



Institute for Scientific Computing Research

Fiscal Year 1999 Seminar Series Abstracts

(in reverse chronological order)



September 17, 1999

Dv: A Toolkit for Building Remote Visualization Services

David O'Hallaron

droh@cs.cmu.edu

Carnegie Mellon University

Abstract

The talk describes a toolkit being developed at Carnegie Mellon, called Dv, for building remote visualization services on computational grids. The long-range goal is develop a general solution to the problem of providing heavyweight Internet services—such as visualization, datamining and advanced search engines—that require significant computation in order to satisfy requests. Initially, Dv is being used to produce a grid-enabled version of vtk for visualizing the output of earthquake ground motion simulations produced by the CMU Quake project.

The Dv toolkit is based on a form of mobile object, called an active frame, that consists of Java and native application data, a Java application program that describes the visualization flowgraph, and a scheduler. Each node in the flowgraph corresponds to a vtk routine. Active frames are processed by active frame servers running on the grid. Each server waits for an active frame to arrive, invokes the scheduler to determine the portion of the flowgraph to execute and the identity of the next host, executes that part of the flowgraph using native vtk routines, constructs an output frame, and sends it to the next host.

A remote Dv visualization consists of a series of active frames that hop from server to server until they reach the client for display. The idea is to provide a flexible and uniform framework for partitioning applications between remote and local systems based on available computing and communication resources, and for studying and deploying application-level schedulers.

Speaker's web page: <http://www.cs.cmu.edu/~droh>

September 13, 1999

Algorithms and Software for Sensitivity Analysis and Optimal Control of Large-Scale Differential–Algebraic Systems

Linda Petzold

petzold@engineering.ucsb.edu

University of California, Santa Barbara

Abstract

In recent years, as computers and algorithms for simulation have become more efficient and reliable, an increasing amount of attention has focused on the more computationally intensive tasks of sensitivity analysis and optimal control. In this lecture we describe algorithms and software for sensitivity analysis and optimal control of large-scale differential–algebraic systems, focusing on the computational challenges. We introduce a new software package, DASPK3.0, for sensitivity analysis, and discuss our progress to date on the COOPT software and algorithms for optimal control which has been used successfully for a wide range of applications including spacecraft trajectory optimization, manufacture of superconducting thin films, and optimization of tissue engineering processes. Results for applications will be described as time permits.

Speaker's web page: <http://eci2.ucsb.edu/mee/petzold>

Institution web page: <http://www.ucsb.edu>

Friday, September 10, 1999

A Graph Partitioning Approach for Irregular Block Decomposition

Jarmo Rantakokko

jarmo@cs.ucsd.edu

University of California, San Diego

Abstract

Irregular block decompositions arise in parallel implementations of structured adaptive mesh refinement (SAMR) as well as in multiblock methods. In this talk we consider the problem of how to efficiently cover and partition an irregular domain arising in ocean modeling. The boundary, i.e., the coastline, is irregular and there may be disconnected islands within the domain. Our objective is to eliminate as many land points as possible without introducing high overhead costs, and to balance the workloads evenly. A similar problem also arises in SAMR where the flagged points cover regions of high error. We have developed a partitioning algorithm based on spectral graph decomposition, which includes tuning parameters for adjusting the granularity of partitioning, and procedures for fine-tuning the partitioning with respect to overhead costs and load imbalance.

We will discuss our algorithm in detail and compare it with two basic partitioning methods, commonly used in ocean modeling. Partitioning with our approach exhibit low communication overheads and load imbalance. We discuss performance results on ASCI Blue-Pacific and the IBM-SP2 and derive an analytic performance model which accurately predicts the running time of the application.

Institution web page: <http://www.ucsd.edu/>

September 7, 1999

Scientific Computing Tools for 3D Field Problems

Gundolf Haase and Ulrich Langer

ghaase@numa.uni-linz.ac.at

ulanger@numa.uni-linz.ac.at

Johannes Kepler University, Linz

Abstract

The 3D magnetic field problems are challenging not only for their interesting applications in industry, but also from the mathematical point of view. Usually, technical 3D magnetic field problems are characterized by complicated interface geometries with potentially moving parts (e.g. rotating parts), large coefficient jumps, non-linearities, singularities, and the necessity of calculating the exterior magnetic field. In practice, the aim of the simulation is often the optimization of the magnetic device we are dealing with. In order to handle such kind of technical magnetic field problems, it is not sufficient to have a fast Maxwell solver and an optimizer; a geometry modeler, an advanced 3D mesh generator, mesh-handling and refinement methods, parallelization tools, and postprocessing tools including advanced visualization techniques also are required.

We will present such pre- and postprocessing tools, specially adapted to the adaptive multilevel methods used in the solver and the optimizer. Of course, the heart of our magnetic field problem solver environment FEPP is an adaptive multilevel Maxwell solver. In the magnetostatic case, the Maxwell solver is based on special mixed variational formulations of the Maxwell equations and their discretization by the Nédélec and Lagrange finite elements. Combining this with an adaptive nested multilevel preconditioned iteration approach, we obtain an

optimal solver with respect to the complexity. This is confirmed by the results of numerical experiments for academic problems and real-life applications as well. We also propose a concept for coupling finite elements with boundary elements. Coupled finite and boundary element schemes are well suited for problems where it is necessary to take into account the exterior magnetic field.

For the parallelization, two different strategies have been developed. The first approach uses a thread-based implementation that is especially suited for shared memory parallel computers such as the ORIGIN 2000. It is highly efficient if small numbers of processors are used. The second concept is based on distributed data algorithms and has been developed for massively parallel computers and workstation clusters.

The algorithms developed are used for solving more challenging engineering problems including 3D transient magnetomechanical problems.

More information on the problem solver environment FEPP and on the presented results can be obtained from the homepage <http://www.sfb013.uni-linz.ac.at> of the special research project (SFB) "Numerical and Symbolic Scientific Computing" that is supported by the Austrian Science Foundation (FWF).

September 3, 1999

Intelligent Miner and the Data Mining Challenge

Balakrishna Iyer

balaiyer@us.ibm.com

IBM Santa Theresa Labs

Abstract

The decreasing cost of computation allows for significant amounts of computing to solve problems previously unaddressed or addressed by other means. The decreased cost of communications bandwidth, and standardization of communications protocols, have made it possible to transport vast amounts of remotely stored/generated data. The decreased cost of storage makes it possible to collect/store/replicate more data.

The defeat of the world's best chess player by a computer, although controversial, is a landmark in data analysis. Intelligent Miner is a toolkit to help data analysts. It has been used successfully to spot fraud in the health sector and to predict financial trends. The talk will describe the toolkit, the kernel algorithms in the toolkit, and the data transformation function it offers.

August 26, 1999

Recent Dynamic Scheduling Strategies for Scientific Computing

Ioana Banicescu
ioana@cs.msstate.edu

Mississippi State University

Abstract

Scientific problems are large, irregular, and result in computationally intensive applications. Although many scientific applications are amenable to parallel execution, high performance is difficult to obtain due to load imbalance, one of the essential performance degradation factors. Load imbalances result from problem, algorithmic, and systemic characteristics. Recently, a number of dynamic scheduling schemes have been proposed and implemented in scientific applications. Some of these schemes (factoring, fractiling, weighted factoring, factoring with adaptive weights) are based on a probabilistic analysis and thus, they accommodate load imbalances caused by predictable phenomena, such as irregular data, as well as unpredictable phenomena, such as data-access latency and operating system interference.

Factoring and fractiling were successfully implemented in Monte-Carlo simulations and N-Body simulations, respectively. Weighted factoring and factoring with

adaptive weights also proved to be cost-effective in computational field simulations using unstructured and adaptive grids. In this talk, I will report on our experiments with these techniques on IBM-SP2 and a SuperMSPARC. Performance of simulation codes was improved by as much as 52%. This improvement underscores the need for scheduling schemes that accommodate load imbalances due to application as well as system induced variances.

Future work will be dedicated to extensions of these methods, and to new techniques that could improve performance of scientific applications on parallel and distributed architectures.

Speaker's web page: <http://www.cs.msstate.edu/~ioana>

Research web page:
<http://www.erc.msstate.edu/~bilderba/research.html>

August 25, 1999

Scientific Visualization in ... Python?

