U.S. Department of Energy
Office of Advanced Scientific Computing Research

Applied Mathematics Research Program
Annual PI Meeting

May 22 – 24, 2007

ABSTRACTS

Plenary Talk Abstracts

Optimization Research in 2007: Incremental, Transformational, or Recycled?

Margaret H. Wright
Courant Institute of Mathematical Sciences
New York University

Several federal agencies, including the Department of Energy, have recently been urged to support "transformational" and "disruptive" research at the frontiers of innovation, rather than tinkering with incremental changes on the sidelines.

If we assess continuous optimization research during the past few decades in terms of these categories, what do we find? Taking an obvious example, the interior-point revolution in constrained optimization was undoubtedly both disruptive and transformational, but it began in 1984, more than twenty years ago---and much of its theoretical content was known in the 1960s and even earlier.

At least four related questions come to mind in surveying research on continuous optimization today:

- Are there genuinely new ideas with the potential to transform our ability to analyze and solve difficult optimization problems? If so, what are they?
- Is our progress steady but primarily incremental?
- To what extent are we recycling old ideas, shaping them into new forms, and then applying them to new problems?
- How does the interplay between theory and implementation affect the likelihood of transformational research in optimization?

The idea of this talk is to examine selected topics in constrained optimization from these perspectives, in the hope of provoking a lively discussion.

Adaptive Control of Multiscale Modeling Error, With Applications to Large-Scale Molecular Systems

Tinsley Oden
Institute for Computational Engineering and Sciences
The University of Texas at Austin

"If *error is corrected whenever it is recognized as such, the path to error is the path of truth*." So said the philosopher Hans Reichenbach in his 1951 treatise, *The Rise of Scientific Philosophy*. In this lecture, we describe a general approach to multiscale modeling that embraces the

Reichenbach philosophy: estimate the modeling error and then systematically correct it through an adaptive process called the goals algorithm.

Central to this approach is the identification of a so-called *base model,* which involves the finest spatial and temporal scales of physical events important in simulating the phenomenon of interest. The base model is rarely solved and is generally intractable; its purpose is to provide a datum with respect to which other coarser-scale models can be compared. The goal of the computation is to obtain accurate approximations of specific features of the base model solution called *quantities of interest* (or target outputs) $Q(u)$, u being the fine-scale solution to the base model, the level of desired accuracy being specified a priori by setting error to tolerances, γ. The goals algorithm attempts to generate sequences of coarse-scale *surrogate models*, generally obtained by using methods of homogenization or ensemble averaging, but designed so that the error in the quantities of interest, $Q(u) - Q(u_o)$ ($u^{(k)}_o$ being a solution of the k^{th} surrogate problem) is brought when the preset tolerance: $|Q(u) - Q(u^{(k)}_o)| < \gamma$.

Our methods of estimating modeling error is based on the theory in [1] (see also [2,3]), and makes use of adjoint problems and residual a posteriori error estimates. We describe ongoing work on the application of these ideas to large-scale molecular models of polymer densification encountered in nanomanufacturing of semiconductor devices.

References
1. J.T. Oden and S. Prudhomme, "Estimation of Modeling Error in Computational Mechanics", *J. Computational Physics*, **182**, 496-515 (2002).
2. J.T. Oden, S. Prudhomme, A. Romkes, and P.T. Bauman, "Multiscale Modeling of Physical Phenomena: Adaptive Control of Models", *SIAM J. Scientific Computing*, **28** (6), 2359-2389 (2006).
3. S. Prudhomme, P.T. Bauman, and J.T. Oden, "Error Control for Molecular Statics Problems", *Int. J. Multiscale Comp. Eng.*, **6** (5) (2007).

Predictive Capability in Computational Science and Engineering

William L. Oberkampf
Validation and Uncertainty Quantification Department
Sandia National Laboratories

Since the beginning of the Advanced Simulation and Computing Program in 1997, high-performance computing capability has increased by roughly three orders of magnitude. More difficult than determining improvements in hardware performance, however, is determining how predictive capability in computational science and engineering has improved. Although several key elements contribute to predictive capability, this presentation will focus on model validation, model calibration, and extrapolation of models to conditions of intended use. These three topics address fundamental issues in mathematical modeling of physics, interpretation of experimental measurements, and physics-based statistical inference. We believe the dominant theme in predictive capability should be the quantification of uncertainties to enable decision makers to better understand the anticipated performance of a system with respect to system requirements and the consequences of system failure. A brief review will be given of the strengths and weaknesses of existing frameworks for uncertainty quantification of system performance. A proposed framework will be presented that sequentially links (a) model accuracy assessment by

comparison with experimental data, (b) calibration of model parameters, and (c) extrapolation of the model to the conditions of intended use. Each of these steps involves the estimation of aleatory and epistemic uncertainties, e.g. those due to experimental measurements, physics modeling deficiencies, calibration of model parameters, the growth of model-form uncertainty due to extrapolation, and uncertainties in the environment in which the system is required to operate. Open questions with the proposed framework will also be discussed.

Numerical modeling of low Mach number flows

John B. Bell
Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory

A number of important application areas in fluid mechanics are low Mach number flows in which the fluid velocity is much smaller than the speed of sound in the fluid. Although it is possible to solve these types of problems using explicit compressible flow methodologies, there is a potential for considerable computational savings if one can effectively exploit the separation of scales between fluid motion and acoustics. The starting point for this type of approach is the use of low Mach number asymptotics to derive effective equations that capture this separation of scales. The simplest model of this type is the incompressible Navier-Stokes equations, which describes a single-component non-reacting isothermal viscous fluid. More sophisticated models can include a variety of additional physics including multicomponent mixtures, chemical reactions, and effects of background stratification. These types of low Mach number models recast the system in a form in which the natural scale of the dynamics is given by the fluid velocity not the sound speed. As a result, this approach offers the possibility of taking much larger time steps than a compressible formulation. However, how Mach number models change the underlying mathematical structure of the system that needs to be solved numerically. Adopting a low Mach number formulation modifies the underlying conservation laws and introduces a constraint on the evolution. Realizing the potential of a low Mach number model requires that one develop effective numerical integration schemes for these more complex constrained systems. Developing efficient numerical algorithms for modeling low Mach number flows involves three basic elements: basic discretization strategy, gridding strategy and parallel implementation. Here we discuss an approach based on combining a projection method for the constrained system with structured adaptive mesh refinement algorithm implemented using a coarse-grained parallelization strategy. We will discuss the interplay between the different components of the methodology and how design decisions for these components relate to the overall parallel performance of the method. Illustrations of the resulting methodology will be presented from several application areas including combustion and astrophysics.

Recent Developments in Systems Solvers

Tom Manteuffel
University of Colorado at Boulder

This talk will begin with a brief survey of current methods for solving large sparse linear and nonlinear systems of equations that arise from numerical approximation to partial differential equations. In particular, Nested Iteration-Newton-Krylov-Multigrid methods will be examined.

In this approach, computation is initiated on the coarsest grid possible. Solution on each grid serves as a starting guess for the next level of grid refinement. In the presence of sharp a-posteriori error measure, the approximate solution guides the refinement process, involving both grid placement and choice of finite element space on the next grid level. Once the new grid level is determined, the solution on that grid is obtained through a Newton-Krylov process that employs multigrid to provide an optimal preconditioning.

Problems with irregular grids, heterogeneous materials, and non-elliptic behavior require more robust solvers. Adaptive algebraic multigrid (aAMG) methods use information obtained during the iteration process to modify the components of the multigrid algorithm and expand the domain of applicability.

Numerical examples, in the context of First-Order System Least-Squares discretization, will be presented. It will be shown that the solution to complex systems, such as coupled fluid/elastic systems, can be achieved in the amount of work corresponding to a relatively small number of work units on the final grid refinement level. Here, a work unit is defined to be the amount of work required to evaluate the residual on the finest grid level.

Architecture Trends and Implications for Algorithms

William D. Gropp
Mathemathics and Computer Science Division
Argonne National Laboratory

Computer architecture has come a long way since floating-point operations were expensive and memory references were relatively cheap, yet much algorithm design optimizes for simple cost metrics based on floating-point operations. The situation has recently become even more complicated as processor clock rates have stagnated and multicore processors become common. This talk reviews some of the trends in architecture and what they mean for algorithm design. Examples of relatively simple models that can be used to guide algorithm development will be presented. Challenges in scalability to millions of processor cores will be covered, both in terms of scalability and fault management.

Towards Uncovering Simplicity from Complexity: Finding the Dots, Connecting the Dots, Understanding the Dots in Petascale Data

Nagiza F. Samatova
Computer Science and Mathematics Division
Computational Biology Institute
Oak Ridge National Laboratory

Biological systems, such as living cells, are inherently complex. This complexity arises from the selective and nonlinear interconnections of functionally diverse components to produce coherent behavior. Computational modelling and simulation that reproduce and predict such behavior form the Holy Grail of systems biology. The key challenge is to reveal underlying simplicity from biological complexity. Unlike the four Maxwell's equations describing all the electro-magnetic phenomena, the fundamental rules (simplicity) that quantify the low dimensional behavior of biological systems are yet to be discovered. The promising approach aims to interrelate emerging disparate and noisy "omics" observations by relying on mathematics, computer science, information technology, and computing.

In this talk, I will highlight the current state-of-the-art in this area, the challenges that still must be addressed, and how the DOE applied mathematics community may be able to contribute towards predictive understanding of biological systems. The difficulty lies in finding informative features (identifying the dots) and linking them (connecting the dots) to formulate fundamental principles governing complex natural phenomena (understanding the dots).

Linear Solver Abstracts

Algebraic Multigrid with Applications

R. D. Falgout, T. V. Kolev and P. S. Vassilevski
Lawrence Livermore National Laboratory

Multigrid methods are referred to as optimal methods because they require O(N) work to solve a problem with N unknowns. This property is necessary for achieving scalability on huge problems on thousands of processors. Algebraic multigrid (AMG) solves linear systems based on multigrid principles, but in a way that depends only on the coefficients in the underlying matrix. This makes them highly suitable for solving complex problems with jumps in coefficients on unstructured grids. Although AMG has proven extremely useful on a variety of problems, there are some applications for which suitable multigrid methods either currently do not exist or have only recently been developed.

In this talk, we will discuss our recent work and accomplishments in developing AMG methods for difficult problems that arise in electromagnetics (EM) and quantum chromodynamics (QCD) applications. We will first give an overview of multigrid and AMG, then discuss our theoretical developments that have led to some of our recent accomplishments. In particular, we will describe our recent accomplishments in developing specialized algebraic multigrid methods that target definite Maxwell equations. More specifically, we will discuss the auxiliary mesh approach which later led to (based on a theoretical result by Hiptmair and Xu) the development of our new Maxwell solver (AMS) that is provably scalable. Another approach of more theoretical value at this point is the construction of new finite element spaces on agglomerated elements suitable for the design of multigrid methods on general unstructured meshes. These spaces provide commuting de Rham complexes similarly to the ones that naturally arise for discretizations on uniformly refined meshes. One of the pair of spaces in the de Rham complex is suitable for solving Maxwell equations and thus provides a tool for constructing element based algebraic multigrid on general unstructured meshes. The approach itself goes beyond Maxwell problems; it provides a whole sequence of spaces that can also be used in discretizing flow problems, as well as offers the potential for upscaling of variety of linear and nonlinear PDEs. The techniques will be illustrated with a number of examples. We will also describe our adaptive AMG methods and recent progress applying these methods to QCD.

Center for Applied Scientific Computing, UC Lawrence Livermore National Laboratory, 7000 East Avenue, Mail stop 560, Livermore, CA 94550, U.S.A. E-mail address: {rfalgout,tzanio,panayot}@llnl.gov

Submitting Author: Rob Falgout

Fast Iterative Solution of Models of Incompressible Flow

Howard C. Elman
University of Maryland

Abstract

We discuss efficient numerical algorithms for solving the algebraic systems of equations that arise from discretization of the incompressible Navier-Stokes equations

$$(1) \qquad \alpha\mathbf{u} - \nu\Delta\mathbf{u} + (\mathbf{u}\cdot\mathrm{grad})\mathbf{u} + \mathrm{grad}\,p = \mathbf{f}\,, \qquad \mathrm{div}\,\mathbf{u} = 0$$

on two-dimensional and three-dimensional domains. The methods under consideration are *preconditioning strategies* for linearized versions of this problem. They apply to both steady ($\alpha = 0$) and transient ($\alpha = 1$) problems, and they can be generalized to handle more complex dynamics (e.g., thermal effects) and the eigenvalue problems arising from stability analysis of steady solutions. We give an overview of these methods, with emphasis on three main points:

1. General properties. Discretization of (1) and solution of the nonlinear algebraic equations by Newton or Picard iteration requires the solution of systems of equations of the form

$$(2) \qquad \begin{pmatrix} F & B^T \\ B & -C \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ g \end{pmatrix},$$

where C is positive-semidefinite and is 0 for div-stable discretizations. The preconditioners have the general form

$$(3) \qquad \begin{pmatrix} Q_F & B^T \\ 0 & -Q_S \end{pmatrix},$$

where Q_F^{-1} is an approximation to the action of the inverse of F, and Q_S^{-1} represents an approximation to the action of the Schur complement $C + BF^{-1}B^T$. We show that with appropriate choices of Q_F and Q_S, this preconditioner can be combined with Krylov subspace methods to produce solution algorithms of nearly optimal complexity and with only mild dependence on Reynolds numbers for solving (2).

2. Performance in a practical computing environment. Using the finite element package MPSalsa developed at Sandia National Laboratories, we demonstrate the effectiveness of these methods for solving models of enclosed flow and flows over obstacles in two and three dimensions. In particular, we demonstrate the effective implementation of the ideas in a large-scale parallel setting, using multigrid for component problems. Moreover, we show improvements in performance (with speedups of factors of five to ten for large problems) achieved in comparison to more traditional methods used for incompressible flow, including the SIMPLE algorithm and domain decomposition strategies.

3. Generalization to other settings The methodology embodied in (3) can be adapted to handle more general problems, including numerical solution of the Boussinesq model, which gives rise to augmented versions of (2) of the form

$$(4) \qquad \begin{pmatrix} F_u & B^T & G \\ B & -C & 0 \\ H & 0 & F_T \end{pmatrix} \begin{pmatrix} \delta \\ \eta \\ \tau \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ 0 \\ g \end{pmatrix},$$

and *stability analysis* of steady solutions of the Navier-Stokes or Boussinesq equations. To elaborate briefly on the latter problem: stability analysis entails the solution of generalized eigenvalue problems

$$\mathcal{A}x = \lambda \mathcal{B}x$$

where $\mathcal{A}$ has the form of the coefficient matrix in (2) or (4) and $\mathcal{B}$ is a symmetric positive-semidefinite matrix. Computation of critical eigenvalues (typically, those with rightmost real parts) by iteration requires the subsidiary computation of the action of the inverse of $\mathcal{A}$, for which the solution algorithms discussed above are designed. This idea can also be combined with specific enhancements for the eigenvalue problem. In particular, it is not necessary to solve the subsidiary systems to high accuracy to obtain information about wanted eigenvalues, and the preconditioner can be modified in a computationally efficient manner to be more effective in the subspace spanned by the corresponding eigenvectors.

Inexact Krylov Subspace Methods for PDEs and Control Problems

Daniel B Szyld
Department of Mathematics, Temple University

In many circumstances, a known good preconditioner is not easily computable. Instead, an approximation to it is available. This is the case, for example, when the preconditioner has an inverse associated with it, such as in Schur complements (e.g., in saddle point problems), or in the reduced Hessian in some control problems. The application of the preconditioner implies then an iterative solution of a linear system. In these cases, the question is how accurately to solve the (inner) iteration. In our work on Inexact Krylov methods, we have shown that the inner iterations can be solved progressively less accurately, as the underlying Krylov method (e.g, GMRES) converges to the overall solution. Computable inner stopping criteria have been developed to guarantee convergence of the overall method. We will discuss these criteria, and illustrate its application to several problems. Currently, we are applying these ideas to parabolic control problems, where the reduced Hessian has two different inverses; and thus two inner iteration criteria are needed.

Submitting Author: Daniel B Szyld

A Newton-Krylov solver for fully implicit 3D extended MHD

L. Chacon
Los Alamos National Laboratory

We present results on the application of Jacobian-free Newton-Krylov (JFNK) methods applied to the time-dependent, primitive-variable, 3D extended magnetohydrodynamics (MHD) equations. MHD is a fluid description of ionized gases (plasmas) in the presence of electromagnetic fields. The standard MHD description treats the plasma as a single fluid, with ions and electrons evolving together. Extended MHD (XMHD) includes nonideal effects such as nonlinear, anisotropic transport and two-fluid (Hall and diamagnetic) effects. XMHD supports so-called dispersive waves (whistler, kinetic Alfven), which feature a quadratic dispersion relation $\omega \sim k^{2}$. In explicit time integration methods, this results in very stringent CFL limits, which severely limits their applicability to the study of long-frequency XMHD phenomena.

A fully implicit implementation promises efficiency without sacrificing numerical accuracy. However, the nonlinear nature of XMHD and the numerical stiffness associated with its fast waves make this endeavor very difficult. Newton-Krylov methods can meet the challenge provided suitable preconditioning is available.

We have developed a successful preconditioning strategy for the 3D primitive-variable XMHD model, which we term ``physics-based.'' Physics-based preconditioning attempts to parabolize an otherwise hyperbolic system of equations (which is diagonally submissive for $\Delta t > \Delta t_{CFL}$) to restore diagonal dominance (and therefore the effectiveness of multilevel methods). The use of approximate multigrid (MG) techniques to invert the ``parabolized'' operator is a crucial step in the effectiveness of the preconditioner and the scalability of the overall algorithm. The parabolization procedure can be properly generalized using the well-known Schur decomposition of a 2x2 block matrix. In the context of XMHD, the resulting Schur complement is a system of

PDE's that couples the three plasma velocity components, and needs to be inverted in a coupled manner. A system MG treatment of this coupled system is possible since the XMHD Schur complement is block diagonally dominant by construction (and therefore block smoothing is effective).

In this presentation, we will discuss the derivation and validity of the physics-based preconditioner for resistive MHD and its generalization to XMHD, the connection with Schur complement analysis, and the system-MG treatment of the associated systems. Grid convergence studies will demonstrate that the algorithm scales optimally under grid refinement, and very favorably with the time step size. We will also discuss parallel scalability of the overall algorithm (which has been parallelized using the PETSc framework). Finally, we will discuss the status of the implementation of this solver in an adaptive mesh refinement framework.

Submitting Author: L. Chacon

Multiresoluton analysis and low separation rank methods for scientific simulations

George Fann, ORNL; Robert Harrison, ORNL/University of Tennessee, Knoxville; Gregory Beylkin, University of Colorado at Boulder

We describe the development of fast, adaptive and accurate O(N) methods for solving three and six dimensional Schrodinger equations.These algorithms are based on multiresolution analysis and low separation rank representations of functions and operators with finite but arbitrary precision. The methods are under development and have been applied to solving Hartree-Fock equations and Density Functional Theory equations to high accuracy (10+ digits). Features of these type of methods permit scalable implementations on massively parallel computers. Extensions of these methods to higher dimensions will also be discussed.

