeling methods as well as the techniques leveraged to achieve excellent scaling on up to 294,912 processor cores.

Amanda Peters
Harvard School of Engineering and Applied Sciences
apeters@fas.harvard.edu

Simon Melchionna
EPFL
Lausanne
simone.melchionna@epfl.ch

Jonas Latt
Ecole Polytechnique Federale de Lausanne
jonas.latt@epfl.ch

Sauro Succi, Massimo Bernaschi
Istituto per le Applicazioni del Calcolo
succi@iac.rm.cnr.it, m.bernaschi@iac.cnr.it

Efthimios Kaxiras
Dept. of Physics and School of Engg. and Applied Sciences
Harvard University
kaxiras@physics.harvard.edu

Mauro Bisson
Dipartimento di Informatica
Universita' La Sapienza
bisson@di.uniroma1.it

MS109

Performance Optimizations for Heterogeneous and Hybrid 3D Lattice Boltzmann Simulations on Highly Parallel On-Chip Architectures

Optimized CFD solvers for GPU computing can be up to

an order of magnitude faster than on current standard x86-type servers. We use a lattice Boltzmann based flow solver kernel, which is proven to perform well both on CPUs and GPUs. The focus is to go beyond a single compute node and show the potential of multinode GPU clusters for our solver as well as the capability of utilizing heterogeneous hardware setups together with hybrid OpenMP+MPI programming techniques.

Johannes Habich
Regional Computing Center Erlangen
Johannes.Habich@rrze.uni-erlangen.de

MS109

GPU Implementation of a Numerical Wave Tank for the Simulation of Non-linear and Turbulent Free Surface Flow

We present an efficient GPU implementation for the numerical simulation of three-dimensional free surface flow on the basis of the Lattice Boltzmann method (LBM) and the nVIDIA CUDA framework. Several validations and applications in the field of civil and environmental engineering will be presented. The runtimes of the numerical simulations are only up to one order of magnitude higher than the time scale of the real world event, for single precision simulations on one GPU.

Christian Janssen
Institute for Computational Modeling in Civil Eng.
Technical University of Braunschweig, Germany
janssen@irmb.tu-bs.de

Martin Schönherr
Institute for Computational Modeling in Civil Engineering
Technische Universität Braunschweig
schoenherr@irmb.tu-bs.de

Manfred Krafczyk
Tech. Univ. of Braunschweig
Germany
kraft@irmb.tu-bs.de

MS109

Comparing Performance and Energy Efficiency of Lightweight Manycore to GPU for LBM methods

The rising power consumption of conventional cluster technology has prompted investigation of architectural alternatives that offer higher computational efficiency. This presentation compares the performance and energy efficiency of highly tuned code on the three architectural alternatives the Intel Nehalem X5530 multicore processor, the NVIDIA Tesla C2050 GPU, and a gate-level architectural model of a manycore chip design called "Green Wave."

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

MS109

The Lattice Boltzmann Method: Basic Performance Characteristics and Performance Modeling

High memory bandwidth requirements and regular data access patterns are common features of many CFD applications. Choosing a 3D Lattice Boltzmann kernel and focusing on single node performance we first demonstrate basic optimization techniques for modern multicore CPUs and GPUs. Concepts and potentials of more advanced multicore-aware temporal blocking techniques will also briefly be covered. Guided by established performance models we compare typical performance characteristics of the latest x86-CPU and GPU generations.

Gerhard Wellein, Georg Hager
Erlangen Regional Computing Center (RRZE), Germany
Gerhard.Wellein@rrze.uni-erlangen.de,
georg.hager@rrze.uni-erlangen.de

Johannes Habich
Regional Computing Center Erlangen
Johannes.Habich@rrze.uni-erlangen.de

MS110

Toward Future Environmental Modeling

Abstract not available at time of publication.

Mohammad Abouali
CSRC
San Diego State University
mabouali@sciences.sdsu.edu

MS110

Isogeometric Analysis in the World of CAD

Isogeometric analysis is a new finite element approach. It unifies the worlds of computer aided design (CAD) and finite element analysis (FEA) and makes it possible to use the same free-form-surfaces model for CAD and FEA.
The presentation shows a full static analysis in the commercial NURBS-based geometry modeling tool "Rhino" using the isogeometric approach. The use of the same geometric description for both the CAD and FEA model has significant advantages in practical application. The presentation shows illustrative examples from structural analysis to demonstrate the practical application of the approach.

Michael Breitenberger
BGCE
Technical University Munich
michael.breitenberger@mytum.de

MS110

Local-isotropy in Direct Numerical Simulation of Turbulent Channel Flow at High-Reynolds Number

Turbulent channel flow (TCF) is one of the most canonical wall-bounded turbulent flows, and there have been therefore extensive studies on TCF by direct numerical simulations (DNS). High-resolution DNS of TCF provides us not only with detailed practical data for modeling of wall-bounded turbulence but also with detailed fundamental data to explore universality in the small-scale statistics of high-Reynolds-number wall-bounded turbulence. In the present study, we developed a DNS code of TCF for the current version of the Earth Simulator (ES), attained a 5.9Tflops computation (11.3% of the peak performance) in the DNS of TCF on 1024x1536x1024 grid points using 64 nodes of ES, and achieved the friction Reynolds number $Re_\tau = 2560$, which is the world's largest Re_τ so far achieved in DNS of TCF. The analysis of the DNS data shows that the one-dimensional longitudinal energy

spectrum which is consistent with the prediction of Kolmogorov's theory (K41) is realized in the so-called log-law layer of high-Reynolds-number wall-bounded turbulence, but shows also that local isotropy, which is a key hypothesis in K41, does not necessarily hold at sufficiently small scales.

Koji Morishita
Dept Computational Science and Engineering
Nagoya University
morishita@fluid.cse.nagoya-u.ac.jp

MS110

Hybrid Discontinuous Galerkin Methods for Incompressible Flow Problems

We propose a hybrid discontinuous Galerkin method for incompressible flow, that allows for nonconforming h-refinements and approximations with locally varying polynomial degrees (hp-adaptivity). We present order optimal a-priori estimates and present a simple, efficient and reliable error estimator. The agreement of the theoretical bounds with numerical results will be demonstrated.

Christian Waluga
AICES Graduate School
RWTH Aachen University
waluga@aices.rwth-aachen.de

Herbert Egger
Institute for Mathematics and Scientific Computing
University of Graz
herbert.egger@uni-graz.at

MS111

Parallel Mesh Generation by Distributing CAD Geometry

Having geometry information is important to make good quality of meshes such as in highly curved geometry or adaptive meshing. Therefore, an approach to parallel mesh generation by distributing assembly CAD geometry will be described in this talk. This method is based on partitioning of geometric surfaces and volumes across processors, with some surfaces shared between processors. Meshing of interface surfaces and volume bodies is partitioned across processors using a graph partitioning method to minimize communication cost and maximize load balance. Model edges are meshed serially by the root processor, interface surfaces meshes are generated by one processor and sent to the other sharing processor, and non-interface surfaces and volume interiors are meshed in parallel. Non-interface surface meshing and asynchronous messages are used to hide the cost of message passing latency. It is implemented using ITAPS interfaces to read/distribute geometry by iGeom, to store/communicate mesh information by iMesh and to relate them for meshing geometry in parallel by iRel.

Hong-Jun Kim
Argonne National Laboratory
hongjun@mcs.anl.gov

Timothy J. Tautges
Argonne National Labory
tautges@mcs.anl.gov

MS111

Parallel Hybrid Mesh Adaptation

A procedure for parallel anisotropic mesh adaptation accounting for mixed element types is presented. The parallel adaptive approach uses local mesh modification procedures in a manner that maintains layered and graded elements near the walls, which are popularly known as boundary layer or semi-structured meshes, with highly anisotropic elements of mixed topologies. The technique developed is well suited for parallel viscous flow applications where the exact knowledge of the mesh resolution is unknown a priori.

Aleksandr Ovcharenko
Rensselaer Polytechnic Institute
shurik@scorec.rpi.edu

MS111

Towards Large-Scale Predictive Flow Simulations using ITAPS Services

This talk will present advances made in massively parallel fluid-dynamics simulations based on tools/services provided by ITAPS. Examples including complex geometries and flow physics will be covered. It will be shown that even with these advances, unknowns associated with real-world problems limit the use of current tools in a comprehensive and predictive manner. Therefore, we will also present basic steps to advance in this regard, for example, based on ITAPS-related manipulation services for geometry, etc.

Onkar Sahni
PECOS/ICES at the Univeristy of Texas at Austin
sahni@ices.utexas.edu

Kenneth Jansen
Rensselaer Polytechnic Institute
jansen@scorec.rpi.edu

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

MS112

Damping of Spurious Reflections off of Coarse-fine Adaptive Mesh Refinement Grid Boundaries

Adaptive mesh refinement (AMR) is an efficient technique for solving systems of partial differential equations numerically. The underlying algorithm determines where and when a base spatial and temporal grid must be resolved further in order to achieve the desired precision and accuracy in the numerical solution. However, systems of PDEs with low dissipation prove problematic for AMR. In such a system, a wave traveling from a finely resolved region into a coarsely resolved region encounters a numerical impedance mismatch, resulting in spurious reflections off of the coarse-fine grid boundary. Here, we present a scheme for damping these spurious reflections and apply it to Maxwell's Equations.

Sven Chilton
University of California, Berkeley
chilton@berkeley.edu

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

MS112

Fast Summation Method for Electro-Magnetics using the Yukawa Screening Potential

Abstract not available at time of publication.

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

MS112

Asymptotic-Based Numerical Methods for Plasmas in the Quasineutral and Strong Magnetic Confinement Regimes

An efficient and accurate numerical scheme for the solution of highly anisotropic elliptic equations is presented. The anisotropy is driven by a magnetic field. Hence, it can be strong in some regions and weak in the others. Moreover, the direction may vary. Our method is based on the Asymptotic Preserving reformulation of the original problem, permitting an accurate resolution independly of the anisotropy strength and direction without the need of mesh adapted to the anisotropy.

Pierre Degond
CNRS, Institut de Mathematiques de Toulouse
CNRS et Universite Paul Sabatier
pierre.degond@math.univ-toulouse.fr

Fabrice Deluzet
Institut de Mathematiques
fabrice.deluzet@math.univ-toulouse.fr

Alexei Lozinski, Jacek Narski
Institut de Mathematiques de Toulouse
Universite Paul Sabatier - Toulouse III
alexei.lozinski@math.univ-toulouse.fr,
jacek.narski@gmail.com

Claudia Negulescu
CMI/LATP
Universite de Provence
claudia.negulescu@cmi.univ-mrs.fr

MS112

Design and Preliminary Results for PIC on GPUs with Python

The Particle-in-Cell (PIC) model, popular for many multiscale plasma physics problems, is well suited to parallel architectures. Our goals include mapping PIC to GPUs, and providing a programmatic interface. We exploit multiple levels of parallelism using tools such as PyCUDA/PyOpenCL, IPython, and ZeroMQ. The design, and promising preliminary performance results will be presented. Our GPU code achieved one order of magnitude performance improvement over a single CPU core on a simplified push subproblem.

Min Ragan-Kelley
UC Berkeley AS&T
minrk@berkeley.edu

John Verboncoeur
Dept. Nuclear Engineering
University of California
johnv@nuc.berkeley.edu

MS113

Eigenanalysis of Polynomial Chaos Representations of Uncertain ODE Systems

Ordinary differential equations (ODEs) with uncertain parameters and/or initial conditions are of practical interest in the analysis and reduction of many physical systems. We use Galerkin projection onto polynomial chaos (PC) basis functions to reformulate an uncertain system of ODEs as a deterministic ODE system that describes the evolution of the PC modes. We analyze the eigenstructure of the uncertain Jacobian and that of the Jacobian of the projected PC-system, outlining general statements about the relation between the two solutions. We also discuss the use of eigenvalues and eigenvectors of both systems for the reduction of ODE systems under uncertainty.

Robert D. Berry
Sandia National Laboratories
rdberry@sandia.gov

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

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

MS113

Efficient Computation of Failure Probability

Computing failure probability is a critical step in many applications such as reliability based optimization. The most straightforward method is to sample the response space. This can be prohibitively expensive because each sample requires a full scale simulation of the underlying physical system. An alternative way is to generate an accurate surrogate model for the system and sample the surrogate. This can be extremely efficient as long as one can construct such a surrogate. In this talk we demonstrate that a explicit surrogate approach is fundamentally flawed, no matter how accurate the surrogate is. Furthermore, we present a hybrid algorithm combines the surrogate and sampling approach and address the robust problem described above. Rigorous error estimate will be presented as well as the numerical examples.

Dongbin Xiu, Jing Li
Purdue University
dxiu@purdue.edu, li176@purdue.edu

MS113

Multiscale Methods for Statistical Inference in Elliptic Problems

Estimating the coefficients of an elliptic PDE from observations of the solution is an ill-posed inverse problem. Finite data resolution and the smoothing character of the forward operator limit one's ability to recover fine-scale information, but simultaneously suggest an intrinsically multiscale approach to inversion. We formulate a new Bayesian inference approach in this context. Essential components of the approach are (1) constructing a prior distribution over stiffness matrices at the coarse scale; and (2) generating realizations of the fine scale conditioned on coarse-scale information, via nonlinear constraints. Numerical examples

demonstrate the efficiency and accuracy of the method.

Matthew Parno, Youssef M. Marzouk
Massachusetts Institute of Technology
mparno@mit.edu, ymarz@mit.edu

MS113

Sensitivity-based Reduced Order Modeling of High Dimensional Uncertain Inputs

We present an approach of reducing the dimensionality of the input space of a stochastic simulation using SVD of randomly sampled sensitivity gradients. Such reduction of the dimension of parameter space has deep implication on the number of samples required to build accurate surrogate surfaces. The curse of dimensionality encountered by most uncertainty propagation schemes makes the required samples grow very fast as the dimensionality of the parameter space. Therefore, reduction in the parameter dimensions through our method can dramatically reduce the computational cost for uncertainty quantification with high dimensional stochastic space.

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
doostan@colorado.edu

Qiqi Wang
Massachusetts Institute of Technology
qiqi@mit.edu

Paul Constantine
Sandia National Labs
pconsta@sandia.gov

Gianluca Iaccarino
Stanford University
jops@stanford.edu

MS114

Challenges for Optimization in Future Electric Power Systems

Abstract not available at time of publication.

Jeremy Bloom
IBM
bloomj@us.ibm.com

MS114

Resource Commitment and Dispatch in the PJM Wholesale Electricity Market

This presentation will provide an overview of the size and scope of the PJM wholesale electricity market and the scheduling and operational challenges we expect to encounter in the future with increased penetration of intermittent and distributed resources. It will also provide a brief overview of the development and implementation of the mixed-integer programming based unit commitment and dispatch programs in the PJM market and will describe the benefits achieved to date.

Andrew L. Ott
PJM
ott@pjm.com

MS114

Long-term Planning for the Power Grid

Abstract not available at time of publication.

Jean-Paul Watson
Sandia National Laboratories
Discrete Math and Complex Systems
jwatson@sandia.gov

MS114

High-Performance Computing for Transmission Expansion Planning

We address computational issues arising in transmission expansion planning. The large-scale adoption of smart-grid programs and intermittent renewables coupled to changing market designs require of highly flexible transmission systems. A powerful technique that can be used to address this problem is stochastic mixed-integer optimization. However, the associated network complexity, time scales, number of uncertain scenarios, and number of integer variables arising in transmission planning problems severely limit the scope of applications. In this project, we leverage high-performance computing capabilities through the development of a set of optimization tools that enable the solution of problems of unprecedented complexity. We present scalability studies on BlueGene/P. In addition, we present a novel set of complementarity formulations to handle transmission switching decisions in DC and AC power flow models.

Victor Zavala, Zhen Xie
Argonne National Laboratory
vzavala@mcs.anl.gov, zhenxie@mcs.anl.gov

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

MS115

Performance of a Sparse Hybrid Linear Solver on a Multicore Cluster

Sparse hybrid solvers are a trade-off between direct methods and iterative methods. Part of the computation is first performed with a direct method in order to ensure numerical robustness; the algorithm then switches to an iterative method to alleviate the computational complexity and memory usage. The convergence and number of iterations depend on the amount of computation performed in the direct part. We present the benefits of such a hierarchical approach on a multicore cluster.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

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

Abdou Guermouche
LaBRI-INRIA futurs
abdou.guermouche@labri.fr

Azzam Haidar
Department of Electrical Engineering and Computer

Science
University of Tennessee, Knoxville
haidar@utk.edu

Yohan LEE-TIN-YIEN
INRIA
yohan.lee-tin-yien@inria.fr

Jean Roman
INRIA / LaBRI
jean.roman@inria.fr

MS115

Auto-Tuned Linear Algebra Computations for Krylov Methods on Multicore and GPUs

We discuss performance optimizations for sparse Krylov solvers. By modelizing and then micro-benchmarking matrix and vector computations, we manage to choose a better sparse structure for the matrices and a better data distribution for the vector computations on multicore processors and GPUs. Our experiments on different generations of processors and three generations of GPUs from first CUDA capable GPUs to the Fermi architecture show different optimums, with speed-ups of 2-3x on processors and up to 10x on GPUs versus respective original performance. We conclude that autotuning is a good way to deliver better performance on both hardware.

Serge G. Petiton
CNRS/LIFL and INRIA
serge.petiton@lifl.fr

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

Jerome Dubois
CEA-DEN
jerome.dubois@cea.fr

MS115

Developing a 3-D Sweep Radiation Transport Code for Large-scale GPU Systems

Solution of the seven-dimensional Boltzmann equation is a computationally expensive operation used in multiple science applications, typically solved using 3-D sweep algorithms which are difficult to parallelize efficiently. This talk discusses a 3-D sweep algorithm implemented on NVIDIA GPUs, as part of the effort to port the Denovo radiation transport code to ORNL's next-generation GPU-based petascale system. This talk presents strategies used to expose thread parallelism as well as performance results on NVIDIA Fermi GPUs.