Valerio Pascucci

pascucci@cs.utexas.edu

University of Texas, Austin

Abstract

Scientific visualization systems have a special need for the integration of many heterogeneous modules to perform numerical computations, large data management, graphics and client server communication among diverse platforms over the net.

A scripting language such as Python seem to encompass several interesting characteristics that are useful in the development of a scientific visualization system, especially in an experimental environment where continuous changes to several modules are developed and integrated independently. The missing element that I have been working on is a simple light-weighted visualization framework called "ndVP" in which the integration of independently developed modules becomes a trivial task. The system is designed to take full advantage of the Python flexibility, robustness and portability including the possibility of being modified/extended at runtime. Following the Python philosophy, the visualization framework proposed here involves a small kernel based on few basic classes: Groups, Models and

Display Windows. All the rest will grow as a set of external modules.

After an overview of the ndVP system I will present the details of a new isocontouring algorithm called Progressive Isocontouring that I have been developing as an extension module of ndVP. The Progressive Isocontouring is an isocontouring that incrementally builds a multiresolution representation of an isocontour. Progressive Isocontouring generates a data-structure that is consistent (valid mesh representation, no cracks, ...) at any time so that partial results can be displayed at any intermediate stage during its evaluation. In the ndVP system the Progressive Isocontouring algorithm has been tested using independent threads for contour evaluation and contour visualization. The visualization thread accesses the data-structure that the evaluation thread is building and performs the redraw action at regular intervals. If a request for a new isocontour is issued the main process terminates the isocontour computation thread and starts a new one.

August 24, 1999

A Multigrid Method Enhanced by Krylov Subspace Iteration for Discrete Helmholtz Equations

Howard Elman
elman@cs.umd.edu

University of Maryland

Abstract

Standard multigrid algorithms have proven ineffective for the solution of discretizations of Helmholtz equations. In this work we modify the standard algorithm by adding GMRES iterations at coarse levels and as an outer iteration. We demonstrate the algorithm's effectiveness through theoretical analysis of a model problem and experimental results. In particular, we show that the combined use of GMRES as a smoother and outer iteration produces an algorithm whose performance depends relatively mildly on wave number and is robust for normalized wave numbers as large as 200. For fixed wave numbers, it displays grid-independent convergence rates and has costs proportional to number of unknowns.

Speaker's web page: <http://www.cs.umd.edu/~elman/>

August 20, 1999

Interactive Visualization of Medical and Biological Data Sets

Joerg Meyer

jmeyer@ucdavis.edu

University of California, Davis

Abstract

Interactive visualization of large medical and biological data sets requires advanced techniques in image processing and data reduction. We have designed a platform-independent rendering system that allows the user to interact with an object on a screen with no delay, i.e., in the same way as he or she would handle an object in the real world. The system provides a user-friendly user interface and a pipeline architecture that maintains interactive behavior independent from the amount of data, which must be processed at a given time frame.

We describe an interactive rendering system that uses hierarchical, multiresolution representations of the data set, and scalable rendering algorithms. The system, called InVIS, is designed to always guarantee a certain frame rate by introducing a time-control unit.

The rendering pipeline accepts input data from CT or MRI scanners, confocal laser scanning microscopes, CAD systems, and other imaging techniques. A flexible import module makes it easy to customize the system and adapt it to user-specific needs. Due to hierarchical refinement, the system ensures interactive response at any time. The rendering stage is based on OpenGL, which makes it easily portable to different hardware platforms. Ongoing research focuses on large and time-variant data sets.

August 20, 1999

Topological Vector Field Visualization with Clifford Algebra

Gerik Scheuermann
scheuer@ucdavis.edu

University of California, Davis

Abstract

Topology is an important tool for the analysis of vector fields. In the talk, it will be shown that Clifford algebra can help in finding features such as non-linear critical points. Clifford algebra is a mathematical model extending the usual vector space description of geometry by a multiplication of vectors. A description will be given only to enable understanding the visualization results.

In a second part, the effect of a triangulation on the complexity of interpolated flows is analyzed, as are the ways in which this complexity can be reduced by simple changes in the grid connectivity. The complexity here is just the number of critical points in the flow. Its dependence on the triangulation is used to define a data-dependent triangulation that reduces this number by flipping edges. This can also be used for the triangulation of curvilinear grids.

The last part shows the topological results from using C1-interpolation as vector field model. It will be demonstrated that a higher order interpolation enhances the local topology extraction compared to linear methods.

August 17, 1999

Reduction of NO_x in Lean Exhaust by Selective NO_x-Recirculation (SNR-Technique) and an Outlook to the Progress in Soot Modeling for Engines

Petra Stapf

petra.stapf@daimlerchrysler.com

Daimler Chrysler AG

Abstract

The Selective NO_x-Recirculation Technique (SNR-Technique) is a new NO_x aftertreatment system for lean-burn gasoline and diesel applications. The objective of SNR is NO_x removal from lean exhaust gas by NO_x adsorption and subsequent selective external recirculation and decomposition of NO_x in the combustion process. The SNR project is composed of two major parts. First is the development of NO_x adsorbents that are able to store large quantities of NO_x in lean exhaust gas, and second is the NO_x decomposition by the combustion process.

A recent soot-formation model using detailed chemistry was modified for high pressure applications. Since the required computing times do not allow a direct implementation in engine codes, the soot model was incorporated into a dedicated group combustion model to assess its performance under high-pressure conditions. The combined model is capable of mimicking the microscopic details of combustion and pollutant formation in diesel engines.

July 29, 1999

A Parallel, Block, Jacobi–Davidson Implementation for Solving Large Eigenproblems on Coarse-Grain Environments

Andreas Stathopoulos

andreas@cs.wm.edu

College of William and Mary

Abstract

Iterative methods often provide the only means of solving large eigenvalue problems. Their block variants converge slowly but they are robust, particularly in the presence of multiplicities. Preconditioning is often used to improve convergence. Yet, for large matrices, the demands posed on the available computing resources are huge.

Clusters of workstations and SMPs are becoming the main computational tool for many scientific and engineering groups. In these environments, high communication costs suggest coarse grain implementations of iterative methods.

We have combined the block and preconditioning functionalities into a parallel implementation of a block Jacobi–Davidson method. We combine a fine-grain iterative scheme with the coarse grain capability that each processor can precondition different eigenpairs. We outline the design and present some timings and convergence results on a small workstation cluster and on a SUN Enterprise.

Speaker's web page: <http://www.cs.wm.edu/~andreas/>

July 26, 1999

The Computation of Bounds for the Norm of the Error in the Preconditioned Conjugate Gradient Method

Gerard Meurant

meurant@bruyeres.cea.fr

Commissariat à l'énergie atomique

Abstract

In this talk we will first recall how to compute bounds for the A-norm of the error in the conjugate gradient (CG) method. This involves expressing the norm of the error as a Stieltjes integral. Then, quadrature formulas are used to approximate the integral. The Lanczos method is used to generate the nodes and weights of the quadrature rules. We will show how to use these techniques with CG and illustrate this by numerical experiments. Finally, we will extend these methods to the preconditioned version of CG and give numerical experiments showing that we can obtain very accurate upper and lower bounds for the norm of the error provided we have estimates of the smallest and largest eigenvalues of the matrix. This can be obtained adaptively during the first CG iterations leading to a reliable criterion for stopping PCG iterations.

July 23, 1999

Calculation of the Overall Magnetic Properties of Magnetorheological Fluids

Fernando Reitich
reitich@math.umn.edu

University of Minnesota

Abstract

Magnetorheological fluids (MRF), composed of micron-sized polarizable particles dispersed in a carrier liquid, constitute examples of controllable (“smart”) fluids whose rheological properties vary in response to an applied magnetic field. Understanding the magnetic behavior of MRF is crucial to the development of MRF-based devices and it also provides valuable insight into the character of the microstructure responsible for their field-dependent rheology. In this talk we will present results on the calculation of the overall magnetic response of MRF. We will show that effective medium approximations and particle dynamics simulations deliver numerical results that are in good agreement with experimental data and which can, therefore, be used to assist in the design of improved MRF.