Submitting Author: George Fann

Parallel segregated Schur complement methods for fluid density functional theories

Michael A Heroux, Sandia National Laboratories; Laura J. D. Frink, Sandia National Laboratories; Andrew G. Salinger, Sandia National Laboratories

Density functional theories for inhomogeneous fluids (Fluid-DFTs) can be applied to a variety of problems from simple fluids in small confined regions, to corrosion effects, to biological phenomena. Fluid-DFTs require the solution of large systems of equations with many degrees of freedom (DOFs) per node. Historically, solvers for these problems have used simple Picard iterations across DOFs or, more recently, fully-coupled general algebraic techniques.

Multiple degrees of freedom (DOFs) per node properties are common to many numerical applications. Segregated solvers, which attempt to view each DOF across the grid as a sub-problem within the larger fully-coupled problem, have been successfully used in many problem domains. Similarly Schur complement methods, which formally eliminate variables by block

Gaussian elimination, can reduce the complexity and cost of solution. In this paper we present a combination of these two classes of algorithms applied to Fluid-DFTs. We will show that viewing Fluid-DFTs from a segregated variable perspective yields a rich structure that can be exploited in the development of robust, scalable solution methods. Furthermore, the algorithms we develop can be applied in a similar fashion to all types of problems generated by the target application, showing that our basic approach is a useful general-purpose technique in this problem domain.

By viewing Fluid-DFTs from this perspective, we develop a mathematical framework and a collection of solution algorithms that have a dramatic impact on the robustness, performance and scalability of the implicit equations generated by Fluid-DFTs. The resulting solution algorithms have several desirable properties: minimal or no user tuning parameters, convergence rates that are invariant under increasing processor counts, and slow growth in cost vs. problem complexity.

In this presentation we provide an overview of the Fluid-DFTs application Tramonto, and a description of the discrete problem formulation. We then characterize the equations in ways that allow formation of a Schur complement that is easy to solve and amenable to scalable parallel execution. Next we discuss the solution methods and the parallel implementation of the solvers. Finally we analyze the cost of these algorithms as a function of problem complexity, and parallel execution and show results from problems that have been impacted by these new solvers.

Submitting Author: Mike Heroux

Multiscale Methods/Mathematics Abstracts

Efficient Solution Methods, Error Estimation, and Atomistic-Continuum Adaptivity for the Quasicontinuum Approximation

Mitchell Luskin, Marcel Arndt, Matthew Dobson, Ryan Elliott, and Ellad Tadmor
University of Minnesota

The quasicontinuum (QC) method has been successfully used to efficiently couple atomistic and continuum models for crystalline solids. The atomistic model is used in regions with highly non-uniform deformations such as around dislocations, whereas the continuum model is used in regions with nearly uniform deformations to reduce the number of degrees of freedom.

We give a convergence rate for an iterative method to solve the nonlinear equilibrium equations that can be used to give an efficient solution method for a quasi-static process.

The quasicontinuum methodology involves the application of the Cauchy-Born rule to the underlying lattice in continuum regions. The validity of this approximation is dependent on the utilization of a unit cell that does not restrict possible lattice instabilities. At the same time, the computational efficiency of the method relies on the use of a minimal cell size. We describe recent work on the development and analysis of an adaptive algorithm to change the element cell size as the element strain evolves during a quasi-static process.

We also develop an a posteriori error estimator which quantifies the modeling error for a goal function and allows for an adaptive decision about which regions should be modeled as a continuum and which regions should be modeled atomistically.

Submitting Author: Mitchell Luskin

Analysis of Atomistic-to-Continuum (AtC) Coupling Methods

Rich Lehoucq, SNL; Pavel Bochev, SNL; Mike Parks, SNL; Max Gunzburger, FSU; Jacob Fish, RPI

My presentation describes project work for "A Mathematical Analysis of Atomistic-to-Continuum (AtC) Coupling Methods" funded under the MICS Multiscale Mathematics program. This work is collaborative with Pavel Bochev and Mike Parks (SNL), with Don Estep (CSU), Jacob Fish and Mark Shephard (RPI), and Max Gunzburger (FSU).

Material behavior is governed by processes that are often controlled by structures and dynamics spanning many length and time scales. Theoretical and computational treatment of multi-scale

phenomena, including the development of predictive capabilities, is therefore of fundamental importance. Synthesizing, or coupling, atomistic and continuum descriptions of physical phenomena is an attempt to ameliorate the overwhelming computational costs associated with an all atomistic simulation.

Atomistic-to-Continuum (AtC) coupling enables a continuum calculation to be performed over the majority of a domain of interest while limiting the more expensive atomistic simulation over a subset of the domain. Unfortunately, combining atomistic and continuum calculations is challenging because the former is based on individual non-local force interactions between atoms or molecules while continuum calculations deal with bulk quantities that represent the average behavior of millions of atoms or molecules.

My presentation describes two related efforts. The first considers blending non-local and local force models across a subdomain, or bridging region. This leads to a class of methods where the synthesis of the two models is imposed by constraint operators using Lagrange multipliers. We employ variational methods leading to an abstract framework suitable for analysis, algorithm design and mechanical interpretation. The second effort considers inserting a microcontinuum model, peridynamics, between the atomistic and continuum models. Peridynamics is a continuum theory employing a non local force interaction that describes long-range material interaction. This is precisely the same mechanism used by atomistics, and preliminary work demonstrates that peridynamics can be implemented within molecular dynamics software. Hence, peridynamics can be employed for the study of both microscopic phenomena that are beyond the realm of atomistics and classical continuum mechanics. The former is prohibitive for all but the smallest material domains while the latter, and its resulting mesh- based discretization, assumes a local model of force interaction that is difficult to reconcile with atomistics.

Submitting Author: Rich Lehoucq

Multiscale Modeling of Spatially Distributed Biological Systems

Dion Vlachos1*, Markos Katsoulakis2, Jeremy Edwards3, and James Faeder4
1Department of Chemical Engineering, University of Delaware
2Department of Mathematics and Statistics, University of Massachusetts
3 Molecular Genetics and Microbiology, Cancer Research and Treatment Center,
University of
New Mexico Health Sciences Center
4 Theoretical Biology and Biophysics Group, Los Alamos National Laboratory
* Contact information: vlachos@udel.edu; Tel. (302)-831-2830; Fax (302)-831-1048

Most biological systems exhibit a vast disparity of length and time scales and are inherently stochastic in nature. Currently, stochastic Monte Carlo (MC) simulation is plagued by the multiscale nature of these systems. The overall objective of our work is to develop the necessary mathematical and computational framework that can handle the full range of time and length scales required to model complex biological systems, with emphasis on spatially distributed systems, and apply this framework to an important biological system, namely the epidermal growth factor receptor (EGFR), whose signal disregulation is implicated in a number of cancers. The intrinsic fluctuations present in intracellular networks, especially at the genetic level, cannot be captured using deterministic models. The potential impact of these fluctuations on the cell

physiology necessitates use of stochastic algorithms. These algorithms lead to tremendous computational load due to (i) the execution of one event at a time, (ii) the presence of some large populations, and (iii) stiffness.

We have recently devised the binomial tau-leap method for the first two challenges. Drawing analogies from singular perturbation and low dimensional manifold methods in deterministic systems, we have also developed a multiscale MC algorithm to handle stiffness. Its essence lies in the relaxation of the fast network to a stable quasi-equilibrium manifold in a short relaxation time compared to the time scales of the slow network. Unlike deterministic models, this stochastic manifold is represented by a time-invariant distribution of the fast states that controls the rates of slow variables. Novel relaxation criteria, invoking hypothesis testing principles and computational singular perturbation concepts, have been developed. We have finally developed a hybrid algorithm that seamlessly integrates the above methods to handle simultaneous separation in time scales and populations of species.

The large disparity in length scales within a cell or in communication between compartments of a cell or in a population of cells in tissues demands development of suitable mathematical and computational tools. In this regard, we have been developing algorithms that couple spatially distributed MC models for the plasma membrane with well-mixed stochastic models representing intracellular processes for a single cell. In addition, we have developed a prototype hybrid model for communication between subsystems coupled by transport, modeled in the simplest case with diffusion. We model each compartment stochastically and their coupling using either the deterministic PDE (diffusion) model or our corresponding, recently developed coarse-grained MC model. The former way of coupling is the only one available in the literature but its correctness has never been assessed. The latter is a fully stochastic (more expensive but still tractable) way of coupling that provides exact results. Parametric studies followed by autocorrelation analysis have been conducted to identify the strength of coupling between submodels and the regime of applicability of hybrid PDE/stochastic models.

Submitting Author: Dion Vlachos

Multiscale energy constraints in rotating and stratified flows

Susan Kurien, Los Alamos National Laboratory; Beth Wingate, Los Alamos National Laboratory; Mark Taylor, Sandia National Laboratories

We find that in strongly rotating and stratified three dimensional flows in the 'quasi-geostrophic' parameter regime, the forward cascade of potential enstrophy implies anisotropic constraints on the wavenumber distribution of kinetic and potential energy. We show that in the tall, narrow scales in physical space, the horizontal kinetic energy is suppressed and correspondingly scales with the horizontal wavenumber as k_h^{-3}. In the limit of flat, wide scales in physical space, the potential energy is suppressed and scales with the vertical wavenumber as k_z^{-3}. These constraints differ from the classical scaling estimates of Charney 1971 for quasi-geostrophic dynamics wherein both kinetic and potential energies are expected to scale with total wavenumber $k = (k_h^2 + k_z^2)^{1/2}$ as k^{-3}. We propose that our scaling estimates are more appropriate as diagnostics for realistic flows which rarely exactly satisfy quasi-geostrophy.

As an extension of the work in the quasi-geostrophic regime, we consider the cases where rotation dominates over stratification the two pieces of linear potential enstrophy force the kinetic and potential energies into anisotropic cascades and scaling regimes. We verify our predictions using resolved numerical simulations of the Boussinesq equations computed in the relevant parameter regimes.

Submitting Author: Susan Kurien

Dimensional reduction for multiscale problems

Alexander Chorin
Lawrence Berkeley National Laboratory

I will present methods for the reduction of the complexity of computational problems, both time-dependent and stationary, and their connections to probability, renormalization, scaling, and statistical mechanics, together with examples. The main points, are: (i) in time dependent problems, it is not legitimate to average equations without taking into account memory effects and noise; (ii) mathematical tools developed in physics for carrying out renormalization group transformations yield effective block Monte-Carlo methods; (iii) the Mori-Zwanzig formalism, which in principle yields exact reduction methods but is often hard to use, can be tamed by approximation; and (iv) more generally, problem reduction is a search for hidden similarities. In the examples I will emphasize the "t-model" for problems where the memory (=autocorrelation of the noise) has a large support; this is important in applications to hydrodynamics.

Submitting Author: Alexander Chorin

Multiscale Strategies for Critical Phenomena in Heterogeneous Materials

Kim F. Ferris[1], Alexandre Panchenko[2], and Alexandre M. Tartakovsky[1]
[1]Pacific Northwest National Laboratory, Richland, WA; [2]Washington State University; Pullman, WA

Strategies to reconcile phenomena over broad spatial and temporal scales have been bounded and restricted by conventional single scale formalisms. In the simplest sense, the physical properties of materials at the different size scales form an 'information' hierarchy, requiring new mathematical and computational methods which allow us to seamlessly join the scales. Complicating the development of these models are the common aspects of materials: 1) heterogeneity – materials can be multiphase and exhibit complex microstructures, and 2) materials properties are often determined by critical phenomena, not the straightforward statistical averages. This latter situation is particularly vexing in materials modeling as the majority of the computational effort is spent on building complex simulation frameworks, while the controlling aspects lie at the extremes of the feature probability distributions (ie. Cracks and fractures in a material breaking a thermal and/or conduction pathway).

In this presentation we present recent developments in hybrid methodologies as a multiscale approach to system heterogeneity; and specifically the case of reactive flow through porous

media. Because of the difficulties associated with geometrically complex pore-scale boundaries that are dynamically changing as a result of precipitation and the non-linearity of the multiphase flow and reactive transport equations, traditional grid-based methods have not been successfully applied to pore-scale processes. An alternative approach is to use mesh-free methods such as spherical particle hydrodynamics (SPH). Because of its Lagrangian particle nature, SPH has several advantages for modeling pore-scale flow and reactive transport: 1) using a Lagrangian framework there is no non-linear term in the momentum conservation equation; SPH allows accurate solution of momentum dominated flows; and 2) complicated physical and chemical processes associated with changes in solid boundaries due to dissolution or precipitation and chemical reactions are more directly accessible to simulation.

Submitting Author: Kim Ferris

Multiscale Approach to Self-Organization of Microtubules and Actin Filaments

Dmitry Karpeev, Igor Aronson, Lev Tsimring, Hans Kaper
Argonne National Laboratory

Recent experiments have shown that a system of microtubules and molecular motors is capable of sustaining a variety of large-scale two-dimensional structures (asters, vortices, and other topological defects). Underlying this phenomenon is a multiscale process, where nonlinear interactions on a microscopic scale result in the emergence of coherent structures on the macroscopic scale. For the study of this process, we propose a mathematical framework based on a master equation on the mesoscopic scale. The equation describes the evolution of the density of microtubules as a function of position and angular orientation. The interaction kernel is determined by the molecular interactions on the microscopic scale. Macroscopic spatiotemporal structures are identified through dimension reduction from the mesoscopic scale.   All three scales interact strongly, resulting in macroscopically observable patterns of self-organization.

Recently we have developed mathematical analysis and simulations to address different aspects of the problem. At the mesoscopic level we have developed an approach at analytically solving the master equation in the homogeneous 2D case.

At the microscopic case we have investigated the effects of flexibility of filaments on their collision properties that control the kernel of the master equation at the mesoscopic level. We have shown that flexibility enhances the inelasticity of interaction and that motor dwelling may be responsible for complete alignment of tubules. Furthermore, we have investigated interaction of much softer filaments, such as actin, which is responsible for the cytoskeleton formation. Interacting actin pairs exhibited an Euler-type buckling instability for sufficient motor stengths. It is conjectured that the buckled states are responsible for cross-linking of actin into cytoskeletal networks and, thereby, for the rheological properties of these networks.

Submitting Author: Dmitry Karpeev

Algorithms and Software for Multiscale Simulation of Chemically Reacting Systems

Linda Petzold
Dept. of Computer Science, University of Santa Barbara

Traditional deterministic approaches for simulation of chemically reacting systems fail to capture the randomness inherent in such systems at scales common in intracellular biochemical processes. In these systems, the small population of a few critical reactant species can result in discrete stochastic behavior. But for many realistic biochemical systems, the computational cost of a fully discrete stochastic simulation can be prohibitively high. We will describe our efforts to develop adaptive multiscale algorithms and the STOCHKIT software package for high-performance simulation of biochemical reacting systems.

Submitting Author: Linda Petzold

Mortar multiscale methods for flow in porous media

Ivan Yotov, University of Pittsburgh; Todd Arbogast, The University of Texas at Austin; Gergina Pencheva, The University of Texas at Austin; Mary F. Wheeler, The University of Texas at Austin

Fully resolving flow in highly heterogeneous porous media is often a computationally intractable task. Several upscaling and variational multiscale methods have been developed to address this problem. In this talk we discuss the relationship between mortar finite element methods and multiscale finite element methods. The latter represent the solution as a sum of a coarse scale and a fine scale (subgrid) component and require solving local fine scale problems to compute the multiscale basis functions. Mortar methods with coarse mortar spaces also resolve the solution on the fine scale in each subdomain, but impose continuity conditions on the coarse scale. The mortar formulation is more flexible than existing variational multiscale methods, since it allows for locally varying the interface degrees of freedom if higher resolution is needed. A domain decomposition algorithm reduces the multiscale algebraic system to a coarse scale interface problem. We study the accuracy and efficiency of the mortar multiscale approach in the context of mixed finite element and related methods.

Submitting Author: Ivan Yotov

Multiscale Numerical Methods for Plasma Kinetics

Russel Caflisch, UCLA; Bruce Cohen, LLNL; Andris Dimits, LLNL; CM Wang, UCLA

Plasmas exhibit behavior over an enormous range of temporal and spatial scales.

These multiscale phenomena are both difficult to simulate and essential for applications ranging from astrophysics to controlled fusion. For Coulomb collisions between charged particles, a single plasma volume may include both long (kinetic) and short (fluid-like) mean-free-path

regions. Effective computation of these combined processes has proved to be a significant challenge, for example in edge plasma simulations. We are developing a hybrid computational method for Coulomb collisions that combines a Monte Carlo particle simulation and a fluid dynamic solver. This approach is distinct from domain decomposition, since it is a single uniform method throughout phase space.

The new method is based on a hybrid representation of the velocity distribution function f(v), as a combination of a Maxwellian equilibrium M(v) and a collection of discrete particles g(v). The Maxwellian M evolves in space and time through fluid-like equations, and the particles in g convect and collide through a Monte Carlo particle method. Interactions between M and g are represented by a thermalization process that removes particles from g and includes them in M and a dethermalization process that samples particles from M and inserts them into g. As a guide to this development and an alternative hybrid method, we are also formulating a related delta-f method.

For rarefied gas dynamics (RGD), a related hybrid method has been developed, combining the Direct Simulation Monte Carlo (DSMC) method for particle dynamics and a fluid dynamic solver [1,2]. This method is based on a on a power-series expansion of the solution to the kinetic equation, and it has shown to be both fast and accurate for simulation of a variety of RGD flows, such as shocks and channel flows.

Extension of the hybrid method from RGD to plasmas is challenging because of the long range of Coulomb interactions and simultaneous presence of kinetic and continuum regions in both velocity and position space. We are using   binary collision method for Coulomb interactions that was originally developed by Takizuka & Abe [3] and modified by Nanbu [4]. As a possible alternative to the binary collision method for collisions of the kinetic plasma with the Maxwellian fluid, we are also implementing a particle/field method [5] that is more efficient than binary collisions. We have developed a new thermalization/dethermalization method that has a mathematical basis and appropriately discriminates between different regions of velocity space. We have derived a generalized set of continuum equations that includes sources and sinks from the thermalization/dethermalization.

As test cases for the hybrid method, we have used spatially-homogeneous problems, including relaxation of an anisotropic Maxwellian and evolution of a bump-on-tail, for the collional process alone. For testing the combination of collisions and convection, we are using a plasma sheath problem.

This talk will describe the derivation and implementation of the hybrid method, and illustrate its performance of the test problems described above. As background, we also describe the relation between the hybrid methods for RGD and the extension to plasmas.

[1] L. Pareschi and R.E. Caflisch, "An implicit Monte Carlo method for rarefied gas dynamics," J. Comput. Phys., 154, 90-116, (1999).

[2] R.E. Caflisch, H. Chen, E. Luo and L. Pareschi. "A Hybrid Method that Interpolates between DSMC and CFD" AAIA Proceddings, 2006.

[3] T. Takizuka and H. Abe "A Binary collision model for plasma simulation with a particle code" J. Comp. Phys. 25 (1977) 205-219.

[4] K. Nanbu. "Theory of cumulative small-angle collisions in plasmas" Phys. Rev. E. 55 (1997) 4642-4652.