Wayne Joubert
Oak Ridge National Laboratory
joubert@ornl.gov

MS115

Parallel Preconditioning Methods for Ill-Conditioned Problems on Multicore Clusters

In this work, authors developed a method for parallel preconditioning based on HID (Hierarchical Interface Decomposition) for finite-element applications with ill-conditioned coefficient matrices, and implemented the developed method on to multi-core/multi-socket clusters using OpenMP/MPI hybrid parallel programming model. In this method, HID is also applied to computations on each domain. Robustness and efficiency of the developed method have been demonstrated on T2K Open Supercomputer (Tokyo) and Cray XT4 (LBNL).

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

Satoshi Ohshima
Univ. of Tokyo
ohshima@cc.u-tokyo.ac.jp

MS116

Viscosity based Shock Capturing for High Order Space and Time Adaptive Discontinuous Galerkin Schemes

This talk deals with shock capturing strategies using the explicit space-time expansion discontinuous Galerkin (STE-DG) scheme. With its high order local time-stepping functionality, the increasing time-step restrictions, introduced by using viscosity for capturing shocks, can be overcome. Starting with ideas by Persson and Peraire, we adopted and refined their approach, while searching for alternative, less parameter-dependent strategies. The talk will introduce the scheme, shock sensors and capturing mechanisms, together with applications for Navier-Stokes and Magnetohydrodynamics.

Christoph Altmann
Institut fuer Aerodynamik und Gasdynamik
Universitaet Stuttgart
altmann@iag.uni-stuttgart.de

Gregor Gassner
Institute for Aerodynamics and Gasdynamics
Universitaet Stuttgart
gassner@iag.uni-stuttgart.de

Claus-Dieter Munz
Institut für Aerodynamik und Gasdynamik (IAG)
munz@iag.uni-stuttgart.de

MS116

Shock Capturing in a Time-Explicit Discontinuous Galerkin Method on the GPU

Having recently shown that high-order unstructured discontinuous Galerkin (DG) methods form a spatial discretization that maps well onto graphics processing units (GPUs), I will report on ongoing work aiming to design a GPU-capable (i.e. data-local) shock capturing scheme based on DG. I will discuss the mathematical and algorithmic motivation for a shock sensor and the design of a GPU-suited smoother that uses the sensor's findings. In addition, I will touch upon accuracy results for the resulting schemes on a variety of nonlinear conservation laws.

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

Timothy Warburton
Department of Computational and Applied Mathematics
Rice University
timwar@caam.rice.edu

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

MS116

A h-p Adaptive Non-hydrostatic Atmospheric Flow Solver

We present a framework aimed at problems arising in the geosciences: it can solve the non-hydrostatic equations of the atmosphere, shallow-water problems, tsunami prediction and advection tests. The DG method is employed, however, any element based discretization can be supported. A mesh database enables parallel non-conforming mesh refinements in both h-p. Modern coding techniques are employed. They enable the seamless optimization of compute kernels: SSE vector operations are supported but could be extended to GPUs.

Amik St-Cyr
National Center for Atmospheric Research
Institute for Mathematics Applied to the Geosciences
amik@ucar.edu

Sébastien Blaise
National Center for Atmospheric Research
sblaise@ucar.edu

David Hall
National Center for Atmospheric Research
1850 Table Mesa Drive, Boulder CO 80305
halld@ucar.edu

MS116

Exponential Integrators: Construction, Implementation and Performance

Exponential integrators offer an efficient alternative to explicit and implicit methods for integration of large stiff systems of ODEs. While first exponential schemes were introduced as early as 1958, they did not attain wide popularity since the algorithms were deemed too computationally expensive. However, recent research on exponential integrators allowed construction of efficient methods which are competitive with commonly used implicit and explicit schemes. In this talk we will provide an overview of the structure and performance of exponential integrators. We will describe what design principles allow for construction of efficient exponential schemes and quantitatively illustrate their computational advantages compared to Krylov-based implicit and explicit integrators using a set of test problems. We will present a new class of exponential propagation iterative schemes (EPI) and discuss how these methods can be derived using the B-series theory. The major building blocks of the new EPI methods, such as their structure, adaptivity and performance will be discussed.

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

John Loffeld
University of California
Merced
jloffeld@ucmerced.edu

MS117

Designing Optimal Filters for Ill-posed Inverse Problems

Ill-posed inverse problems arise in many scientific and engineering applications. Regularization via filtering of the singular value decomposition can be used to compute reasonable solutions. In this talk, a general framework for designing optimal filters is developed, where techniques from stochastical and numerical optimization are utilized. We employ prior knowledge of the application and investigate error metrics, such as p-norms. Numerical examples from image deblurring illustrate better performance compared to well-established filtering methods.

Julianne Chung
University of Maryland
jmchung@cs.umd.edu

Matthias Chung
Department of Mathematics
Texas State University
mc85@txstate.edu

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

MS117

A Fast GPU-based Method for Image Segmentation

The goal of image segmentation is to partition an image into two or more regions that are characterized by, e.g., similar intensity or texture in order to simplify image analysis. Applications range from location of tumors to face recognition. In many of these applications realtime or close to realtime performance is required. GPUs offer high computational performance at low cost and are therefore an interesting architecture especially for data parallel algorithms found in imaging. Therefore, we have implemented a variational shape optimization approach for image segmentation based on a Mumford-Shah functional on GPU using the Open Computing Language (OpenCL). We present performance results on different GPUs and CPUs and additional scaling results for an MPI-parallel version on multi-GPU.

Harald Koestler
University of Erlangen-Nuremberg
harald.koestler@informatik.uni-erlangen.de

MS117

Computational Challenges in SPECT Reconstruction

This talk presents a comprehensive reconstruction framework for the task of Single Photon Emission Computed Tomography (SPECT) reconstruction. Different configurations of this framework lead to different reconstruction problems. Selected problems are discussed in detail and the challenges with respect to the computational complexity

are highlighted. The talk conludes with basic approaches to master the aforementioned challenges and the corresponding reconstruction results are discussed.

Jan Modersitzki
University of Lübeck
Institute of Mathematics and Image Computing
jan.modersitzki@mic.uni-luebeck.de

Sven Barendt
Institute of Mathematics and Image Processing
sven.barendt@mic.uni-luebeck.de

MS117

Open Curve Evolution for Mumford-Shah Segmentation Models

In two dimensions, the Mumford-Shah functional allows for piecewise-smooth minimizers u with the edge set K made of open or closed curves. Prior level set formulations for Mumford-Shah only allow for segmentation of images with closed edge sets. We propose here an efficient level set based algorithm for segmenting images with edges made of open curves or crack tips, adapting Smereka's approach combined with M-S energy minimization. Numerical results on synthetic and real images will be presented.

Rami Mohieddine
UCLA Math department
ramim@math.ucla.edu

Luminita A. Vese
University of California, Los Angeles
Department of Mathematics
lvese@math.ucla.edu

MS118

Adaptive Mesh Refinement for Time-Harmonic Inverse Scattering Problems

Abstract not available at time of publication.

George Biros
Georgia Institute of Technology
biros@gatech.edu

MS118

Shape Optimization for Free Boundary Problems — Analysis and Numerics

In this talk the solution of a Bernoulli type free boundary problem by means of shape optimization is considered. Four different formulations are compared from an analytical and numerical point of view. By analyzing the shape Hessian in case of matching data it is distinguished between well-posed and ill-posed formulations. A nonlinear Ritz-Galerkin method is applied for discretizing the shape optimization problem. In case of well-posedness existence and convergence of the approximate shapes is proven. In combination with a fast boundary element method efficient first and second order shape optimization algorithms are obtained.

Karsten Eppler
Technische Universitaet Dresden
karsten.eppler@mailbox.tu-dresden.de

Helmut Harbrecht
Universitaet Stuttgart
Institut fuer Angewandte Analysis und Numerische Simulation
harbrecht@ians.uni-stuttgart.de

MS118

Optimization of Shell Structure Acoustics

This work analyzes a mathematical model for shell structure acoustics based on the Naghdi shell equations, and thin boundary integral equations, with full coupling at the shell mid-surface. The use of adjoint equations allows the computation of derivatives with respect to large parameter sets in shape optimization problems where the thickness and mid-surface of the shell are computed so as to generate a radiated sound field subject to broad-band design requirements.

Sean Hardesty
Rice University
hardesty@caam.rice.edu

MS118

Boundary Element Methods for Dirichlet Control Problems

We present a boundary integral approach for the solution of Dirichlet boundary control problems with box constraints. The reduced minimization problem results in a variational inequality for which we analyse equivalent boundary integral formulations involving Bi-Laplace boundary integral operators. In addition to the numerical analysis we also comment on the solution of the resulting discrete variational inequality.

Olaf Steinbach
TU Graz, Institute of Computational Mathematics
o.steinbach@tugraz.at

Guenther Of, Thanh Phan Xuan
TU Graz
of@tugraz.at, thanh.phanxuan@tugraz.at

MS119

Sensitivity Analysis Concerning Geometry and Modeling for Aneurysms

In this work examples of cerebral aneurysms will be used to discuss effects of uncertainties in the model reconstruction from medical images and the rheological models to describe blood. Preliminary findings indicate an acute sensitivity to regional variations in the geometry definition. Furthermore, within the aneurysm, a change of rheological models incurs noticeably on the wall shear stress up to comparable levels as the uncertainty in geometry definition.

Alberto Gambaruto
Departamento de Matematica
Instituto Superior Tecnico
agambar@math.ist.utl.pt

Alexandra Moura, Joao Janela, Adelia Sequeira
Department of Mathematics
Instituto Superior Tecnico
alexandra.moura@math.ist.utl.pt, jjanela@iseg.utl.pt, adelia.sequeira@math.ist.utl.pt

MS119

Hemodynamics and Morphology in the Parent Vessels could Predict the Aneurysm Phenomenology

In the cerebral circulation, the morphological features of the feeding arteries could determine a hemodynamics environment which is more or less protective from the formation and development of cerebral aneurysms. This research aims at defining a set of parameters describing both the morphological features and hemodynamics of cerebral arteries, which significantly and jointly correlate with the location of the disease and possibly with its tendency to rupture.

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

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

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

MS119

Stress Analysis of Cerebral Aneurysms

The rupture of intracranial aneurysms is a mechanical phenomenon caused by underlying biological events. The distribution of pressure induced tissue wall tension may provide some insights into this phenomenon and perhaps help assess rupture risk. Patient-specific aneurysms models reconstructed from diagnostic imaging data have been considered. We employed anisotropic finite elastic material models and principal curvature-based material fiber directions. Strain energy stored during deformation is minimized when material fiber directions conform to the principal curvatures.

Madhavan Raghavan
Department of Biomedical Engineering
University of Iowa
ml-raghavan@uiowa.edu

MS119

Functional Data Analysis of Three-dimensional Cerebral Vascular Geometries for the Study of Aneurysms Pathogenesis

We perform statistical analysis of three-dimensional cerebral vascular geometries, obtained from reconstructions of angiographic images. This exploratory study highlights the role of vascular morphology on the pathogenesis of cerebral aneurysms. Advanced techniques are developed for the statistical analysis of these functional data, including methods for multidimensional curve fitting, dimension reduction, registration and classification.

Laura M. Sangalli
MOX Department of Mathematics
Politecnico Milano, Italy
laura.sangalli@polimi.it

Piercesare Secchi, Simone Vantini
MOX - Department of Mathematics
Politecnico di Milano, ITALY
piercesare.secchi@polimi.it, simone.vantini@polimi.it

MS120

Viscoelastic Flows with Microscopic Evaluation of Kramers Rod Forces

Abstract not available at time of publication.

Gregory H. Miller
UC Davis & LBNL
grgmiller@ucdavis.edu

MS120

Time-parallel Continuum-kinetic-molecular Computation of Polymer Rod Models

Abstract not available at time of publication.

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

MS121

Heterogeneous Simulation of Particulate Flows on GPU Clusters

Particulate flows are crucial for various industrial processes, however understanding the underlying transport processes is still ongoing research. For our numerical approach, we couple a rigid body dynamics and a lattice Boltzmann flow solver, which fully resolves the particles. Hence, the simulation of real-world scenarios is very compute intensive. In this talk we present a performance study of our approach on GPU clusters including heterogeneous simulations, and a comparison between CPU and GPU.

Christian Feichtinger
Chiar for System Simulation
University of Erlangen-Nuremberg, Germany
christian.feichtinger@informatik.uni-erlangen.de

MS121

LBM Simulation of Multi-phase Flow — Large-scale Parallel Implementation on the Mole-8.5 GPGPU Supercomputer

Mole-8.5 is the first GPGPU supercomputer (Rpeak of about 1100 Tflops) using NVIDIA Tesla C2050 in the world, designed and established in April 2010 by Institute of Process Engineering (IPE), Chinese Academy of Sciences. It holds the No. 19 spot on the 35th TOP500 list of worldwide supercomputers and debuts at No.8 for the most energy-efficient supercomputers in the Green500 list of July 2010. Mole-8.5 system has already carried out many CFD applications covering such areas as chemical engineering, oil exploitation and recovery, metallurgy, based on some discrete particle method, among which Lattice Boltzmann method is representative due to its natural parallelism and high efficiency. Here we present an application of the direct numerical simulation (DNS) of particle-fluid systems. Particle-fluid systems exist widely in both natural and engineer, such as sandstorm, debris flow, sediment transport, blood flow, fluidized bed reactors, pneumatic conveying, etc. Understanding the physical mechanisms underlying the complex multi-scale behaviors of these systems requires detailed physical information at high resolution below the particle scale, which can not be provided adequately by

experiments so far. We have implemented a coupled numerical method for DNS of particle-fluid both in 2D and 3D. In our proposed coupling scheme, the particle motion is described by the particle method, while the hydrodynamic equations governing fluid flow are solved by LBM. Particle-fluid coupling is realized by an immersed boundary method (IBM). As a result of the fast and efficient simulation of the present scheme and computational capability of Mole-8.5, the scale that DNS can be performed with for particle-fluid system can even be expanded to an engineering scale, that is, meters in magnitude. This implementation makes DNS an attractive alternative to explore the complexity in particle-fluid systems and a potential tool for industrial application.

Limin wang, Xiaowei Wang, Qingang Xiong, Guofeng Zhou
State Key Laboratory of Multiphase Complex Systems, Institute of Process Engineering, Chinese Academy of Science
lmwang@home.ipe.ac.cn, xwwang@home.ipe.ac.cn, qgxiong@home.ipe.ac.cn, gfzhou@home.ipe.ac.cn

Wei Ge
State Key Laboratory of Multiphase Complex Systems
Chinese Academy of Sciences, China
wge@home.ipe.ac.cn

MS121

Highly Interactive Computational Steering of CFD Simulations Utilizing Multiple GPUs

Traditionally, computational fluid dynamics (CFD) is done in a cyclic sequence of independent steps. As CFD simulations are computationally intensive they are usually executed on high performance systems with rather limited user interaction. Not surprisingly, it is a long term wish of scientists and engineers to closely interact with their running simulations. We will show that the convergence of massive parallel computational power of GPUs and a steering environment into a single system significantly improves the usability, application quality and user-friendliness.

Jan Linxweiler
Institute for Computational Modeling in Civil Engineering
Technical University of Braunschweig, Braunschweig, Germany
linxweiler@irmb.tu-bs.de

Manfred Krafczyk
Tech. Univ. of Braunschweig
Germany
kraft@irmb.tu-bs.de

MS121

Subject-sepcific Pulmonary Airflow Simulation by Lattice Boltzmann Method on GPU Cluster

We propose a novel meshing method for multi-GPU LBM computation of subject-specific pulmonary airflows. Since GPU computation cannot bring out its performance under unoptimized memory access, there is little GPU computation of complex geometries including pulmonary airway. In this method, we overcome the issue by decomposing computational domain by small cubes consisting of 4^3-size Cartesian mesh. This method can compute airflow 40 times faster than a Core i7-930 by using Geforce GTX 480.

Takahito Miki
Department of Biomedical Engineering
Tohoku University, Japan
takahito@pfsl.mech.tohoku.ac.jp

Xian Wang
Tokyo Institute of Technology
Japan
wang@sim.gsic.titech.ac.jp

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

Yohsuke Imai, Takuji Ishikawa
Dept. of Bioengineering and Robotics
Tohoku University, Japan
yimai@pfsl.mech.tohoku.ac.jp, ishikawa@pfsl.mech.tohoku.ac.jp

Kei Takase
Dept. of Diagnostic Radiology
Tohoku University, Japan
ktakase@rad.med.tohoku.ac.jp

Takami Yamaguchi
Dept. of Biomedical Engineering
Tohoku University, Japan
takami@pfsl.mech.tohoku.ac.jp

MS122

Establishing a Consortium of CSE Programs (Open Discussion)

Abstract not available at time of publication.

Max Gunzburger
Florida State University
School for Computational Sciences
gunzburg@csit.fsu.edu

Martin Ruess
Technische Universität München
Computation in Engineering
ruess@tum.de

Michael Hanke
KTH Royal Institute of Technology
School of Computer Science and Communication
hanke@nada.kth.se

MS122

Toward a Predictive Model of Tumor Growth

A general continuum theory of mixtures is used as a basis for diffuse-interface models of multiple interacting constituents. The resulting models are of Cahn-Hilliard type, governed by systems of evolution equations in the volume fractions of species. We develop a framework of statistical inverse analysis based on Bayesian methods and employ methods of statistical calibration and validation of various tumor models that lead to the quantification of uncertainties in specific quantities of interest. Virtual medical imaging data are used to inform representative tumor growth models for the investigation of the initial robustness of the

scheme.