Speaker's web page: <http://www.math.umn.edu/~reitich/>

Research web page:
<http://www4.ncsu.edu/~hvly/dynam.html>

July 16, 1999

COMPOSE: An Object-Oriented Framework for PDE Solvers

Krister Ahlander

krister@tdb.uu.se

Uppsala University

Abstract

The issue of how to develop reusable software in scientific computing is addressed. With object-oriented analysis and design an extendable set of collaborating objects, a framework named COMPOSE, has been developed for the numerical solution of partial differential equations.

Speaker's web page: <http://www.tdb.uu.se/~krister/>

Research web page:
<http://www.tdb.uu.se/research/swtools/>

July 12, 1999

Coupling Discretization Methods on Non-Matching Grids

Raytcho Lazarov
lazarov@rieman.llnl.gov

Texas A&M University

Abstract

The mortar method for coupling various approximation techniques has become an important tool in the construction and analysis of discretization schemes on non-matching grids.

First, we discuss briefly the main ingredients, issues, advantages, and the price we have to pay when using this approach. Second, on a differential level we introduce two hybrid formulations, which are the bases for the mortar and non-mortar approximations on non-matching grids.

Third, we discuss mortar finite volume methods and mortar mixed finite element discretizations and present the main steps in the stability and error analysis. Finally, we discuss iterative methods for solving the corresponding linear system and present some numerical experiments on model second-order problems.

July 9, 1999

An Extensible Wrapper Construction System for Internet Information Sources

Calton Pu

calton@cse.ogi.edu

Oregon Graduate Institute

Abstract

The amount of useful semi-structured data on the web continues to grow at a stunning pace. Often interesting web data are not in database systems but in HTML pages, XML pages, or text files. The data in these formats is not directly usable by standard SQL-like query processing engines that support sophisticated querying and reporting beyond keyword-based information retrieval. Hence, the web users or applications need a smart way of extracting data from these web sources. One of the popular approaches is to write wrappers around the sources, either manually or with software assistance, to bring the web data within the reach of more sophisticated query tools and general mediator-based information integration systems.

In this talk, we describe the methodology and the software development of a semi-automatic wrapper construction toolkit XWrap, which provides an XML-enabled, feedback-based, interactive wrapper construction facility

for generating value-added wrappers for Internet information sources. By XML-enabled, we mean that the extraction of information content from the web pages will be captured in an XML-compliant format and the query-based content filtering process is performed against XML documents. By feedback-based, we mean that the wrapper construction process is feedback-driven and the wrapper programs generated by XWrap can be incrementally revised and tuned according to the feedback collected through inductive learning by the wrapper generator. By value-added, we mean that any XWrap-generated wrapper program may accept more sophisticated queries than the corresponding web source can take. XWrap enables such enhanced functional capabilities through a smart composition of data wrapping and function wrapping. We will demonstrate the benefits of using wrappers generated by the XWrap in the context of the Continual Queries project.

Speaker's web page:

<http://www.cse.ogi.edu/~calton/>

July 7, 1999

Automated Feature Extraction from Transient CFD Simulations

Robert Haimes
haimes@orville.mit.edu

Massachusetts Institute of Technology

Abstract

In the past, feature extraction and identification were interesting concepts, but not required to fully understand the underlying physics of a steady-state flow field. This is because the results of the traditional visualization tools like iso-surfaces, cuts and streamlines are interactive and the results are easily represented to the user. These tools worked and properly conveyed the collected information at the expense of a great deal of interaction. For unsteady flow fields, the investigator does not have the luxury of spending time scanning only one “snap-shot” of the simulation. Automated assistance is required in pointing out areas of potential interest contained within the flow. This must not require a heavy compute burden (the extraction should not significantly slow down the solver for co-processing environments). With equal importance, methods must be developed to abstract the feature and display it in a manner that makes sense physically and is easily understood.

This talk presents the current status on locating shocks, vortices, separation surfaces (regions of recirculation) and boundary layers (and wakes) in steady-state and transient regimes.

June 18, 1999

A Particle-Partition of Unity Method for the Solution of Elliptic, Parabolic and Hyperbolic Partial Differential Equations

Alexander Schweitzer
schweitz@iam.uni-bonn.de

Universitaet Bonn

Abstract

In this paper, we present a meshless discretization technique for nonstationary convection–diffusion problems. It is based on operator splitting, the method of characteristics and a generalized partition of unity method. We focus on the discretization process and its quality. The method may be used as an h- or p-version. Even for general particle distributions, the convergence behavior of the different versions corresponds to that of the respective version of the finite element method on a uniform grid. We discuss the implementational aspects of the proposed method. Furthermore, we present the results of numerical examples, where we considered instationary convection-diffusion, instationary diffusion, linear advection and elliptic problems.

June 15, 1999

Tools and Techniques for High-Performance and Distributed Data Mining

Robert Grossman
grossman@uic.edu

University of Illinois at Chicago

Abstract

Data mining is the semi-automatic discovery of patterns, correlations, associations, changes and anomalies in large data sets. During the past two decades, the amount of data has grown explosively, yet the number of statisticians and other data analysts has remained relatively constant. There are only two possibilities: We either ignore larger and larger amounts of data, or we develop tools and techniques to aid in its automated analysis. In this talk, we give an overview data mining from this perspective and discuss some techniques for mining large and distributed data sets. We also provide an overview of some of the infrastructure that can make the mining of these types of data sets more practical.

June 14, 1999

An Adaptive Mesh Refinement Scheme for Hyperbolic Equations with Stiff Source Terms

Donald Schwendeman

Email: schwed@rpi.edu

Rensselaer Polytechnic Institute

Abstract

Multilevel adaptive mesh refinement (AMR) techniques have become popular for the numerical solution of nonlinear hyperbolic equations in which sharp gradients and/or discontinuities (shock, detonations, etc.) in the solution develop and propagate. The basic idea is to refine the computational grid in regions where the solution changes rapidly in order to resolve these structures and to achieve high accuracy.

The talk will focus on current research aimed at the development of an AMR scheme for the numerical solution of hyperbolic equations with or without stiff source terms. The present scheme uses a second-order finite volume method on a structured hierarchical grid. The refinement is done by identifying grid patches where the solution changes rapidly in either space or time. Each

grid patch, in turn, can be refined so that the refinement scheme is recursive. Interpolation is used to generate fine grid patches from coarse grid data and to provide boundary data at fine-coarse grid boundaries. A description of the AMR scheme will be given for the case of unsteady, compressible, inviscid flow of a binary mixture of burnt and unburnt fuel as modeled by the reactive Euler equations.

Several examples will be used to illustrate the method and the special issues related to the numerical treatment of the stiff source term. Of particular interest will be the study of ignition and the evolution to detonation provoked by a small initial temperature gradient in an explosive sample, and the interaction of detonations created by several initial "hot spots."

June 9, 1999

Mesh Generation by Delaunay Refinement

Jonathan Shewchuk
jrs@cs.berkeley.edu

University of California, Berkeley

Abstract

Delaunay refinement is a technique for generating unstructured meshes of triangles or tetrahedra suitable for use in such applications as

- Graphics (especially radiosity).
- Terrain databases and Geographical Information Systems.
- Function interpolation.
- Numerical methods such as the finite element method.

In response to the needs of Carnegie Mellon's Quake project, an interdisciplinary Grand Challenge project whose goal is to simulate earthquake-induced ground motion in the Los Angeles basin, I have produced two general-purpose mesh generators. Triangle and Pyramid are implementations of two- and three-dimensional Delaunay refinement, respectively. Triangle has been released to the public, and has hundreds or thousands of users in both research and industrial settings, who use Triangle for all the applications listed above and more.

In this talk, I discuss the algorithmic underpinnings and current state of the art of Delaunay refinement methods. Popularized by the engineering community in the mid-1980s, Delaunay refinement operates by maintaining a Delaunay triangulation or Delaunay tetrahedralization, which is refined by the insertion of additional vertices. The placement of these vertices is chosen to force the mesh to conform to an object's boundaries, and to improve its quality as a grid for interpolation.