[5] W.M. Mannheimer, M. Lampe and G. Joyce. "Langevin representation of Coulomb collisions in PIC simulations" J. Comp. Phys. 138 (1997) 563-584.

Submitting Author: Russ Caflisch

Multi Spatial and Time Scale Coupling in Rotating and Stratified Flows

Mark Remmel, Jai Sukhatme, Leslie Smith
University of Wisconsin

The dynamics of rotating, stratified fluids include the complex interaction of waves and turbulence, and are the basis for understanding the earth's oceans and atmosphere. The wide range of spatial scales in geophysical flows, from thousands of kilometers to meters, is one reason why they are so rich in behavior, so costly to compute and so difficult to understand. In certain scale regimes, intermediate-scale motions self-organize to generate larger-scale structures such as jets and vortices, while in other regimes, energy is transferred from largescale winds and tides to small-scale turbulent fluctuations. Furthermore, the dispersive nature of the waves introduces multiple time scales, starting first with linear effects, followed by nonlinear contributions from exactly resonant triads, and then on longer time scales, the dynamics of nearly resonant triads become important and can drastically change the nature of the flow.

Dispersive waves are common to many of the mathematical systems used to model atmospheric and oceanic dynamics, including _-plane flow, the rotating shallow-water equations, and the 3D Boussinesq approximation. All these systems are governed by partial differential equations (PDEs) with a linear, skew-hermitian operator and a quadratic nonlinearity. The waves are the eigenmodes of the linear operator, and the wave frequencies are the eigenvalues given by the dispersion relation. The quadratic nonlinearity induces triad interactions between three eigenmodes, leading to energy transfer between modes and to turbulence.

In general, dispersive wave motions tend to reduce the energy transfer between scales that would be achieved in their absence. However, in all of the PDE systems mentioned above, there are resonant wave interactions for which the nonlinear transfer remains strong. These interactions between wavevectors k, p and q satisfy the resonance condition $\sigma(k) + \sigma(p) + \sigma_{-}(q) = 0$, where $\sigma(k)$ is the wave frequency of the mode with wavevector k. Exactly resonant triads dominate for short times, while nearly resonant triads with $|\sigma(k) + \sigma(p) + \sigma(q)| = \varepsilon \ll 1$ play a significant role for long times. The well-known Quasi-Geostrophic (QG) model retains only interactions between eigenmodes with zero frequency (the so-called vortical modes), and so all nonlinear interactions are exactly resonant. Thus the QG model is inherently a short-time approximation.

There are two goals of the present work. First, we show how nearly resonant triads fundamentally change the fluid dynamics by generating large-scale coherent structures on long time scales; those

structures are jets, layers and vortices as observed in the atmosphere and oceans. Near resonances also capture asymmetry between eastward and westward jets, and between cyclonic and anti-cyclonic vortices. Second, we introduce a new, non-perturbative method to include physics beyond QG dynamics. New PDE reduced models are derived by extracting any combination of wave- and vortical-eigenmode triad interactions, which may contain a subset of near-resonant interactions. These PDE models allow for a complete understanding of three-wave and wave-vortical interactions in rotating and stratified flows.

Submitting Author: Leslie Smith

Wavelet-based Spatiotemporal Multiscaling in Diffusion Problems with Chemically Reactive Boundary

S. Pannala, S. Simunovic, C. S. Daw, and P. K. Nukala
Oak Ridge National Laboratory

G. Frantziskonis, S. Mishra, A. Mallik, K. Muralidharan, and P. Deymier
University of Arizona

R. Fox and Z. Gao
Ames Laboratory

T. J. O'Brien and M. Syamlal
National Energy Technology Laboratory

Chemically reacting flows over catalytic and non-catalytic surfaces are one of the elementary operations in chemical processing plants. The underlying physical phenomena span time- and length-scales over several orders of magnitude. The challenge is to account and to bridge (as seamlessly as possible) the length and time scales involved in the problem and the multi-physics phenomena operating at various scales. In this talk, development of a multiphysics and multiscale mathematics framework for coupling various modeling methods over a range of scales will be presented along with a simple example that will illustrate how it will enable accurate modeling of heterogeneous reacting flows over catalytic surfaces. The methods that are coupled are the finite differences for bulk diffusion, to model slow and intermediate meso- and macro-scopic phenomena, and the Kinetic Monte Carlo (KMC), to model the fast chemical reactions occurring on the reactive surface.

Coupling between the scales and methods are developed in the framework of Compound Wavelet Method (CWM). CWM is a wavelet based spatio-temporal operator which can be used to perform either compounding operation or transfer (up-scaling or down-scaling) operation. The talk will provide formal introduction to the CWM steps involved (described below) and the current status of these steps:
- **Decomposition:** Wavelet transforms (WT) of the coarse and fine signals
- **Clipping:** Retain relevant scale information from coarse and fine scales
- **Blending/Interpolation** (for cases with no clear separation of scales): Weighted averages of WT coefficients over the overlapping scales and preserve properties like total energy, smoothness in spatio-temporal correlations, etc.
- **Prolongation:** This is the operation in which we stretch out the fine scale WT coefficients to fill the appropriate level in the compounded matrix.

- **Renormalization:** This step involves renormalization of the WT coefficients if the coarse and fine scales are sampled differently.
- **Compounding:** Compound the results of the above step (union operator)
- **Reconstruction:** In this step we recover the approximate true solution which can be used as is for a steady-state scenario or used as the initial condition for a transient simulation.
- **Transfer:** Use CWM for Up-scaling or Down-scaling

Our recent results on application of this multiscale framework on coupling of two prototype methods for the problem of species generated on a chemically reactive boundary and diffusing through the bulk will be presented. The first method in this example, termed "fine," models the chemical reactions on the reactive boundary stochastically by the KMC and the diffusion in the bulk medium deterministically using the time and space discretization based on the chemical reaction process. The second method, termed "coarse," models both the reaction and the diffusion deterministically and uses drastically larger time and space increments. The two methods are coupled by forming a spatiotemporal CWM that merges the respective time and spatial scales in each method. The developed method allows for considerable speed up of the simulation that incorporates multiple scales and methods.

Submitting Author: S. Pannela

Numerical Methods Abstracts

Mimetic discretizations and what they can do for you.

P. Bochev [1,2]
Sandia National Laboratories

Recent advances in compatible discretizations have enabled impressive gains in computational science and affirmed the key role of homological principles in numerical PDEs. Thanks to homological ideas and tools, we now have a much better understanding of why some discretization methods work so well and why other methods fail spectacularly. More importantly, homological ideas can be used to develop stable and physically consistent discretizations, such as mimetic methods, which translate PDEs to algebraic equations that inherit their fundamental structural properties.

We provide a common framework for mimetic methods using algebraic topology to guide our analysis. The key concept in our approach is the natural inner product on co-chains. This inner product is sufficient to generate a combinatorial Hodge theory on co-chains but avoids complications attendant in the construction of a robust discrete Hodge-star operators. In particular, using a reduction and a reconstruction maps between differential forms and co-chains we define mutually consistent sets of *natural* and *derived* discrete operations that preserve the invariants of the De Rham homology groups and obey a discrete Stokes theorem. By choosing a specific reconstruction operator we obtain well-known mixed FE, mimetic FD and covolume methods and explain when they are equivalent.

To illustrate the potential of compatible discretizations we use the mimetic framework to reformulate the discrete Maxwell's equations into a system that is dominated by discrete Hodge-Laplace operators and has much improved scaling for highly heterogeneous conductors - a setting that arises the modeling and simulation of Z-pinch at Sandia. Using properties of mimetic operators we show that the kernel of the mimetic Hodge-Laplace operator has the same dimension as the kernel of the analytic Hodge-Laplacian. Together with the improved scaling, this means that the discrete Laplacian is equivalent to a standard vector Laplace operator. As a result, the reformulated system can be used with standard "black-box" AMG solvers for the Poisson equation.

We will close with a brief summary of the INTREPID project at Sandia. INTREPID is a software package that represents a radical departure from conventional discretization libraries that are designed to support a single discretization paradigm. In contrast, INTREPID provides a flexible software infrastructure that can support a mixture of FE, FV and FD discretizations, defined on different types of cells, including general polyhedral cells. The core toolset of INTREPID represents a software realization of our mimetic framework and offers intuitive API for restriction and reconstruction operators and the natural and derived discrete operations in the framework. INTREPID follows the successful "package" design principle of Trilinos and leverages specialized tools such as MOAB to provide production level support for mesh management.

This talk is based on joint work with J. Hu, C. Siefert, R. Tuminaro (SNL), M. Gunzburger (CSIT, Florida State University), M. Shashkov and M. Hyman (Theoretical Division, Los Alamos National Laboratory).

Submitting Author: Pavel Bochev

Adaptive finite volume method for distributed non-smooth parameter identification

Eldad Haber, Emory University; Stefan Heldmann, Emory University; Uri Ascher, University of British Columbia

In this work we develop a finite volume adaptive grid refinement method for the solution of distributed parameter estimation problems with almost discontinuous coefficients. We discuss discretization on locally refined grids, as well as optimization and refinement criteria. An OcTree data structure is utilized. We show that local refinement can significantly reduce the computational effort of solving the problem, and that the resulting reconstructions can significantly improve resolution, even for noisy data and diffusive forward problems.

Submitting Author: Eldad Haber

Advances in Tracking Interfaces: Semiconductors, InkJets, Seismic Imaging and Transport

J.A. Sethian

Mathematics Department, LBNL

Department of Mathematics, UC Berkeley

In recent years, a collection of numerical techniques have been developed to track moving interfaces. These techniques include Level Set Methods, Narrow Band Level Set Methods, Fast Marching Methods, Ordered Upwind Methods, and Liouville Escape Equation Methods. The unifying theme in these approaches has been to recast moving interface problems in the language of PDEs and differential geometry, and to solve the associated embedded equations through numerical techniques borrowed from hyperbolic conservation laws. These techniques have been applied to a wide spectrum of topics, including combustion, two-phase fluid flow mixing, semiconductor manufacturing, crystal growth, medical imaging, and optimal structural design.

We will discuss some algorithmic advances in these topics, and then focus on some new and on-going applications, mostly taken from the efficient design of manufacturing processes. As examples, we consider the following:

- Numerical simulations of microjets, and display devices: we present coupled two-phase projection method/level set techniques for Newtonian and viscoelastic flow in the manufacturing of display devices for plasma screens, inkjet devices, and and bioinjections

- Seismic inversion algorithms based on fast Dijkstra-like time-to-depth transformations, multiple arrival escape equation methods and moving interface simulations: we present algorithms and field examples for estimating and imaging in geophysical analysis

- Dijkstra-like methods for Terrain mapping, differential games, logistical transport, and shortest paths for navigation and path planning

- Irregular mesh methods with coupled sub-grid interface solvers for materials problems.

As an example, we developed a numerical model of two-phase axi-symmetric viscoelastic flow to simulate a wide collection of nanoscale devices and flow phenomena. We used the Oldroyd-B viscoelastic fluid model, in which both the dynamic viscosity and relaxation time are constant, to simulate two-phase immiscible incompressible flows with surface tension and both viscosity and density jumps across interfaces separating viscoelastic fluid from air, incorporating a macroscopic slipping contact line model. We model air/wall/fluid interactions and measure interactions between geometry and viscoelastic forces. The approach handles problems in which either of the two fluids is either viscoelastic or Newtonian.

The coupled algorithm incorporates (1) projection methods fot fluid incompressibility, (2) level set methods to capture the moving interface, (3) high order

Godunov schemes for convection terms in the momentum and level set equations, (4) first-order upwind algorithms for convective viscoelastic stress, and higher order central schemes for viscosity, surface tension, and upper-convected derivatives, and (6) a slipping line contact model at air/wall/fluid triple points.

These algorithms were tested in the context of inkjet plotters. Regular dye-based inks used in desktop printers are Newtonian. Pigment-based inks, introduced in the 1990's, improve the color durability of inkjet printouts in industrial manufacturing, and are usually viscoelastic, requiring a history dependence between the stress and deformation tensor. These viscoelastic jets now form core new technology for such varied applications as printing plasma screens and automated drug design.

Newtonian Ejection: Experiment

Newtonian Ejection: Full Two-Phase Flow Simulation

Collapsing Fluid Simulation: Charged Drop-on-Demand

λ=.4: Little viscoelasticity λ=1.0: Medium viscoelasticity λ=3.0: More viscoelasticity

2

High-Order Methods for Petascale Science

Paul Fischer
Mathematics and Computer Science Division, Argonne National Laboratory

High-order numerical discretizations will play an increasingly important role as DOE Science moves towards petascale simulation. In a majority of science areas, the quest for increased computational resources is driven by a need to span a broader range of scales, that is, to capture the interaction of small scales with the large. In transport problems such as electromagnetics and fluid mechanics, this implies a need to propagate small scale features over long times and distances, which is most effectively accomplished using high-order methods, as established by Kreiss and Oliger in the 70s. Here, we present recent advances in spectral element methods designed for the petascale architectures that are soon to be deployed by several national agencies, including the DOE and NSF. In particular, we address stable high-order methods, scalable multilevel iterative solvers, and implementation for efficient single- and multi-node performance. Application areas include the study of magnetorotational turbulence in accretion disk models, thermal hydraulics of advanced recycling reactors, wakefield computations in accelerators, and transition in vascular flows.

Submitting Author: Paul Fischer

Simulation Environments for Electromagnetic Scattering

Leslie Greengard
New York University

We will review the analytic and computational foundations of Green's function-based methods for electromagnetic scattering, including high order integral representations, fast solvers, and quasi-periodicity. We will then discuss the development of easy-to-use numerical simulation environments, and present some applications to photonic crystals, random microstructures, and negative index materials.

Submitting Author: Leslie Greengard

Solving Multi-Material and Multi-Physics Problems on Overlapping Grids

William D. Henshaw
Centre for Applied Scientific Computing
Lawrence Livermore National Laboratory

This talk will describe approaches for solving multi-material and multi-physics problems on overlapping grids with applications to high-speed reactive flow, thermal-hydraulics and electromagnetics. The overlapping grid approach can be used to accurately treat problems with interfaces and moving boundaries. In a typical grid construction, narrow boundary fitted grids are used to represent the moving boundary or interface, and these grids overlap with stationary

background Cartesian grids. In moving grid problems, the boundary fitted grids will evolve over time with the overlapping grid interpolation points updated at each time step. Block structured adaptive mesh refinement is used to resolve multiple scales. Refinement grids are added to both the back-ground grids and the near-body curvilinear grids.

This approach has been used to simulate the motion of rigid bodies in high-speed reactive flow. The Euler equations are solved numerically using a second-order extension of Godunov's method. The stiff source term in the reactive case is handled using a Runge-Kutta error-control scheme. Adaptive mesh refinement is used to resolve fine features such as detonations and is coupled with moving grids.

The approach has also been used to model temperature dependent incompressible flow coupled to heat transfer in solids. Overlapping grids are built for the separate fluid and solid domains. Different governing equations are solved in the different sub-domains. The appropriate jump conditions applied at interfaces to couple the solutions.

The final example considers an efficient high-order accurate algorithm for solving the time-domain Maxwell's equations of electromagnetics. New high-order accurate symmetric difference approximations for curvilinear grids are developed. Compatibility conditions are used to derive stable, high-order accurate centered approximations at material interfaces.

Submitting Author: Bill Henshaw

A Local Corrections Algorithm for Solving Poisson's Equation in Three Dimensions

P. W. McCorquodale, P. Colella, G. T. Balls, S. B. Baden, LBNL (PWM, PC) UCSD (GTB, SBB)

We present a second-order accurate algorithm for solving the free-space Poisson's equation on a locally-refined nested grid hierarchy in three dimensions. Our approach is based on linear superposition of local convolutions of localized charge distribution with the nonlocal coupling represented on coarser grids. The representation of the nonlocal coupling on the local solutions is based on Anderson's Method of Local Corrections and does not require iteration between different resolutions. A distributed-memory parallel implementation of this method is observed to have a computational cost per grid point less than three times that of a standard FFT-based method on a uniform grid of the same resolution, and scales well up to 1024 processors.

Submitting Author: Phil Colella

Numerical Methods of Electromagnetics for Complex Inhomogenous Systems

Wei Cai
Department of Mathematics & Statistics
UNC Charlotte

Fast and accurate computation of electromagnetic phenomena plays an important role in understanding the underlying physics for many complex physical and biological systems, such as lasing in optical fiber lasers, electrostatics forces in solvation model of biomolecules, and irradiation damage in materials under extreme conditions. In this talk, we will present two new algorithm developments with applications in these areas.

Image Charge Approximations of Reaction Fields and FMM for Charges inside a Dielectric Sphere

The reaction field of a charge inside a dielectric sphere, induced by a surrounding dissimilar dielectric medium, has applications in the study of electrostatic forces in the defect evolutions in material under extreme neutron irradiation, and hybrid explicit/implicit solvation models for biomolecules. In both cases, the long range Coulomb interactions have been identified as of primary influence in material's resistance to amorphization under extreme conditions in the first case, and the free energy and the solvation study of biomolecules for the second. We have developed new discrete image charge approximations for the reaction field of a charge inside a dielectric sphere at high accuracy with only 2-3 image charges. Based on this result, we have extended the Fast Multipole Method to calculate the electrostatic interactions of charges inside or outside a dielectric sphere. The resulting O(N) algorithm has applications in computational materials and biology.

A Generalized Discontinuous Galerkin (GDG) Method based on Split Distributions for PDE with Nonsmooth Solutions

To model optical wave propagations in inhomogenous waveguides under the paraxial approximation, we need to solve time dependent Schrödinger equations with nonsmooth solutions as a result of field discontinuities at material interfaces. We will present a new type of discontinuous Galerkin method based on split distributions and their incorporations into the PDEs to account for jumps in solutions and derivatives. Special integration by parts formula for the split distributions is developed. The resulting generalized discontinuous Galerkin (GDG) method will be flexible to handle various types of interface jump conditions (time dependent and nonlinear) with high accuracy and easy to extend to multi-dimensional and other type PDEs with nonsmooth solutions. A full vector GDG-BPM (beam propagation method) will be developed to study gain guided fiber laser for efficient generations of high energy power sources.

Submitting Author: Wei Cai

Systematic Deceleration of Fast Modes in Multi-Scale Systems
Ilya Timofeyev, Univ. of Houston; Eric Vanden-Eijnden, NYU, Courant Institute

Often, the main objective of numerical studies (e.g. geophysical flows) is to analyze the statistical behavior of slowly-evolving large-scale structures (low wavenumbers). On the other hand, typically, small-scale modes are the fastest variables in the system and are the limiting factor for

selecting larger time-step in simulations. We introduce a systematic procedure for "slowing down" the fast degrees of freedom by modifying the triad interactions coefficients in quadratically nonlinear systems. The goal of the methodology is to preserve the statistical properties of large scales while making the

modified equations less stiff and allowing for direct numerical simulations with a larger time-step. This procedure can be carried out consistently with the conservation of energy in the original equations. Truncated Burgers-Hopf system is utilized to illustrate this approach.

Submitting Authors: Ilya Timofeyev

Mimetic Methods for Partial Differential Equations

Mikhail Shashkov
Los Alamos National Laboratory

The need of the office of science programs to approximate the solutions of strongly nonlinear, coupled partial differential equations in complex domains has been a continuous driver for the dual development of supercomputing platforms and for more accurate and efficient numerical algorithms.