Andrea J. Hawkins-Daarud
Institute for Computational Engineering and Sciences
The University of Texas at Austin
andjhawkins@gmail.com

J. Tinsley Oden
The University of Texas at Austin
ICES
oden@ices.utexas.edu

Kris van Der Zee
Institute for Computational Engineering and Sciences
The University of Texas at Austin
vanderzee@ices.utexas.edu

Serge Prudhomme
ICES
The University of Texas at Austin
serge@ices.utexas.edu

MS122

Continuum Mechanical Modeling and Numerical Simulation of Finite Damage in Multi-phasic Materials within the Framework of the Extendend Finite Element Method

In this contribution, fluid-saturated materials are modeled by use of the Theory of Porous Media (TPM) which allows the description of multi-component continua with internal interaction. Multiple numerical models for the computation of tearing mechanisms are discussed in the context of the TPM. For the numerical treatment of strong discontinuities, the Extended Finite Element Method (XFEM) is applied. Computational examples show the numerical realization of the presented damage mechanisms.

Hans-Uwe Rempler
Institute of Applied Mechanics (Civil Engineering)
University Stuttgart
rempler@mechbau.uni-stuttgart.de

MS122

Mathematical Models for Inverse Design Problem of Painless Electrode and Automatic Detection of Brain Metastases

This talk presents two mathematical models; inverse design problem of painless electrode and automatic detection of brain metastases. Inverse design problem for finding an optimal geometry of electrodes aims to deal with edge singularity problem causing painful sensation or skin burn around the perimeter of the electrode. In the second model of automatic brain metastases detection, we develop a robust detection method using a special filtering function which is designed to pick out tumor-like anomalies having a certain size. Theoretical results of the two models are supported by numerical simulations.

Yizhuang Song
Department of Computational Science & Engineering
Yonsei University
yizhuangsong@hotmail.com

Hyeuknam Kwon
Department of Computational Science & Engineering
Yonsei University
3c273-85@hanmail.net

MS123

Generic and Adaptive PDE Solvers on Multicore and Distributed Machines

Implementations of parallel algorithms to solve partial differential equations using adaptively changing meshes are large and complex. At the same time, if implemented intelligently, they can efficiently make use of parallelism both within a single machine using multiple cores as well as across nodes of a cluster computer. We will review our experiences in extending the widely used finite element library deal.II to use both of these paradigms, along with the difficulties one encounters and how they can be resolved.

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

MS123

Petascale Simulations of Reacting Flows

We investigate local and parallel performance issues for direct numerical simulation of turbulent combustion using the spectral element code Nek5000. We analyze single-node multi-core performance across the range of kernels encountered in combustion with detailed chemistry and also discuss a variety of scaling issues, including an AMG-based coarse grid solve, a parallel communication framework, and I/O. We present simulation results for several applications, including reacting and non-reacting turbulent flows and demonstrate scaling for up to 290,000 cores of the IBM BG/P.

Stefan Kerkemeier, Paul F. Fischer
Argonne National Laboratory
stefanke@mcs.anl.gov, fischer@mcs.anl.gov

James Lottes
Oxford University
lottes@maths.ox.ac.uk

MS123

Trading Memory Usage for Improved Communication/compute Overlap in a Peta-scale Reacting Flow Solver

Effective scaling of massively parallel PDE solvers using finite-difference techniques with a physical domain decomposition is limited by the inter-process and inter-node communication necessary to communicate 'halo' data between neighboring tasks. We have reorganized a peta-scale reacting follow solver, S3D, to improve overlap between the communication and computation resulting in improved scaling and increased efficiency when running at the full machine size of jaguar, the XT5 at Oak Ridge National Laboratory (224, 000 MPI tasks) . The reorganization is enabled by increasing the amount of scratch memory used during the computation to break the conventional data dependency limitations. Specifically, derivative operands are computed and halo zone communication begun asynchronously before evaluation of the chemical reaction rates. Traditionally, each operand is computed immediately before evaluation of the derivative to minimize working memory requirements, and reaction rates are evaluated after computing the derivatives. Reacting flow solvers such as S3D tend to be compute and memory-bandwidth bound and are rarely

limited by available memory; with the proposed reorganization we have significant flexibility to vary the amount of memory used in order to minimize work array space while fully overlapping communication with computation.

Ray W. Grout
National Renewable Energy Laboratory
ray.grout@nrel.gov

John M. Levesque
Cray
levesque@cray.com

Ramanan Sankaran
National Centre for Computational Sciences
Oakridge National Laboratories
sankaranr@ornl.gov

Steven Hammond
National Renewable Energy Laboratory
steven.hammond@nrel.gov

Jacqueline Chen
Sandia National Laboratories
jhchen@sandia.gov

MS123

MAESTRO and CASTRO - Petascale AMR Codes for Astrophysical Applications

We present a suite of AMR hydrodynamics codes for astrophysical applications developed at the Center for Computational Sciences and Engineering at LBNL. MAESTRO is suitable for low Mach number flows and CASTRO is a general compressible code. Both codes scale to 100k-200k cores using a hybrid MPI/OpenMP approach on the Jaguar XT5 supercomputer at OLCF. We are currently studying a variety of astrophysical phenomena including Type Ia supernovae.

Andy Nonaka, Ann S. Almgren
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov, ASAlmgren@lbl.gov

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

Michael Lijewski
Lawrence Berkeley National Laboratory
mjlijewski@lbl.gov

Michael Zingale
Dept. of Physics and Astronomy
Stony Brook University
mzingale@scotty.ess.sunysb.edu

MS124

Progress in Large-Scale Differential Variational Inequalities for Heterogeneous Materials

Modeling the mesoscale behavior of irradiated materials is an essential aspect of developing a computationally predictive, experimentally validated, multiscale understanding of the thermo-mechanical behavior of nuclear fuel. Phase field models provide a flexible representation of time-dependent heterogeneous materials. We explain how differential variational inequalities (DVIs) naturally arise in phase field models, and we discuss recent work in developing advanced numerical techniques and scalable software for DVIs as applied to large-scale, heterogeneous materials problems.

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

Anter El-Azab
Florida State University
aelazab@fsu.edu

Lois Curfman McInnes, Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
mcinnes@mcs.anl.gov, tmunson@mcs.anl.gov

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

Lei Wang
University of Michigan
olivewl@umich.edu

MS124

Scalable Solvers for the Cahn-Hilliard Equation with Applications in Nuclear Fuel Modeling

The Cahn-Hilliard (CH) equation is often used for the simulation of phase separation. In order to use large time steps, we consider some implicit methods, which are traditionally unconditionally stable. But for the CH equation, standard fully implicit methods are not stable due to the anti-diffusion term in the equation. We propose and test some stabilized domain decomposed fully implicit methods. The parallel scalability of a PETSc-based implementation will be reported for 2D and 3D problems.

Chao Yang
University of Colorado, Boulder
chao.yang@colorado.edu

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

Michael Pernice
Idaho National Laboratory
Michael.Pernice@inl.gov

MS124

A Posteriori Error Analysis for a Cut Cell Finite Volume Method

We consider elliptic problems where the diffusion coefficient changes discontinuously across a smooth curved interface. The discontinuity has a strong impact on the accuracy of numerical methods. We derive goal-oriented a posteriori error estimates for the numerical solutions by describing a systematic approach to discretizing a cut-cell problem that handles complex geometry in the interface in a natural fashion. Our approach reduces to the well-known Ghost

Fluid Method in simple cases.

Donald Estep
Colorado State University
estep@math.colostate.edu or don.estep@gmail.com

Michael Pernice
Idaho National Laboratory
Michael.Pernice@inl.gov

Simon Tavener
Colorado State University
tavener@math.colostate.edu

Haiying Wang
Michigan Technical University
estep@math.colostate.edu

MS124

Phase-Field Modeling for Nuclear Materials Applications

In nuclear materials, the characteristics of irradiation-induced gas bubble structures are dependent on the defect production rates as well as microstructure. Here, we implement a phase-field model capable of capturing multi-component defect diffusion, bubble nucleation and growth, and bubble/grain boundary interactions to investigate these processes throughout time and for varying irradiation conditions. Furthermore, we have utilized this simulation capability to develop models of bubble percolation and the effective thermal transport across heterogeneous microstructures.

Paul Millett
Idaho National Laboratory
Nuclear Fuels and Materials
paul.millett@inl.gov

Anter El-Azab
Florida State University
aelazab@fsu.edu

Michael Tonks
Idaho National Laboratory
michael.tonks@inl.gov

MS125

Two-Phase Flow Simulation on GPU cluster using an AMG Preconditioned Sparse Matrix Solver

Multi-GPU computing is applied to a gas-liquid two-phase flow simulation which is one of the challenging themes in CFD applications. The volume of fluid should be locally conserved and the interface between gas and liquid is captured by level-set function and the WENO scheme. The pressure Poisson equation including coefficients with the high density ratio is solved by a AMG preconditioning and BiCGStab solver. A high performance is achieved on an infiniband-connected GPU cluster.

Takayuki Aoki, Kenta Sugihara
Tokyo Institute of Technology
taoki@gsic.titech.ac.jp, skenta@sim.gsic.titech.ac.jp

MS125

Engineering a Kernel-agnostic Distributed Linear Algebra Library for Multi/many-core in Trilinos

In a shared-memory parallel library, user-authored kernels are required to be parallelized to avoid the introduction of serial bottlenecks. In the context of multiple programming models, this is a significant burden for users, especially those unaccustomed to shared-memory programming. Conversely, vendors/heroes may wish to replace generic library kernels with tuned versions. However, this can be an intrusive and complicated substitution. We discuss Trilinos efforts to decouple distributed objects and their kernels to address these difficulties.

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

MS125

Preparing Multi-physics, Multi-scale Codes for Hybrid HPC

Effective scientific code teams create development environments that allow them to focus on their scientific goals. Here we examine how one such team is preparing for gpu-accelerated multicore computers for the shock hydrodynamics on an unstructured mesh capability within their multi-scale multi-physics application program. The focus is on the organization of data structures and computations so that they execute effectively in this new environment as well as existing homogeneous multi-core environments.

Richard Barrett, Richard R. Drake, Allen C. Robinson
Sandia National Laboratories
rfbarre@sandia.gov, rrdrake@sandia.gov, acrobin@sandia.gov

MS125

A Sequential Programming Framework for Large-Scale GPU-Accelerated Structured Grids

Although structured-grid applications have demonstrated significant speedups with GPUs, the complexities of memory system involved in large-scale GPU clusters makes programming such machines very challenging. We propose a programming framework that allows the user to express structured grid applications in a concise and declarative way by automatically translating the user-written sequential program code to parallel programs in MPI and CUDA. We report performance studies using several benchmarks on the TSUBAME2 supercomputer.

Tatsuo Nomura, Naoya Maruyama, Toshio Endo
Tokyo Institute of Technology
tatsuo@matsulab.is.titech.ac.jp,
naoya@matsulab.is.titech.ac.jp, endo@is.titech.ac.jp

Satoshi Matsuoka
Tokyo Insitute of Technology
matsu@is.titech.ac.jp

MS126

Optimum Experimental Design for Nonlinear Dynamic Processes

Modeling, simulation and optimization are a powerful 'enabling technology' for mastering the challenges of today's science and engineering. An important precondition is the validation and calibration of the models by means of experimental data. We present new algorithms to optimize

the expensive yet often ambiguous experiments to provide significant data, solve the corresponding nonlinear mixed-integer control problems and incorporate model uncertainties. Applications in chemical industry and cell biology demonstrate a truly amazing potential of Nonlinear Optimum Experimental Design.

Hans Georg Bock
IWR, University of Heidelberg
bock@iwr.uni-heidelberg.de

Ekaterina Kostina
Fachbereich Mathematik und Informatik
Philipps-Universitaet Marburg
kostina@mathematik.uni-marburg.de

Stefan Koerkel, Johannes Schloeder
IWR, University of Heidelberg, Germany
stefan.koerkel@iwr.uni-heidelberg.de, schloeder@iwr.uni-heidelberg.de

MS126

Designing Maximally Informative Experiments in Systems Biology

Biological measurements are technically challenging and resource demanding. In this talk, we show how to identify maximally informative experimental protocols for the selection between conflicting nonlinear dynamical hypotheses. The expected information gain is optimized by specifying configurations of crucial subsets of chemical species, sampling schedules and interventions. Our approach is validated with complex mechanistic biochemical networks. Automatically cycling through hypothesis generation, design and testing, this computational framework enables the feedback loop between modeling and experimentation.

AlbertoGiovanni Busetto
ETH Zurich
Department of Computer Science and CC-SPMD
busettoa@inf.ethz.ch

Mikael Sunnaker, Sotiris Dimopoulos, Jörg Stelling
ETH Zurich
mikael.sunnaker@bsse.ethz.ch,
sotiris.dimopoulos@bsse.ethz.ch,
joerg.stelling@bsse.ethz.ch

Joachim M. Buhmann
ETH Zurich
Departement Informatik and CC-SPMD
jbuhmann@inf.ethz.ch

MS126

Approximation Algorithms for Bayesian Experimental Design

Many sensor management and experimental design problems require us to adaptively select observations to obtain the most useful information. These problems involve sequential stochastic optimization under partial observability – a fundamental but notoriously difficult challenge. Fortunately, many observation selection problems have a structural property that makes them easier than general sequential stochastic optimization. In this talk, I will introduce this structural property – a new concept that we call *adaptive submodularity* – which generalizes submodular set functions to adaptive policies. In many respects adaptive submodularity plays the same role for adaptive problems such as sequential experimental design as submodularity plays for nonadaptive problems (such as placing a fixed set of sensors). Specifically, just as many nonadaptive problems with submodular objectives have efficient algorithms with good approximation guarantees, so too do adaptive problems with adaptive submodular objectives. I will illustrate the usefulness of the concept by giving several examples of adaptive submodular objectives arising in diverse applications including sensor selection, viral marketing and active learning. Proving adaptive submodularity for these problems allows us to recover existing results in these applications as special cases and handle natural generalizations. In an application to Bayesian experimental design, we show how greedy optimization of a novel adaptive submodular criterion outperforms standard myopic heuristics such as information gain and value of information.

Andreas Krause, Daniel Golovin
California Institute of Technology
krausea@caltech.edu, dgolovin@caltech.edu

MS126

An Optimal Simultaneous Source for Inverse Problems with Multiple Sources

PDE-constrained parameter estimation problems typically involve multiple right-hand sides. For large-scale problems in 3D, the computational cost and memory requirements to solve these problems increases exponentially. In this work, we introduce techniques from stochastic optimization to reduce the number of PDE solves drastically. This reduction is made possible by a superpostion principle where data from a multitude of sources is replaced by data from a single designed source. We demonstrate the viability of this principle on realistic examples.

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

MS126

Simulation-Based Optimal Bayesian Experimental Design

We propose a Bayesian framework for optimal experimental design with nonlinear simulation-based models. The formulation accounts for uncertainty in model parameters, experimental conditions, and observables, capturing their interdependence using polynomial chaos expansions. The objective function is constructed from information theoretic measures, reflecting expected information gain from potential sequences of experiments. Stochastic approximation algorithms are then used to make optimization feasible in computationally intensive and high-dimensional applications. The setup is demonstrated on a stiff combustion kinetics problem.

Xun Huan, Youssef M. Marzouk
Massachusetts Institute of Technology
xunhuan@mit.edu, ymarz@mit.edu

MS127

Exploration of a Cell-Centered Lagrangian Hydro-

dynamics Method

Abstract not available at time of publication.

Donald Burton
LANL
burton@lanl.gov

Scott Runnels, Theodore Carney, Mikhail Shashkov
Los Alamos National Laboratory
srunnels@lanl.gov, tedc@lanl.gov, shashkov@lanl.gov

MS127

High Order Finite Elements for Lagrangian Hydrodynamics, Part I: General Framework

This talk presents a general Lagrangian framework for discretization of compressible shock hydrodynamics using high order finite elements. The novelty of our approach is in the use of high order polynomial spaces to define both the mapping and the reference basis functions. This leads to improved robustness and symmetry preservation properties, better representation of the mesh curvature that naturally develops with the material motion, significant reduction in mesh imprinting, and high-order convergence for smooth problems.

Veselin Dobrev, Truman Ellis
Lawrence Livermore National Laboratory
dobrev1@llnl.gov, ellis35@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

MS127

High Order Finite Elements for Lagrangian Hydrodynamics, Part II: Numerical Results

This talk focuses on several practical considerations of the high order finite element discretization from Part I, including the generalization of traditional concepts such as corner forces and artificial viscosity. We consider an extensive set of test problems to examine shock wave propagation over unstructured/distorted meshes and symmetry preservation for radial flows in 2D, 3D and axisymmetric geometry. In each case we demonstrate robust performance of the high order FEM implementation in our research code BLAST.

Veselin Dobrev, Truman Ellis
Lawrence Livermore National Laboratory
dobrev1@llnl.gov, ellis35@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

MS127

Shock Hydrodynamics on Tetrahedral Meshes

A new, variational multiscale stabilized formulation Lagrangian shock hydrodynamics is presented. To the author's knowledge, it is the only hydrocode that can accurately compute highly unsteady shock hydrodynamics transients on triangular/tetrahedral meshes in two/three dimensions, as well as the more commonly used quadrilateral/hexahedral meshes. Piecewise linear, equal-order interpolation is adopted for velocities, displacements, and thermodynamic variables. This last aspect makes the current formulation insensitive to the typical pathologies affecting standard hydrocodes (namely hourglass on quadrilateral/hexahedral meshes, and artificial stiffness on triangular/tetrahedral meshes). Numerical tests for the unsteady Euler equations of gas dynamics are presented in two and three dimensions.