I present a new three-dimensional Delaunay refinement algorithm, which builds on the pioneering two-dimensional work by L. Paul Chew and Jim Ruppert. My algorithm produces tetrahedral meshes that are nicely graded (in a provable sense) and whose tetrahedra have circumradius-to-shortest edge ratios bounded below 1.63. This theoretical guarantee ensures that all poor quality tetrahedra except slivers (a particular type of poor tetrahedron) are removed. The slivers that remain are easily removed in practice, although there is no theoretical guarantee.

May 27, 1999

Iterative Methods for Solving Linear Systems

Gene Golub

golub@sccm.stanford.edu

Stanford University

Abstract

We discuss several problems in connection with solving linear systems. First, we consider various pre-conditioners for solving a variety of problems. This includes pre-conditioners for solving indefinite systems, and non-symmetric problems, especially when the skew-symmetric part of the matrix is dominant. Next, we describe the use of inner and outer iteration methods for solving problems where it is not possible to solve the equations using the pre-conditioner exactly. We attempt to restore the rate of convergence of the problem. Finally, we analyze the convergence properties of a method when the initial vector is considered to be a random vector. We give some numerical examples.

May 26, 1999

The Blob Projection Method for Immersed Boundary Problems

Michael Minion

minion@amath.unc.edu

University of North Carolina

Abstract

This talk will examine the problem of numerically approximating the evolution of a thin elastic membrane immersed in a constant density incompressible fluid. Immersed boundaries are currently used in modeling biological systems such as the beating of the heart and the swimming of flagellated organisms. A numerical method developed in collaboration with R. Cortez of Tulane will be presented which represents the membrane by a collection of regularized point forces while solving the incompressible Navier–Stokes equations on a regular Cartesian grid. The effects of the membrane on the fluid are calculated at grid points via a smoothed dipole potential and fast summation techniques, while the fluid equations are solved using a projection method. Comparisons between this method and the traditional immersed boundary method due to Peskin will also be presented.

May 24, 1999

Multigrid on 3D Adaptive Sparse Grids

Pieter Hemker

pieth@cwi.nl

Centrum voor Wiskunde en Informatica

Abstract

We show some aspects of adaptive multigrid (MG) for the solution of 3D partial differential equations (PDEs). In particular, we discuss different possibilities of using partially ordered sets of auxiliary grids in multigrid algorithms. Because, for the classical sequence of grids, with 3D regular grids the number of degrees of freedom grows much faster with the refinement level than for 2D, it is more difficult to find sufficiently effective relaxation procedures and it makes sense to study the possibility of using larger sets of (regular rectangular) auxiliary grids.

Semi-coarsening is one technique in which a partially ordered set of grids is used. In this case still a unique fine-grid discrete problem is solved. On the other hand, sparse grid techniques may be more efficient if we compare the accuracy obtained with the number of degrees of freedom used. In the latter case, however, it is not straightforward to identify a discrete equation to be solved. Different approaches are compared.

We show different multigrid strategies, and results are given for transonic Euler-flow over the ONERA M6-wing and for a singular perturbation problem.

May 20, 1999

Theory, Numerics and Applications of Dynamic Sensitivity Analysis

Paul Barton
pib@mit.edu

Massachusetts Institute of Technology

Abstract

Dynamic sensitivity analysis has many applications in model reduction, model parameter estimation, controller design and numerical optimal control. The original existence and uniqueness results for dynamic sensitivities date back to Gronwall's work early this century. In this work, we extend Gronwall's results to include nonlinear ODEs with general consistent initial conditions, and linear time invariant DAEs.

Furthermore, we develop a framework, and existence and uniqueness results, for sensitivity analysis of hybrid discrete/continuous dynamic systems with these differential systems embedded. Highlights of these analyses include the qualitative "jumping" of sensitivities at implicit discontinuities, and nonexistence of sensitivities at critical parameter values related to qualitative changes in the solution of the original dynamic system.

From the numerical standpoint, we discuss a novel and efficient algorithm for computing sensitivities simultaneously with BDF integration of the ODEs/DAEs. This algorithm is then extended to compute the sensitivities of hybrid systems. We present ABACUSS, a high level equation-based modeling environment for simulation and sensitivity analysis. The above theoretical results have some important practical implications for the development of robust general-purpose sensitivity software.

Finally, we discuss some practical applications of sensitivity analysis using ABACUSS, including model development and reduction for a citric acid process, and numerical optimal control. We conclude with some speculation concerning the role of sensitivity analysis in numerical technologies for the global solution of mixed-integer dynamic optimization problems.

May 18, 1999

Building Really Huge Scalable Storage Systems

Ethan Miller

elm@csee.umbc.edu

University of Maryland, Baltimore County

Abstract

In recent years, we have seen rapid increases in processor speed, memory capacity, and storage capacity. However, file system bandwidth has not kept up because of various limitations, both hardware and software. This talk will describe the storage challenges posed by slow hardware and inefficient hardware, and the research in parallel file systems and mass storage that we are doing to address these issues.

The first part of the talk will describe RAMA, our massively parallel file system. We are building a prototype whose bandwidth goal is 1 GB/sec using 100–200 commodity disks on commodity PC's connected by a Fast Ethernet switch. We can achieve this bandwidth by eliminating all centralized operations on file reads and writes, and by guaranteeing file system consistency even when file metadata is not written out immediately, reducing the number of disk I/Os per file block written. In addition to providing high bandwidth, this file system will parity-pro-

tect individual files, allowing a per-file tradeoff between performance and protection against disk failure. Moreover, this file system can accommodate large numbers of small files along side gigabyte (and larger) files without reducing performance on large files. We are currently testing a small prototype of this file system, and hope to deploy a larger prototype over the summer.

The second part of the talk will discuss the issues facing builders of petabyte (and larger) file systems. In particular, designers of such systems must deal with bandwidth that will inevitably lag behind increases in storage density. This problem affects access bandwidth, file system backup, and migration to newer storage technologies. I will quantify the problem and describe possible approaches to addressing it. Because this part of the talk is more speculative, feedback on both issues and directions will be welcome.

May 14, 1999

Integrating Archival Storage into Information Management Systems

Chaitanya Baru
baru@SDSC.EDU

University of California, San Diego

Abstract

Organizations are increasingly having to deal with archiving of large amounts of data over long periods of time. This has been a well known requirement in computational science and supercomputing applications where much effort is expended in creating large data collections (e.g. observational data, simulation outputs), and where it is typical for many applications to routinely generate large outputs.

We will discuss issues in combining the functionalities of archival storage systems with database management systems to efficiently support search and management of very large data collections. Archival storage systems such as the High Performance Storage System (HPSS) are designed to efficiently manage relatively small numbers (e.g. millions) of large data sets (e.g. hundreds of MB and higher). Combining DBMS technology with archival systems can allow efficient management of collections con-

taining large numbers (e.g. tens to hundreds of millions) of relatively small data sets (e.g. a few KB, or less, each). An example of such a system is the DB2/HPSS project, which is a joint effort between the IBM T.J. Watson Research Center and San Diego Supercomputer Center (SDSC), where the DB2 UDB system has been “integrated” with HPSS. We will provide details of this system and also describe other approaches to handling the “small data problem” in HPSS, including the “container” concept being developed for the Storage Resource Broker middleware at SDSC.

Finally, we will describe some projects at SDSC that are employing the above technologies. These include a project with the USPTO where the data collection is 2 million US patents, and a project with the California Digital Library (CDL), where the collection is about 100,000 art images from the Art Museum Image Consortium (AMICO).

May 10, 1999

Adaptive Operating Systems: An Architecture for Evolving Systems

Barton Miller
bart@cs.wisc.edu

University of Wisconsin

Abstract

Operating systems formerly were viewed as static entities, changing almost as slowly as the underlying hardware. Recent commercial systems have provided for a small amount of run-time change by allowing device drivers to be installed while the system is running. We are developing "adaptive operating systems" whose code can change and evolve while the system is running. This adaptation can be used to instrument the code for profiling or debugging purposes, or to modify and extend the operating system to adapt to changing work loads, application-demands, and configuration.

Our work differs from other efforts in this area in two ways. First, we can instrument and modify a stock, commercial operating system (Solaris) as it was delivered to the customer. We operate directly on the executable code while it is running. Second, we can modify the operating system at almost any point in its code. We are not constrained to system call or procedure call replacements.