Despite the years and magnitude of the effort that has been put into computational science, in many ways the construction of new algorithms for complex problems remains as much of an art than a science. While the accuracy and efficiency of algorithms for idealized problems can be studied with mathematical tools of numerical analysis, increased predictiveness of realistic simulations more typically requires the incorporation of physical principles into the algorithm.

The goal of "Mimetic Methods for Partial Differential Equations" project at LANL (M. Shashkov, PI) is to develop methods that are constructed to embed physical principles. The physical principles that might be incorporated include causality, coordinate invariance, conservation laws, symmetries, asymptotics, and well-posedness.

The basis for mimetic finite difference (MFD) methods is discrete vector and tensor analysis (DVTA). The DVTA addresses discretizations of scalar, vector and tensor functions on different types of the grids, including the discrete analogs of differential operators div, grad, and curl. Many of the physical properties of the solutions, e.g., conservation, symmetry and coordinate invariance, follow directly from the properties of these elementary operators or from the relationships among these operators. We have constructed discrete analogs of divergence, gradient and curl, and have proved discrete analogs of the basic theorems of vector analysis: Gauss', Stokes', orthogonal decomposition (Hodge). One of the examples of practical application of DVTA is construction and investigation of MFD methods for diffusion equation on 3D generalized polyhedral meshes with non-flat faces as well as for meshes with local refinement and non-matching meshes. We have suggested several approaches to enforce discrete maximum principle that includes new non-linear discretizations, approach based on constrained quadratic optimization and postprocessing repair paradigm.

We have made important advances in development of solid mathematical foundation of Arbitrary Lagrangian-Eulerian (ALE) method in which a grid motion could be determined as an

independent degree of freedom. A typical ALE method consist of the following stages: a Lagrangian phase --- we have developed discretization of Lagrangian equations on general polyhedral meshes, developed new artificial viscosity, created new approach for dealing with artificial grid motion. We investigated new closure models needed for the computation with mixed cells containing several materials; a rezoning phase in which mesh is modified --- new methods for rezoning stage include mesh untangling algorithms, volume and surface mesh smoothing methods for general polyhedral meshes, as well error based mesh adaptation; a remapping phase in which the Lagrangian solution is interpolated onto the rezoned grid --- we have developed new conservative interpolation methods on general polyhedral meshes, as well as conservative repair process for enforcing local bounds of the discrete solution.

Submitting Author: Mikhail Shashkov

Adaptive Hybrid Mesh Optimization and Refinement

Ahmed Khamayseh, Oak Ridge National Laboratory ; Computer Science and Mathematics Division, ; Oak Ridge National Laboratory

The accuracy and convergence of computational solutions of mesh-based methods is strongly dependent on the quality of the mesh used. We have developed methods for optimizing meshes that are comprised of elements of arbitrary polygonal and polyhedral type. We present in this research the development of r-h hybrid adaptive meshing technology tailored to application areas relevant to multi-physics modeling and simulation. Solution-based adaptation methods are used to reposition mesh nodes (r-adaptation) or to refine the mesh cells (h-adaptation) to minimize solution error. The numerical methods perform either the r-adaptive mesh optimization or the h-adaptive mesh refinement method on the initial isotropic or anisotropic meshes to maximize the equidistribution of a weighted geometric and/or solution function. We have successfully introduced r-h adaptivity to a least-squares method with spherical harmonics basis functions for the solution of the spherical shallow atmosphere model used in climate forecasting. In addition, application of this technology also covers a wide range of disciplines in computational sciences, most notably, time-dependent multi-physics, multi-scale modeling and simulation.

Submitting Author: Ahmed Khamayseh

Models and Computational Algorithms for Multiphase Magnetohydrodynamics

Roman Samulyak, Brookhaven National Laboratory; Tianshi Lu, Brookhaven National Laboratory; Jian Du, Stony Brook University; James Glimm, Brookhaven National Laboratory / Stony Brook University; Paul Parks, General Atomics

New mathematical models, numerical algorithms, and computational software for the study of magnetohydrodynamics (MHD) of 3D multiphase flows at low magnetic Reynolds numbers in the presence of phase transitions and external energy sources have been developed. The governing system of equations includes a coupled hyperbolic -- elliptic system in geometrically complex, evolving domains and equations for phase transitions and external sources. Numerical algorithms are based on the method of front tracking for material interfaces, high resolution

hyperbolic solvers, the embedded boundary method for the elliptic problem on a dynamic mesh conforming to the interface, new equation of state (EOS) models, and kinetic models for external sources. New algorithms for phase transitions couple balance equations at the phase boundary with nonlinear waves in the interior, enabling the simulation of shock (rarefaction) wave induced mass transfer. Algorithms are multiscale as they use subgrid models and grid patches for processes occurring at different scales. They have been implemented as an MHD extension of FronTier, a hydrodynamic code with free interface support. Meshing components of the algorithms have been developed in collaboration with the SciDAC ITAPS Center.

The models and software have been used for the simulation of mercury targets for future particle accelerators and pellet fueling of thermonuclear reactors with the magnetic confinement of burning plasma (tokamaks). Main results of the accelerator target simulations will be summarized and thermonuclear applications will be discussed in more details. A magnetohydrodynamic numerical model for the pellet ablation in tokamaks has been developed and validated. The main features of the model are the explicit tracking of interfaces, a surface ablation model, a kinetic model for the electron heat flux, an equation of state accounting for atomic processes in the ablation cloud, and new models for conductivity and electric potential. The interaction of the pellet ablation flow with the magnetic field is studied in our work systematically for the first time. This study discovered new features of the pellet ablation and disproved some prevailing expectations on the role of the directional heating on the pellet ablation rate. Current simulations are oriented towards the study of the International Thermonuclear Experimental Reactor (ITER) fueling, pellet induced striation instabilities, and "killer-pellets" for the plasma disruption mitigation.

Submitting Author: Roman Samulyak

Efficient Nonlinear/Non-Gaussian Hydrological Estimation in Composite Media

Juan M Restrepo, University of Arizona; Daniel Tartakovsky, UCSD; Michael Holst, UCSD

Machine learning tools, applied to available data of the subsurface material, enables us to find a partitioning of the hydrological domain with the aim of producing a set of smaller near-linear, near-Gaussian hydrology estimation problems: a path-integral sampling of the posterior probability density is used to monitor and propose partitions. Partitioning is also used in the weak-formulation solver to propose efficient, accurate, and stable solvers. Efficient partitions are found via optimization of the computational parameter space. This talk will present the algorithm and preliminary results of this data assimilation strategy, aimed at improving estimates of large-scale hydrology in stratified domains by combining field data and models in a forward approach.

Submitting Author: Juan Restrepo

Lagrangian Simulations of Combustion

Ahmed F Ghoniem, Fabrice Schegel and Daehyun Wee
Massachusetts Institute of Technology

A Lagrangian methodology is developed for multiscale direct simulations of high Reynolds number reacting flow simulations. New concepts, including the transport element method to deal with transport-chemistry interactions and a multi-purpose tree-code for the evaluation of N-body interaction, are added to our previously developed distribution-based simulation of diffusion and an efficient k-means clustering based strategy for parallel domain decomposition.

In the transport element method, when solving a reactive transport problem of a scalar variable, one discretizes the gradient of the scalar field instead of discretizing the scalar field itself. The evolution of the scalar field is hence determined by solving the corresponding transport equation for its gradients. The advantages of this approach are identical to those ascribed in the vortex methods. Since the gradients are discretized by computational elements confined to a small fraction of the computational domain, one can utilize the computational elements more efficiently. The concept of transport elements is applied to the three-dimensional simulations of multi-physics problems, where the flow field and the scalar field are coupled by various interaction mechanisms, i.e., baroclinicity, thermal expansion, and chemical reactions. The scheme is equipped with a multi-purpose adaptive tree-code optimized to run on a massively parallel architecture.

Lagrangian simulations require the evaluation of various N-body interactions: One needs to compute the vortical velocity, $\mathbf{u}_\omega$, the expansion velocity, and the conserved scalar field, s, from information on the vorticity, expansion sources and the scale gradient, $\mathbf{g}$, respectively. Moreover, the evolution of vorticity and the scalar gradients relies on the velocity gradient field, and it is also necessary to be able to compute $\nabla\mathbf{u}_\omega$ and $\nabla\mathbf{u}_\varepsilon$ simultaneously. All of these interactions are effectively computed from one single Taylor series expansion of a potential, and recombination of its Taylor coefficients for the quantity of interest. Particle-particle interactions are replaced by particle-cluster interactions, which can be efficiently computed from the Taylor coefficients obtained. Finally, the extension of the current tree-code from a low-order algebraic kernel to a high-order algebraic kernel is presented. The use of a high-order algebraic kernel guarantees efficiency and better convergence characteristics.

Particular attention has been paid to the development of Lagrangian tools for multi-physics problems, including flows with buoyancy and transport-controlled combustion. Two examples are shown next. In Figure 1, the evolution of the reaction surface of a diffusion-controlled combustion of a methane sphere is presented. The vortex ring generated through the baroclinic interaction deforms the reaction surface into a mushroom shape. Figure 2 shows the vorticity field of a jet issuing in a cross flow, demonstrating the complex transition of the vorticity structure as the dynamics undergoes several transitions from large coherent scales to small well distributed scales.

t = 3.0   t = 4.5

Figure 1. Evolution of the reaction surface of a diffusion-controlled combustion of a methane

sphere.
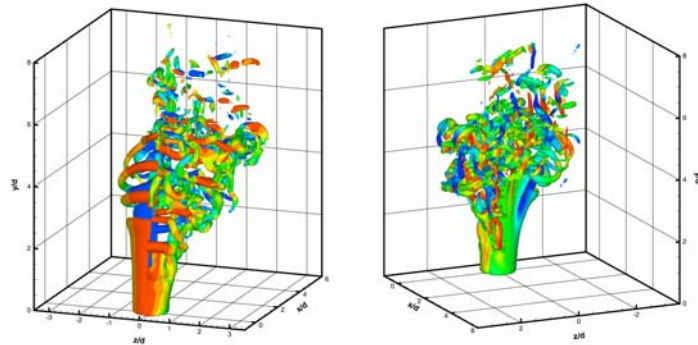
Figure 2. Vorticity fields of transverse jets. Two jets with different nozzle boundary conditions

are simultaneously shown to emphasize the impact of boundary conditions.

Submitting Author: A. Ghoniem

Optimization Abstracts

What's New in Mixed-Integer Nonlinear Optimization

Sven Leyffer, Argonne National Laboratory

Many important design applications can be formulated as mixed-integer nonlinear programs (MINLPs). Applications of interest to DOE include the process design for FutureGen (the planned zero emission fossil fuel plant), blackout prevention in the national power grid, and the optimization of the reloading operation of nuclear reactor cores. We outline these and other applications, present a survey of recent advances in the solution of MINLPs based on branch-and-cut, and outline future challenges arising out of applications of interest to DOE.

Submitting Author: Sven Leyffer

Optimization: Preconditioners, PDE Constraints and Equilibrium

Jorge Nocedal, Northwestern University

This talk surveys recent advances in large-scale nonlinear optimization. We begin by considering the design of preconditioners for iterative methods used in constrained optimization. Although constraint preconditioners have been highly successful, important open questions remain, particularly in the context of interior-point methods and when many equality constraints are present in the problem. We also discuss general (non-constraint) preconditioners, the effects of roundoff errors and iterative refinement.

In the second part of the talk, we consider PDE-constrained optimization and some of its main challenges from the point of view of (optimization) algorithm design. We study the advantages and disadvantages of full-space primal-dual approaches compared with step decomposition techniques. Then we consider the issue of inexactness in the step computation. After reviewing advances made by other research teams, we give some answers to the following question: when applying an iterative solver to the full KKT system, when is an inexact solution adequate in the sense that it guarantees progress toward a minimizer, even in the nonconvex case? Surprisingly, it is possible to provide satisfactory answers even when negative curvature is present in the Hessian of the Lagrangian. This research is motivated by PDE optimization problems, but is also relevant to classification algorithms.

We conclude the talk by discussing the interplay between equilibrium and optimization problems. We outline some promising optimization approaches for the solution of a strategic bidding problem in electricity markets. We give numerical results for both deterministic and a stochastic formulations. Although we are able to produce satisfactory results quite rapidly, guaranteeing that the problem has been solved to optimality remains a formidable challenge in the general case.

Submitting Author: Jorge Nocedal

Speeding the Training of Support Vector Machines and Solution of Quadratic Programs

Jin Hyuk Jung, University of Maryland; Dianne P. O'Leary, University of Maryland; Andre Tits, University of Maryland

Data classification is a fundamental task in science and engineering. For example, given data gathered about a patient's tumor, we might need to decide whether the tumor is malignant or benign. Ideally, we would like to determine a mathematical function whose evaluation would indicate the classification of the tumor. Linear discriminant analysis provides one such function, but functions more general than a separating hyperplane are needed in many applications.

Support vector machines (SVMs) provide a means to classify data into two groups (positive and negative) using criteria more descriptive than separating hyperplanes. SVMs are trained using a large set of positive and negative examples. Classifiers such as neural networks are trained by an iterative process of presenting examples and adjusting network weights until convergence. In contrast, the training of an SVM is accomplished by solving a single quadratic programming problem whose size is determined by the number of examples and the number of parameters in the classifier. This simple training regime is a major advantage of the SVM framework.

These quadratic programming problems can, however, be quite large. In this work we use two approaches to improve computational efficiency. First, we apply an adaptive constraint reduction method in an interior point method for solving the quadratic programming problem. This has the effect of allowing later iterations to focus on only a few of the example datapoints. Second, specific to SVM training, we consider clustering the data and initially training on a small number of examples drawn from each cluster. This potentially has the effect of reducing computation time in the early iterations.

We discuss our algorithm and its convergence theory and illustrate its performance on a variety of examples.

Submitting Author: Dianne O'Leary

Inverse problems for systems constrained by reaction diffusion equations

Santi S. Adavani, University of Pennsylvania; George Biros, University of Pennsylvania; Shravan K. Veerapaneni, University of Pennsylvania

In the first part of this talk we will discuss a method for the fast and accurate computation of distributed (or volume) heat potentials. The distributed source is assumed to be given in terms of piecewise space-time Chebyshev polynomials. We discretize uniformly in time, whereas in space the polynomials are defined on the leaf nodes of a quadtree data structure. The quadtree can vary at each time step. We combine a product integration rule with fast algorithms (fast heat potentials, nonuniform FFT, fast Gauss transform) to obtain a high-order accurate method with optimal complexity. If the input contains $q3$ polynomial coefficients at M leaf nodes and N time steps, our method requires $O(M N \log M)$ work to evaluate the heat potential at arbitrary MN space-time target locations. The overall convergence rate of the method is of order q. We present numerical

experiments for q = 4, 8, and 16, and we verify the theoretical convergence rate of the method. When the solution is sufficiently smooth, the 16th-order variant results in significant computational savings, even in the case in which we require only a few digits of accuracy.

In the second part of my talk will present a multigrid algorithm for solution of distributed parameter estimation problems with variable coefficients. The main feature of the method is that it is mesh-independent even in the case of zero regularization (when the data is in the range of the inversion operator). This makes the method algorithmically robust to the value of the regularization parameter. The method is based on a reduced space formulation in which we iterate in the inversion parameter space. We use a full multigrid scheme with a spectrally filtered stationary approximate-Hessian stationary smoother and standard intergrid transfer operators. We use Fourier analysis to estimate the overall performance of the scheme, and we provide numerical experiments that demonstrate the effectiveness of the method for different diffusion coefficients and regularization parameters. We have observed mesh-independent convergence factors resulting in O(N) complexity, where N is the number of state variables. We present and compare two smoothers, one based on simple spectral thresholding and one based on reduced Hessian approximation, based on inexact domain-decomposition approximations of the forward solvers. The latter approach is found to be robust, i.e., it results in a mesh independent convergence for all diffusion coefficient and regularization parameter values. Although the spectral filtering preconditioner has negligible computational cost it does not converge in all the cases unlike pointwise preconditioner, e.g., it only works for large values of the regularization parameter. We also discuss the overall multigrid by using a 2-step iterative smoother, and Fourier transforms to do the intergrid transfers. We have shown that this scheme is extends to nonlinear reaction problems. Even in case of partial observations, these schemes result in near optimal algorithmic complexity.

Submitting Author: George Biros

New Heuristic Techniques for General Mixed-Integer Programs

Robert Carr, Sandia National Laboratories; Jonathan Eckstein, Rutgers University; Cynthia Phillips, Sandia National Laboratories; Jean-Paul Watson, Sandia National Laboratories

Mixed-integer linear programming (MIP) is a general technology for solving combinatorial optimization problems exactly. MIP is the optimization of a linear function subject to linear and integrality constraints. MIPs naturally represent resource allocation problems. Thus they are a workhorse technology for decision support, such as scheduling, logistics, manufacturing optimization, and sensor placement, and for the study of natural systems such as bioinformatics. Ali Pinar of LBNL has identified several new potential science applications of MIP include scheduling telescope time, groundwater flow analysis, and analysis of particle collisions in supercolliders.

PICO (Parallel Integer and Combinatorial Optimizer) is a massively-parallel MIP solver developed, and under active refinement, at Sandia National Laboratories and Rutgers. Like all current general-purpose MIP solvers, PICO computes (approximately) optimal solutions by intelligent search based on branch-and-bound and branch-and-cut. Finding good feasible solutions early in the process provides a reasonable approximation if the user or the system terminates the search early. It also speeds the search by enabling stronger search-space pruning.

In this talk, we will describe new heuristic methods for finding feasible solutions for general MIPs. In particular, we will describe methods developed by members of the PICO team that are particularly well suited to exploit the large number of extra processors available early in the search.

Mixed-integer programming is NP-complete (formally intractable) in the worst case, but in practice, general solvers can find (near) optimal solutions to large instances. Many MIP heuristics begin by relaxing the integrality contraints to obtain a linear program (LP). This problem is theoretically tractable and usually practically solvable. The solution to this LP relaxation provides a numerical lower bound for a minimization problem, and the structure of the fractional solution can sometimes guide heuristics.

This year we developed the fractional decomposition tree (FDT) heuristic. The FDT heuristic runs in time polynomial in the input size and it is guaranteed to find a feasible solution to an integer programming instance provided the instance has bounded integrality gap. For a minimization problem, the integrality gap is the largest ratio of the value of the optimal MIP solution to its linear-programming (LP) relaxation. At first glance, it might appear that we are claiming to solve an NP-hard problem, but we will explain in the talk why this is not the case. The FDT method is based on a previous theoretical result relating convex combinations of feasible MIP solutions to a scaled solution to the LP relaxation. Previous provably-good approximation algorithms for specific combinatorial optimization problems have used problem-specific methods for decomposing solutions to LP relaxations. The FDT is a first attempt to automate this process for general MIPs. Because the FDT uses many linear-programming computations, it can exploit extra processors. In fact there are two variants that exploit parallelism in different ways. We will compare the FDT to other MIP heuristics including randomized rounding, feasibility pump, and the parallel cut-pivot-dive heuristic due to Nediak and Eckstein.This talk will be self-contained; we assume no prior knowledge of MIP.