Guglielmo Scovazzi
Sandia National Laboratories
gscovaz@sandia.gov

MS128

Fully Implicit Methods for Kinetic Simulation of Plasmas

Kinetic plasma simulation presents formidable challenges due to its strongly nonlinear nature, the many dimensions of phase space (up to 6D), and the time-scale disparity. Current algorithmic approaches rely on inefficient explicit time stepping. Here, we explore fully implicit methods, employing Newton-Krylov solvers, for particle-based kinetic methods. Key to the viability of the approach is a suitable formulation of the nonlinear residual. We present proof-of-principle results in 1D+1D electrostatic PIC to demonstrate the potential of the approach.

Luis Chacon
Oak Ridge National Laboratory
chaconl@ornl.gov

Guangye Chen
ORNL
cheng2@ornl.gov

Daniel Barnes
Coronado Consulting
coronadocon@msn.com

MS128

3D Penning Tap Simulations using Space-time Parallel Particle Solvers on GPGPUs

Abstract not available at time of publication.

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

MS128

High Order Hybrid Semi-Lagrangian Scheme for Kinetic Equations

In my talk, I will discuss the coupling between semi-Lagrangian framework with high order finite volume/difference/element methods. We aim to combine the

advantages of no CFL time step restriction from semi-Lagrangian schemes with the state-of-art high order finite volume/difference WENO scheme and discontinuous Galerkin method. Vlasov simulation results from proposed methodology will be presented to demonstrate the effectiveness and efficiency of the method.

Jing-Mei Qiu
Mathematical and Computer Sciences
Colorado School of Mines
jingqiu@mines.edu

Wei Guo
The Department of Mathematical & Computer Sciences
Colorado School of Mines
wguo@mines.edu

MS128

Constrained Transport Schemes for Ideal MHD on Unstructured Grids

Standard shock-capturing numerical methods fail to give accurate solutions to the equations of magnetohydrodynamics (MHD). The essential reason for this failure is that by ignoring the divergence-free constraint on the magnetic field, these methods can be shown to be entropy unstable. In this talk we will briefly review the entropy stability theorem for ideal MHD. We will then present a class of discontinuous Galerkin constrained transport (DG-CT) methods on unstructured grids that give both stable and accurate results. The proposed CT approach can be viewed as a predictor-corrector method, where an approximate magnetic field is first predicted by a standard DG method, and then corrected through the use of a magnetic potential.

James A. Rossmanith
University of Wisconsin
Department of Mathematics
rossmani@math.wisc.edu

MS129

Investigation of Air Quality Forecasting Biases through Mass Budget and Process Analysis of Science Algorithms in the Model

Predictability of an air quality model is determined by the accuracy of model algorithms representing atmospheric processes as well as the initial and boundary conditions. As the present forecasting models rely on previous forecasting results as initial conditions, errors in meteorological and air quality predictions may accumulate and propagate different places for multiple days. To understand the causal relations of biases in ozone and particulate matter predictions, mass budget and process analysis of individual science algorithms is applied to the Community Multiscale Air Quality (CMAQ) model. By analyzing differences in the process contributions between results of sensitivity simulations, we bound how much of the original forecasting error can be reduced by the improved initial and boundary conditions.

Daewon Byun
NOAA Air Resources Laboratory
daewon.byun@noaa.go

Yunsoo Choi, Richard Saylor, Tianfeng Chai
Air Resources Laboratory, NOAA
yunsoo.choi@noaa.gov, rick.saylor@noaa.gov, tianfeng.chai@noaa.gov

MS129

Sensitivity and Process Analysis of Daytime and Nighttime Atmospheric Chemistry

Process analysis is used to quantify changes in air pollutant mixing ratios attributable to specific processes while sensitivity analysis is used to calculate the response of mixing ratios to small perturbations in parameters or to the initial state. Process and sensitivity analysis are critical for the development of chemical data assimilation methods. An analysis for a number of modeling cases with models employing the Regional Atmospheric Chemistry Mechanism, version 2 (RACM2) will be presented.

Charlene Lawson
Department of Chemistry, Howard University
ipurdiichari@aol.com

William R. Stockwell
Howard University
William.R.Stockwell@gmail.com

Wendy Goliff
CE-CERT University of California, Riverside
wendyg@engr.ucr.edu

John Lewis
NOAA National Severe Storms Laboratory and DRI
john.lewis@dri.edu

MS129

The Mathematics of the Air Quality Modeling System

Deterministic air quality models link meteorology with atmospheric chemistry. The three-dimensional Eulerian approach is the most commonly used modeling method. Eulerian air quality models typically solve tens of thousands of continuity equations. These are partial differential equations that describe the sum total of processes affecting the mixing ratio of a chemical species within each grid box. An overview of the processes, equations that are employed by the models and their solutions will be presented.

William R. Stockwell
Howard University
William.R.Stockwell@gmail.com

Charlene Lawson
Department of Chemistry, Howard University
ipurdiichari@aol.com

John Lewis
NOAA National Severe Storms Laboratory and DRI
john.lewis@dri.edu

MS130

Multiscale Modeling for Stochastic Forest Dynamics

Individual-based models are widely employed to represent and simulate complex systems. Those descriptions however come with a high computational cost and perhaps unnecessary degrees of detail. In this work, we start with a spatially explicit representation of the interacting agents and attempt to explain the systems dynamics at multiple scales by means of successive coarse-graining steps. We apply our technique to a forest model subjected to different

disturbance regimes.

Maud Comboul
University of Southern California
comboul@usc.edu

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

MS130

Length-scale Dependent Active Microrheology of Biological Materials

Abstract not available at time of publication.

M. Gregory Forest
University of North Carolina at Chapel Hill
Dept of Math & Biomedical Engr.
forest@amath.unc.edu

MS130

Correlated Brownian Motions and the Depletion Effect in Colloids

We first review the model of correlated Brownian motions as derived from deterministic dynamics (Kotelenez 1995, 2005). We then describe the qualitative behavior of correlated Brownian motions at short distances. In particular, we obtain that at short distances and for random times two correlated Brownian motions are attracted to each other (K., Leitman and Mann 2008). This attractive behavior is in good agreement with the depletion phenomena, experimentally observed in colloids (Asakura and Oosawa (1954)). [Based on joint work with Marshall Leitman (CWRU) and Jay Mann (CWRU)]

Peter Kotelenez
Case Western Reserve University
pxk4@case.edu

MS130

A Continuum Model for Moving Contact Lines and the Spreading of Liquid Thin Films

PWe will discuss a continuum model for the moving contact line problem derived based on thermodynamics principles and molecular dynamics simulations. Macroscopic thermodynamic argument is used to place constraints on the form of the boundary conditions; molecular dynamics is then used to measure the detailed functional dependence of the boundary conditions. A new continuum model is obtained for the case of partial wetting as well as the case of complete wetting in a unified form.

Weiqing Ren
Courant Institute of Mathamatical Sciences
weiqing@cims.nyu.edu

MS131

A Novel Sparse Preconditioner for High-Order Finite Element Problems

In the 1980s, Orszag [J. Comp. Phys., 37, 70 (1980)] introduced a method for preconditioning high-order finite element problems with a low-order discretization. The low-order discretization is performed on a finer mesh defined by the high-order finite element nodes. We describe an approach for constructing a low-order preconditioner that only needs the high-order element stiffness matrices. New results on 2D and 3D meshes for various diffusion-type systems will be considered.

Travis M. Austin
Tech-X Corporation
austin@txcorp.com

Ben Jamroz
Tech-X
jamroz@txcorp.com

Chetan Jhurani, Srinath Vadlamani
Tech-X Corporation
jhurani@txcorp.com, srinath@txcorp.com

Marian Brezina
U. of Colorado, Boulder
Applied Math Department
marian.brezina@colorado.edu

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

John Ruge
University of Colorado at Boulder
ruge@colorado.edu

MS131

Application of C1 Finite Elements to Two-Fluid Magnetohydrodynamics

The vector potential/stream function representation of the magnetic and velocity fields can be advantageous when computing solutions of the two-fluid magnetohydrodynamic equations, but results in fourth-order derivatives. In applications where shocks are not expected to form, the use of finite-elements having C1-continuity is particularly useful. The M3D-C1 code, which employs reduced-quintic C1 elements on an unstructured mesh, is discussed, and results primarily regarding tokamak fusion applications are presented.

Nathaniel Ferraro
Oak Ridge Institute for Science and Education
General Atomics
nathaniel.ferraro@gmail.com

Stephen Jardin
Princeton Plasma Physics Laboratory
sjardin@pppl.gov

MS131

Smoothers for High Order $H(curl)$ Bases

We consider algebraic multigrid for systems discretized by high order compatible finite elements. In particular, we focus on $H(curl)$ discretizations. Algebraic multigrid methods require efficient smoothers to damp high frequency error. However, the nullspace of purely $H(curl)$ discretizations is large, hence standard smoothers are insufficient for these types of problems. Hybrid smoothers targeting the nullspace have been introduced for lowest order elements, and we show how to generalize this smoother to high order

elements for hierarchical and interpolatory bases.

James Lai, Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
jhlai2@illinois.edu, lukeo@illinois.edu

MS131

High Order Elements in FOSLS; Or Why Linears Suck

We will talk about the use of high order elements in the context of first-order system least-squares (FOSLS) discretization. Numerical results show that high order elements perform better than standard analysis predicts. That is, in general, FOSLS yields locally-sharp a-posteriori error estimates equivalent to H^1-seminorms. The L^2-error is bounded through coercivity constants that depend on domain shape and boundary conditions. It suffers when using linears; however, the accuracy is greatly enhanced when high order elements are used. We will present results that show dramatic improvement in the accuracy per degree of freedom (DOF) for high order elements, and discuss the theoretical analysis.

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

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

John Ruge
University of Colorado
jruge@colorado.edu

Lei Tang
Department of Applied Mathematics
University of Colorado-Boulder
l6tang@gmail.com

MS132

Apex to Base Heterogeneous Electrophysiology in a Ventricular Dog Model of the Heart

We created a transmurally and longitudinally heterogeneous, fully-coupled electromechanical model of the dog left ventricle. With pacing the spread of the Action Potential restitution (AP90) was larger in the longitudinal heterogeneous cases compared with the transmural (AP90 longitudinal = 25.9 10 ms; AP90 transmural = 6 2 ms). Longitudinal plus transmural heterogeneity improved stroke volume from 6.7 0.2 ml to 8.7 0.2 ml compared with a transmurally heterogeneous only model.

Jazmin Aguado-Sierra
UCSD
Dept of Bioengineering
jaguadosierra@ucsd.edu

Roy Kerckhoffs
UCSD
Department of Bioengineering
roy@bioeng.ucsd.edu

MS132

Finite Element Analysis of the Mitral Valve with Active Muscle Fibres

The mitral valve located between the left atrium and left ventricle of the heart prevents blood from flowing back into the atrium when the ventricle contracts. Several experimental findings showed that this valve contains contractile cells, meaning that the mitral valve is not only a passive structure. We present a transversely isotropic hyperelastic material model for the leaflets in which we add a contractile element to simulate the mechanical function of the muscle cells present in the leaflets.

Victorien Prot, Bjørn Skallerud
Norwegian University of Science and Technology
mailto:victorien.prot@ntnu.no, bjorn.skallerud@ntnu.no

MS132

Electromechanical Models of the Heart

Over the last decade multiscale models of cardiac electrical behavior that incorporate biophysical models of myocyte subcellular processes as well as continuum models incorporating anatomy and passive mechanics have been developed. Advancements in modeling of these processes have occurred largely independently. However, in the last few years coupled models of cardiac electromechanics have emerged. We present the state-of-the-art in ventricular electromechanical modeling and discuss areas of clinical significance where these models could make important contributions.

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

MS132

Modeling the Infarct Injured Heart, Insights into Mechanical Dysfunction

As direct stress measurements in the beating heart are not possible, computational methods are required to estimate mechanical stress that the heart experiences during the cardiac cycle. We present a fully-coupled electromechanical model which was built to match the in vivo measurements of myocardial strain in an infarct injured ovine left ventricle. This validated model was used to evaluate stress fields as the ventricle beats and to asses function / dysfunction of the injured heart.

Samuel Wall
Simula Research Laboratory
Center for Biomedical Computing
samwall@simula.no

MS133

Lattice Boltzmann Approach to the Numerical Modeling of Electrowetting Phenomena

A 3D lattice Boltzmann approach for the modeling of single-component multi- phase (SCMP) fluid is applied to the modeling of basic transport and splitting processes in digital microfluidics. The liquid free surface is modeled by using a Shan-Chen interaction potential and the physical properties of this surface set to the case of water-air interface. Contact angle at the walls and electrowetting actuation at the electrodes is modeled by considering an

anisotropy in the interaction of fluid particles with the solid obstacles. Different contact angles and different geometries of the device are then used in order to simulate basic processes consisting in displacements of droplets from one electrode to another as well as splitting of single droplets in a three-electrode approach. Comparisons of numerical results to both analytical models and experimental measurements reveal good agreement and point out strength and limitations of both numerical model and electrowetting actuation.

Liviu Clime, Daniel Brassard, Teodor Veres
National Research Council Canada
clime.liviu@imi.cnrc-nrc.gc.ca, daniel.brassard@imi.cnrc-nrc.gc.ca, teodor.veres@imi.cnrc-nrc.gc.ca

MS133

Capillary Flows with Parallel Free-surface Lattice Boltzmann Method

Many interesting flow phenomena in micro scales such as porous media and nano tubes are influenced by capillary effects. Our free surface lattice Boltzmann method (FSLBM) simulates gas-liquid flows, including surface tension effects and has been incorporated in waLBerla, which is a framework combining several lattice Boltzmann applications in one efficient and parallel code. A recently developed extension of the FSLBM introduces capillary forces to enable the simulation of capillary effects for gas-liquid systems.

Stefan Donath
Computer Science 10 - System Simulation
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
stefan.donath@informatik.uni-erlangen.de

Ulrich Rüde
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
ulrich.ruede@informatik.uni-erlangen.de

MS133

Drop-on-demand Drop Impact on Patterned Surfaces

Lattice Boltzmann simulations of micron-scale water drop impact on dry patterned surfaces are carried out over a wide range of impact velocities and equilibrium contact angles. Minimization of the total free energy subject to the polynomial wall free energy determines the contact angle and the density profile at solid surfaces. Time evolution of dimensionless kinetic energy and surface energy of an impacting drop at various contact angles will be discussed.

Taehun Lee
Department of Mechanical Engineering
The City College of New York
thlee@ccny.cuny.edu

MS133

Lattice Boltzmann Simulation of Isothermal Vaporization in a Porous Material

We apply a lattice Boltzmann multiphase approach to model the drying process at the pore scale and compare the simulated liquid distribution in a soil sample with laboratory measurements to study effects of soil and transport properties on the evaporation behavior. The parallel efficiency and simplicity of lattice Boltzmann method pave the way for such a computational challenge. Quantitative comparisons with measured water retention curves and transient drying front depth will be presented.

Ying Wang
Technische Universität Braunschweig
Germany
yingwang@irmb.tu-bs.de

Manfred Krafczyk
Tech. Univ. of Braunschweig
Germany
kraft@irmb.tu-bs.de

Benjamin Ahrenholz
Technische Universität Braunschweig
Germany
ahrenholz@irmb.tu-bs.de

Peter Lehmann, Dani Or
Eidgenössische Technische Hochschule Zürich
Switzerland
peter.lehmann@env.ethz.ch, dani.or@env.ethz.ch

MS134

Uncertainty Assessment for Dynamical Systems, with Applications to Chemical Plants

We discuss the issue of material tracking under uncertainty in a dynamical system as a model of a chemical process. As opposed to typical state estimation approaches, we are interested in using current information to reduce the uncertainty of past states. Due to the high dimensionality and complexity of the distribution, sequential Monte Carlo (SMC) methods are used to estimate the material amounts in the production phases. We discuss sufficient conditions on observed quantities that would result in slowly increasing variance in time and thus in a tractable problem. We prove that SMC methods converge when the proposal density at each time step is normal if the distribution is defined on compact space. We demonstrate our findings using SISTOS, a parallel code that was developed for this purpose.

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

Xiaoyan Zeng
Argonne National Laboratory
zeng@mcs.anl.gov

MS134

A Scalable Algorithm for Solutions of Large-scale Statistical Inversions

We present an attempts to reduce the *cost* of evaluating the probability density by employing a Bayesian response surface method based on a Gaussian process model. Our method incorporates gradient and Hessian information into the prior Gaussian process, which is expected to better explore variations of the posterior and reduce the number of sampling points. We apply the proposed method to various synthesized problems and inverse shape scattering problems governed by Maxwell's equations in the

time domain. The shape gradient and shape Hessian-vector products are computed via adjoint equations and continuous shape derivatives. Scatterer shape is parametrized via Fourier bases, yielding increasingly higher dimensional inverse problems.

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

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

MS134

A Discretized Law of Total Probability Approach to Inverse Problems

Given a set of physical measurements with a known error model, the goal of model calibration is to find the input parameters of a simulation that bring the output into agreement with measurements. We examine a procedure based on the law of total probability, which formulates the inverse problem as a standard convex optimization problem. In contrast to MCMC-based inversion, this procedure naturally separates the measurement error from the model error.

Paul Constantine
Sandia National Labs
pconsta@sandia.gov

Christopher Maes, Peter Glynn
Stanford University
cmaes@stanford.edu, glynn@stanford.edu

MS134

Sparse Approximation of SPDEs

We propose a method for the approximation of solutions of PDEs with stochastic coefficients based on the direct, i.e., non-adapted, sampling of solutions. This sampling can be done by using any legacy code for the deterministic problem as a black box. The method converges in probability (with probabilistic error bounds) as a consequence of sparsity and a concentration of measure phenomenon on the empirical correlation between samples. We show that the method is well suited for truly high-dimensional problems (with slow decay in the spectrum).