Our research is embodied in a facility called KernInst. We will describe the basic KernInst mechanism and several uses of KernInst, including performance profiling and dynamically modifying and customizing the operating system in response to its work load. We will also present a case study using KernInst to profile the Solaris kernel under a Web proxy server (Squid) workload.

May 7, 1999

Lowes, Highs, Jets & Storms: A Mathematical View of the Weather

David Muraki
muraki@cims.nyu.edu

New York University

Abstract

The most prominent features of the North American weather pattern, as often highlighted in TV weather segments, are the (west-to-east) jetstream and pressure systems of relative lows and highs. These behaviors are also reproduced in computations of simple weather models. The mathematics of the atmosphere is really a problem in fluid mechanics, but specialized to the situation where density/temperature effects and Coriolis force (due to the Earth's rotation) play dominant roles.

The basic equations for the midlatitude atmosphere are best understood in the limit of small Rossby number, an important theory known as "quasi-geostrophy." As a leading-order asymptotic theory, quasi-geostrophy captures much of the dynamics of the atmosphere, including the fundamental instability of the jetstream

to pressure cells (falling pressure means bad weather). This theory, however, fails to explain the preferential development of low-pressure cells that lead to the asymmetric spirals characteristic of baroclinic storm systems.

By resolving a degeneracy inherent in this small Rossby number limit, a systematic asymptotic extension to the theory of quasi-geostrophy has been developed. Our computations of the next-order corrections are in excellent agreement with simulations of the full model equations, and provide a basis for explaining the asymmetric development of storms. Applications to other atmospheric phenomena, in collaboration with meteorologists at the National Center for Atmospheric Research (NCAR Boulder), are also discussed.

May 4, 1999

Cache-Based Multigrid Algorithms

Craig Douglas
douglas@ccs.uky.edu

University of Kentucky

Abstract

Multigrid methods combine a number of standard sparse matrix techniques. Usual implementations separate the individual components (e.g., an iterative methods, residual computation, and interpolation between grids) into nicely structured routines. Today, however, many computers employ quite sophisticated and potentially large caches whose correct use is instrumental in gaining much of the peak performance of the processors. This is true independent of how many processors are used in a computation.

We investigate when it makes sense to combine several of the multigrid components into one, using block-oriented algorithms. We determine how large (or small) the blocks must be in order for the data in the block to just fit into the processor's primary cache. By re-using the data in cache several times, a potential savings in run time can be pre-

dicted. This is analyzed for a set of examples. It is surprising to see how large a subdomain can fit into a relatively small, well designed cache (e.g., 256KB). As caches continue to increase in size, the ideas here will extend quite nicely to three-dimensional problems. In particular, machines with 1–16MB caches are under design which will make three-dimensional cache-based multigrid a reality.

While most of the savings in time are with respect to the approximate solver, using a multiplicative or additive domain decomposition method saves even more time and allows us to use theoretical convergence rates from that field for our problems.

An automatic, low overhead software tool for determining a good cache based ordering for the multigrid components for general problems will also be discussed.

May 3, 1999

On the Accuracy and Efficiency of Coupling Different Grids and Discretizations in Flow in Porous Media Simulations

Ivan Yotov

yotov@math.pitt.edu

University of Pittsburgh

Abstract

Combining different grids and discretizations in a single numerical simulation has gained much popularity in recent years. It provides means for efficient modeling of irregular geometries, internal features, and local solution behavior. Critical for the success of this approach is imposing physically meaningful interface matching conditions in a numerically stable and accurate way. In this talk we consider flow in multiblock porous media domains discretized by mixed finite element methods on non-matching grids. We propose several mortar schemes (using specially chosen interface finite elements) and a non-mortar scheme for coupling along interfaces. The accuracy and efficiency of these approaches are compared theoretically and numerically. Efficient parallel implementation using a non-overlapping domain decomposition is discussed. The algorithms are illustrated by numerical simulations of multiphase flow in irregular heterogeneous domains.

April 29, 1999

Recent Advances in Morphological Methods for Image Segmentation and Feature Extraction

Luc Vincent

lucv@adoc.xerox.com

Xerox Corporation

Abstract

Mathematical morphology is an image processing methodology that was born in France in the late Sixties and has been gaining increasing importance and popularity since then. Today, it is broadly accepted as a powerful alternative to traditional linear techniques. Its rich set of operators is especially well suited for complex segmentation and feature extraction problems. Morphology has been successfully applied in many areas, including medical imaging, biology, radar, sonar, infra-red, and remote sensing images, industrial inspection, material science, fingerprints, identification, document recognition, etc.

The morphological approach to image analysis is natural and attractive: Binary images are considered as sets, whereas grayscale images are viewed as functions or topographic reliefs. These sets and functions are then transformed—in the spatial domain—via morphological operators, whose definitions are usually based on structuring elements, i.e., particular shapes that are translated in images and used as probes.

The presentation will start with a quick introduction to morphology's most well-known operations, namely erosions, dilations, openings, and closings, which are quite

useful for simple noise filtering and feature extraction tasks. We then move on to “granulometries.” These operations, based on openings and closings, are particularly interesting for extracting size information in images, for feature extraction, or for texture characterization. In addition, recent algorithmic advances have cut down their computation time by orders of magnitude, thereby opening up a range of new potential applications for these operators. Finally, we focus on segmentation. The watershed transformation forms the cornerstone of the morphological approach to segmentation. By constraining watersheds using previously extracted object markers (or “seeds”), a particularly powerful segmentation methodology is derived. Analyzing cases where watersheds do not perform well, However, leads us to propose alternative methods based on gray-values distance transforms and minimal paths. Among others, these methods are very efficient for the extraction of faint linear structure in grayscale images.

Though the talk covers a broad range of topics, it does not require any particular background in morphology from attendees. Operations will be described in an intuitive manner, and a large number of examples of applications will be shown.

April 28, 1999

GFS: A New File System Architecture for Network Devices

Matthew O'Keefe
okeefe@ee.umn.edu

University of Minnesota

Abstract

In computer systems today, speed and responsiveness are determined by network and storage subsystem performance. Faster, more scalable networking interfaces such as Fibre Channel and Gigabit Ethernet provide the scaffolding upon which higher-performance implementations may be constructed, but new thinking is required about how machines interact with smart network and storage devices

To this end, the Global File System (GFS) group's goal is to develop GFS, a scalable server-less file system that integrates IP-based network attached storage (NAS) and Fibre-Channel-based storage area networks (SAN). We call this new architecture Storage Area InterNeTworking (SAINT). It exploits the speed and device scalability of SAN clusters, and provides the client with the scalability and network interoperability of NAS appliances. GFS works under Linux and IRIX, though 95% of our efforts are now focused on Linux. GFS is freely-available, open source code released under the GNU General Public License. In this talk I will describe how GFS can provide the underlying parallel file system in large Linux clusters.

More information on GFS can be found at
<http://gfs.lcse.umn.edu>.

April 27, 1999

Java Grande: Prospects for High-Performance Computing Using Java

Dennis Gannon
gannon@cs.indiana.edu

NASA Ames and Indiana University,
Dept. of Computer Science

Abstract

This talk will present an overview of the Java Grande Forum and its valiant efforts to make Java safe for high-performance computing. Java Grande consists of two working groups. The forum is led by Geoffrey Fox. One working group is focused on the issue of numerics and the associated language features deemed necessary to make it a productive tool for large scale scientific computation. The other working group is focused around the area of parallelism and concurrency. This subgroup is working on ways to improve the efficiency of Java serialization and remote method invocation. Members are also working on proposals for Java MPI, though the talk will not dwell on this topic.

April 20, 1999

Alpha-Linux-ServerNet HPC Clusters and Visualization

Alan Heirich

alan.heirich@compaq.com

Compaq Tandem Laboratories

Abstract

This talk will present a strategy being executed by the Tandem division and Custom Systems Section of Compaq to develop a "Beowulf-class" cluster computing technology, and a related effort in scalable visualization. The clusters are built using the ServerNet-II SAN, a commodity HPC network technology with 350 MB/s bandwidth and application-to-application software latencies below 10 microseconds.