Submitting Author: Cindy Phillips

Orbital Branching

JEFF LINDEROTH, JAMES OSTROWSKI

Department of Industrial and Systems Engineering,
Lehigh University
200 W. Packer Ave. Bethlehem, PA 18015, USA

`jtl3@lehigh.edu` · `jao204@lehigh.edu`

FABRIZIO ROSSI, STEFANO SMRIGLIO

Dipartimento di Informatica,
Università di L'Aquila
Via Vetoio I-67010 Coppito (AQ), Italy

`rossi@di.univaq.it` · `smriglio@di.univaq.it`

March 30, 2007

Many important optimization models involve a combination of discrete and continuous decision variables and logical relations between these decision variables. These models are typically called (linear) integer programs (IP), and can be stated mathematically as

$$\min_{x \in \mathbb{R}^n_+ \times \mathbb{Z}^p_+} \{c^T x \mid Ax \geq b\}. \tag{IP}$$

(IP) is NP-Hard, but large instances are successfully solved in practice by a branch-and-bound tree search method that relies on linear programing relaxations of (IP) to perform the bounding. In fact, powerful and reliable commercial software systems, such as CPLEX and XPRESS-MP, have led the way toward broad-based acceptance of integer programs as an important tool in the decision-making and design process. However, many important (as yet unsolved) integer programs suffer from "the curse of symmetry," a concept that can be formalized as follows. Let Π^n be the set of permutations of $\{1, 2, \ldots, n\}$, and $A(\pi, \sigma)$ be the matrix A when the columns are permuted by $\pi \in \Pi^n$ and the rows are permuted by $\sigma \in \Pi^m$. The *symmetry group* of (IP) is

$$\mathcal{G}(\text{IP}) \stackrel{\text{def}}{=} \{\pi \in \Pi^n \mid \pi(c) = c \text{ and } \exists \sigma \in \Pi^m \text{ s.t. } \sigma(b) = b, A(\pi, \sigma) = A\}.$$

If $\pi \in \mathcal{G}(\text{IP})$, then π establishes a relabeling of the columns that leaves the set of feasible solutions to (IP) unchanged. If $|\mathcal{G}(\text{IP})|$ is large, then (IP) is said to be *highly-symmetric*, and the branch-

1

ing operation performed by standard algorithm has little or no effect on the relaxation of (IP), and branch-and-bound is unable to make progress towards a solution. Integer programs coming from important applications such as error correcting codes, covering codes, covering designs, and Steiner triple systems, are all highly symmetric.

Recently, we have shown how to *exploit* the symmetries present in IP by explicitly computing the groups of variables that are equivalent with respect to the symmetry remaining in the problem after branching These groups of equivalent variables, called *orbits*, are used to create a valid partitioning of the feasible region which significantly reduces the effects of symmetry while still allowing flexibility in the branching scheme. We also show how to exploit the symmetries present in $\mathcal{G}(\text{IP})$ to fix variables throughout the branch-and-bound tree. Our new method, *orbital branching*, can easily be incorporated into standard IP software. Through an empirical study on a test suite of symmetric integer programs, we show that the method is orders of magnitude more effective than a state-of-the-art commercial solver. Table 1 highlights the effectiveness of the method by reporting the CPU time (in seconds) and number of nodes in the enumeration tree to solve symmetric instances of (IP).

		Orbital Branching		CPLEX v10.1	
Application	Size	Time	Nodes	Time	Nodes
Error Correcting Code	$\{0,1\}^8, d = 3$	2	25	391	32077
Error Correcting Code	$\{0,1\}^9, d = 3$	176	539	**fail**	488136
Error Correcting Code	$\{0,1\}^{10}, d = 5$	306	11	1245	1584
Covering Code	$\{0,1,2\}^5, d = 1$	66	935	1150	54018
Covering Design	(9,5,4)	22	401	9	1514
Covering Design	(10,5,3)	50	745	937	99145
Covering Design	(10,5,4)	2	11	**fail**	239266
Covering Design	(10,7,5)	292	377	141	10278
Steiner Triple System	27	1	71	0	1647
Steiner Triple System	45	12	5559	24	51078
Steiner Triple System	81	202	10691	**fail**	12572533

Table 1: Comparison of Orbital Branching and CPLEX v10.1

2

Tuesday Poster Abstracts

New Eigensolvers and Preconditioners for Large Scale Nanoscience Simulations

A. Canning, O. Marques, C. Voemel, L-W Wang , Lawrence Berkeley National Laboratory ; S. Tomov, J. Langou and J. Dongarra, University of Tennessee; Knoxville

We present results for new iterative eigensolvers based on conjugate gradients and Jacobi-Davidson in the context of semi-empirical plane wave electronic structure calculations. These new methods give significant speedup over existing conjugate gradient methods used in electronic structure calculations. The new methods will be demonstrated for CdSe quantum dots as well as quantum wires (single electron devices) constructed from layers of InP and InAs. These systems are studied in the context of a semi-empirical potential where we typically solve for a few states around the gap allowing us to study large scale nanosystems. The parallelization of this approach will also be discussed as well as scaling results to large processor counts. We will also present results for a new preconditioner based on bulk-band states. This new preconditioner is based on the observation that in the interior of large quantum dots the eigenstates are well approximated by the bulk eigenstates. Therefore the bulk eigenstates which are cheap to calculate can be used as a preconditioner for the Quantum Dot as well as for the starting vector for the eigensolver.

Submitting Author: Andrew Canning

Density matrix methods as eigensolver replacements in density functional theory

Richard P. Muller, Sandia National Laboratories; Richard B. Lehoucq, Sandia National Laboratories

Density functional theory (DFT) is a quantum mechanical technique for computing energies and related properties of molecular and material systems. The rate determining step in these calculations for large systems is solving the eigenvalues/vectors of the Hamiltonian matrix, from which the density matrix may be formed. Our work has investigated techniques that form the density matrix directly from the Hamiltonian, which promises to reduce the overall scaling, and make sparsity or parallel computer architectures easier to utilize. I will present preliminary results of our work.

Submitting Author: Richard Muller

Inverse Problem in Seismic Imaging

Maria Cameron*[1], Sergey Fomel[2], James Sethian[1]
[1] Department of Mathematics, Lawrence Berkeley National Laboratory,
[2] Bureau of Economic Geology, U of Texas, Austin

Imaging of Earth regions with nonhorisontal subsurface structures and laterally varying sound speed (seismic velocity) is of great importance for the oil industry. The ultimate goal of producing an accurate map of the Earth's interior from the external soundings and measurements is exceedingly difficult due to noisy data, artifacts, shadow zones and ill-posedness, and currently out of reach.

Time migration is a robust and efficient type of seismic imaging which does not require the knowledge of seismic velocities. However, it is adequate only for regions with mild lateral velocity variation. Moreover, time migrated images are obtained in somewhat unintuitive time coordinates.

Depth migration is a type of seismic imaging which is adequate for regions with arbitrary lateral velocity variation and produces images in regular Cartesian (depth) coordinates. However, its implementation requires knowledge of seismic velocity.

Time migration has an additional output: time migration velocities given in time coordinates.

The goal of this work is to build an efficient algorithm to
(a) estimate seismic velocities from time migration velocities (then one can use these seismic velocities for obtaining depth migrated earth images);
(b) convert the time migrated images directly to depth coordinates and correct the distortions due to the lateral velocity variations.

We established theoretical relations between time migration velocities and seismic velocities in 2D and 3D using paraxial ray tracing theory: the seismic velocities relate to the time migration velocities through the geometrical spreading of the so called image rays.

We formulated an inverse problem of finding seismic velocities from time migration velocities and developed two numerical approaches which involve three algorithms for solving it.

Our approaches include Dijkstra-like Hamilton-Jacobi solvers for first arrival Eikonal equations and techniques for data smoothing.

The first algorithm is based on the ray tracing. It produces seismic velocities in time coordinates from time migration velocities. Time evolution equations of the geometrical spreading involve the second derivatives of the unknown velocity, which makes the numerical solution of the problem very difficult.

The second algorithm is time-to-depth conversion. It converts velocities given in time coordinates to depth coordinates. This algorithm is based on J. Sethian's fast marching method designed for solving of Eikonal equation with known right-hand side. Our algorithm solves Eikonal equation with unknown right-hand side, which makes this extension nontrivial, coupled with an orthogonality relation.

These two algorithms constitute the first approach.

The third algorithm is based on the level set method. It computes seismic velocities in depth coordinates directly from time migration velocities. It involves the time-to-depth conversion as an essential part of its time cycle, and implements the techniques for data smoothing and second derivatives estimating used in the ray tracing algorithm. Then one can use the fast marching method to compute the transition matrices from time to depth coordinates. This is our second approach.

We tested these approaches on synthetic data examples and applied them to a field data example. We demonstrated that our algorithms give a significantly better estimate of seismic velocities than the Dix inversion which is the standard approach.

Submitting Author: Maria Cameron

Spectral Methods for the Analysis of Stochastic Dynamical Systems

Bert Debusschere*, Habib Najm*, Dongjin Kim[†], Olivier Le Maître[‡]
* Sandia National Laboratories, Livermore, CA
[†] **Georgetown University, Washington, DC**
[‡] **Université d'Evry Val d'Essonne, Evry, France**

Due to recent advances in biophysics and nanoscale science, many molecular-level system models have now become available. Given the small scales and limited numbers of particles involved, these systems generally exhibit stochastic dynamics. Examples are found in the stochastic biochemical reaction networks that underlie Gene Regulatory Networks or in surface reaction chemistry. Even though stochastic behavior is an integral feature of the system dynamics, there is currently a lack of robust analysis methods for such systems. This talk presents our development of spectral methods for sensitivity analysis and reduced order modeling of stochastic dynamical systems that are governed by the chemical master equation. The sensitivity analysis relies on polynomial chaos expansions, which capture the nonlinear behavior of the system dynamics in response to finite-sized parametric perturbations. Reduced order models are obtained from Karhunen-Loève decompositions. We show applications of these methods to a bi-stable model system and a viral infection model.

Submitting Author: Bert Debusschere

A network-particle computational method for dense suspensions

Alexander Panchenko, Department of Mathematics, Washington State University;
Alexander Tartakovsky, PNNL)

We propose a new method for computer simulation of dense suspensions. In the regions of high particle concentration, fluid flow is dominated by the local flows in the narrow inter-particle gaps. Asymptotic analysis of these local flows produces good approximate formulas for the inter-particle forces. The forces are calculated analytically as functions of particle velocities. Our approach leads to an efficient discretization of the problem, since there is no need to simulate

fluid motion directly. As a results, our method is capable of handling millions of particles while Stokesian dynamics (the most popular direct simulation method) typically deals with no more than a few thousands. In the regions of lower concentrations, we supplement asymptotic formulas with the smooth particle hydrodynamics . This is done in order to resolve large scale fluid motions which are not captured by our lubrication-type asymptotics.

Submitting Author: Alexander Panchenko

New Results from "VEGA" Vulnerability of Electric Grids Analyzer

Javier Salmeron and Kevin Wood
Operations Research Department,
Naval Postgraduate School, Monterey, CA 93943-5001

This talk describes new optimization methods for analyzing the vulnerability of an electric power grid to attack, and for planning defenses. We focus on coordinated, terrorist attacks on sets of grid components, but our technology also applies to more standard security analyses (e.g., N−1, N−k). All models and algorithms are integrated into our decision-support system called "VEGA" (Vulnerability of Electric Grids Analyzer), which includes a database, a graphical user interface and optimization modules; the talk includes a demonstration.

Our ultimate goal is to identify sets of components in a power grid which, if protected, would minimize the damage from a coordinated attack on unprotected components. But, to protect a grid, we must first understand its vulnerabilities, so we begin by describing an interdiction model for electric power grids. This "attacker-defender" model (AD) is formulated as a special type of bilevel optimization problem: Using limited offensive resources, an "attacker" (e.g., a terrorist group) seeks to destroy a set of grid components that inflicts maximum damage; damage is measured in terms of load shedding or its cost, which is minimized after the attack by an optimizing "defender" (e.g., the system operator). In fact, the defender solves a set of optimal power flow models, accounting for load-duration curves and restoration time of destroyed components. (Our current models cover medium- and long-term disruptions, although collaborative research efforts are underway to incorporate cascading failures.)

Using AD as a submodel, we then present a trilevel model for active defense of the grid; this is a "defender-attacker-defender" model (DAD). In the first level of DAD, the defender applies limited defensive resources to protect components against attack. The solution to an AD evaluates the effectiveness of that defensive plan.

AD and DAD cannot be solved as standard mathematical programs, at least not if we wish to analyze real-word grids. Therefore, we describe and demonstrate specialized, primal decomposition methods for their solution. We can solve AD on a U.S. grid, to near-optimality, in 1.5 hours on a personal computer. We cannot yet solve DAD on such large problems, however, so we present computational results for smaller test grids, and describe new research results that are getting us closer to the ultimate goal. For instance, we show that approximating the ADsubmodel in DAD as a minimum-cost-flow network-interdiction problem gives very fast and accurate "heuristic guidance" to the exact decomposition.

We also briefly describe collaborative work with Dr. Ross Baldick at the University of Texas at Austin. He and his students are developing a cascading-analysis tool that will be incorporated into

VEGA. This module will enable a more accurate assessment of the short-term effects that attacks (or even random events) can have on a grid.

Submitting Author: Kevin Wood

A multiscale coupling of continuum elasto-viscoplasticity and discrete dislocation dynamic

Firas Akasheh, Washington State University; Hussein M. Zbib, Washington State University

It is known that the origin of plasticity in crystals lies in the motion and interaction of dislocations (crystalline line defects). Although dislocation mechanics is well understood, the complexity and multiplicity of dislocation interactions among themselves, applied stresses, other defects, as well as surfaces and interfaces makes dislocation motion intractable in a continuum crystal plasticity framework leading to current models being mostly phenomenological in nature. The success of such models in addressing different boundary value problems has been limited and problem-dependant. To overcome this drawback, microscale discrete dislocation dynamics (DD) analysis is coupled with finite element continuum analysis of elasto-viscoplasticity problems. Based on the actual dislocation content and motion of the underlying dislocation structure, explicit calculations of the internal stress field and plastic strain are made and homogenized over the finite elements. Those quantities replace any phenomenological constitutive relationships for the plastic strain and back stress used at the continuum scale. Furthermore, through the superposition principle, this coupled framework allows the treatment of different boundary conditions like free surfaces. In this work, the DD framework is presented along with the coupling scheme and the numerical solution algorithms. Applications pertaining to the deformation in nanoscale multilayered metallic composites and bending of micro-sized beams are also presented to demonstrate the capability of the coupled model.

Submitting Author: Firas Akasheh

Multilevel Upscaling in Multiscale Modeling and Parameter Estimation

(1)David Moulton (Co-PI), (1,2)Daniel Tartakovsky (Co-PI), (3)Scott MacLachlan, (4)Ethan Coon

(1)Mathematical Modeling and Analysis Los Alamos National Laboratory, Los Alamos, NM (2)Department of Mechanical and AeroSpace Engineering University of San Diego, San Diego, CA (3)Applied Physics and Applied Mathematics, Columbia University, New York, NY (4)Delft Institute of Applied Mathematics Delft University of Technology, The Netherlands

The predictive capability that is needed for multiscale simulations to drive scientific discovery and guide policy decisions remains elusive for many DOE applications, including carbon sequestration, aquifer assessment and protection, and nuclear waste disposal. This problem arises because fully resolved simulations are intractable, the mathematical model may be inaccurate or

incomplete, and its parameters may be poorly characterized by sparse noisy data at disparate lengths scales. Developing efficient and robust techniques that facilitate the accurate quantification of uncertainty in these situations is the objective of the "Predictability with Stochastic Partial Differential Equations" project at LANL (D. Tartakovsky, Co-PI).

In this discussion we focus on the class of problems in which the fine-scale continuum model is well understood, but the medium may be highly heterogeneous and it may be represented stochastically. Here, homogenization (or upscaling) techniques are necessary to develop computationally feasible models on scales coarser than the variation of the coefficients of the continuum model. The accuracy of such techniques depends dramatically on assumptions that underlie the particular upscaling methodology used. For example, decoupling of fine- and coarse-scale effects in the underlying medium may utilize artificial internal boundary conditions on sub-cell problems. Such assumptions, however, may be at odds with the true, fine-scale solution, yielding coarse-scale errors that may be unbounded.

To address this problem we step beyond the traditional view of multigrid methods as iterative solvers or preconditioners and explore their potential as a rich multilevel variational setting for the practical treatment of multiscale problems. In particular, the development of robust efficient multilevel solvers has naturally led to the development of general multiscale concepts, such as operator-induced variational coarsening. Such an approach implicitly treats multiscale aspects of the fine-scale model as it generates a sequence of coarser representations and, thus, implicitly creates multiscale basis functions and coarse-scale closures. We review results for the upscaling of models and parameters in single-phase saturated flow through porous media that clearly demonstrate the potential flexibility and efficiency of this approach. We then highlight recent work on multilevel upscaling of two-phase flow (oil and water).

To further demonstrate the flexibility of developing a robust multilevel modeling framework we examine a Bayesian approach to the related inverse problem of impedance tomography. In particular, we leverage a recent advance in Markov Chain Monte Carlo (MCMC) sampling that rigorously accommodates approximations in the forward map through a delayed acceptance strategy. Since the overall efficiency of the inversion is strongly influenced by the accuracy of the approximation, this is particularly well suited for our multilevel approach.

Submitting Author: David Moulton

Coarse-graining and microscopic reconstruction in simulations of stochastic systems

Petr Plechac, Oak Ridge National Laboratory

We discuss some mathematical and computational issues arising in problems where the microscopic Markov process is approximated by a hierarchy of coarse-grained processes. We present mathematical tools developed for error control in microscopic simulations using the coarse-grained stochastic processes and stochastic reconstruction of microscopic scales. Derived a posteriori error control allows us to design adaptive coarse-graining of the configuration space. We shall briefly discuss how the methods of statistical mechanics (e.g., cluster expansions) lead to improved sampling schemes with compressed long-range interaction kernels. Such compression of the underlying equilibrium measure makes parallel implementation of the simulation algorithm easier. On specific examples of lattice as well as off-lattice dynamics

(simulations of spin systems or polymers) we demonstrate that computational implementation of constructed coarse-grained approximation leads to significant CPU speed up of simulations.

Submitting Author: Petr Plechac

Convergence Analysis for Takizuka & Abel and Nanbu's Collision Models

CM Wang, University of California, Los Angeles; Tungyou Lin, University of California, Los Angeles; Russel Caflisch, University of California, Los Angeles; Bruce Cohen, Lawrence Livermore National Laboratory; Andris Dimits, Lawrence Livermore National Laboratory

A plasma consists of a large number of charged particles. If a plasma is highly collisional, the plasma kinetics can be approximated by a fluid description. On the other hand, if a plasma is collisionless, each particle interacts with the rest of the plasma through long-range electromagnetic fields. In the intermediate regime, collisional effects have to be included specifically to provide an adequate description of plasma kinetics. One significant example is the edge plasmas in a confinement fusion device [2]. A kinetic approach is essential for satisfactory physical modeling and numerical simulations.