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
doostan@colorado.edu

Houman Owhadi
Applied Mathematics
Caltech
owhadi@caltech.edu

MS135

Large Sparse Matrix Problems in Ab-initio Nuclear Physics

A microscopic theory for the structure of light nuclei poses formidable challenges for high-performance computing. The ab-initio no-core full configuration method frames this quantum many-body problem as a large sparse matrix eigenvalue problem which is solved for the lowest eigenvalues (binding energies) and their associated eigenvectors. The eigenvectors are employed to evaluate experimental quantities. We discuss different strategies for distributing and solving this large sparse matrix on current multicore computer architectures.

Pieter Maris
Iowa State University
pmaris@iastate.edu

MS135

An Interpolatory Parallel Method for Large-scale Nonlinear Eigenvalue Problems

We present a course-grained parallel method for computing interior eigenvalues and their corresponding eigenvectors of large-scale nonlinear eigenvalue problems. Our method consists of a rational interpolation that is constructed from solutions of independent systems of linear equations with different shift points. This enables us to avoid communication between computing nodes assigned for each linear solver, and provides a good scalability on many-core processors with massively parallel computing resources.

Tetsuya Sakurai, Hiroto Tadano
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp, tadano@cs.tsukuba.ac.jp

Tsutomu Ikegami
Information Technology Research Institute
AIST
t-ikegami@aist.go.jp

Ichitaro Yamazaki
Computational Research Division
Lawrence Berkeley National Laboratory
iyamazaki@lbl.gov

MS135

Sparse Matrix Techniques in a Parallel Hybrid Solver for Large-scale Linear Systems

A parallel hybrid linear solver based on the Schur complement method has a great potential to utilize thousands of processors for solving large-scale linear systems that are becoming increasingly difficult to solve using standard techniques. In this talk, we outline the algorithm implemented in our parallel hybrid solver, particularly focusing on the sparse matrix techniques for achieving high-performance. We also present numerical results of solving highly-indefinite linear systems from real applications.

Ichitaro Yamazaki, Xiaoye Sherry Li
Computational Research Division
Lawrence Berkeley National Laboratory
iyamazaki@lbl.gov, xsli@lbl.gov

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

MS135

Sparse Matrix Techniques in X-ray Diffractive Imaging

An emerging technique in X-ray diffractive imaging is Pty-

chography. In a Ptychography experiment, a large number of small and overlapping diffraction patterns of an unknown object are collected. These diffraction images are used in an iterative procedure to recover both the phase and amplitude of the object. An efficient implementation of this algorithm relies on using sparse matrix vector multiplications to update the approximate reconstruction. We will describe the use of such a technique on high performance computers in this talk.

Chao Yang, Stefano Marchesini
Lawrence Berkeley National Lab
CYang@lbl.gov, smarchesini@lbl.gov

Filipe Maia
NERSC
Lawrence Berkeley National Laboratory
frmaia@lbl.gov

Andre Schirotzek
Advanced Light Source
Lawrence Berkeley National Laboratory
aschirotzek@lbl.gov

MS136

Inexact Newton Methods for Large-Scale Nonlinear Optimization

InexactNewtonmethods play a fundamental role in the solution of large-scale unconstrained optimization problems and nonlinear equations. The key advantage of these approaches is that they can be made to emulate the properties ofNewton's method while allowing flexibility in the computational cost per iteration. Due to the multi-objective nature of *constrained* optimization problems, however, that require an algorithm to find both a feasible and optimal point, it has not been known how to successfully apply aninexact Newtonmethod within a globally convergent framework. In this talk, we present a new methodology for applying inexactness to the most fundamental iteration in constrained optimization: a line-search primal-dualNewtonalgorithm. We illustrate that the choice of merit function is crucial for ensuring global convergence and discuss novel techniques for handling non-convexity, ill-conditioning, and the presence of inequality constraints in such an environment. Numerical results are presented for PDE-constrained optimization problems.

Frank E. Curtis
Industrial and Systems Engineering
Lehigh University
frank.e.curtis@gmail.com

MS136

A Stochastic Newton Method for Large-scale Statistical Inverse Problems with Application to Geophysical Inverse Problems

We present a Langevin-accelerated MCMC method for sampling high-dimensional, expensive-to-evaluate probability densities that characterize the solution to PDE-based statistical inverse problems. The method builds on previous work in Metropolized Langevin dynamics, which uses gradient information to guide the sampling in useful directions, improving acceptance probabilities and convergence rates. We extend the Langevin idea to exploit local Hessian information, leading to what is effectively a stochastic version of Newton's method. We apply the method to the Bayesian solution of a seismic inverse problem, for which we observe several orders of magnitude faster convergence over a reference blackbox MCMC method.

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

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, Lucas Wilcox
University of Texas at Austin
georgst@ices.utexas.edu, lucasw@ices.utexas.edu

MS136

Proper Orthogonal Decomposition (POD) Basis Design by Means of Non-linear Inversion

A framework for the design of Proper Orthogonal Decomposition (POD) bases in introduced in this study. While basis selection for PDE-based problems has been extensively studied, limited research has addressed the problem form non-linear inversion standpoint. In this study, we formulate the basis design problem in terms of stochastic optimization, based on the Stochastic Sampling Average (SSA) approach. We demonstrate the utility and effectiveness of the proposed approach over a set of PDE-based problems.

Lior Horesh
Business Analytics and Mathematical Sciences
IBM TJ Watson Research Center
lhoresh@us.ibm.com

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

Marta D'Elia
MathCS, Emory University
mdelia2@mathcs.emory.edu

MS136

A Control-theoretic Approach to Inference for Prediction

Design of systems with uncertain parameters is ubiquitous in engineering. Many approaches proceed sequentially with experimentation, inference, and design. These methods may waste resources by conducting experiments not significantly contributing to the final control objective or constraints. We propose a goal-oriented, control-theoretic approach to inference for prediction. The algorithm is demonstrated on a 2-D contamination mitigation problem.

Chad E. Lieberman
MIT
celieber@mit.edu

Karen Willcox
Massachusetts Institute of Technology

kwillcox@mit.edu

MS137

Career Preparation of Undergraduate and Graduate Students through an Interdisciplinary Consulting Approach

It is vital for Mathematicians and Statisticians to interact effectively with colleagues from other fields. We have implemented a training approach through consulting with clients from application areas to provide our students with demonstrated experience of their skills in this area. On the graduate level, we use a consulting center, and on the undergraduate level, we have created an REU Site using the same philosophy. This talk will show both the ideas and their implementation in more detail.

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

Nagaraj Neerchal
Department of Mathematics and Statistics
University of Maryland, Baltimore County
nagaraj@umbc.edu

MS137

Undergraduate Foundations of Applied Mathematics

In this talk we explore some of the fundamental concepts in applied mathematics and demonstrate how they can be taught to undergraduates. What should be the foundational core of applied mathematics at that level? We are motivated by the desire to create an undergraduate degree program an applied and computational mathematics that will provide students with a strong background in computation, without sacrificing the mathematical rigor of a traditional mathematics degree. Audience participation is encouraged.

Jeffrey Humpherys
Brigham Young University
jeffh@math.byu.edu

MS137

Undergraduate Research Experiences in Computational Mathematics

Abstract not available at time of publication.

Eric J. Kostelich
Arizona State University
Department of Mathematics
kostelich@asu.edu

MS137

Undergraduate CSE Programs in the U.S.

This talk will discuss the current state of Computaitonal Science programs in the US (and abroad) with reference to the updated SIAM Undergraduate CSE report. The new SIAM Undergraduate Research Online publication will also be presented as a suitable outlet for publishing quality undergraduate research in applied and computational mathematics.

Peter R. Turner
Clarkson University
School of Arts & Sciences
pturner@clarkson.edu

MS138

Compressive Sensing: A New Approach to Image Acquisition

Imaging systems are under increasing pressure to accommodate larger data sets. The foundation of today's digital data acquisition systems is the Shannon/Nyquist sampling theorem. However, the physical limitations and inherently high Nyquist rates of current systems impose a performance barrier. We will overview our work on compressive sensing, which digitizes analog signals via measurements using more general, even random, test functions. The implications of compressive sensing are promising for medical image reconstruction.

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

MS138

Hardware Acceleration of Medical Imaging Applications

Medical imaging applications often involve a pipeline of algorithms that include image restoration, registration, and analysis (e.g., segmentation, feature extraction). However, most algorithms are constrained to the research environment due to a lack of computational power. This talk discusses the challenges of deploying these algorithms in clinical practice and focuses on how domain-specific computing accelerates pipeline algorithms using a customizable heterogeneous platform that is tuned to the specific needs of each image processing algorithm.

William Hsu, Gene Auyeung, Sugeun Chae, Yi Zou, Alex Bui
University of California, Los Angeles
willhsu@mii.ucla.edu, gauyeung@ucla.edu, sugeun@ucla.edu, zouyi@cs.ucla.edu, buia@mii.ucla.edu

MS138

Numerical Challenges in Constrained Image Registration

This talk introduces image registration and explains the necessity of application conform constraints such as volume and mass preservation. The talk presents a variational framework for constrained image registration and its numerical implementation. The presented implementation is based on a sequence of coupled discretizations, where for each discretization a finite dimensional constrained optimization problem is to be solved. Central challenges are the size of the problems which can reach billions of unknowns and constraints and the nature of the non-linear constraints.

Jan Modersitzki
University of Lübeck
Institute of Mathematics and Image Computing
jan.modersitzki@mic.uni-luebeck.de

MS138

Modeling Image Processing Algorithms Using the Concurrent Collections Coordination Language

Concurrent Collections (CnC) is a graphical parallel coordination language in which a node corresponds to a step, data or item collection, and a directed edge corresponds to a put, get, or step-creation operation. A CnC program execution is guaranteed to be deterministic and data-race-free. We summarize our experience with modeling image-processing algorithms using CnC, and discuss how CnC can be used to naturally uncover multiple levels of parallelism intrinsic to an algorithm or application.

Vivek Sarkar
Rice University
vsarkar@rice.edu

MS138

Expectation Maximization and Total Variation Based Model for Computed Tomography Reconstruction from Undersampled Data

Computerized tomography plays an important role in medical imaging. However, the higher dose of radiation of CT will result in increasing of radiation exposure in the population. We propose a method combining expectation maximization and total variation regularization, called EM+TV, to preserve the quality of the image while reducing the exposure. This method can reconstruct the image using undersampled views while providing good result, then reduce the overall dose of radiation.

Ming Yan, Luminita A. Vese
University of California, Los Angeles
Department of Mathematics
yanm@math.ucla.edu, lvese@math.ucla.edu

MS139

A High Order Cell Centred Lagrangian Godunov Scheme for Shock Hydrodynamics

A new cell centred Lagrangian Godunov scheme for shock hydrodynamics is proposed. The new method uses a transient dual grid to define the motion of the vertices in a way that is consistent with the geometric conservation law. The extension of the scheme to second order accuracy in space is considered and an initial assessment of the performance of the method is made by comparison with results obtained with a compatible staggered grid scheme.

Andrew J. Barlow
Atomic Weapons Esatblishment, United Kingdom
Andy.Barlow@awe.co.uk

MS139

Staggered Lagrangian Discretization based on Cell-centered Riemann Solver - A Bridge from Staggered to Cell-centered Lagrangian Schemes

In this presentation we provide a new formalism to bridge well-known staggered Lagrangian and newly developed cell-centered Lagrangian schemes. We will present this new approach that leads to formally second-order accurate scheme in space and time. A 3D implementation of this scheme will provide numerical results on classical test cases of hydrodynamics.

Raphael Loubere
University of Toulouse, France
raphael.loubere@math.univ-toulouse.fr

Pierre-Henri Maire
UMR CELIA, Université Bordeaux I
maire@celia.u-bordeaux1.fr

Pavel Vachal
Czech Technical University in Prague
vachal@galileo.fjfi.cvut.cz

MS139

A General Formalism to Derive Cell-centered Schemes for Two-dimensional Lagrangian Hydrodynamics on Unstructured Grids

The aim of this work is to develop a general formalism to derive cell-centered schemes for 2D Lagrangian hydrodynamics on unstructured grids that meet the compatibility GCL requirement. The high-order extension of this general cell-centered scheme is constructed using the two-dimensional extension of the Generalized Riemann Problem methodology in its acoustic version. Various numerical results on representative compressible fluid flows are presented to demonstrate the accuracy and the robustness of these schemes.

Pierre-Henri Maire
UMR CELIA, Université Bordeaux I
maire@celia.u-bordeaux1.fr

MS139

Issues in Designing a 2D Cylindrically Symmetric Conservative Lagrange Hydro Scheme

Work by Burton, Caramana, et al, developed practical energy conserving Lagrangian hydrodynamics schemes in the late 1990's. The Wilkins area-weighted discretization, which preserves spherical symmetry in cylindrical geometry, can be made energy conserving, but the change requires Lagrangian corner masses, which in turn raise some issues about the definitions of corner volumes. We present one such scheme, some test results, and discuss the impact of our definitions on other packages in a multi-physics code.

Douglas S. Miller
Lawrence Livermore Nat. Lab
dougmiller@llnl.gov

MS140

Energy Stable Space-Time Discontinuous Galerkin Approximations of the 2-Fluid Plasma Equations

Energy stable variants of the space-time discontinuous Galerkin (DG) finite element method are developed that approximate the ideal two-fluid plasma equations. Using standard symmetrization techniques, the two-fluid plasma equations are symmeterized via convex entropy function and the introduction of entropy variables. Anaysis results for the DG formulation assuming general unstructured meshes in space and arbitrary order polynomial approximation include

- a cell entropy bound for the semi-discrete formulation,
- a global two-sided entropy bound and L_2 stability for the space-time formulation,
- a modification of the DG method with provable stability when entropy variables are not used.

Numerical results of the 2-fluid system including GEM magnetic reconnection are presented verifying the analysis and assessing properties of the formulation.

Timothy J. Barth
NASA Ames Research Center
Timothy.J.Barth@nasa.gov

James A. Rossmanith
University of Wisconsin
Department of Mathematics
rossmani@math.wisc.edu

MS140

Discontinuous Galerkin Method Applied to the Multi-fluid Plasma Model

The multi-fluid plasma model only assumes local thermodynamic equilibrium within each fluid. Physical parameters indicate the importance of the two-fluid effects. The algorithm implements a discontinuous Galerkin method with approximate Riemann fluxes for the fluids and electromagnetic fields. The two-fluid plasma model has time scales on the order of the electron and ion cyclotron frequencies, the electron and ion plasma frequencies, the electron and ion sound speeds, and the speed of light.

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

MS140

Simulation Challenges in Industrial Plasma Applications

Abstract not available at time of publication.

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

MS140

Robust and Efficient Schemes for Highly Compressible Magnetohydrodynamics with Applications to Interstellar Turbulence and Stratified Magneto-atmospheres

Ideal magnetohydrodynamics is a widely used fluid model for astrophysical plasma. The dynamics often contain shocks, large density fluctuations and vast scale ranges, making numerically stable simulations challenging. We present stable finite volume schemes based on techniques of entropy stability, positivity and well-balancing. Simulation examples include chromospheric waves and interstellar turbulence.

Knut Waagan
Center for Scientific Computation and Mathematical Modeling
University of Maryland
kwaagan@cscamm.umd.edu

MS141

Analysis of Nudging as a Method for Dynamic Data Assimilation

While nudging as a method for dynamic data assimilation has been around since early 1970's, much of the analysis related to the effectiveness of nudging is largely based on empirical testing. In this talk, for the first time, using a simple linear scalar model, we prove that the nudged forecast is closer to the data than the non-nudged model. In particular, we derive a closed expression for the optimal value for the nudging parameter.

S. Lakshmivarahan
University of Oklahoma
varahan@ou.edu

John Lewis
NOAA National Severe Storms Laboratory and DRI
john.lewis@dri.edu

MS141

Forward Sensitivity Approach to Dynamic Data Assimilation

Abstract not available at time of publication.

John Lewis
NOAA National Severe Storms Laboratory and DRI
john.lewis@dri.edu

Ramesh Vellore
Desert Research Institut
ramesh.vellore@dri.edu

MS141

Inverse Modeling of Black Carbon and NOx Emissions over North America Using CMAQ: Influence of Data Sets and Types Used

Abstract not available at time of publication.

Armistead Russell
Georgia Institute of Technology
ted.russell@ce.gatech.edu

MS141

New Computational Tools for Chemical Data Assimilation

Abstract not available at time of publication.

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

MS142

Variance Reduction as a Multiscale Tool

Fluctuations are one of the defining characteristics of molecular simulation methods. Although sometimes responsible for a number of interesting physics, in many cases of practical interest they amount to an inconvenience, due to the considerable computational cost required for their elimination through statistical sampling. In this talk we present variance reduction methods for drastically reducing the statistical uncertainty associated with Monte Carlo methods for solving the Boltzmann transport equation. The variance reduction, achieved by simulating the devia-

tion from equilibrium, provides a speedup which increases quadratically as the deviation from equilibrium goes to zero, thus enabling the simulation of arbitrarily small deviations from equilibrium. In addition to enabling classical hybrid methods based on domain decomposition by facilitating the coupling process, these control variate methods can themselves be thought of as multiscale methods in which the continuum/molecular decomposition occurs throughout the computational domain, namely by having part of the solution described by a deterministic component (equilibrium) and the remainder described by a molecular simulation.

Nicolas Hadjiconstantinou
Massachussets Institute of Technology
ngh@mit.edu

MS142

Accelerated Kinetic Monte Carlo Methods: Hierarchical Parallel Algorithms and Coarse-graining

In this talk we present two intimately related approaches in speeding-up molecular simulations via Monte Carlo simulations. First, we discuss coarse-graining algorithms for systems with complex, and often competing particle interactions, both in the equilibrium and non-equilibrium settings, which rely on multilevel sampling and communication. Second, we address mathematical, numerical and algorithmic issues arising in the parallelization of spatially distributed Kinetic Monte Carlo simulations, by developing a new hierarchical operator splitting of the underlying high-dimensional generator, as a means of decomposing efficiently and systematically the computational load and communication between multiple processors. The common theme in both methods is the desire to identify and decompose the particle system in components that communicate minimally and thus local information can be either described by suitable coarse-variables (coarse-graining), or computed locally on a individual processors within a parallel architecture.