As part of this effort we are developing a scalable image combining architecture using 100% COTS components. The architecture computes the pixel-by-pixel merge of a set of images. It achieves low latency and high frame rates by using data compression and aggressive pipelining. The architecture uses gigabit crossbar switching in a fat tree topology to implement operator associativity, and uses an in-band control protocol to support high-speed mode changes in multi-pass rendering algorithms. The pixel combining function is user-specifiable and may involve pixel components (r,g,b,z,alpha,stencil) both as operands and as results. One use of the architecture is for interactive visualization of large-scale problems in computational science running on Beowulf-class workstation clusters. These clusters may contain dozens or hundreds of computational nodes and this architecture is designed to achieve interactive latency at these scales. This talk will discuss some requirements of algorithms for volume and vector field visualization and for photo-realistic lighting calculations.

April 14, 1999

Automatic High-Performance Fluid Computations in Complex Geometry

Marsha Berger

berger@nyu.edu

New York University

Abstract

We review aspects of our work on efficient and robust computational methods for computing moderate to high Mach number flows in complicated geometries. Our approach uses regular Cartesian meshes where the geometry intersects the mesh in an essentially arbitrary way. We draw on a variety of techniques in developing this approach. Adaptive mesh refinement is used to concentrate the computational effort in regions requiring greater resolution. The fluid flow equations need to be discretized in a special way to handle the cut cells where the geometry intersects the grid. Efficient algorithms from computational geometry allow us to find the cut cells and set up the discretization directly from a triangulation of the surface. We have taken great care to structure the geometric and flow computations so as to make the best use of cache-based moderately parallel computers.

April 1, 1999

pV3 — parallel Visual3k

Robert Haimes

haimes@orville.mit.edu

Massachusetts Institute of Technology

Abstract

pV3 builds heavily on the technology developed for Visual3. Visual3 is an interactive environment for visualization of 3D structured or unstructured data, on workstations with graphics hardware. The data may be steady or time-varying. Using Visual3, one can examine and probe a computational grid, as well as, display isosurfaces, cutting planes, and number of vector visualization functions such as streamlines and ribbons. Visual3 was originally developed for CFD applications.

pV3 is specifically designed for visualization of distributed data using a coprocessing paradigm. To create the visualization, pV3 runs simultaneously on a graphics workstation and an MPP, communicating by pvm. The workstation controls and displays the visualization. On the MPP, pV3 is coupled with the simulation and extracts the desired data for visualization and sends it to the workstation. pV3 has been designed to minimize network traffic. The client-side library extracts lower dimensional data required by the requested visualization tool from the volume of data in place.

March 25, 1999

Experimenting with ZPL

Larry Snyder

snyder@cs.washington.edu

University of Washington

Abstract

ZPL is a recently developed programming language designed to be fast, portable and convenient for scientific and engineering computations. ZPL has achieved these goals — experimental evidence will be presented in the lecture — because it was developed from first principles. Specifics of ZPL's machine and performance models will also be covered. The message is that programmers can write efficient, machine independent parallel programs without resorting to message passing.

ZPL was released in July 1997 and has been in regular use by a modest, but growing user community. Some users are simply solving their research problems. Others are experimenting with it because ZPL is the data parallel subset of a fully general language, Advanced ZPL. They want to become familiar with ZPL's approach in order to influence A-ZPL's design. The lecture will describe the evolutionary path to Advanced ZPL, with emphasis on its direct support for sparse arrays.

March 22, 1999

Caching Dynamic Web Objects and Measuring Network Performance

Mei-Ling Liu
mliu@csc.calpoly.edu

California Polytechnic State University

Abstract

In this talk, the presenter describes two of her projects at Cal Poly where, she conducts research with the assistance of Master's thesis students.

The first project studies the known techniques for reducing the latency on the web by using "Optimistic Delta", an approach developed by Lucent Technologies. A version of the approach is being implemented on the Cal Poly systems, and the performance impact of the technique is being measured. Our hope is to come up with improvements on the existing techniques.

In another project, a group of students performs measurements on the latency, throughput, and CPU utilization of the TCP/IP protocol stack on the Microsoft NT platform. The goal of the project is to develop a model to characterize the behavior of the protocol stack, in an attempt to maximize the network throughput on NT systems.

March 17, 1999

A Investigation of Myrinet, Linux, and Cost for High-Performance Scientific Computation

Peter Beckman

beckman@lanl.gov

Los Alamos National Laboratory

Abstract

Linux, the free, Open Source operating system, has seen explosive growth the last couple of years. According to a study from International Data Corporation, in the past year, the market share of Linux leapt from 6.8% to an estimated 17.2% of server operating system shipments. At Los Alamos National Laboratory, Linux has made similar inroads. Linux is being used lab-wide for desktop workstations, software development platforms, file and computation servers, and web and email servers.

Furthermore, Linux has proved to be extremely versatile for high-performance parallel computation on clusters of workstations. In 1994, NASA pioneered the use of Linux for building extremely cheap clusters with the Beowulf project. Since that time, dozens of laboratories, universities, and companies have been exploring Linux for high-performance computation. At Los Alamos, we have been exploring the performance and stability of our "Little Blue Penguin" cluster, a 64 dual-node (128 CPUs) Myrinet-connected parallel computer, across several different scientific applications. We have also looked at the cost effectiveness of high-speed interconnects for Linux clusters, and the current software available for parallel scientific programming.

March 16, 1999

A Framework for the Real-Time Walkthrough of Massive Models

Dinesh Manocha

dm@cs.unc.edu

University of North Carolina, Chapel Hill

Abstract

Computer-aided design (CAD) applications and scientific visualizations often need user-steered interactive displays or walkthroughs of very complex environments. Structural and mechanical designers often create models of ships, oil platforms, spacecraft, process plants and urban environments whose complexity exceeds the interactive visualization capabilities of current graphics systems. Yet for such structures the design process (and especially the multi-disciplinary design review process) benefits greatly from interactive walkthroughs.

In this talk we give an overview of our recent work on interactive walkthrough of massive models. These include new algorithmic approaches for accelerating rendering based on visibility culling, model simplification and image-based representations. We also present a framework for rendering very large models at nearly interactive rates. The framework scales with model size. Our framework can integrate multiple-rendering acceleration techniques, including visibility culling, geometric levels-of-detail and image-based approaches. We describe the database representation scheme for massive models used by the framework and a pipeline to manage the allocation of system resources among different techniques. We demonstrate the system of a coal-fired power plant composed of more than 15 million triangles.

March 5, 1999

Methods for Hyperbolic Systems with Stiff Relaxation

Robert Lowrie
lowrie@lanl.gov

Los Alamos National Laboratory

Abstract

A major challenge for numerical methods is the solution of hyperbolic systems with stiff relaxation terms. Applications include chemically reacting flows, multi-phase flows, elastic-plastic solid dynamics, gas kinetics, and radiation transport. It is usually straightforward to develop methods that are accurate when all of the time and length scales are resolved. When the time scale of the relaxation is unresolved, however, it is well known that many popular methods fail in that they are inconsistent with the correct asymptotic behavior (the Chapman–Enskog, continuum, or diffusion limit). In applications such as radiation hydrodynamics, it is critical that the correct limit be represented, including an accurate resolution of the diffusion term.

In this talk I will focus on shock-capturing, high-resolution (“Godunov-based”) schemes. Even if the relaxation term is treated implicitly, unless an intricate time-stepping scheme is used (such as that due to Jin), conventional finite-volume methods have the wrong asymptotic behavior. However, using a simple time-integration method, the discontinuous Galerkin method obtains the correct diffusion limit for many problems. Results will be shown for radiation hydrodynamics, a simple elastic-plastic model, and the Broadwell model of gas kinetics. I will also discuss the requisite modification of the flux solver.

March 4, 1999

High-Order Spectral Elements on Triangles and Other Domains

Beth Wingate
wingate@lanl.gov

Los Alamos National Laboratory

Abstract

In this work we explore high-order spectral finite elements on domains such as triangles and tetrahedra. In the quadrilateral and cube, both the space (the tensor products of Legendre polynomials), and the points (the Gauss-Lobatto points), are well known. In the simplex, however, things are not as well understood. Current methods often give ill-conditioned matrices. Preconditioning can fix some of this at moderate order, but at higher-degree this fails. Part of the problem lies in the choice of space, and its basis. The other difficulty lies in the choice of points. For example, it doesn't appear that there exists a general formula for optimal (same number of points as degree of space) Gaussian quadrature in the triangle.