One of the earliest and most influential Monte Carlo binary collision model was proposed by T. Takizuka & H. Abe in 1977 [5]. In their method, randomly chosen pairs of particles undergo binary collisions. The resulting scattering angle is sampled through a Gaussian distribution to compute the change in velocities.

Nanbu proposed a new Monte Carlo binary collision model in 1997 [1], [4]. Nanbu's method computes the cumulative scattering angle for many small binary deflections. Successive small angles are grouped into one single collision angle.

The two methods have been widely used in the plasma physics community. We believe their models have potential to be used to extend a hybrid method for rarified gas to collisional plasmas [3]. For this reason, we are interested in performing convergence analysis to compute the errors and derive the orders of accuracy for both methods to quantify their performance.

Our poster will present computational results for a spatially homogeneous plasma. We simulate the relaxation of anisotropic temperatures over time due to collisions [4], [5] and [6], using the results to evaluate the accuracy and efficiency of these two methods. We test both electron-electron and electron-ion collision cases. We will present the deterministic errors and random errors at different number of particles and time step for each method. We will also compare the accuracy of the two methods.

[1] A. V. Bobylev and K. Nanbu, Physical Review E, Vol 61, No. 4, p. 4576-4582 (2000)

[2] O.V. Batishchev, M. M. Shoucri, A. A. Batishcheva and I. P. Shkarofsky. J. Plasma Physics, Vol 61, part 2, p. 347-364 (1999)

[3] R. E. Caflisch and L. Pareschi. J. Compt. Phys. 154, 96 (1999)

[4] K. Nanbu, Phys. Rev. E. 55 (1997).

[5] T. Takizuka and H. Abe, J. Comp. Phys. 25 (1977)

[6] B. A. Trubnikov, Rev. Plasma Physics (Consultant Bureau, New York, 1965), Vol. 1, p. 105

Submitting Author: Chiaming Wang

Multiscale Model of Morphogenesis

Mark Alber, Department of Mathematics, University of Notre Dame
Scott Christley, Department of Computer Science and Engineering,
University of Notre Dame

Cells of the embryonic vertebrate limb in high-density culture undergo chondrogenic pattern formation, which results in the production of regularly spaced islands of cartilage similar to the cartilage primordia of the developing limb skeleton. In this talk we describe a discrete, stochastic model for the behavior of limb bud precartilage mesenchymal cells in vitro [1]. The model uses a biologically motivated reactiondiffusion process and cell-matrix adhesion (haptotaxis) as the bases of chondrogenic pattern formation, whereby the biochemically distinct condensing cells, as well as the size, number, and arrangement of the multicellular condensations, are generated in a self-organizing fashion. Improving on an earlier lattice-gas representation of the same process, it is multiscale (i.e., cell and molecular dynamics occur on distinct scales), and the cells are represented as spatially extended objects that can change their shape.

In the development of multiscale biological models it is crucial to establish a connection between discrete microscopic or mesoscopic stochastic models and macroscopic continuous descriptions based on cellular density. In this talk a continuous limit of a two-dimensional Cellular Potts Model (CPM) with excluded volume will be demonstrated, describing cells moving in a medium and reacting to each other through both direct contact and long range chemotaxis [2]. The continuous macroscopic model is obtained as a Fokker-Planck equation describing evolution of the cell probability density function. All coefficients of the general macroscopic model are derived from parameters of the CPM and a very good agreement is demonstrated between CPM Monte Carlo simulations and numerical solution of the macroscopic model.

1. Christley, S., Alber, M.S. , Newman, S.A. [2007], Patterns of Mesenchymal Condensation in a Multiscale, Discrete Stochastic Model, PLoS Computational Biology, 3, 4, e76.

2. Alber, M., Chen, N., Glimm, T., and P. Lushnikov [2006], Multiscale dynamics of biological cells with chemotactic interactions: From a discrete stochastic model to a continuous description, Phys. Rev. E. 73 051901.

Submitting Author: Mark Alber

Fast Marching Methods for the continuous Traveling Salesman Problem

June Andrews, Department of Mathematics, University of California Berkeley; James Sethian, Department of Mathematics, University of California Berkeley

We consider a problem in which we are given a domain, a cost function which depends on position at each point in the domain, and a subset of points ("cities") in the domain. The goal is to determine the cheapest closed path that visits each city in the domain once. This can be thought of as a version of the traveling salesman problem, in which an underlying known metric determines the cost of moving through each point of the domain, but in which the actual shortest path between cities is unknown at the outset. We describe algorithms for both a heuristic and an optimal solution to this problem. The complexity of the heuristic algorithm is at worst case M.N log N, where M is the number of cities, and N the size of the computational mesh used to approximate the solutions to the shortest paths problems. The average runtime of the heuristic algorithm is linear in the number of cities and $O(N \log N)$ in the size N of the mesh.

Submitting Author: June Andrews

Statistical Mechanics and Complex Dynamics in Rapidly Rotating Bounded Flows

Chjan Lim
Rensselaer Polytechnic Institute

This is a short review of preliminary results in a new nonlocal / non-Hamiltonian approach to the statistical mechanics of stationary, nearly inviscid quasi-2D bounded flows. Early rigorous and numerical analysis of quasi-2D turbulence focused on periodic and unbounded 3d flows in a rapidly rotating setting in the vanishing viscosity and Rossby number limits. Recent experiments - Eindhoven, Sommeria, Tabeling, Swinney - clearly showed that boundaries and angular momentum transfer play important roles in quasi-2D turbulence. The dual cascades - K41 and KLB - first introduced for periodic flows have to be modified to fit bounded flows - (i) impose finite enstrophy since 2D flows generally do not allow entrophy blow-up in finite time; (ii) this implies zero enstrophy dissipation and thus no forward cascade; (iii) with boundaries, the energy cascade to large scales can be realized but is modified to stop either at the largest (domain-size) scale OR at a smaller scale before the ground state- in both cases Ekman dissipation remains $O(1)$ for nearly inviscid flows; (iv) boundaries further change the transfer of fluid angular momentum by a complex torque mechanism that need not be dissipative.

While rigorous analysis - Chemin, BMN, Kato, Vishik, Kelliher - is making some headway, it is dependent on Kato-like conditions which does not allow anomalous BL energy dissipation AND it cannot answer the BASIC ?: at which large scale (and under ?conditions) does the quasi-2D turbulent bounded flow stop its inverse energy cascade? Only cascades to the largest scale will result in angular momentum transfer.

First steps to answer this Question, we introduce a statistical mechanics with the properties (a) it is based on path-integrals where the energy functional is explicitly non-Hamiltonian to allow nonconservation of fluid angular momentum; (b) this action is based on a non-local or weak formulation corresponding to the weak solns of the NSE that support energy cascades; (c) to model a stationary energy cascade – interior injection and Ekman dissipation - and angular

momentum transfer, we propose a canonical ensemble in energy; (d) zero enstrophy dissipation / fixed finite enstrophy in (i) leads to a microcanonical constraint on enstrophy.

Exact solns and MC results based on spin glass technique show that angular momentum transfer accompany the inverse energy cascade only at the lowest and highest energies, with interesting dichotomy between the prograde and retrograde states in relation to the planetary spin and enstrophy. For all intermediate energies, the ground mode or largest scale is not occupied. This phenomena occur in the form of a BEC phase transition. Our work provides an answer to the BASIC ? and also suggests conditions for the Ekman layer instability as a mechanism for intermittent dissolution of the BECondensed state.

Submitting Author C. Lim

Towards 1,000,000 Atom First Principles Electronic Structure Simulation Code

William A. Shelton, Oak Ridge National Laboratory

A new screened multiple scattering method has been developed that can possibly treat nanoscale and bio-systems consisting of hundreds of thousands of atoms. This new method leads to an extremely sparse representation, which makes use of both non-symmetric direct and iterative solvers. A new aspect of this work is to develop fast sparse techniques for addressing full potential Poisson-type equation and to extend to solving full potential scattering based equations. As part of this presentation we highlight our current developments using discontinuous-Garlekin type formulations for addressing Poisson type equation.

Submitting Author: William Shelton

A Linear Scaling 3-Dimensional Fragment Method for Petascale Nanoscience Simulations

Zhengji Zhao, Lin-Wang Wang, Juan Meza
Computational Research Division, Lawrence Berkeley National Laboratory

Material simulations require ab initio methods and codes that scale linearly both with the system size and with the number of processors. Unfortunately, Density functional theory (DFT), which is the most widely used ab initio method in material simulations, scales as O(N3) with the size N of the system. Furthermore, the existing codes don't scale well with the number of processors used due to the large amount of communication required between processors. We have developed a linear scaling three-dimensional fragment (LS3DF) method to overcome this $O(N^3)$ scaling and also addresses the computer parallelization issue. This method divides a large system into small fragments, and solves the electronic wavefunctions of each fragment independently, then patches the results together. The novel patching scheme assures that the artificial boundary effects of the spatial division will cancel out. As a result, the LS3DF results are numerically the same as the original full system DFT results, while the required floating point operations are reduced from $O(N^3)$ to O(N). In addition, each fragment is solved by a small number of processors independently, without communication with other groups of processors resulting in a code that

scales almost linearly with the number of processors used. We have tested the scalability of our code using over 2000 processors on system sizes up to 15,000 atoms (Silicon quantum dots). We have also tested the accuracy of the LS3DF method, which compares favorably with the original full system LDA results.

Submitting Author: Zhengji Zhao

Analytic Jacobians in the Solution of Nonlinear PDEs

Paul D Hovland, Argonne National Laboratory; Boyana Norris, Argonne National Laboratory

We discuss the use of automatic differentiation in the solution of nonlinear PDEs. Automatic differentiation makes the computation of cheap and accurate Jacobian-vector and transposed-Jacobian-vector products possible, which can alter the choice of the best linear solver in a Newton-Krylov or Newton-multigrid method. We present experimental results suggesting that while GMRES is often the clear winner for finite difference based Jacobian-free Newton-Krylov, recurrence-based linear solvers may provide superior performance when analytic Jacobian-vector products are available. We describe the use of coloring to reduce the cost of computing sparse Jacobians for use in preconditioning or multigrid methods, including recent advances in the use of two-stage coloring heuristics to reduce the cost of Jacobian computations for nonlinear PDEs discretized using regular grids.

Submitting Author: Paul Hovland

Parallel Optimization Using TAO

Jorge More, ANL/MCS Division; Todd Munson, ANL/MCS Division; Jason Sarich, ANL/MCS Division

In this talk, we discuss the parallel optimization capabilities in TAO, the Toolkit for Advanced Optimization. Particular attention is given to the preconditioned limited-memory, variable-metric methods for unconstrained and bound-constrained optimization problems. We discuss their implementation and provide numerical results demonstrating their effectiveness on standard test sets and scalability on parallel architectures.

Submitting Author: Todd Munson

A polynomial-time interior-point method for conic optimization with inexact barrier evaluations

Simon P. Schurr, Dianne P. O'Leary, Andre Tits
University of Maryland

Interior point methods (IPMs) are the algorithms of choice for solving many convex optimization problems, including semidefinite programming problems and second-order cone problems. The basic ingredients of efficient IPMs are the construction of a self-concordant barrier function and the use of Newton's method to determine the sequence of iterates. The bulk of the work in each iteration is the evaluation of the first and second derivatives of the barrier function and the solution of the linear system of equations formed by Newton's method.

The success of IPMs for these problems has not been matched for a more general class of problems, conic optimization problems. The primary hindrance is the difficulty of explicitly evaluating the derivatives of the barrier functions, particularly those for the dual problem. The construction of the universal barrier function provides a definition for the derivatives of the primal barrier but does not necessarily prescribe a polynomial-time evaluation procedure.

Even for less general problem classes, evaluation of the derivatives can be quite expensive. For example, in semidefinite optimization, where the variables are symmetric positive definite matrices, derivatives of the barrier function are usually evaluated by computing inverses or Cholesky decompositions of matrices. Since these linear algebra tasks can be expensive when the matrices are large, it might be preferable to instead compute an approximate inverse or approximate Cholesky decomposition, and hence an approximate gradient and Hessian of the barrier function, if doing so does not greatly increase the number of iterations.

Therefore, in this work, we consider applying a simple primal-dual short-step IPM to convex conic optimization problems for which exact evaluation of the gradient and Hessian is either impossible or too expensive, even if rounding errors are ignored. As our main contribution, we show that if the relative errors in our estimates of the gradient and Hessian are not too large, our algorithm has the desirable property of polynomial worst-case iteration complexity.

Our work is directly applicable to problems involving a large number of (conic) constraints. Barrier functions for such problems are often expensive to evaluate; however, they (or rather their gradients and Hessians) can be approximately evaluated by ignoring all but a small, appropriately selected subset of the constraints. Indeed, barrier functions associated with individual constraints are often inexpensive to evaluate. A special case of interest is that of finite discretized semi-infinite linear programming. Problems involving a large number of (possibly nonlinear) convex constraints with similar structure are also prime candidates for application of our method.

Submitting Author: Andre Tits

Convergence of a Constraint-Reduced Variant of Mehrotra's Predictor-Corrector Algorithm

Luke Winternitz, Stacey Nicholls, Andre L. Tits, Dianne P. O'Leary
University of Maryland

In a recent paper [SIOPT, Vol. 17, pp. 119-146, 2006], Tits et al. [TAW] presented an interior point algorithm for solving unbalanced linear programs. Specifically, they considered problems in dual standard form with many more (dual) inequality constraints than variables (i.e., with n>>m, where A is mxn). For example, such problems arise in the discretization of semi-infinite problems. Their algorithm, termed the "reduced Primal Dual Affine Scaling" (rPDAS) algorithm, uses a modified affine scaling direction (no centering), and requires a dual feasible initial point. At each iteration, the search direction is based on solving an "approximate normal equation" where the normal matrix is formed using only a small subset of the constraints. Since forming the normal matrix is the most computationally costly step of the algorithm the savings are potentially significant.

In [TAW], Tits et al. proved global and local quadratic convergence of rPDAS under certain non-degeneracy assumptions. They reported promising numerical results on a collection of random problems and on a few Netlib LP problems, even when only a small fraction of constraints are used in computing the search directions. In addition to the rPDAS algorithm, Tits et al. proposed a constraint reduced version of the Mehrotra Predictor Corrector algorithm (rMPC), the current algorithm of choice for implementations of interior point LP solvers. They reported promising performance of rMPC on the test problems, but provided not convergence analysis.

In the present work, the analysis of [TAW] is extended to prove global and local quadratic convergence of a modified rMPC algorithm. Additional numerical tests are performed, focusing in particular on discretized semi-infinite problems, such as the Chebyshev approximation problem. Also, a strategy for adapting the working set of constraints as the algorithm proceeds is proposed and is shown to generally outperform the simple rule used in TAW.

Further efforts on this work aim to weaken the non-degeneracy assumptions made in the convergence analysis, extend the idea to problems more general than linear programs, and to precisely characterize when this type of constraint reduction will and will not be effective.

Submitting Author: Luke Winternitz

Convex Duality and Entropy-Based Moment Equations: Characterizing Degenerate Densities

Cory D. Hauck, Los Alamos National Laboratory, Andre L. Tits, University of Maryland, College Park, C. David Levermore, University of Maryland, College Park

A common method for constructing a function from a finite set of moments is to solve a constrained minimization problem. The idea is to find, among all functions with the given moments, that function which minimizes a physically motivated, strictly convex functional. In the kinetic theory of gases, this convex functional is the kinetic entropy; the given moments are a vector of macroscopic densities; and the solution to the constrained minimization problem is used to formally derive PDE which describe how the macroscopic densities evolve in time. Known as

moment equations, such PDE are important because they simplify the kinetic, phase-space description of a gas; and when derived with entropy-based methods, they retain many of the fundamental properties of kinetic transport.

Unfortunately, there often exist physically relevant densities for which the constrained minimization problem does not have a solution; and in such cases, the entropy-based derivation is not well-defined. In this work, we provide a geometric description of these so-called degenerate densities. Our key tool is a complementary slackness condition that is derived from a dual formulation of the minimization problem with relaxed constraints. We show that, under reasonable assumptions, the set of degenerate densities is small in both a topological and measure theoretic sense. Such a result is important for further assessment and implementation of entropy-based moment equations.

Submitting Author: Cory Hauck

A Fitness Landscape Analysis of the Capacitated Vehicle Routing Problem

Jean-Paul Watson, Sandia National Laboratories

Metaheuristics are a family of generalized local search strategies for escaping local minima in (typically combinatorial) optimization problems; well-known examples of such techniques include simulated annealing, tabu search, and variable neighborhood search. Although metaheuristics fail to provide quality bounds on the obtained solutions, they are pervasive in a number of practical contexts - including logistics, routing, and scheduling problems - due to exceptional performance in terms of run-time and solution quality. Over the last two decades, metaheuristics researchers have concentrated almost exclusively on developing high-performance algorithms in an ad-hoc fashion, at the expense of principled algorithm engineering and scientific model-building.

Moving against this historical trend, we carefully analyze metaheuristics for a fundamental problem in logistics - the Capacitated Vehicle Routing Problem or CVRP. Innumerable metaheuristics have been introduced for the CVRP over the last 15 years, yet very little is known about why these algorithms work so well in practice. We approach this question by analyzing the fitness landscape of the CVRP. A fitness landscape is a 3-tuple consisting of a search space, an objective function, and a move operator. The structure of the fitness landscape is theoretically known to be central to understanding why metaheuristics work well, and in what contexts.

We first explore the impact of feasible versus infeasible solutions on the fitness landscape structure, and examine the run-time dynamics of state-of-the-art metaheuristics in terms of searching this landscape. Our analysis indicates that these algorithms - despite their state-of-the-art status - can be modeled with high fidelity as a simple Ehrenfest diffusion model of a random walk with a central restoring force. This model illustrates that these algorithms spend the vast majority of their time searching far away from globally optimal or near-optimal solutions, and suggests a concrete mechanism by which newly proposed algorithms can be assessed. Next, we consider the distribution of solution quality relative to their distance to a globally optimal solution; this is historically known as fitness-distance correlation (FDC) analysis. We expose a number of central problems with prior FDC analyses of the CVRP, and demonstrate that this correlation is in fact far lower than previously thought. This finding has implications for the design of long-term memory mechanisms for metaheuristics, which we detail.

Overall, we demonstrate that careful empirical analysis of metaheuristics can lead to both accurate behavioral models and insight into understanding the "why" of metaheuristic performance. Ultimately, we believe such analysis will lead to scientific models of metaheuristics, which in turn can eliminate the "iterative hacking" approach to algorithm design so widely employed in the metaheuristics community.

Submitting Author: Jean-Paul Watson

Recent Advances in High Order methods for Simulation of Richtmyer-Meskhov Instability

David Gottlieb and Wai-Sun Don
Brown University

The Richtmyer-Meshkov instability (RMI) is a fundamental fluid instability that occurs when a shock wave passes through an interface separating gases with different properties and amplifies perturbations. This instability is typically studied in shock tube experiments, in which an incident shock passes through an initially perturbed interface separating the gases. Following the passage of the shock, the interface is set in motion along the direction of shock propagation and a transmitted shock enters the second gas. The misalignment of the density and pressure gradients cause a deposition of vorticity $\omega \equiv \nabla x u$ on the interface through baroclinic vorticity production.