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

MS142

Multiscale Computation of Rolie-Poly Flow

Abstract not available at time of publication.

Paula A. Vasquez
Department of Mathematics
University of North Carolina at Chapel Hill
pvasquez@email.unc.edu

MS143

A Full Navier-Stokes Solver on Irregular Domains coupled with a Poisson-Boltzmann Solver with Neumann or Robin Boundary Conditions on Non-Graded Adaptive Grid

Second-order solver for the full Navier-Stokes equations coupled with the Poisson-Boltzmann equation on irregular domains is introduced. The Poisson-Boltzmann solver is on Quadtee/Octree grids. The fluid solver is an improved projection method where Neumann boundary conditions for pressure are easily enforced at the irregular domain. The projection matrix is positive definite. The Poisson-Boltzmann matrix is an invertible M matrix, leading to a simple and robust second-order accurate solver.

Asdis Helgadottir
Dept. Mechanical Engineering
UCSB
asdis@engineering.ucsb.edu

MS143

Computational Modeling of Porous Supercapacitors Using Non-graded Adaptive Cartesian Grids

An efficient finite difference discretization of the non-linear Poisson-Boltzmann equation is presented for complex geometries. The level-set method is adopted to represent the object while Octree/Quadtree data stuructures are used to generate adaptive grids, required in the singular limit of thin double layers. Several numerical experiments indicate second order accuracy in L_1 and L_∞ norms. Finally, we use our method to model porous geometries as electrochemical supercapacitors.

Mohammad Mirzadeh
Dept. Mechanical Engineering
UCSB
m.mirzadeh@engineering.ucsb.edu

MS143

A Level Set Approach for Diffusion and Stefan Problems with Robin Boundary Conditions on Non-graded Adaptive Grids

Abstract not available at time of publication.

Joe Papac
Dept. Mechanical Engineering
UCSB
jepapac@gmail.com

MS143

Multigrid Method on Octree Data Structures

We present a new numerical method for solving the equations of linear elasticity on irregular domains in two and three spatial dimensions. We combine a nite volume and a nite dierence approach to derive second-order discretizations. The domain's boundary is represented implicitly by a level set function. Our model is a sharp model in the sense that we solve for the solution inside the domain only, without smearing of the solution near the interface. The approach is well suited for handling shape optimization problems, in which the boundary of the domain is evolving in time with a velocity that depends on the solution of the elasticity problem.

Maxime Theillard
Dept. Mechanical Engineering
UCSB
maxime_theillard@umail.ucsb.edu

MS144

A High-order Finite-volume Method for Hyperbolic Conservation Laws on Locally-refined Grids

We present a fourth-order-accurate finite-volume method for solving time-dependent hyperbolic systems of conservation laws on Cartesian grids with multiple levels of refinement. From coarser to finer levels, we interpolate in

time using the fourth-order Runge–Kutta method, and in space by solving a least-squares problem over a neighborhood of each target cell. The method uses slope limiters, slope flattening, and artificial viscosity. We show results demonstrating fourth-order convergence for some smooth problems in 2D and 3D.

Peter Mccorquodale, Phillip Colella
Lawrence Berkeley National Laboratory
PWMcCorquodale@lbl.gov, PColella@lbl.gov

MS144

Higher Order Finite Volume Methods on Mapped Grids, with Application to Gyrokinetic Plasma Modeling

We describe the development and application of high-order, mapped-grid, finite-volume methods for the solution of gyrokinetic models of magnetically confined fusion plasmas. Phase space advection of plasma species distribution functions and solution of the associated field equations are performed using fourth-order discretizations facilitated by a general formalism. Coordinate mapping is employed to solve the gyrokinetic system on locally rectangular grids aligned with magnetic flux surfaces. Performance of the approach on benchmark problems will be demonstrated.

Milo Dorr
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
dorr1@llnl.gov

Ronald Cohen, Mikhail Dorf
Lawrence Livermore National Laboratory
rcohen@llnl.gov, dorf1@llnl.gov

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

Phillip Colella, Daniel Martin
Lawrence Berkeley National Laboratory
PColella@lbl.gov, dfmartin@lbl.gov

MS144

Another Look at H-box Methods

We consider an explicit finite volume method for the approximation of hyperbolic conservation laws on embedded boundary grids. By using a method of lines approach with a strong stability preserving Runge Kutta method in time, the complexity of our previously introduced h-box method is greatly reduced. Our method is second order accurate and stable for time steps that are appropriate for the regular part of the computational mesh.

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

Marsha Berger
Courant Institute of Mathematical Sciences
New York University
berger@cims.nyu.edu

MS144

Fourth-order Finite Volume Cut Cell Approach for Elliptic and Parabolic PDE's

Abstract not available at time of publication.

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

MS145

Microfluidic Simulation of Dewetting over Solid Plane with Micro-structrured Array

DNA molecules in a solution can be immobilized and stretched into a highly ordered array on a solid surface containing micro-features by molecular combing technique. In this paper, the microfluidic dynamics of dewetting over substrates covered with microwell and micropillar array was simulated based on a front capturing approach with the deforming body-fitted grid system. The simulation results provide insights for explaining the stretching, immobilizing and patterning of DNA nanowire observed in the experiments.

Chun-Hsun Lin
Chung Yuan University
tba@tba.edu

Jingjiao Guan
Florida State University
guan@eng.fsu.edu

Shiu-Wu Chau
National Taiwan University of Science and Technology
chausw@mail.ntust.edu.tw

L. James Lee
Ohio State University
USA
leelj@chbmeng.ohio-state.edu

MS145

Lattice Boltzmann Simulations of Capillary Flows in Patterned Channels

Recent years have seen rapid progress in the technology of fabricating channels at micron length scales. As narrower channels are used, to conserve space and reagents, surface effects will have an increasing influence on the fluid flow. In particular, it may be possible to exploit surface patterning to control the flow within the channels. In this talk, I will discuss how lattice Boltzmann simulations may be used to guide smart designs of such channels. Firstly, I will show how the capillary filling of microchannels is affected by posts or ridges on the sides of the channels. Secondly, I will consider the equilibrium behaviour and dynamics of liquid drops on superhydrophobic surfaces patterned with sawtooth ridges or posts. On these surfaces, liquid drops exhibit complex anisotropic behaviors. Interestingly, this observation allows us to interpret recent experiments describing the motion of water drops on butterfly wings.

Halim Kusumaatmaja
Max Planck Institute of Colloids and Interfaces
Golm, Germany
halim.kusumaatmaja@mpikg.mpg.de

Julia Yeomans
The Rudolf Peierls Centre for Theoretical Physics
Oxford, UK
j.yeomans1@physics.ox.ac.uk

MS145

Lattice Boltzmann Modeling of Microchannel Flow

We present the lattice Boltzmann equation (LBE) with multiple relaxation times to simulate pressure-driven flows in a long microchannel. We use the first-order slip boundary conditions at the walls. The LBE results are validated against those of the compressible Navier-Stokes equations, the information-preservation direct simulation Monte Carlo (IP-DSMC) and DSMC methods. The LBE results agree very well with IP-DSMC and DSMC results in the slip-flow regime, but not in the transition-flow regime in part due to the inadequacy of the slip velocity model. We also compare the LBE simulations of high-Kn flows with molecular dynamics simulations.

Li-Shi Luo
Old Dominion University
lluo@odu.edu

MS145

Monte Carlo Study on Water Cluster Formation and Degradation Mechanisms in PEMFC GDLs

The water household within the gas diffusion layer (GDL) plays a prominent role both for the performance and degradation of PEM fuel cells. To improve understanding of the significant processes and investigate their dependence on the relevant parameters, a Monte Carlo (MC) model working on the μm scale has been developed to simulate the water distribution within a GDL. A description of the model will be given together with results highlighting water cluster analysis capabilities.

Florian Wilhelm
Zentrum für Sonnenenergie- und Wasserstoff-Forschung
Baden-Württemberg, Helmholtzstrae 8, 89081 Ulm,
Germany
florian.wilhelm@zsw-bw.de

Katrin Seidenberger, Tanja Schmitt, Joachim Scholta
Zentrum für Sonnenenergie- und Wasserstoff-Forschung
Germany
katrin.seidenberger@zsw-bw.de, tanja.schmitt@zsw-bw.de,
joachim.scholta@zsw-bw.de

MS146

Uncertainty Quantification for Fluid Mixing

Uncertainty Quantification (UQ) for fluid mixing has to start with lengths scales for observation: macro, meso and micro, each with its own UQ requirements. New results are presented for each. For the micro observables, recent theories argue that convergence in the LES regime should be governed by pdfs (Young measures) which satisfy the Euler equation. VV results for Rayleigh-Taylor mixing, jet breakup and chemical processing are presented to support this point of view.

James G. Glimm
Stony Brook University
Brookhaven National Laboratory
glimm@ams.sunysb.edu or glimm@bnl.gov

MS146

Probabilistic Reduced-order Models for Uncertainty Quantification

This paper is concerned with the development of probabilistic reduced-order models for complex systems in view of applications in uncertainty quantification. We consider systems characterized by a high-dimensional vector of input parameters which are random variables. Given a density for the input uncertainties, the proposed reduced-order models attempt to learn the induced density on the solution vector. We discuss hierarchical mixture models and parallelizable, adaptive inference strategies for learning such probabilistic models.

Phaedon-Stelios Koutsourelakis
School of Civil and Environmental Engineering &
Center for Applied Mathematics, Cornell University
pk285@cornell.edu

MS146

Generalised Polynomial Chaos for Differential Algebraic Equations with Random Parameters

We consider dynamical systems consisting of differential algebraic equations, where some physical parameters include uncertainties. Thus the parameters are replaced by random variables. We resolve the stochastic model by the generalised polynomial chaos. A stochastic Galerkin method yields a larger coupled system of differential algebraic equations, which is satisfied by an approximation of the unknown random process. We analyse the properties of the larger coupled system. In particular, the index of the system of differential algebraic equations is investigated. Numerical simulations of an illustrative example are presented.

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

MS146

A Framework for Managing the Combined Effect of Aleatoric Uncertainty, Epistemic Uncertainty, and Numerical Error

Abstract not available at time of publication.

Gianluca Iaccarino
Stanford University
jops@stanford.edu

Jeroen Witteveen
Stanford University
Center for Turbulence Research
jasw@stanford.edu

MS147

Title Not Available at Time of Publication

Abstract not available at time of publication.

George Biros
Georgia Institute of Technology
biros@gatech.edu

MS147

An Effective Method for Parameter Estimation

with PDE Constraints with Multiple Right Hand Sides

Many parameter estimation problems involve with a parameter-dependant PDEs with multiple right hand sides. The computational cost and memory requirements of such problems increases linearly with the number of right hand sides. For many applications this is the main bottleneck of the computation. In this paper we show that problems with multiple right hand sides can be reformulated as stochastic optimization problems that are much cheaper to solve. We discuss the solution methodology and use the direct current resistivity as a model problem to show the effectiveness of our approach.

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

MS147

"Map-based" Bayesian Inference for PDE-constrained Inverse Problems

Computational expense and convergence challenges associated with MCMC can hinder the application of Bayesian methods to large-scale inverse problems. We present a new technique that entirely avoids Markov chain-based simulation, by constructing a map under which the posterior becomes the pushforward measure of the prior. Existence and uniqueness of a suitable map is established by casting our algorithm in the context of optimal transport theory. Functional descriptions of the posterior distribution also facilitate subsequent stages of uncertainty propagation, such as posterior prediction and sequential analysis.

Youssef M. Marzouk, Tarek El Moselhy
Massachusetts Institute of Technology
ymarz@mit.edu, tmoselhy@mit.edu

MS147

Compressive Sensing and Underwater Acoustics

We will discuss the application of the theory of compressive sensing to two problems encountered in acoustics: source localization/tracking and communication over an uncertain channel. For the first problem, we show how a randomized group testing can significantly reduce the number of acoustic simulations needed to locate and then track a source. For the second problem, we treat the scenario of communicating over an unknown channel as a blind deconvolution problem, and show that if the (unknown) channel is sparse and a random encoding scheme with a small amount of redundancy is used by the source, then the receiver can discover both the message and the channel response by solving a well-posed optimization program

William Mantzel, Salman Asif
Georgia Institute of Technology
willem@gatech.edu, sasif@gatech.edu

Justin Romberg
School of ECE
Georgia Tech
jrom@ece.gatech.edu

Karim Sabra
Georgia Institute of Technology
karim.sabra@me.gatech.edu

MS148

Reproducible Research, Lessons from the Madagascar Project

The Madagascar open-source project is a community effort, which implements reproducible research practices, as envisioned by Jon Claerbout. More than 100 geophysical papers have been published, together with open software code and data, and are maintained by the community. We have learned that continuous maintenance and repeated testing are necessary for enabling long-term reproducibility. As noted by Claerbout and others, the main beneficiary of the reproducible research discipline is the author.

Sergey Fomel
University of Texas at Austin
sergey.fomel@beg.utexas.edu

MS148

Top 10 Reasons to NOT Share your Code and Why you Should Anyway

The research codes used to produce results (tables, plots, etc.) in publications are rarely made available, limiting the readers' ability to understand the algorithms that are actually implemented. Many objections are typically raised to doing so. Although there are some valid concerns, my view is that there are good counter-arguments or ways to address most of these issues. In this talk I will discuss what may be the top 10 reasons.

Randall J. LeVeque
Applied Mathematics
University of Washington (Seattle)
rjl@uw.edu

MS148

The Challenge of Reproducible Research in the Computer Age

Computing is increasingly central to the practice of mathematical and scientific research. This has provided many new opportunities as well as new challenges. In particular, modern scientific computing has strained the ability of researchers to reproduce their own (as well as their colleagues') work. In this talk, I will outline some of the obstacles to reproducible research as well as some potential solutions and opportunities.

Kenneth J. Millman
University of California, Berkeley
millman@berkeley.edu

MS148

Intellectual Contributions to Digitized Science: Implementing the Scientific Method

Our stock of scientific knowledge is now accumulating in digital form, and the underlying reasoning is often in the code that generated the findings, which is often never published. The case for open data is being made but open code must be recognized as equally important in a principled approach, that of reproducibility of computational results. Issues involved with code and data disclosure are presented, along with possible solutions.

Victoria Stodden
Columbia University

Statistics
vcs@stanford.edu

MS149
Entropy Viscosity for Lagrangian Hydrodynamics

A new technique for approximating nonlinear conservation equations is described (entropy viscosity method). The novelty is that a nonlinear viscosity based on the local size of an entropy production is added to the numerical discretization at hand. The methodology initially introduced in Eulerian coordinates is adapted to Lagrangian Hydrodynamics. The methodology is numerically illustrated on standard benchmark problems.

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

MS149
A Mimetic Tensor Artificial Viscosity Method for Arbitrary Polygonal and Polyhedral Meshes

We construct a new mimetic tensor artificial viscosity on general polygonal and polyhedral meshes. The tensor artificial viscosity is based on discretization of coordinate invariant operators, divergence of a tensor and gradient of a vector. We consider both non-symmetric and symmetric forms of the tensor artificial viscosity. We demonstrate performance of the new viscosity for the Noh implosion, Sedov explosion and Saltzman piston problems on a set of various meshes.

Konstantin Lipnikov, Mikhail Shashkov
Los Alamos National Laboratory
lipnikov@lanl.gov, shashkov@lanl.gov

MS149
A Corner and Dual Mesh ALE Remapping Algorithm for use with the Compatible Energy Lagrangian Discretization

The energy conserving Lagrangian hydrodynamic discretization introduced by Caramana et al. in the late 90's has proven to be quite successful. However, the introduction of subzonal Lagrangian elements complicates ALE remapping algorithms for such methods, both in remapping the subzonal masses as well as properties on the dual mesh. We will discuss new ALE algorithms for computing such remapped properties which obey important goals such as consistency with the primary mesh remapping, conservation, and monotonicity.

J. Michael Owen
Lawrence Livermore Nat. Lab
mikeowen@llnl.gov

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

MS149
Recent Advances in Lagrangian Hydro Methods

I will describe recent advances in Lagrangian hydro methods: the symmetry preservations, shock resolution and energy conservation; novel finite element and cell-centered discretization techniques, new forms of artificial viscosity and methods for axisymmetric problems.

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

MS150
A Massively Parallel FMM Algorithm

We present new scalable algorithms and implementations of the kernel-independent fast multiple method (KIFMM), employing hybrid distributed memory message passing (via MPI) and graphics processing unit (GPU) acceleration to rapidly evaluate two-body non-oscillatory potentials.

George Biros
Georgia Institute of Technology
biros@gatech.edu

MS150
A Domain Decomposition and Load Balancing Algorithm for Large Scale Molecular Dynamics

After reviewing some general aspects of large scale molecular dynamics on modern supercomputers, we will present some of the new computational algorithms implemented in ddcMD to achieve unprecedented performances. In particular, we will focus on a new automatic load balancing algorithm which swaps particles between nearest neighbor processors to continuously balance work. We will also present a heterogeneous decomposition of forces computation to improve parallel scaling for applications with long-range Coulomb interactions.

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

David F. Richards, Jim N. Glosli, Erik W. Draeger
Lawrence Livermore Nat. lab.
richards12@llnl.gov, glosli1@llnl.gov, draeger1@llnl.gov

Milo Dorr
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
dorr1@llnl.gov

Fred H. Streitz
Lawrence Livermore Nat. lab.
streitz1@llnl.gov

MS150
Multilevel Summation Method for Large-scale Simulations of Biomolecules

An ongoing challenge in the simulation of biomolecules is to extend computational techniques to larger systems and longer timescales. However, the use of particle–mesh Ewald (PME) to account for long-range electrostatic interactions has become a restrictive bottleneck for large-scale parallel computation. The multilevel summation method is being investigated as an alternative approach, offering comparable accuracy to PME while permitting better parallel scalability by avoiding the need to calculate 3D FFTs.