We present the multi-dimensional extension of the Legendre (and Tchebyshev) polynomials over the simplex. In the triangle these are the Koornwinder polynomials (1974). We are also able to compute points which give well-behaved Lagrange polynomials for the triangle and tetrahedra, the Fekete points, which are analogous in some ways to the Gauss-Lobatto points. These points also give a Gauss-like quadrature which has desirable symmetry properties.

This work has applications not only to the spectral element method, but to accurate interpolation in general domains, such as on the sphere, and to accurate and symmetric integration over the triangle and tetrahedra.

March 4, 1999

Simulations of Compressible MHD Problems Using Higher-Order Godunov Methods

Andrea Malagoli

a-malagoli@uchicago.edu

University of Chicago

Abstract

Higher-order Godunov methods are very popular as the shock-capturing methods of choice for the pure fluid Euler equations. In the first part of this talk I will discuss issues related to the extension of Godunov methods to solve the equations of compressible ideal MHD, and more in general issues related to the finite volume discretization of these equations (e.g., the Div(B) problem). In the second part, I will present results from simulations performed with a higher-order Godunov method for the MHD equations. In particular, I will discuss the effects of these numerical methods on the physics of the system, and their possible interpretation as a subgrid-model for the MHD equations.

March 3, 1999

Preconditioning Constrained Systems

Andy Wathen

wathen@comlab.ox.ac.uk

Oxford University

Abstract

The general importance of preconditioning in combination with an appropriate iterative technique for solving large scale linear(ised) systems is widely appreciated. For definite problems (where the eigenvalues lie in a half-plane) there are a number of preconditioning techniques with a range of applicability, though there remain many difficult problems. For indefinite systems (where there are eigenvalues in both half-planes), techniques are generally not so well developed.

Constraints arise in many physical and mathematical problems and invariably give rise to indefinite linear(ised) systems: The incompressible Navier–Stokes equations describe conservation of momentum in the presence of viscous dissipation subject to the constraint of conservation of mass, for transmission problems the solution on an interior domain is often solved subject to a boundary integral which imposes the exterior field, in optimisation the appearance of constraints is ubiquitous.

We will describe two approaches to preconditioning such constrained systems and will present analysis and numerical results for each. In particular, we will describe the applicability of these techniques to approximations of incompressible Navier–Stokes problems using mixed finite elements and the MAC finite difference scheme. We will also present the results of computations on various matrices from the CUTE optimisation test set.

February 25, 1999

Large-Scale Parallel Unstructured Multigrid Computations for Steady-State Aerodynamic Problems

Dimitri Mavriplis
dimitri@icase.edu

NASA Langley Research Center

Abstract

Recent experiences in the development and execution of a parallel unstructured multigrid solver are discussed. The target applications consist of steady-state viscous turbulent flow over external aircraft configurations. The discretization is based on the vertices of an unstructured mesh that may contain arbitrary combinations of tetrahedra, prisms, pyramids and hexahedra. An FAS agglomeration multigrid algorithm has been developed for accelerating the solution to steady-state. This approach uses a graph algorithm to construct the coarse multigrid levels from the given fine grid, similar to an algebraic multigrid approach, but operates directly on the non-linear system using the FAS approach. Two preconditioning techniques are employed to relieve stiffness due to high-aspect ratio grid cells in the boundary layer regions, and the stiffness associated with regions of nearly incompressible flow.

The preconditioned multigrid solver is parallelized using the domain-decomposition approach with the MPI message passing interface. Good scalability for

medium size problems is demonstrated on the SGI Origin 2000 using up to 128 processors, and the Cray T3E using up to 512 processors. A large scale case using 25 million grid points (150 million degrees of freedom) is solved on a dedicated Cray T3E-1200e provided by Cray Research Inc. using up to 1450 processors. Near perfect scalability is demonstrated in going from 256 to 1450 processors on this machine, with the 1450-processor case producing a solution in about one hour of wall clock time in 500 multigrid cycles.

Future work will concentrate on parallelizing the pre-processing operations such as coarse multigrid level construction and grid partitioning which currently are performed sequentially on an SGI ORIGIN 2000. A case involving 100 million grid points (600 million degrees of freedom) should be feasible from a memory and cpu-time perspective on the large T3E configuration, or other ASCI machines, once the pre-processing bottlenecks are addressed.

February 23, 1999

Overlapping Grids Applied to Naval Hydromechanics

Anders Petersson

andersp@na.chalmers.se

Chalmers University of Technology

Abstract

The three-dimensional overlapping grid generator “Chalmesh” is presented. The code implements a combination of previous approaches implemented in, for example, PEGSUS and CMPGRD together with some new ideas.

The algorithm starts by forming a boundary description based on the overlapping surface patches. The union of the overlapping surface patches is used to cut the holes in the component grids by using a mark-and-fill technique. The grid cells intersected by the boundary are identified by an octree-based search technique and the grid points in those cells are marked as being inside or outside the boundary by employing a modified ray method. The holes are then filled starting from the outside points in the marked cells.

After the holes have been cut, the remainder of the grid points are classified by the CMPGRD approach as either interpolation, discretization, or hole points. Thereafter, the grid is trimmed to remove all unnecessary interpolation points to minimise the overlap between the components. Finally, the interpolation data in viscous boundary layer grids is compensated for boundary mismatch.

Examples will be given for two marine applications: a ship and a propeller, where these grids have been used to solve the Reynolds averaged Navier–Stokes equations.

February 19, 1999

The Role of Nonlinear Wavelet Approximation in Image Processing and Partial Differential Equation Solvers

Ronald DeVore
devore@math.sc.edu

University of South Carolina

Abstract

Wavelets and multiresolution give efficient decompositions of functions into simple building blocks. These decompositions can be utilized in the design of numerical algorithms for image processing or solving partial differential equations (PDEs). Much emphasis has recently been placed on nonlinear algorithms whose goal is to use adaptively selected terms of the wavelet decomposition to minimize computational cost. We shall discuss optimal methods for adaptively choosing the terms in a wavelet decomposition and then discuss some of their implications in image processing and PDE solvers.

February 10, 1999

Automatic Differentiation: A Tool for Computational Science

Paul Hovland

hovland@mcs.anl.gov

Argonne National Laboratory

Abstract

Derivatives play an important role in computational science. They are used in optimization, sensitivity analysis, the solution of nonlinear equations, and a variety of inverse problems.

Automatic differentiation (AD) provides a mechanism for transforming code for computing a function into code for computing that function and its derivatives. Unlike finite differences, the derivatives computed by AD are analytic, and do not suffer from truncation error. AD also avoids the time-consuming and error-prone task of coding derivatives by hand. We discuss how AD works, the tools available for its application, and techniques that enable a knowledgeable programmer to achieve improved performance.

February 9, 1999

A Negative Norm Least-Squares Approach to Div-Curl Systems in Two and Three Dimensions

James Bramble

bramble@math.tamu.edu

Texas A&M University

Abstract

In this talk, we will present a direct finite element approximation of div-curl systems. The discretization is based on a reformulation of the div-curl systems using a negative norm least-squares approach. The method is optimal with respect to both the order of the approximation of the finite element space as well as the regularity of the solution. The corresponding algebraic system is symmetric positive definite and, moreover, it is well conditioned and hence can be solved effectively by the conjugate gradient method, for example. The computation of the discrete negative norms only involves the preconditioners of second order finite element problems. In particular a multigrid or multilevel preconditioner can be used.