We will review some of the research e_orts and results on the applications of high order methods for the simulation of the RMI including the Spectral methods and high order (up to 9th order) weighted essentially non-oscillatory (WENO) finite difference methods. We have systematically and self-consistently explore and quantify the sensitivity of a broad array of quantities characterizing single-mode Richtmyer-Meshkov instability-induced mixing with grid resolutions 128, 256 and 512 points per initial perturbation wavelength λ via the 5th and 9th order WENO schemes 2. The mixing layer widths, circulation, mixing profiles, energy spectra and other statistics and probability distribution functions of the flow, before and after reshock by reflecting the shock from the end of the shock tube, were obtained. Comparison of the results with some existing models before reshock is also studied. Some preliminary results on three dimensional RMI with membrane with noise will also be presented. The high order methods are shown to be effective and accurate and therefore suitable for the class of highly unstable ows that feature both discontinuous function (shocks) and fine scale structures (interfacial mixing and = vorticity generations).

We will also present a new high order hybrid Spectral and WENO finite difference scheme for hyperbolic= conservation laws 3. The new spatial and temporal adaptive algorithm combines two very high order high resolution numerical methods, namely, Spectral methods (Spectral) and high order WENO finite difference scheme (WENO), for hyperbolic conservation laws in a multi-domain framework (Hybrid). The Hybrid algorithm employs the high order Multi-Resolution analysis (MR) by Harten to determine the local smoothness of the solution in each sub-domain. The subdomains adapt to the change of the smoothness of the solution by switching between the Spectral and WENO methods as required spatially and temporally. It is effective in avoiding the appearance of the Gibbs phenomenon if shock appears in the spectral subdomains while maintaining high resolution and efficient when the solution is smooth. We will demonstrate the efficacy and accuracy of the Hybrid scheme for two-dimensional high Mach shocked ows, such

as the Mach 3 and 6 shock-vortex interactions and the Mach 4.5 and 8 Richtmyer-Meshkov instabilities (RMI).

Submitting Author: David Gottleib

Wednesday Poster Abstracts

An Eigensolver with Lowrank Updates for Spin-Fermion Models

E. D'Azevedo*, P. Nukala, G. Alvarez
Computer Science and Mathematics Division
Oak Ridge National Laboratory

Monte Carlo (MC) simulation on a two dimensional periodic lattice is used in the investigation of strongly correlated electron or spinfermion systems where the fluctuations are adiabatically slow. The MC algorithm visits each location of the lattice and a local change is proposed. The probability that the change is accepted is based on the Boltzmann weight that is a function of all the eigenvalues obtained from diagonalization of the Hamiltonian matrix H. If a change is accepted, the matrix H is updated by a low rank modification and all the eigenvalues are recomputed. On a two dimensional n by n rectangular attice system with two spins, the matrix H is N by N, ($N = 2 n^2$) and has a block sparsity pattern resulting from nearest neighbor connections and periodic boundary condition. A key property is that only a 2 by 2 diagonal block of the Hamiltonian matrix I modified after each transition. Moreover, matrix H may be considered the sum of a banded matrix with bandwidth O(n) and a low rank matrix with O(n) off diagonal entries corresponding to the edge connections of the periodic boundary. Recomputation of eigenvalues after a symmetric rank one update is performed in the divide and conquer algorithm for computing eigenvalues. The standard algorithm is adapted for complex Rank two updates. We present our work in developing an efficient eigensolver with low rank updates for large-scale lattice simulations.

Submitting Author: Ed D'Azevedo

A simple, fast, high-order Helmholtz solver

Leonid Kunyansky, University of Arizona

The Helmholtz equation is frequently solved by a reduction to the Liepmann-Schwinger integral equation, whose kernel is a free-space Green's function. Accurate computation of convolutions with such a kernel can usually be achieved only by application of sophisticated numerical techniques.

We present an extremely simple Helmholtz solver that does not require manipulations with special functions or high-order interpolations between the computational grids, and is very easy to implement. In spite of its simplicity, this method yields super-algebraic convergence for smooth scatterers. For non-smooth scatterers an improved order of convergence is obtained by using the band-limited approximation of the scatterer.

The main computational expense of the solver is two FFT's per iteration, each twice the size of the discretization grid. In practical terms, computation of one iteration on the grid 256 by 256 by 256 (about 17 million unknowns) requires about 2 minutes on a desktop computer.

Submitting Author: Leonid Kunyansky

Robust and Scalable Factorization-based Preconditioners Suitable for Petascale Simulations

Ming Gu, UC Berkeley
Xiaoye Sherry Li, Lawrence Berkeley National Lab
Panayot S. Vassilevski, Lawrence Livermore National Lab

The multiscale, multiphysics components in petascale scientific simulations require development of scalable and reliable mathematical algorithms and software. Such simulations typically result in systems (discretizations of PDEs) that are extremely large and often poorly conditioned. The solution of the resulting sparse linear equations or eigenvalue equations presents a major simulation bottleneck. Much progress for solving such systems has been achieved in recent years by developing new approaches, in particular, domain decomposition and multigrid methods. We are investigating alternatives to these solvers in areas where they are less successful.

One new study is primarily focused on fast nearly exact sparse matrix factorizations exploiting highly accurate low-rank approximations existing for matrices that exhibit certain semi-separable structure. Such structure arises naturally for matrices coming from elliptic type discretization problems. Our new factorization algorithms possess nearly linear complexity in both time and memory for solving both 2D and 3D elliptic PDE problems. This preserves the robustness of the direct methods, but significantly reduces the computational cost of the traditional (direct) factorization methods.

Another study is focused on the development and analysis of a new class of block factorization preconditioners for more general problems. Here, we develop a semiseparable matrix strategy that allows, for any a priori given directions, the resulting approximate factorization matrix to be SPD and have the same action as the original SPD matrix. The direction preserving property, in its simplest form, generalizes the popular MILU methods. We are investigating the new block-factorization technique for both elliptic and more general problems. In the latter case, the new technique can be employed to construct vector preserving interpolation matrices as well as kernel preserving non-Galerkin coarse matrices in algebraic multigrid (or AMG). The AMG is perhaps the only class of methods that has the potential to produce scalable and robust preconditioners for systems of PDEs like elasticity problems, that comprise our main test bed example.

Submitting Author: Sherry Li

Adaptive Multimethod Linear Solvers for Nonlinear PDEs

Lois Curfman McInnes,[1] Boyana Norris,[1] Sanjukta Bhowmick,[1,2]
Dinesh Kaushik,[1] and Padma Raghavan[3]

[1] Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Ave., Argonne, IL 60439

[2] Department of Applied Physics and Applied Mathematics, Columbia University, 200 S.W. Mudd Building, 500 W. 120th Street, New York, NY 10027

[3] Department of Computer Science and Engineering, The Pennsylvania State University, 343K IST Building, University Park, PA 16802

[mcinnes,norris,bhowmick,kaushik]@mcs.anl.gov, raghavan@cse.psu.edu

Abstract

As computational science progresses toward ever more realistic multiphysics and multiscale applications, the complexity of linear and nonlinear solvers for partial differential equations (PDEs) is becoming such that no single research group can effectively select or tune all of the components in a given application, and no single tool or solver can seamlessly span the entire spectrum *efficiently*. In addition, overall simulation times often depend to a large extent on the robustness and efficiency of sparse PDE solvers. This poster describes initial work on (1) parallel multimethod (non)linear solvers that incorporate algorithmic adaptivity to enhance robustness and decrease overall application runtimes, and (2) software infrastructure to support this (as well as more general) adaptivity during long-running simulations.

Consider a multi-component application A, which is composed of components $C_1, C_2, ...C_n$, where multiple implementation choices are available for each component. Each implementation $C_{i,k}$ of component C_k has a cost function $P(C_{i,k}) = F_{i,k}(x,p)$, where $P(C_{i,k})$ is the execution time of component C_k with implementation $C_{i,k}$, and $F_{i,k}(x,p)$ is a mathematical model based on a representation of the simulation's state x and parameters p (considering both software and hardware parameters). While in some cases an analytical model exists for $P(C_{i,k})$, in most cases we only have a few datapoints indicating performance of prior runs, though each datapoint contains a wealth of information. Our goals are twofold: (1) to derive cost functions $P(C_{i,k})$ for all component implementations $C_{i,k}$ in an application (possibly from results obtained in other related applications), and (2) given all $P(C_{i,k})$, to derive a model for an overall application cost function, $P(A)$, where A represents the set of all components in the application. The ultimate goal is to determine the set of $C_{i,k}$ that minimizes $P(A)$ within application-dependent constraints, thus providing an initial application configuration, and to do so cheaply enough to enable dynamic adaptivity during runtime.

This poster presents initial research on these challenges for three motivating parallel applications involving nonlinear PDEs: compressible Euler flow, driven cavity flow, and radiation-diffusion. Within the family of preconditioned inexact Newton methods, we explore algorithms and software for adaptive linear solvers, which combine more robust, but more costly, methods when needed in particularly challenging phases of solution, with cheaper, though less powerful, methods in other phases. We demonstrate that this adaptive multimethod approach can improve robustness and reduce overall simulation time for long-running time-dependent and/or nonlinear applications, which have properties that may significantly change during the coarse of a given simulation.

We also describe prototype software infrastructure for implementing multimethod solvers using the concept of *computational quality of service* (CQoS), or the automatic selection and configuration of components to suit a particular computational purpose. CQoS embodies the familiar concept of quality of service in networking as well as the ability to specify and manage characteristics of the application in a way that adapts to the changing (computational) environment. In this context, CQoS focuses on making appropriate algorithmic choices for linear solvers and the parameters that configure them.

AMS: A Scalable Maxwell Solver

T. V. Kolev and P. S. Vassilevski
Lawrence Livermore National Laboratory

This poster begins with a brief overview of our research on scalable solution techniques for large electromagnetic problems. We highlight several topics, ranging from basic research to parallel implementation with impact in ASC simulation codes. Our focus is on algebraic multigrid (AMG) algorithms for finite element discretizations on 3D unstructured meshes. The AMG is perhaps the only feasible approach that can provide scalable solution of PDEs that describe a number of DOE applications.

We first give a brief description of the so-called auxiliary mesh approach, which uses an auxiliary problem posed on a related mesh that is more amenable for constructing optimal order (geometric) multigrid methods. In combination with a domain embedding technique, this method can precondition a problem defined on a very complicated grid by a more standard geometric multigrid (on uniform meshes). This approach was studied in a collaboration with Prof. Joseph Pasciak from Texas A&M University; we have developed its theory and demonstrated its numerical performance in a recent journal paper.

The main tool in the theory and the implementation of the auxiliary mesh approach is the construction of an interpolation operator that maps auxiliary functions into the (edge) finite element space. Recently, Hiptmair and Xu showed that in contrast to the auxiliary mesh method, one could instead use an auxiliary space that is conforming and is defined on the same original mesh. Motivated by their result and based on our experience with the auxiliary mesh method we developed a new unstructured mesh Maxwell solver that is provably scalable.

The poster focuses on a parallel version of the solver and its implementation in the scalable linear solvers library hypre under the name AMS (Auxiliary Maxwell Solver). We present numerical experiments that demonstrate the scalability of AMS on unstructured 3D meshes for variety of test problems including ones with highly variable coefficients as well as with zero conductivity. We also show results from the application of the AMS to a real MHD simulation from an ALE3D application code.

The results of our AMS computational research have been documented in three technical reports and a conference proceeding paper.

Center for Applied Scientific Computing, UC Lawrence Livermore National Laboratory, 7000 East Avenue, Mail stop 560, Livermore, CA 94550, U.S.A.

E-mail address: ftzanio,panayot@llnl.gov

Submitting Author: Panayot Vassilevski

3D Boundary Integral Analysis by a Fast Spectral Method

S. Nintcheu Fata*, L. J. Gray†, T. Kaplan‡
Computer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA.
*nintcheufats@ornl.gov, †graylj1@ornl.gov, ‡kaplant@ornl.gov

Spurred by the need to solve large scale engineering problems involving complicated geometries, fast boundary integral methods have been developed. In this work, an acceleration of a three-dimensional Galerkin Boundary Element Method (BEM) for solving the Laplace equation is investigated in the context of the Precorrected-FFT (PFFT). Originally proposed for capacitance extraction problems, the PFFT technique is an algorithm for rapid computation of dense matrix-vector multiplications arising in an iterative solution of discretized integral equations. Unfortunately, the original fast spectral method is (i) formulated solely in terms of weakly-singular kernels associated with single-layer potentials (indirect method), and (ii) entirely characterized by the utilization of piecewise constant interpolation functions.

To remove these limitations, an implementation of the PFFT to accelerate the solution of direct boundary integral equations involving weakly-singular, strongly-singular and hyper-singular potential kernels has been developed. Moreover, this algorithm is capable of dealing with arbitrarily varying piecewise polynomial interpolation functions. In the fast spectral method, the problem domain is overlaid with a regular Cartesian grid that serves as an auxiliary platform for computation. With the aid of the Fast Fourier Transform, the necessary influence matrices of the discretized problem are rapidly evaluated on the Fourier grid in a sparse manner. Compared to traditional boundary integral techniques, the sparse representation of the coefficient matrices results in a significant reduction in resolution time and computer memory requirements. The computational cost associated with the sparse approximation of influence matrices is asymptotically lower than that of conventional methods. For a fast solution of direct boundary integral equations, a Krylov-subspace (BiCGSTAB) iterative method is employed, and the sparse influences are used to rapidly compute the matrix-vector products involved at each iteration. Several key features of the formulation, including the mapping of density distributions to/from the regular Fourier grid, are highlighted. Numerical experiments are presented to illustrate the performance of the fast spectral method.

For a robust general-purpose elliptic solver, new analytic integration techniques are being developed for both collocation and Galerkin boundary integral equation methods. The new exact integration method, based on piecewise linear approximations, is expected to significantly improve the accuracy of both the conventional Galerkin BEM and the PFFT method. Preliminary results indicate that the usual inaccuracies associated with quasi-singular integrals are completely eliminated. This latter behavior will be discussed using some numerical examples.

Submitting Author: Nintcheu Fata

Multiscale numerical methods for flows in heterogeneous porous media using limited global

Yalchin Efendiev, Texas A&M University; Tom Hou, Caltech; Victor Ginting, Colorado State University

In this talk, I will describe multiscale numerical methods for flows in heterogeneous porous media. The main idea of these methods is to construct local basis functions that can capture the small scale information when the basis functions are coupled via some global formulation. It is known that the use of purely local basis functions leads to resonance errors which depend on the ratio between the mesh size and the characteristic length scale of the problem. To improve the accuracy of subgrid capturing errors, some type of limited global information is needed. I will discuss how using limited global information, one can remove the resonance errors. Applications of these methods to uncertainty quantification and stochastic porous media equations will be also presented.

Submitting Author: Yalchin Efendiev

Low Mach Number Modeling of Type Ia Supernovae

Ann S. Almgren
John B. Bell
Mike Zingale

Type Ia supernovae (SNe Ia), the largest thermonuclear explosions in the universe, are of considerable astrophysical and cosmological significance. Numerical modeling of SNe Ia has traditionally relied on discretizations of the fully compressible equations for fluid dynamics, supplemented by equations describing the thermonuclear reactions and heat release. The low Mach number formulation, by contrast, analytically removes acoustic wave propagation while retaining the compressibility effects resulting from nuclear burning, radial stratification, and large-scale heating. This formulation generalizes low Mach number models used in terrestrial atmospheric modeling and combustion to arbitrary equations of state such as those describing the degenerate matter found in stellar material. A further generalization allows the background state to evolve in response to large-scale heating of the stellar atmosphere. While a compressible formulation is required for simulation of the final stage of SNe Ia, the low Mach number formulation is well-suited to investigation of the long-time evolution and subsequent ignition of the pre-supernova white dwarf. The low Mach number formulation results in a substantial improvement in computational efficiency compared to a compressible formulation, enabling calculations of supernova ignition that have previously been intractable.

Submitting Author: Ann Almgren

A Stable Finite Difference Method for the Elastic Wave Equation on Complex Domains with Free Surface Boundary Conditions

Daniel Appelo, LLNL; Anders Petersson, LLNL; Bjorn Sjogreen, LLNL

In this talk we describe a new stable finite difference method for the elastic wave equation in second order form on complex domains with stress free boundary conditions.

Finite difference discretizations of the elastic wave equation in second order form are highly efficient and easy to implement (compared to staggered discretizations). Due to these favorable properties the methods were popular in the early days of seismic modeling. However, it was soon discovered that the early methods suffered from two main flaws: The lack of stable discretization of the free surface boundary condition at high ratios between the P and S-wave speed, and the handling of complex geometries needed for realistic topography.

Recently a remedy to the free surface instabilities was presented in a paper by Nilsson, Petersson, Sjogreen and Kreiss. They introduced a summation-by-parts discretization were, at a free surface, one sided difference operators where used for mixed derivatives in the equations. For their discretization they were able to derive a guaranteed stable discretization of the stress free boundary condition, even when the material properties vary rapidly on the grid. The method was proved to be second order accurate and stable for all ratios of the speeds of the P and S-wave.

Here we generalize the results of Nilsson et al. to curvilinear grids. Using summation by parts techniques we show that there exist a corresponding discretization of the free surface boundary condition on a curvilinear grid. We prove that the discretization is stable and energy conserving both in semi-discrete and fully discrete form. We also establish that the method is second order accurate.

Finally, we present computations on overset grids, where we combine the new curvilinear and the Cartesian discretization into an efficient and accurate numerical method for simulations of elastic waves on complex domains.

Submitting Author: Daniel Appelo

Characterization of Implicit LES Methods

A. .J. Aspden and J. B. Bell
Lawrence Berkeley National Laboratory

The broad range of time and length scales present in high Reynolds number turbulent flows is often prohibitively expensive for direct numerical simulations (DNS) to capture completely. Large eddy simulation (LES) attempts to circum- vent this issue by filtering out the small scale motions in the flow, replacing their effects with a subgrid model. High-order finite-volume

schemes can accurately capture the inviscid cascade of kinetic energy, and the inherent truncation error acts as an implicit subgrid model, forming a natural form of LES. However, the absence of a physical viscosity prohibits conventional characterization of these methods, specifically how kinetic energy is dissipated at the grid scale and how to define a relevant Reynolds number.

Kolmogorov's 1941 papers achieve this characterization for real-world viscous fluids in terms of a universal equilibrium range determined uniquely by the rate of energy dissipation and physical viscosity. Analogously, this paper proposes than an implicit LES method results in behavior that can be characterized by a universal equilibrium range determined uniquely by the energy dissipation rate and computational cell width.

Implicit LES simulations of maintained homogeneous isotropic turbulence are presented to support this proposal and highlight similarities and differences with real-world viscous fluids. Direct comparison with data from high resolution DNS calculations provides a basis for deriving an effective viscosity and an effective Kolmogorov length scale.

