

Bay Area Scientific Computing Day 2006

Contributed Posters

Jamylle Carter

Mathematical Sciences Research Institute

Co-authors: Tony F. Chan – UC Los Angeles; Ke Chen – University of Liverpool

“A Multigrid Method for Dual Total Variation-Based Image Restoration”

This talk will describe a computational method for the inverse problem of edge preserving image restoration. Total Variation (TV) regularization removes noise from an image while retaining its edges. Earlier primal TV methods require a user-chosen smoothing parameter which, if chosen too large, will lead to blurred edges. We solve an equivalent dual version of the TV problem independent of a smoothing parameter. Our multigrid method yields faster convergence than the Chan-Golub-Mulet primal-dual method.

Ichitaro Yamazaki

University of California, Davis

Co-authors: Zhaojun Bai – University of California Davis; Richard Scalettar – University of California Davis; Wenbin Chen – Fudan University

“Robust and Efficient Preconditioned Iterative Solver for Multi-scale Quantum Simulation”

Multi-scale quantum simulation of electron interaction provides a powerful technique to investigate many of the fundamental properties of solid-state materials, such as their magnetism, superconductivity, and metal-insulator transition. However, the quantum simulation is currently limited to a few hundred electrons. The primary bottleneck is in the computational kernel involving the solutions of multi-scale linear system of equations. The computational costs of the state-of-the-art linear solver scales cubically in terms of the number of electrons.

We present a robust and efficient preconditioned iterative solver. Using a specially designed data structure to accommodate the data access pattern of the underlying problem, the computational cost of the new iterative solver scales linearly in terms of the number of electrons when the scale of potential energy of the quantum system is moderate. Linear systems of equations associated with thousands of electrons can now be solved effectively, which provides the promise to simulate physical phenomena of technological importance.

Hui Jiang

Stanford Linear Accelerator Center

Co-authors: Volkan Akcelik, Rich Lee – Stanford Linear Accelerator Center

“Adjoint-Based Shape Determination for the ILC Cavity Design”

In the International Linear Collider (ILC), the accelerating cavity shape is different from the ideal shape due to loose machining tolerances and the tuning process. As a result, the damping of the Higher-Order-Modes (HOMs) in the cavity can be degraded leading to beam instability. The goal of this project is to numerically determine the true cavity shape from measured data and to identify the sensitivity of critical dimensions that affect HOM damping. We will formulate the problem as a PDE constrained optimization problem and use a parallel adjoint based method suitable for solving such large scale systems. In the initial development, we will adopt a simplified 2D cavity model to validate the approach.

Sheng Chen

Stanford Linear Accelerator Center

Co-author: Rich Lee – Stanford Linear Accelerator Center

“Gradient-Based Higher Order Error Estimator in Accelerator Design”

In simulating large, complex geometries in accelerator design, adaptive refinement can improve accuracy and speed up convergence while using optimal computing resources. A posterior error estimator based on the energy gradient evaluated in domains with higher-order surface geometry is developed to form the basis of an adaptive refinement scheme for the parallel finite element eigensolver Omega3P. The energy gradient computes the variation of both the electric and magnetic field components. A higher-order error estimator is then used to apply local refinement by reducing the mesh size or increasing the order of the basis function or a combination of both.

Stefan Lang

University of Heidelberg

Co-authors: Christian Engwar – University of Heidelberg; Oliver Sander – Free University of Berlin

“The DUNE Framework with Application to Unstructured Meshes”

The Distributed Unified Numerical Environment (DUNE) is a framework that enables design and implementation of numerical algorithms completely separated from the type of the used mesh (structured, unstructured, block structured). This algorithmic flexibility is realized via methods from generic programming and is shown to cost only little overhead. The design is based on a formal specification of grid functionality. Several mesh instances are ready for use including structured, unstructured, and parallel meshes.

Further we present an implementation of the DUNE grid interface to enable unstructured mesh support. The interface is realized as a thin software layer on top of the existing powerful unstructured mesh module DUGE. We present an application in biomechanics where multiple instances of this mesh entity are used in a single simulation.

The Distributed Unstructured Grid Engine (DUGE) is capable of handling unstructured and hierarchical meshes of mixed element type. Topological entities include elements, faces, edges and nodes. Via local grid adaptation the hierarchical meshes can be manipulated. The mesh management is parallelized using DDD. Dynamic Distributed Data (DDD) is a parallel programming model on object basis, which has been reworked in the last years for efficiency purposes and enhanced stability. Thus parallel adaptivity with dynamic load balancing is realized in a hierarchical mesh context.

Complex domains which are often a prerequisite for problems from engineering and life sciences are now supported directly on CAD level. Therefore the unstructured mesh module has been integrated with the OpenCascade CAD kernel. This allows a domain description in BREP, IGES or STEP format. Initial meshes for CAD domains can be e.g. generated via the CUBIT meshing environment. CUBIT-generated meshes are after translation valid input meshes for the mesh management module.