David J. Hardy
Theoretical Biophysics Group, Beckman Institute

University of Illinois at Urbana-Champaign
dhardy@ks.uiuc.edu

MS150

A Metascalable Approach to Large Reactive Molecular-dynamics Simulations for Energy and Nanoscience Applications

We are developing a metascalable algorithmic framework that is likely to scale on future many-core clusters. The embedded divide-and-conquer algorithms combine globally informed local solutions into a global solution conforming to correct symmetry. The framework has achieved parallel efficiency over 0.95 on 212,992 IBM BlueGene/L processors for 1.68 trillion electronic degrees-of-freedom quantum molecular dynamics. I will discuss atomistic mechanisms of rapid hydrogen production from water using nanocatalysts and mechanically enhanced reaction kinetics in nanoenergetics-on-a-chip.

Aiichiro Nakano, Rajiv K. Kalia, Priya Vashishta
University of Southern California
anakano@usc.edu, rkalia@usc.edu, priyav@usc.edu

Fuyuki Shimojo
Kumamoto University
shimojo@kumamoto-u.ac.jp

MS151

Multi-level Optimization Algorithms for the Design of Nano-porous Materials

Multi-level optimization problems in nano-porous materials are characterized by the facts that the objective function is different at each level and that it is necessary to get the physics correct at all levels to construct a correct design. We discuss extensions to multi-grid optimization algorithms that allow us to address nano-porous problems that span a wide ranges of scales, include general constraints, and have different objectives at different scales. Some numerical results will be given.

Paul Boggs
Sandia National Lab
ptboggs@sandia.gov

David M. Gay
AMPL Optimization LLC
University of New Mexico
dmg@acm.org

Robert M. Lewis
College of William and Mary
rmlewi@wm.edu

Stephen G. Nash
George Mason University
Systems Engineering & Operations Research Dept.
snash@gmu.edu

MS151

Anomalous Diffusion in Soft Matter Materials and Passive Microrheology

A relatively new field, passive microrheology, has emerged to infer viscoelastic properties of complex media from stochastic fluctuations of passive tracers. It is a generalization of fluctuation-dissipation to complex fluids. The seminal paper was in 1995 by Tom Mason and Dave Weitz; since then the field has evolved extensively. Our theory group focuses on experimental data of our collaborator David Hill in the UNC Cystic Fibrosis Center on human lung mucus and various synthetic simulants. Passive tracer beads from nms to microns are tracked using advanced microscopy, revealing normal and anomalous mean squared displacement statistics. The lecture will survey experimental data, inference methods, models, and simulations. The potential for this methodology to be applied to assess nano-porous materials design is a motivation for this topic in this session.

M. Gregory Forest
University of North Carolina at Chapel Hill
Dept of Math & Biomedical Engr.
forest@amath.unc.edu

Scott McKinley
University of Florida
scott.mckinley@ufl.edu

John Fricks
Dept of Statistics
Penn State University
fricks@stat.psu.edu

Lingxing Yao
Utah University
yaol@math.utah.edu

Christel Hohenegger
University of Utah
Department of Mathematics
choheneg@math.utah.edu

MS151

Pore Networks Optimizing Transport in Permeable Energy Storage Materials

The physics and advantages of gas and electrical energy storage in nanoporous materials are briefly described. Although nanoscale pores enhance storage density, small pores impede transport requiring introduction of wider transport channels to achieve rapid charge and discharge. To this end, the aperture and spacing of hierarchical channel networks are optimized to obtain maximum discharge from a fixed volume of storage material within a specified discharge period.

Robert Nilson, Stewart Griffiths
Sandia National Laboratories
rhnilso@sandia.gov, skgriff@sandia.gov

MS151

Hydrodynamic Theory and Simulations of Nano-Rod and Nano-Platelet Dispersions

We present a new hydrodynamic theory for nano-rod and nano-plate dispersions in the polymer matrix, which are the intermediate phases during processing of most high performance nanocomposite materials. In this theory, we model the surface contact interaction between the nano-inclusion and the polymer matrix explicitly along with the semi-flexibility of the nanorods or nano-platelets. We then explore the phase behavior of the flowing nano-composites in equilibrium and in shear flows. Using a bifurcation study, we explore the various possible phases in simple

shear geometry along with their corresponding rheological consequences.

Qi Wang
University of South Carolina
qwang@math.sc.edu

M. Gregory Forest
University of North Carolina at Chapel Hill
Dept of Math & Biomedical Engr.
forest@amath.unc.edu

Jun Li
Nankai University
nkjunlimail.com@g

MS152

A GPU/Multi-core Accelerated Multigrid Preconditioned Conjugate Gradient Method for Adaptive Mesh Refinement

Abstract not available at time of publication.

Austen C. Duffy
Department of Mathematics
Florida State University
aduffy@math.fsu.edu

MS152

A Robust Symmetric Second-Order Discretization of the Poisson Equation with Mixed Dirichlet-Neumann Boundary Conditions on Arbitrary Geometry

The Poisson equation has countless applications in important engineering problems. Many different approaches have been proposed for solving the Poisson problem subjected to different boundary conditions. However, there is a lack of a straightforward, unified method for dealing with mixed boundary conditions. We propose an approach for imposing mixed Dirichlet and Neumann boundary conditions on irregular domains, which can be encountered for example in the simulation of free surface flows on an arbitrarily shaped topography.

Yen Ting Ng
Dept. Computer Science
UCSB
yenting@cs.ucsb.edu

Chohong Min
kyungHee University
chohong@khu.ac.kr

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

MS152

Computation of Three-dimensional Standing Water Waves

Abstract not available at time of publication.

Chris Rycroft
Department of Mathematics
Lawrence Berkeley National Laboratory
chr@math.berkeley.edu

MS152

A Second Order Virtual Node Algorithm for Elliptic Interface Problems on Irregular Domains

Abstract not available at time of publication.

Joseph Teran
UCLA
jteran@math.ucla.edu

MS153

Adaptive p-refinement Approaches in Stochastic Expansion Methods

This presentation describes automated refinement of stochastic expansions, including uniform and adaptive p-refinement within nonintrusive polynomial chaos expansion and stochastic collocation methods. Adaptive p-refinement techniques include anisotropic tensor and sparse grids, with anisotropy detected from online variance-based decomposition or online or offline spectral coefficient decay rates. These techniques employ general-purpose refinement controls based on response covariance. Alternatively, generalized sparse grids admit refinement controls in statistical quantities of interest, resulting in refinement approaches that are goal-oriented.

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS153

Computational Strategies for UQ: Fault Tolerant Collocation and Input Model Generation

This work focuses on two aspects of scalable UQ. The first part showcases a fault tolerant computational framework for adaptive sparse grid collocation. The second aspect deals with developing a scalable, parallel, computational framework for data-driven model reduction. This framework utilizes various linear and non-linear dimensionality reduction methods to extract the most important features of the given high-dimensional data. This framework allows estimating quantitative information about the data including linearity, compactness and convexity.

Baskar Ganapathysubramanian, Saikiran Samudrala, Jaroslaw Zola, Yu Xie, Srinivas Aluru
Iowa State University
baskarg@iastate.edu, saikiran@iastate.edu, zola@iastate.edu, yuxie@iastate.edu, aluru@iastate.edu

MS153

Optimal Multi-domain Stochastic Collocation

Traditionally stochastic multi-element collocation methods decompose the parametric space into multi-dimensional rectangles. We propose a method based upon high-dimensional edge detection to efficiently decompose the parametric space into arbitrarily shaped regions of high-regularity. A high-degree polynomial interpolant is then constructed in each element using a modified version of De Bore's algorithm. The use of De Bore's 'least interpolant' removes the need for hyper-rectangular elements by allowing construction of polynomial interpolants based upon arbitrarily distributed collocation points.

Dongbin Xiu

Purdue University
dxiu@purdue.edu

John Jakeman
ANU
john.jakeman@anu.edu.au

MS153

Data Analysis for Uncertainty Quantification of Inverse Problems

We present exploratory data analysis methods to assess inversion estimates using examples based on ℓ^2- and ℓ^1-regularization. These methods can be used to reveal the presence of systematic errors such as bias and discretization effects, or to validate assumptions made on the statistical model used in the analysis. The methods include: confidence intervals and bounds for the bias, resampling methods for model validation, and construction of training sets of functions with controlled local regularity.

Luis Tenorio, Christian Lucero
Colorado School of Mines
ltenorio@mines.edu, chlucero@mines.edu

MS154

Fast Schur-complement Approximations for KKT Systems

Abstract not available at time of publication.

George Biros
Georgia Institute of Technology
biros@gatech.edu

MS154

Preconditioning for Optimization Problems with Partial Integrodifferential Equations

We consider an optimization problem with a partial integrodifferential equation. Examples are calibration problems for PIDE models which occur in finance and biology. The discretized versions of these problems lead to dense systems, which however exhibit a particular structure. We use this structure in the design of preconditioners for solving the original system and the resulting necessary optimality conditions. We analyze the preconditioners in a framework that includes aspects of mesh independence. Numerical results are presented for a calibration problem using Levy models in option price modeling.

Ekkehard W. Sachs
University of Trier
Virginia Tech
sachs@uni-trier.de

MS154

All-at-once Solution of Time-dependent PDE-constrained Optimization Problems

One-shot methods for time-dependent PDEs approximate the solution in a single iteration that solves for all time-steps at once. Here, we look at one-shot approaches for the optimal control of time-dependent PDEs and focus on fast solution methods using Krylov solvers and efficient preconditioners. We solve only approximate time-evolutions and compute accurate solutions of the control problem only at convergence of the overall iteration. We show that our approach can give competitive results.

Martin Stoll
Computational Methods in Systems and Control Theory
Max Planck Institute for Dynamics of Complex Technical Syste
martin.stoll80@gmail.com

Andrew J Wathen
Oxford University
Numerical Analysis Group
wathen@maths.ox.ac.uk

MS154

Robust Preconditioners for Distributed Optimal Control of the Stokes Equations

The velocity tracking problem for Stokes flows with distributed control is considered in the steady-state case. The formulation of this optimal control problem involves a regularization parameter, say α, in the cost functional. We will present a preconditioner for the discretized optimality system. If used in Krylov subspace methods like the minimal residual method, we obtain an iterative method for this problem whose convergence rate can be shown to be uniformly bounded in α.

Walter Zulehner
University of Linz, Austria
zulehner@numa.uni-linz.ac.at

MS155

Publishing Reproducible Results with VisTrails

VisTrails is an open-source provenance management and scientific workflow system designed to support scientific discovery. It combines and substantially extends useful features of visualization and scientific workflow systems. Similar to visualization systems, VisTrails makes advanced scientific visualization techniques available to users allowing them to explore and compare different visual representations of their data; and similar to scientific workflow systems, VisTrails enables the composition of workflows that combine specialized libraries, distributed computing infrastructure, and Web services.

Juliana Freire
Scientific Computing and Imaging Institute
School of Computing University of Utah
juliana@cs.utah.edu

Claudio Silva
Scientific Computing and Imaging Institute
University of Utah
csilva@sci.utah.edu

MS155

Reproducible Research: Lessons from the Open Source World

Why are the practices of open source software development often more consistent with our ideas of openness and reproducibility in science than science itself? Today's scientific praxis falls short of our ideals of reproducibility, and these problems are particularly acute in computational domains where they should be less prevalent. I will compare the incentive structure in both cultures to show how this plays an important role, since many of the incentives for

scientific career advancement are directly at odds with the expectations of reproducibility in research. In addition, we can also draw from the everyday toolchain and practices of open source software development to improve our scientific workflows. I will highlight some key tools that can be easily integrated into our everyday research workflows and that provide real benefits without imposing undue burdens.

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

MS155

Reproducible Models and Reliable Simulations: Current Trends in Computational Neuroscience

Computational neuroscientists simulate models of neuronal networks to further our understanding of brain dynamics. Unfortunately, the validity of models of neuronal dynamics and of the simulation software implementing the models is difficult to ascertain, challenging the validity of computational neuroscience. We will describe how the computational neuroscience community is addressing validity through software reviews, best practices, increasing use of established software packages, meta-simulators, systematic testing, and simulator-independent model-specification languages.

Hans E. Plesser
Department of Mathematical Sciences and Technology
Norwegian University of Life Sciences
hans.ekkehard.plesser@umb.no

Sharon M. Crook
Dept of Mathematics and Statistics & School of Life Sciences
Arizona State University
sharon.crook@asu.edu

Andrew P. Davison
Unité de Neuroscience, Information et Complexité
CNRS, Gif sur Yvette, France
andrew.davison@unic.cnrs-gif.fr

MS155

FEMhub, a Free Distribution of Open Source Finite Element Codes

FEMhub (http://femhub.org) is an open source distribution of finite element codes with a unified Python interface. The goal of the project is to reduce heterogeneity in installation and usage of open source finite element codes, facilitate their interoperability and comparisons, and improve reproducibility of results. FEMhub is available for download as desktop application, but all codes are also automatically available in the Online Numerical Methods Laboratory (http://lab.femhub.org).

Pavel Solin, Ondrej Certik
Department of Mathematics and Statistics
University of Nevada, Reno
solin@unr.edu, ondrej@certik.cz

Aayush Poudel, Sameer Regmi
University of Nevada, Reno
aayushpoudel@gmail.com, regmisk@gmail.com

Mateusz Paprock
Technical University of Wroclaw, Wroclaw, Poland
mattpap@gmail.com

PP1

Block Preconditioning for Multi-Physics Problems

The coupling of multi-physics systems is essential for a variety of problems in engineering and computational physiology. Due to their size and complexity, these systems require robust and efficient solution methods. We focus on enhancing monolithic solvers by using a linearly scalable block preconditioner, which is shown to result in mesh-independent convergence of GMRES. This approach enables the extension of the accurate and stable monolithic methods to complex large-scale engineering problems.

Liya Asner
University of Oxford
Computing Laboratory
liya.asner@comlab.ox.ac.uk

David Nordsletten
Massachusetts Institute of Technology
nordslet@mit.edu

David Kay
University of Oxford
Computing Laboratory
david.kay@comlab.ox.ac.uk

PP1

Rabies Epizootic Modeling

I propose an updated model of the rabies epizootic. In this model, a system of partial differential equations represents changes in susceptible, exposed, and infected compartments for both skunk and bat species. The model expands upon one developed by Steinhaus, Kuang, and Gardner (2010). The predictive ability of the model is tested by comparing simulations with the confirmed case data from Texas.

Rebecca Borchering
Arizona State University
School of Mathematical and Statistical Sciences
rborcher@asu.edu

PP1

Bachelor Course Scientific Programming

For the Bachelor course *scientific programming* at Aachen University of Applied Sciences students are also required to train as *mathematical technical software developers* (MATSE) based at Forschungszentrum Jülich or at external companies. The theoretical part of the MATSE qualification (almost equally from mathematics and computer science) is organised as lecture courses at Jülich Supercomputing Centre and is part of the bachelor course. Both the vocational and academic training are designed to last three years.

Oliver Buecker
Research Centre Juelich
Germany
o.buecker@fz-juelich.de

PP1

Spinal Biomechanics Mathematical Model for Lumbar Intervertebral Ligaments

Lumbar Spinal Surgery constitutes a Surgical Specialization that deals with rather complicated operations because of the anatomical and neurophysiological difficulties of the operation field. Both in Lumbar Spinal Surgery and Orthopedic Lumbar Surgery, the vertebrae distraction and/or forced position changes are usual maneuvres carried out during many interventions at the surgical theatre. Vertebrae are tightly joined by a number of strong ligaments to keep the Biomechanical functionality of the Spine. Each ligament has specific mechanical characteristics and its proper functionality. We present a Mathematical-Biomechanical Model of the Lumbar Intervertebral Ligaments, related mainly to the necessary distraction forces that have to be exerted during the operations to separate two adjacent vertebrae (for any specific spinal surgery purpose, e.g., Intervertebral Disck Replacement, Scoliosis, Lumbar Spinal Fusion, Spine Trauma, Surgical Spinal Decompression, etc). The Model has been theoretically fitted to the proper stress and elastic parameters data for each type of Ligament, and simulated computational/numerical results are shown. Applications of fundamental algorithms in Approximation Theory are used to optimize this research. An overview of the Mathematical and Biomechanical calculations, together with the designed software is also presented in this contribution. F Casesnoves MSc (Physics) MD (MPhil Medicine and Surgery).

Francisco Casesnoves
SIAM Activity Group on Geometric Design
NG7 4DW UK
casesnoves@engineering.com

PP1

A Multi-Numerics Scheme for a Multi-Physics Coupling of Free Flow with Porous Media Flow

We present a coupled continuous finite element method with discontinuous Galerkin method scheme for a coupled free flow with porous media flow model. We prove existence and uniqueness of both the weak and numerical solution. Convergence rates are verified for the fully coupled model using grid studies. A two grid method which allows the coupled model to be decoupled into two smaller problems is also presented. The two grid method is tested numerically and compared to solving the fully coupled problem.

Prince Chidyagwai
Temple University
Department of Mathematics
prince.chidyagwai@temple.edu

Beatrice Riviere
Rice University
Houston, Texas, USA
riviere@caam.rice.edu

PP1

Recent Advances in `bout++`

Challenges in turbulent edge simulations in plasma physics include scalability and coupling to other codes for different computational regions. This poster presents recent advances in the solution of time-dependent, nonlinear partial differential equations arising in `BOUT++`, with emphasis on exploration of various preconditioners. We also discuss improvements in interfaces that enable `BOUT++` to function within the `FACETS` framework, which has the goal of providing modeling of the core, edge, and wall regions of fusion devices.