Submitting Author: A. J. Aspden

Comparison of Flux-Corrected-Transport and High Resolution Godunov Methods on Overlapping Grids

J. W. Banks, Sandia National Laboratory; J. N. Shadid, Sandia National Laboratory

Numerical methods for the simulation of conservation law systems in "difficult" regimes, such as fluid flows for Z-pinch, NIF, or even condensed phase explosives applications, are plagued by difficulties which can be traced to the non-lindear stability properties of the method. This talk will focus on the development and comparison of two schemes which directly enforce a kind of non-linear stability (i.e. TVD, monotonicity) and as such have favorable properties in these difficult regimes. Specifically we investigate the development of a flux-corrected-transport (FCT) method and compare simulation results for this method with a high resolution Godunov method. To facilitate simulation in regions of complicated, possibly moving, geometry as well as the inclusion of adaptive mesh refinement, both algorithms are implemented in an overlapping grid framework. A description of the methods in this context is provided and a series of computational examples is given to illustrate their differences. These computational examples range from rather simple (such as an isolated contact wave) to very complicated (such as the impulsive motion of rigid bodies by an incident shock). To conclude we introduce an idealized model of Z-pinch type implosions which is used as a benchmark test for the hydrodynamic portion of algorithms. This difficult problem highlights the strength of methods which are capable of direct enforcement of non-linear stability. In particular the FCT method proves to perform very well even for extreme cases such as density jumps on the order of seven orders of magnitude.

Submitting Authors: Jeff Banks

Domain-Adaptive High-order Accurate Algorithms for PDEs in Moving Geometry

William D. Henshaw, LLNL; Kyle K. Chand, LLNL

Multiscale and multi-physics applications require algorithms capable of adapting the PDE model, order of accuracy or discretization to the simulation requirements in a subset of the domain. For example, in complex 3D applications adaptive grid and accuracy refinements help to resolve local features without enforcing the cost of high resolution on the entire domain. In multi-physics problems the actual PDE under investigation may change throughout the domain and even in time. The domain adaptive algorithms developed by this project will provide an efficient way to solve these problems by treating separate parts of the domain in different ways. A subset of the domain can use a local PDE, moving grid, discretization, mesh and even timestep. With this approach each part of the problem can be treated in a locally optimal manner leading to significant efficiency and accuracy improvements.

This project focuses in particular on the mathematics required to develop domain adaptive algorithms for wave propagation problems involving multi-material and multi-physics applications. Our research includes stable and high order accurate (HOA) adaptive methods for wave equations in complex moving geometry, material interfaces and domain dependent solution methods. A significant feature of this work is a mathematical framework for the treatment of interfaces between grids, different materials and different governing equations. To accurately and efficiently discretize wave propagation problems in complex geometry we are developing compact Hermite-Taylor approximations. High order accurate (HOA) methods for structured adaptive mesh refinement are also being studied along with mollified equation approaches for the associated problem of smooth wave propagation through coarse-fine grid boundaries. Accuracy and stability analysis using energy methods, normal mode and GKS stability analysis provides a theoretical foundation for these methods.

Our poster will present some early results of this research applied to problems in electromagnetics, gas dynamics and moving body problems.

Submitting Author: Kyle Chand

A Cartesian AMR framework for detonation- and shock-driven fluid-structure interaction simulations

Ralf Deiterding, Oak Ridge National Laboratory

Fluid-structure interaction simulations of strong shock and detonation waves in inviscid compressible fluids impinging on heavily deforming solid structures require numerical methods that can cope with large topology changes. A generic Eulerian-Lagrangian approach is presented that allows the incorporation of time-accurate computational solid dynamics solvers with large deformation, fracture and fragmentation capability into time-explicit Cartesian CFD upwind methods. The geometrically complex solid boundary is represented implicitly on the Eulerian mesh with a level set distance function that is evaluated on-the-fly from the evolving solid

dynamics mesh. Fluid boundary conditions are constructed according to location and velocity of the zero isocontour of the level set function. After a fluid update, the structural loading is derived from the hydrodynamic pressure and communicated back to the solid dynamics solver in a straightforward splitting approach.

Structured adaptive mesh refinement (SAMR) is used in the fluid to dynamically capture the near-body interaction and incoming waves at minimal computational costs. The algorithmic integration of fluid-structure coupling calls into the hierarchical time step refinement scheme of the SAMR method will be discussed. All components have been parallelized for distributed memory systems and the scalable implementation of the hierarchical mesh refinement method, the efficient parallel inter-solver communication routines, and the effective transformation of the embedded triangulated boundary into a Cartesian level set function will be sketched briefly.

In order to demonstrate versatility and efficiency of the approach, several computational examples, that utilize second-order accurate approximate Riemann solvers for the hydrodynamic transport, will be considered. Beside simple validation computations, in which shock waves in polytropic gases impinge on elastically deforming panel and propagating rigid objects, miscellaneous three-dimensional multi-physics simulations will be presented. Among those are the fracture of thin aluminum tubes induced by interior gaseous detonations in an ethylene-oxygen mixture, the plastic deformation and fracture of copper plates from strong piston-induced pressure waves in a water shocktube and the elastic wall response of a mercury filled vessel from shock and rarefaction waves created by high energy proton beam.

[1] R. Deiterding, R. Radovitzky, S. P. Mauch, L. Noels, J. C. Cummings, and D. I. Meiron. A virtual test facility for the efficient simulation of solid materials under high energy shock-wave loading. Engineering with Computers, 22(3-4):325–347, 2006.

[2] F. Cirak, R. Deiterding, and S. P. Mauch. Large-scale fluid-structure interaction simulation of viscoplastic and fracturing thin shells subjected to shocks and detonations. Computers & Structures, 2007, in press.

Submitting Author: Ralf Deiterding

High-Resolution Discretization of the Gyrokinetic Vlasov-Poisson Equation on Mapped Grids

M. Dorr, Lawrence Livermore National Laboratory; J. Hittinger, Lawrence Livermore National Laboratory; P. Colella, Lawrence Berkeley National Laboratory; D. Martin, Lawrence Berkeley National Laboratory

We describe our approach and progress in the development of high-resolution methods for the numerical solution of the gyrokinetic Vlasov-Poisson equations on mapped grids. The primary scientific application of interest to the Office of Science that motivates this work is the kinetic simulation of plasmas in the edge of a tokamak fusion reactor. Here, the preponderant magnetic field defines a unique and challenging phase space geometry upon which these equations, as well as further generalizations including full electromagnetics, collisions and other physical processes, must be solved.

The gyrokinetic Vlasov-Poisson equations describe the evolution of gyrocenter distribution functions (ion and electron) and an electrostatic potential under the influence of a large, imposed magnetic field. The use of gyrocenter coordinates provides two major advantages: particle gyromotion, i.e., the gyration of particles about the magnetic field lines, decouples from the remaining particle dynamics and the gyrocenter coordinate system is derived as a volume-preserving mapping from laboratory coordinates. A consequence of the former is that it reduces the dimension of the phase space domain on which the distribution function is defined and eliminates unimportant high-frequency modes from the simulation. The latter ensures that phase space volumes, and hence conservation properties, are preserved.

We are applying a combination of numerical methodologies to address the myriad complications associated with the solution of the gyrokinetic Vlasov-Poisson system in realistic geometries. In particular, we seek a discretely conservative algorithm that effectively addresses the issues of long-time fidelity, large anisotropies, embedded boundaries, and high dimensionality. Our core discretization is a new class of multidimensional higher-order finite volume methods (at least fourth-order in space and time). The use of a finite volume method guarantees discrete conservation of the distribution function, while the use of higher-order methods provides a substantial improvement in the long-time accuracy over more traditional second-order finite-volume methods. We are developing these methods for mapped multiblock grids, thus providing us with sufficient flexibility in the gridding to accommodate the geometrical complexity of the physical space domain, and allow us to use grids that are aligned with the magnetic field directions. The use of a semi-structured computational grid offers additional advantages, including the availability of efficient elliptic solvers, natural extensions to local mesh refinement, and efficient parallel implementation. Since the computational demands of this problem require us to deal with performance on high-end parallel systems in order to make any progress on even the most fundamental research issues, from the outset, we are implementing our algorithms as scalable parallel code, leveraging off of existing software frameworks and tools developed under the SciDAC program.

Submitting Author: Milo Dorr

Implicit Interface Finite Element Method for Elliptic Interface Problems

Jae-Seok Huh, Math. Dept., LBNL and Dept. of Math., UC Berkeley; James A. Sethian, Math. Dept., LBNL and Dept. of Math., UC Berkeley

We establish a finite element method to solve elliptic interface problems on fixed background meshes which are non-conforming to the interface. The solution of an elliptic equation, $-\nabla \cdot (a \nabla u) + Lu = f$, on a domain partitioned by an interface typically experiences singularies on the interface; the solution and/or its normal derivative has jump discontinuity across the interface. In general, such interfacial singularities result from (i) jump in the coefficient, $[[a]] \neq 0$, (ii) given jump of the flux, $[[a \partial u / \partial \nu]] = g$, and/or (iii) given jump of the solution, $[[u]] = h$.

Insufficient regularity of the solution results in unbounded interpolation errors across the interface, which implies a plain application of a finite difference/element method can not work, unless the mesh is conforming to the interface. Moreover, in most applications involving elliptic interface problems, the interface is not fixed but propagating; a conforming mesh, if employed, should be accompanied by a painful re-meshing which introduces a collection of severe numerical difficulties. This naturally demands numerical schemes based on fixed non-conforming meshes, which have been evolved into two distinctive approaches; (i) obtain a smeared solution by regularizing underlying distribution sources, or (ii) achieve sharp resolution by imposing given jump conditions in the discretized operator sacrificing simplicity and robustness in the implementation.

Our approach follows the idea of *regularization by singular correction functions;* first, a jump discontinuity of *a*, if present, is treated by normalizing the equation, which introduces an additional unknown on the interface. Then, a correction function is constructed to satisfy given jump conditions. By decomposing a singular solution into a regular portion and the correction function, the original problem can be transformed into a regular one with an additional source involving the correction function on the right hand side. A regular background method is employed to solve the resulting problem on a fixed non-conforming mesh. Thus, our approach, where the jump conditions are enforced in the correction function, can achieve a sharp resolution of the solution, maintaining the same stable discrete operator as that of the regular background method.

The most crucial algorithm in this method is the construction of the correction function. From given jumps, utilizing a signed distance representation for the interface, the presented algorithm constructs a correction function which is easy to evaluate and whose variation is bounded independently of the relative location of the interface with respect to the background mesh.

Mathematics Department, LBNL and Department of Mathematics, UC Berkeley
E-mail address: jshuh@math.berkeley.edu

Mathematics Department, LBNL and Department of Mathematics, UC Berkeley
E-mail address: sethian@math.berkeley.edu

Submitting Author: Jae-Seok Huh

Analysis of Particle Migration in Dilute Suspensions Using a TC-BEM

Marc Ingber, University of New Mexico; Alan Graham, Los Alamos National Laboratory; Shihai Feng, Los Alamos National Laboratory; Lisa Mondy, Sandia National Laboratories

The rheology of suspensions of rigid particles is still not well understood, and in particular, the origins of the experimentally observed irreversibility in what is nominally a reversible flow are unclear. While there is general agreement that irreversibility is caused by short-range interactions, the exact nature and magnitude of these interactions is unknown, and repulsive or attractive forces are applied more or less arbitrarily. Further, it is not clear whether two-particle, three-particle or many-particle interactions are the mechanisms through which irreversibility is manifested at the macroscale. In this research, the phenomena of self-diffusion and migration of rough spheres in nonlinear shear flows are investigated using a new traction-corrected boundary element method (TC-BEM) in which the near-field asymptotics for the traction solution in the

interstitial region between two nearly touching spheres is seamlessly coupled with a traditional direct boundary element method. The TC-BEM is extremely accurate in predicting particle trajectories, and hence, can be used to calculate both the particle self-diffusivity and a newly defined migration diffusivity for dilute suspensions. The migration diffusivity is a function of a nonlinearity parameter characterizing the shear flow and arises from the net displacement of the center of gravity of particle pairs. This net displacement of the center of gravity of particle pairs does not occur for smooth particles and does not occur for rough particles in a linear shear flow. An explanation is provided as to why two particle interactions of rough spheres in a nonlinear shear flow result in particle migration.

Submitting Author: Marc Ingber

Mimetic Discretizations

Konstantin Lipnikov, Los Alamos National Laboratory; Mikhail Shashkov, Los Alamos National Laboratory

Development of discretizations that mimic important properties of continuous models is a part of the project "Mimetic Methods for Partial Differential Equations". We pursue two research directions: (I) accurate discretizations on generalized polyhedral meshes and (II) monotone schemes for triangular and quadrilateral meshes. The ultimate research goal is to merge two directions to find an accurate and monotone scheme for general meshes.

I. The new mimetic finite difference (MFD) method provides symmetric and second-order accurate discretizations of diffusion problems with heterogeneous anisotropic full tensor coefficients on generalized polyhedral meshes. These meshes appear in different porous media and CFD applications. The MFD method does NOT require subdivision of elements into simplexes. It can be deployed on AMR meshes with degenerate elements and on Lagrangian meshes with non-convex element in exactly the same manner as on tetrahedral meshes. On random hexahedral meshes, the MFD method is superior to the Raviart-Thomas finite element method b/c the latter may not converge. The MFD method is computationally cheaper than the Kuznetsov-Repin finite element method. On tetrahedral meshes, a multi-flux version of the MFD method results in a symmetric scheme with a local expression for the flux. On polyhedral meshes, the MFD and the multi-point flux approximation (MPFA) methods have many common features which have not been yet studied to improve both methods. In general, a rich family of MFD methods does exist (e.g., a 6-parameter family for hexahedral meshes). This family can be analyzed to tackle other computational problems such as monotonicity and optimal accuracy for a fixed mesh. The new MFD method has been developed in 2005-2006 and has been already implemented in ASC codes.

We developed a solid mathematical foundation for the new MFD method. The extension of the new MFD method to (a) nodal discretizations on polyhedral meshes, (b) other PDEs (linear elasticity, Maxwell), and (c) high-order schemes is the research in progress. This is the joint work with F.Brezzi (Italy), and V.Simoncini (Italy).

II. Monotonicity is the most difficult property of continuum models to mimic even on triangular and quadrilateral meshes. The classical finite volume (FV) and finite element (FE) schemes do not guarantee the discrete maximum principle on general meshes and for full diffusion tensors.

Linear two-point flux approximation schemes are usually monotone but very inaccurate for a full diffusion tensor. This may results in unphysical Darcy velocities in porous media applications and therefore to the wrong prediction of a contaminant transport. We present results our research on a nonlinear FV scheme which is monotone for unstructured triangulations and arbitrary heterogeneous full diffusion tensors. The scheme uses a nonlinear two-point flux approximation method and therefore has a compact stencil. For triangular meshes, the schemes gives the second-order convergence rate for the scalar unknown and the first-order convergence rate for the flux unknown.

We developed a theoretical analysis of the nonlinear FV scheme. We have also extended the scheme to shape regular quadrilateral meshes and heterogeneous isotropic diffusion tensors. This is a joint work with D.Svyatskiy (LANL) and Yu.Vassilevski (Russia).

Submitting Author: Konstantin Lipnikov

Computational Nanophotonics and NEKCEM

Misun Min, Argonne National Laboratory; Paul Fischer, Mathematics and Computer Science Division

We have developed a numerical software package, NEKCEM, for computational electromagnetics. The software package employs high-order numerical techniques, specifically a discontinuous Galerkin method based on body-fitted spectral element meshes. We will demonstrate various simulations to study dynamic nature of electromagnetic waves interacting with nanosystems and cavities.

High performance of the code with some built-in examples will be also demonstrated with spectrally accurate convergence and efficient parallel performance.

Submitting Author: Misun Min

Solving PDEs on Irregular Domains with Moving Interfaces

Ying Shan, Lawrence Berkeley National Laboratory; J.A. Sethian, Lawrence Berkeley National Laboratory

We have designed a numerical algorithm for solving partial differential equations on irregular domains with moving interfaces. This approach is faster than existing ones by doing most calculations only in the desired domain instead of a larger rectangular domain. To do so efficiently, we have developed a one-sided multigrid method to solve the corresponding large sparse linear systems.

To test the algorithm, we focus on the simulation of the electrodeposition process in semiconductor manufacturing. Our goal is to track the position of the interface between the metal and the electrolyte as the features are filling and to determine numerically what initial configurations lead to superfill.

We begin by motivating the set of equations which model the electrodeposition process. After that, we show how to solve the level set equation, then the conservation laws, followed by the diffusion equations on moving, irregular domains. We will also explain how to use the multigrid

method to solve the linear systems which are discretizations of the one-sided diffusion equations. Finally some possible future research directions will be proposed.

Submitting Author: Ying Shan

A Subgrid-Scale Turbulence Model for Simulating Compressible Flows

Paul R. Woodward
University of Minnesota

My research team in the Laboratory for Computational Science & Engineering (LCSE) at the University of Minnesota has been simulating compressible, turbulent flows for many years using our PPM gas dynamics codes. These codes solve the Euler equations of inviscid flow. They capture turbulent phenomena automatically, since for very high Reynolds number flows, at least in astrophysics in the absence of walls, the development of turbulence and the nature of the turbulent cascade at longer wavelengths is not affected by viscosity. However, just as one finds from solving the Navier-Stokes equations for such flows, energy tends to accumulate in the region of the velocity power spectrum just above the very short wavelength region where viscous dissipation, either real or numerical in origin, sets in. In this near dissipation range, the velocity power spectrum tends to have a k^{-1} rather than a $k^{-5/3}$ behavior. For our PPM codes, this near dissipation range is from wavelengths of about 8 to 30 grid cell widths, and for Navier-Stokes simulations it tends to extend to somewhat longer wavelengths. This pile up of energy just before the dissipation range is a real physical effect, but it is undesirable when we are trying to get the best possible approximation to the infinite Reynolds number limit of viscous flows. Over the last few years, David Porter and I have been using very highly resolved PPM Euler simulation data for turbulent flows in the place of experimental data. We are able to filter these flows to yield high quality information on the various terms involving sub-filter-scale information that appear in the fluid equations for the filtered variables. We have compared this data with subgrid-scale turbulence modeling ideas, and we have developed a model that when added to the PPM Euler scheme eliminates the accumulation of energy in the near dissipation range of the spectrum. This model includes terms that mimic the form of some of the numerical truncation error terms, but in turbulent flow regions these terms are about an order of magnitude larger than their numerical error counterparts. This new model will be presented along with data that supports it. Although the model was developed in the PPM family of codes, it should work equally well with any modern, dissipative numerical scheme for compressible flow by means of a readjustment of a single, dimensionless free parameter.

Submitting Author: Paul Woodward

Implementing an Interior Point Method for Linear Programs on a CPU-GPU System

Jin Hyuk Jung, University of Maryland; Dianne P. O'Leary, University of Maryland

Graphics processing units (GPUs), present in every laptop and desktop computer, are potentially powerful computational engines for solving numerical problems. We present a CPU-GPU

algorithm for solving linear programming problems using interior point methods. This algorithm, based on the rectangular-packed matrix storage scheme of Gunnels and Gustavson, uses the GPU for computationally intensive tasks such as matrix assembly, Cholesky factorization, and forward and back substitution. Comparisons with a CPU implementation demonstrate that we can improve performance by using the GPU for sufficiently large problems. Since GPU architectures and programming languages are rapidly evolving, we expect that GPUs will be an increasingly attractive tool for matrix computation in the future.

Submitting Author: Dianne P. O'Leary