Kyle K. Chand

Lawrence Livermore National Laboratory

“Simulation Tools for Complex Geometry”

The Agile Solvers project in the Center for Applied Scientific Computing develops methods for solving complex partial differential equation (PDE) systems in complex geometry. Our approach uses composite grids to discretize complex, possibly moving, domains; and high-order accurate methods for the effective and reliable discretization of PDEs. One result of this research is the Overture framework, a collection of software tools that support composite grid generation and discretization technologies. Overture packages these tools into a library that uses C++ for the management of high-level abstractions, such as the geometry, mesh and fields; while employing Fortran kernels for low level operators. At the highest level, a C++ interface eases the prototyping of simulation tools by providing a domain-specific "language" for expressing a PDE discretization. Lower level interfaces can be used to obtain higher performance at the cost of runtime generality. The framework's geometry and meshing tools provide solvers with fast primitives to query and modify the domain which are particularly useful for adaptive and high-order methods. Our poster illustrates Overture's geometry, mesh and discretization components.

UCRL-POST-215863

Tony Drummond and Osni Marques

Lawrence Berkeley National Laboratory

“The DOE Advanced Computational Software Collection (ACTS)”

The ACTS Collection is a set of software tools that help programmers write high performance scientific codes for parallel computers. It differs from other "parallel tools" projects in that it focuses primarily on software used inside an application, instead of on software used to develop an application. The tools were mostly developed at DOE laboratories and Universities. ACTS is an umbrella project that has brought the tools together and is funding developers to provide interoperability. In this poster we highlight some applications of tools that have benefited from the use of ACTS Technologies, and this poster is joined by 5 other poster presentations on new developments of ACTS tools and applications (See Reidy, Yozo, Chand, Kolev)

Tzanio V. Kolev

Lawrence Livermore National Laboratory

“hypre: A Scalable Linear Solvers Library”

The hypre software library provides high performance pre-conditioners and solvers for the solution of linear systems of equations on massively parallel computers. An attractive feature of the library is the availability of a variety of interfaces that allow users to describe their problems in a natural way, such as in terms of grids and stencils, finite elements, or a matrix and right hand side. These interfaces are both easy to use and scalable on machines with tens of thousands of processors, such as BlueGene/L. Furthermore, the interfaces provide access to a large collection of state-of-the-art preconditioners and linear solvers that have been proven scalable on large numbers of processors. For example, a number of geometric multigrid solvers are available for structured-grid problems, while algebraic multigrid is a good choice for many unstructured problems. This poster provides an overview of the many interfaces and solvers available in hypre and highlights some recent additions to the library, as well some results from runs on BlueGene/L.

UCRL-ABS-219272

Jason Riedy

University of California, Berkeley

“Making Static Pivoting Dependable”

For parallel, sparse, unsymmetric LU factorization, static pivoting is an attractive alternative to dynamic pivoting methods like partial pivoting. Static pivoting selects all pivots before numerical factorization and completely decouples the factorization phase from the earlier symbolic manipulations. This decoupling exposes greater parallelism in the numerical phase, but factorization may encounter tiny or zero pivots. Perturbation heuristics to cope with such pivots potentially lose the hallmark stability of a direct method; such problems have been reported with the default heuristics in SuperLU. Experiments to determine static pivoting's dependability address the following questions:

- Do different static pivoting heuristics affect the frequency of tiny pivots?
- How do perturbation heuristics affect backward stability and the forward errors?
- Can iterative refinement restore any lost stability?

We establish that the static pivoting heuristic of maximizing the product of the pivots' magnitudes is somewhat insensitive to approximation, leading to faster pivot selection. We propose a different perturbation heuristic and set of parameters that performs well across a wide selection of test systems. Also, we demonstrate that extended-precision refinement works as well for sparse systems as it does for dense ones and can restore lost stability when the system is not horribly conditioned.

Yozo Hida

University of California, Berkeley

“Extra Precise Iterative Refinement on Linear Systems and Least Squares Problems”

We present the design and testing of an algorithm for iterative refinement of the solution of linear equations, where the residual is computed with extra precision. We report extensive test results on general, symmetric indefinite, and SPD systems. As long as appropriate condition numbers (normwise or componentwise) computed by the algorithm are less than $O(1/\epsilon)$, the computed solution is accurate to $O(\epsilon)$ in the appropriate measure. Some preliminary results applying this algorithm to least squares problems (via augmented linear systems) are also presented.

Ben-Shan Liao

University of California, Davis

“An Algebraic Multilevel Krylov Substructuring Method for Very Large-Scale Eigenvalue Computation”

Algebraic Multilevel substructuring (AMLS) techniques have demonstrated great success in computing eigen-analysis of extremely large scale matrices arising from structural dynamics applications, such as the vibration of car bodies. In these applications, it is typical that a large number of eigenvalues with relatively low accuracy are required. However, such a lack of accuracy is a cause for concern in some applications.

In this talk, we present an algebraic multilevel Krylov substructuring (AMLKS) method, which preserves the mechanism of AMLS techniques, such as efficiency and parallelism, and meanwhile significantly improves the accuracy of computed eigenvalues. The gist of AMLKS method is to replace the eigenmodes of interior substructures by proper Krylov modes of the substructures, which take the coupling among substructures into the account. The accuracy improvement of the AMLKS method are demonstrated by numerical results from electromagnetic and MEMS simulations.