Sean Farley
Mathematics and Computer Science Division
Argonne National Laboratory
sean@mcs.anl.gov

Ben Dudson
Physics Department
University of York
bd512@york.ac.uk

PP1

A Numerical Algorithm for the Solution of a Phase-Field Model of Polycrystalline Alloys

We describe an algorithm for the numerical solution of a phase-field model of microstructure evolution in alloys using physical parameters from thermodynamic and kinetic databases. The system of equations includes a local order parameter, a quaternion representation of local crystal orientation and a species composition parameter. The implicit time integration of the system uses a backward difference formula combined with a preconditioned Newton-Krylov algorithm to solve the nonlinear system at each time step.

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

Milo Dorr
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
dorr1@llnl.gov

Michael E. Wickett, James F. Belak, Patrice E. A. Turchi
Lawrence Livermore Nat. Lab.
wickett@llnl.gov, belak1@llnl.gov, turchi1@llnl.gov

PP1

Scalable Methods for Large-Scale Statistical Inverse Problems, with Applications to Subsurface Flow and Transport

We consider the problem of estimating uncertainty in large-scale nonlinear statistical inverse problems with high-dimensional parameter spaces within the framework of Bayesian inference. The solution is a posterior probability distribution function (pdf) over the desired parameter. Standard Markov chain Monte Carlo approaches are intractable for such high-dimensional problems. We approximate the pdf by constructing a Gaussian process response surface approximation, using the structure of the Hessian matrix and optimization algorithms to achieve scalability.

H. Pearl Flath
Institute for Computational Engineering and Sciences
The University of Texas at Austin
pflath@ices.utexas.edu

Tan Bui, Omar Ghattas
University of Texas at Austin
tanbui@ices.utexas.edu, omar@ices.utexas.edu

PP1

Graph Coloring Software for Sparse Derivative Computation and Beyond

We present a software package, called ColPack, comprising implementations of fast and effective algorithms for a variety of graph coloring and related problems arising in efficient computation of sparse derivative matrices using automatic differentiation. The capabilities of ColPack cover Jacobian as well as Hessian computation using direct as well as indirect methods. Several of the coloring problems supported by ColPack also find important applications in many areas outside derivative computation. In this poster, we will give an overview of the functionalities available in ColPack, highlight its key algorithms, and as an example of an application enabled by ColPack, present results from an optimization problem in chemical engineering.

Assefaw H. Gebremedhin
Purdue University
agebreme@purdue.edu

PP1

Bayesian Markov Chain Monte Carlo Optimization of a Physiologically-Based Pharmacokinetic Model of Nicotine

In an effort to reduce harm caused by tobacco products, the Family Smoking Prevention and Tobacco Act was enacted by Congress and the Food and Drug Administration (FDA) was tasked with regulating tobacco. The FDA is now charged to use the best available science to guide the development and implementation of effective public health strategies to reduce the burden of illness and death caused by tobacco products. Therefore, one significant challenge is to reduce toxicant exposure in tobacco products such as cigarettes. Nicotine is the active drug within tobacco that is delivered during cigarette smoking. Smoking behavior is driven by the need for nicotine and will influence delivery of other tobacco toxicants into the lungs. To better understand nicotine dosimetry a physiologically-based pharmacokinetics (PBPK) model which captures the essential nicotine metabolism and takes into account the variability of kinetic parameters among individuals was developed.

Rudy Gunawan, Chuck Timchalk, Justin Teeguarden
Pacific Northwest National Laboratory
rudy.gunawan@pnl.gov, charles.timchalk@pnl.gov, jt@pnl.gov

PP1

High-Dimensional Adaptive Grids for Time-Dependent PDEs

Accurate solution of time-dependent, high-dimensional PDEs requires massive-scale parallel computing. In particular, the memory requirements for uniform grids with fine enough resolution will become prohibitively large. We discuss an on-going project where a parallel framework for block-adaptive discretization on Cartesian grids of well localized Hamiltonian systems is developed. Numerical simulations from quantum chemistry are presented, where we solve the time-dependent Schrdinger equation using exponential propagators based on the Lanczos algorithm.

Magnus Gustafsson
Uppsala University, Sweden
magnus.gustafsson@it.uu.se

Katharina Kormann, Anna Nissen
Department of Information Technology
Uppsala University
katharina.kormann@it.uu.se, anna.nissen@it.uu.se

Sverker Holmgren
Uppsala University
Department of Scientific Computing
sverker.holmgren@it.uu.se

PP1

H-P Spectral Element Methods for Three Dimensional Elliptic Problems on Non-Smooth Domains

We propose a nonconforming $h-p$ spectral element method to solve three dimensional elliptic boundary value problems on non-smooth domains to exponential accuracy. To overcome the singularities which arise in the neighborhoods of the vertices, vertex-edges and edges we use local systems of coordinates. Away from these neighborhoods standard Cartesian coordinates are used. In each of these neighborhoods we use a geometrical mesh which becomes finer near the corners and edges. We then derive differentiability estimates in these new set of variables and a stability estimate on which our method is based. A parallel preconditioner to solve the normal equations is defined, using the stability estimates. The Sobolev spaces in vertex-edge and edge neighborhoods are anisotropic and become singular at the corners and edges.

Akhlaq Husain
Indian Institute of Technology Kanpur
India
akhlaqiitk@gmail.com

Pravir Dutt
Department of Mathematics,
IIT Kanpur, India
pravir@iitk.ac.in

A S Vasudeva Murthy
TIFR Centre for Applicable Mathematics
TIFR Bangalore, India
vasu@math.tifrbng.res.in

Chandra Shekhar Upadhyay
Indian Institute of Technology Kanpur India
shekhar@iitk.ac.in

PP1

Constructing Sparse Preconditioners for High-Order Finite Element Problems

We present a method for constructing sparse preconditioners for elliptic problems discretized using higher-order polynomial basis functions. The method computes local sparse preconditioners from dense element stiffness matrices. This is done by solving a sparse least-squares problem on an element. The sparse matrices are assembled to be used in an algebraic multigrid method. We implement a fast and memory efficient method for these local problems. The results also show that solving the global problem requires less time and memory.

Chetan Jhurani, Travis M. Austin
Tech-X Corporation
jhurani@txcorp.com, austin@txcorp.com

Ben Jamroz
Tech-X
jamroz@txcorp.com

Srinath Vadlamani, Scott Kruger
Tech-X Corporation
srinath@txcorp.com, kruger@txcorp.com

PP1

Variance Reduction Techniques for Stochastic Differential Equations

Stochastic differential equations are essential to modeling many physical phenomena including option prices in mathematical finance. The stochastic simulation of the quantities of interest is done using repeated realizations of Monte Carlo methods. However these methods are well known for converging slowly with an order of only $\frac{1}{2}$ in terms of the number of realizations. Therefore there is an increasing need for variance reduction in order to improve efficiency of Monte Carlo simulations. Here we use the variance reduction technique known as control variates to improve the computational efficiency. The results presented demonstrate the magnitude of the achieved variance reduction.

Albi Kavo, Todd Caskey, Mandeep Singh
New Jersey Institute of Technology
ak266@njit.edu, tc62@njit.edu, ms247@njit.edu

PP1

Markov Order Estimation and Forecasting

We develop an exact Bayes factor (EBF) classifier to distinguish finite sequences that have Markov order k=0 versus k¿0. Through extensive simulation and testing against a classifier based on the Bayesian Information Criterion (BIC), we find that EBF classifiers can be significantly more accurate, especially when prior knowledge is available. Applying these classifiers to five real-world data sets, we find that EBF-based models have higher out-of-sample predictive accuracy than BIC-based models.

Harish S. Bhat, Nitesh Kumar
University of California, Merced
Applied Mathematics
hbhat@ucmerced.edu, nkumar4@ucmerced.edu

PP1

GPU-Accelerated Preconditioned Iterative Linear Solvers

Sparse linear algebra computations can benefit from the fine-grained massive parallelism of Graphic Processing Units (GPUs). Performance of sparse matrix-vector product kernels can achieve up to 12 GFLOPS in double precision for unstructured matrices. This work focuses on developing high-performance iterative linear solvers accelerated by GPUs. Incomplete Cholesky factorization preconditioned CG method and incomplete LU factorization preconditioned GMRES method are adapted to a GPU environment. Other parallel preconditioners appropriate for GPUs are also explored.

Ruipeng Li
Department of Computer Science & Engineering
University of Minnesota
rli@cs.umn.edu

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

PP1

Solar Design Optimization Process Model

The purpose of this research is to develop a simulation model process algorithm used to optimize solar energy system design. This optimization model will be developed to minimize the cost of a solar energy system in order to be more affordable to marginalized communities. The costing analysis will include a levelized cost of energy (LCE) in determining the maximum power and minimized costs of a system and parameters such as site location and solar radiation.

Scott M. Little
SunWize Technologies
Northcentral University
slittle@sunwize.com

PP1

Modeling the Dynamics of An Insect Heart

The tubular heart of the fly, Drosophila melanogaster, contains excitable cells that cause the heart to contract and pump hemolymph throughout the body. In this work, we model the electrical activity of the heart based on a network of connected excitable cells, where each cell is modeled using a system of ordinary differential equations. Computational studies are performed to investigate the roles of specific ion channels and network connectivity in the generation of the heartbeat.

Miles Manning
Student, School of Mathematical and Statistical Sciences
Arizona State University
miles.manning@asu.edu

April Chiu
student, School of Mathematical and Statistical Sciences
Arizona State University
april.chiu@asu.edu

PP1

A Scalable Preconditioner for Edge Plasma Transport with Neutral Gas Species

Simulating multi-component plasma transport in the edge region of a tokamak is computationally intensive; after discretization, a nonlinear system must be solved at each time step. Including neutral gases in the simulation magnifies the challenge and has limited the scalability of solvers. This work describes numerical experiments with UEDGE, placing emphasis on the development of a component-wise preconditioner that has improved the speed and scalability of the nonlinear solve in the presence of neutral gases.

Michael Mccourt
Argonne National Laboratory
mccomic@mcs.anl.gov

Tom Rognlien
Lawrence Livermore National Laboratory
trognlien@llnl.gov

Lois Curfman McInnes
Argonne National Laboratory
Mathematics and Computer Science Division
mcinnes@mcs.anl.gov

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

PP1

Efficient Implementation of Smoothness-Increasing Accuracy-Conserving Filters for Discontinuous Galerkin Solutions

Smoothness
increasing accuracy conserving post-processors when applied to a discontinuous Galerkin solution can increase the order of accuracy from k+1 to 2k+1. Computationally, this involves solving geometric intersection problems and evaluating several numerical quadratures per mesh element which in turn makes this technique an expensive tool to use. However, as post-processor is a suitable candidate for parallelization, we demonstrate how filtering an entire domain can be achieved in a fast and efficient manner using different parallelization techniques.

Hanieh Mirzaee
Scientific Computing and Imaging Institute
School of Computing, University of Utah
mirzaee@cs.utah.edu

Jennifer K. Ryan
Delft University of Technology
j.k.ryan@tudelft.nl

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

PP1

Local Time-Stepping Adaptive Mesh Refinement Techniques for Discontinuous Galerkin Methods

Adaptive mesh refinement is the process of dynamically adding and deleting mesh elements based on some sort of local error estimator. Time stepping will be achieved using local time-stepping techniques, which allows each mesh element to have its own time step. This differs from the Berger-Oliger method which uses a hierarchy of increasingly refined grid patches, with uniform time steps taken on each grid patch. Various error estimators will be evaluated based on a comparison to the actual known error for both smooth and discontinuous examples. One possible error estimator will involve associating a large gradient with a large error; others will originate from more mathematically rigorous derivations. The LTS-AMR method will be applied to solving the advection equation in one dimension.

Scott A. Moe
University of Wisconsin Madison
smoe@wisc.edu

James A. Rossmanith
University of Wisconsin
Department of Mathematics
rossmani@math.wisc.edu

PP1

Implementation of SpMV Using Indexed Segmented Scan Method for CUDA

We show a new SpMV implementation using Indexed Segmented Scan method for CUDA. This implementation is based on Segmented Scan method, and there are some kind of improvement in pre-calculation and GPU's reduction. In this talk, we show the detail of implementation, performance, and application for numerical library.

Satoshi Ohshima
Univ. of Tokyo
ohshima@cc.u-tokyo.ac.jp

Takao Sakurai
Central Research Laboratory
Hitachi Ltd.
takao.sakurai.ju@hitachi.com

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

Hisayasu Kuroda
Univ. of Tokyo / Ehime U.
kuroda@cs.ehime-u.ac.jp

Ken Naono
HITACHI Ltd.
ken.naono.aw@hitachi.com

Mitsuyoshi Igai
Hitachi ULSI Systems Corporation
mitsuyoshi.igai.bf@hitachi.com

Shoji Itoh
Univ. of Tokyo
itosho@cc.u-tokyo.ac.jp

PP1

Adjoint Methods for Inversion of Rheology Parameters in Ice Sheet Flows

Modeling the dynamics of polar ice sheets is critical for projections of future sea level rise. Yet, there remain large uncertainties in the non-Newtonian constitutive relationships employed within ice sheet models. Here we formulate an inverse problem to infer the rheological parameters that minimize the misfit between observed and modeled surface flow velocities and surface elevation. The inverse problem is solved using an adjoint-based Gauss-Newton method.

Noemi Petra, Hongyu Zhu
Institute for Computational Engineering and Sciences
University of Texas at Austin
noemi@ices.utexas.edu, zhuhongyu919@gmail.com

Georg Stadler, Omar Ghattas
University of Texas at Austin
georgst@ices.utexas.edu, omar@ices.utexas.edu

PP1

Building a Pde Code from Components, Including Shape Optimization and Embedded Uq

A general-purpose PDE code is being written to drive the development of independent yet interoperable software components. This code makes use of dozens of libraries from the Trilinos, Dakota, Cubit, and Sierra projects. To avoid a monolithic framework, we use abstract interfaces at many levels. The heavy use of Automatic Differentiation technology means that any new physics can be instantly be driven by Newton-based solvers, sensitivity analysis, shape optimization, non-intrusive UQ, or embedded UQ.

Andy Salinger
Sandia National Laboratories
agsalin@sandia.gov

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

Steve J. Owen
Sandia National Laboratories
Albuquerque, NM
sjowen@sandia.gov

Matt Staten, Jake Ostien
SNL
mlstate@sandia.gov, jtostie@sandia.gov

PP1

Understanding Memory Effects of Loop Fusion for Linear Algebra Kernels

The performance of scientific applications is often limited by the cost of memory accesses inside linear algebra kernels. We developed a compiler that tunes such kernels for memory efficiency using loop fusion. In this poster, we use statistical methods to determine when fusing vector operations in the presence of matrix operations significantly impacts performance. We present circumstances where the fusion of vector operations can be disregarded when searching for an optimal routine within our compiler.

Erik Silkensen
University of Colorado at Boulder
erik.silkensen@colorado.edu

Ian Karlin
Department of Computer Science
University of Colorado at Boulder
ian.karlin@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

Thomas Nelson
University of Colorado at Boulder
Argonne National Laboratory
thomas.nelson@colorado.edu

Pavel Zelinksy
University of Colorado at Boulder
pavel.zelinksy@colorado.edu

Jeremy Siek
Department of Electrical and Computer Engineering
University of Colorado at Boulder
jeremy.siek@colorado.edu

PP1

Measuring Glioma Tumor Volumes with the Talairach Transformation

Glioblastoma multiforme is a high-grade glioma that significantly lowers patient survival to under two years. Serial magnetic resonance images have been processed to obtain volumetric measurements of brain structures that surround the tumor based on the Talairach Transformation. These measured volumes of the surrounding structures provide estimation of the tumor size. Based on these tumor estimated values, the glioma growth rate can be calculated for forecast models and implemented in MATLAB treatment simulations.

Maple So
Arizona State University
School of Mathematical and Statistical Sciences
mapleso@asu.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 an adjoint-based inexact Newton-CG method.

Jennifer A. Worthen, Georg Stadler
University of Texas at Austin
jworthen@ices.utexas.edu, georgst@ices.utexas.edu

Michael Gurnis
Caltech
gurnis@caltech.edu

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

PP1

A Shape Hessian-Based Analysis of Roughness Effects on Viscous Flow

We study the influence of boundary roughness characteristics on the rate of dissipation in a viscous fluid using first- and second-order shape derivatives. We show that the flat boundary is a stationary point for the dissipation functional, and thus its local behavior is governed by the shape Hessian, whose eigenvectors are shown to be the the Fourier modes. We show numerical examples for laminar

flow and present preliminary results for turbulent flow.

Shan Yang, Georg Stadler, Omar Ghattas, Robert Moser
University of Texas at Austin
syang@ices.utexas.edu, georgst@ices.utexas.edu, omar@ices.utexas.edu, rmoser@ices.utexas.edu

PP1

Comparison of Model Reduction Techniques on High-Fidelity Linear and Nonlinear Electrical, Mechanical, and Biological Systems

An impartial and rigorous comparative study of model order reduction (MOR) techniques is presented. Some of these techniques are classical, while others have been recently developed. Large-scale linear and nonlinear, static and dynamical systems are considered, with a particular focus on nonlinear systems. The study emphasizes accuracy and computational speedups, and considers systems arising in a wide variety of engineering fields. All MOR methods exhibit a trade-off between accuracy and speedup.

Matthew J. Zahr
University of California, Berkeley
Stanford University
bokie89@sbcglobal.net

Charbel Farhat, Kevin T. Carlberg, David Amsallem
Stanford University
CFarhat@stanford.edu, carlberg@stanford.edu, amsallem@stanford.edu

PP1

Adjoint Methods for Inversion of Basal Boundary Conditions in Ice Sheet Flows

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, which represent slip along the basal surface. Here we study mathematical and computational issues in the inverse problem of inverting for basal slip parameters from observations of surface flow velocities and surface elevation. We employ adjoint-based Gauss-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

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