February 8, 1999

SailFlow: Computing Air Flow Around Yacht Sails with Overture

Cheryl Fillekes
sail@iconz.co.nz

Doyle Bouzaid Sailmakers, New Zealand

Abstract

SailFlow is a program to optimize sails for racing yachts in a commercial sail design environment now in its sixth month of development. The Overture framework was selected as the basis for the gridding and aeroelastic computations. Overture's object model allows the coherent and consistent derivation of grids specific to yacht sails and rigging from more general utility grids. Overture's overlapping grid technology allows multiple sails to be trimmed and re-trimmed relative to each other without having to re-grid the entire computational domain. Overture's support for new partial differential equations solver implementation allows SailFlow the flexibility of doing elastic computations on the same grids as the flow solvers are implemented on. This also allows the implementation of new turbulence models, such as the SST $k-\omega$ model. Finally, because Overture is based on A++/P++, a clear path for migration to parallel platforms is pre-established.

February 1, 1999

Viscous Fingering in Non-Newtonian Fluids

Petri Fast

fast@cims.nyu.edu

New York University

Abstract

Thin gap flows of non-Newtonian liquids are important technologically, and can exhibit a rich variety of dynamical behaviors. To study such problems, I use the thin gap limit to reduce a continuum viscoelastic model to a two-dimensional nonlinear generalization of Darcy's Law. This leads to a nonlinear elliptic boundary value problem to be solved in a time-dependent domain. Recent intermediate time-scale simulations of the resulting dynamical system show intriguing consistency with experiments. I present the current state of a project to develop a moving overset grid scheme for the accurate and efficient long-time simulation of this problem.

January 28, 1999

Analytical Performance Modeling of Hierarchical Mass Storage Systems

Yelena Yesha

yeyesha@cs.umbc.edu

University of Maryland and NASA Goddard

Abstract

Mass storage systems are finding greater use in scientific computing research environments for retrieving and archiving the large volumes of data generated and manipulated by scientific computations. This talk presents a queueing network model that can be used to carry out capacity planning studies of hierarchical mass storage systems. Measurements taken on a Unitree mass storage system and a detailed workload characterization provided by the workload intensity and resource demand parameters for the various types of read and write requests. The performance model developed here is based on approximations to multiclass Mean Value Analysis of queueing networks. The approximations were validated through the use of discrete event simulation and the complete model was validated and calibrated through measurements. The resulting model was used to analyze three different scenarios: effect of workload intensity increase, use of file compression at the server and client, and use of file abstractions.

January 25, 1999

Pattern Formation in Non-Newtonian Hele–Shaw Flow

Lou Kondic

kondic@math@duke.edu

Duke University

Abstract:

We explore the morphology of patterns due to the Saffman–Taylor instability in Hele–Shaw cells and find that it can be dramatically altered by the non-Newtonian response of complex fluids such as liquid crystals and polymer solutions. The dense-branching morphology of Newtonian liquids may be replaced by dendritic fingers with stable tips and sidebranches.

Starting from a very general viscoelastic fluid model, we find a distinguished limit where shear thinning effect is dominant. Darcy’s Law leads to the nonlinear boundary value problem for the pressure in the fluid. Full numerical simulations show that shear thinning alone modifies considerably the pattern formation and can produce fingers whose tip-splitting is suppressed, in agreement with experimental results. These fingers grow in an oscillating fashion, shedding “side-branches” from their tips, closely resembling solidification patterns (*Phys. Rev. Lett.*, **80** 1998: 1433).

January 14, 1999

Domain Decomposition Techniques for Localized Space and Time Grid Refinement for Flow Problems

Richard Ewing
ewing@isc.tamu.edu

Texas A&M University

Abstract

We discuss various strategies for solving steady-state and transient problems with localized solution in space and time. The discretization techniques include Galerkin finite element and finite volume methods in space and discontinuous Galerkin and finite difference methods in time on locally refined (matching and non-matching) grids in both space and time. Next, we shall present iterative methods based on domain decomposition technique for solving the composite grid system. These include overlapping and nonoverlapping Schwarz methods (BPS and BEPS iterative algorithms) and domain decomposition methods for mortar approximations. Finally, we discuss numerical results demonstrating the capabilities of the methods for both test problems and applied problems from flows in porous media.

January 12, 1999

Domain Decomposition Preconditioners for High-Order Discretization Methods

Calvin Ribbens
ribbens@vt.edu

Virginia Institute of Technology

Abstract

A widely-used family of techniques for numerically solving elliptic partial differential equations (PDEs) on parallel processors is known as “domain decomposition” (DD). The key step in a DD method is typically the construction and application of a preconditioner for a Krylov subspace iterative linear equation solver. Most DD preconditioners are derived and analyzed in the context of second-order accurate discretization schemes, e.g., “five-point” finite differences, piecewise linear finite elements. In this talk, we describe two preconditioners for linear systems arising from higher-order discretization methods. We first describe HODIEX, a DD-preconditioned iterative solver based on the fourth-order accurate HODIE method (High-Order Difference Approximation with Identity Expansion) of Lynch and Rice. We then focus on ColEVP, a preconditioner for linear systems arising from a piecewise Hermite bicubic collocation discretization method (which also gives $O(h^4)$ accuracy).

The ColEVP preconditioner is defined in terms of a three-level grid and discretization scheme. In the framework of substructuring, we partition the domain into subdomains, edges, and vertices. The three-level discretization scheme has a fine grid $G_{h,h}$ with uniform step h on the first level. The third level is the coarse grid $G_{H,H}$ with step H corresponding to a subdomain width. The middle level is the edge-grid $G_{h,H'}$, a hybrid fine/coarse grid which has step h along each edge and step H' in the perpendicular direction. Here, $H' = H-h$ for all edges except those that are a distance H from a physical boundary—for those edges we use $H' = H$. The edge and vertex subproblems that must be solved in applying the preconditioner are based on $G_{h,H'}$ and $G_{H,H}$, respectively. We show how the coupling between the edge and subdomain subproblems is significantly reduced by using the edge-grid $G_{h,H'}$, and how this can be used to construct an efficient preconditioner. We describe the motivation, construction, and implementation of ColEVP, and illustrate its performance with examples.

January 11, 1999

A Multigrid Approach for the Mortar Finite Element Method

Joseph Pasciak
pasciak@math.tamu.edu

Texas A&M University

Abstract

A multigrid technique for uniformly preconditioning linear systems arising from a mortar finite element discretization of second-order elliptic boundary value problems is described. These problems are posed on domains partitioned into subdomains, each of which is independently triangulated in a multilevel fashion. The multilevel mortar finite element spaces based on such triangulations (which need not align across subdomain interfaces) are in general not nested. Suitable grid transfer operators and smoothers will be discussed which lead to a variable V-cycle preconditioner resulting in uniformly preconditioned algebraic systems. Computational results illustrating the theory are also presented.

January 8, 1999

Spatial Databases and Geographic Information Systems

Hanan Samet

hjs@umiacs.umd.edu

University of Maryland

Abstract

An introduction is given to the spatial database issues involved in the design of geographic information systems (GIS) from the perspective of a computer scientist. Some of the topics to be discussed include the nature of a GIS and the functionalities that are desired in such systems. Representation issues will also be reviewed. The emphasis will be on indexing methods, as well as the integration of spatial and nonspatial data. An example spatial browser for a geographic information system using these concepts will be discussed.

December 18, 1998

Multi-Source Data Analysis in Science and Engineering

Samuel Uselton
uselton@nas.nasa.gov

MRJ Technology Solutions

Abstract

As digital data acquisition becomes easier, cheaper and more pervasive, and computational simulations gain increasing fidelity and detail, many activities can benefit from the combined analysis of data from several sources. Building a useful multi-source analysis system requires solving many problems, some of a pragmatic, engineering nature, and some of a more basic nature. The issues seem remarkably similar whether the application is weather modeling, environmental assessment and remediation planning, oil exploration and production, or engineering design processes. The problems to be addressed range from efficient access to large amounts of data from multiple heterogeneous sources, to design of user interfaces and inventing visualization techniques. Integrated display and comparative analysis of relevant data is interesting and relatively unexplored.

Our Multi-Source Visualization (MSV) project uses a specific problem domain, concurrent design of aircraft, to focus research and development efforts in this area.

Frequent contacts with the Earth Observation System (EOS) project, meso-scale atmospheric modeling researchers, and nanotechnology researchers, among others, keep broader needs in view.

All the work done by the Data Analysis Group of the NAS Division at NASA Ames Research Center is relevant and useful to this project. This group's work ranges from specific visualization techniques, through innovative user interfaces, the software engineering required to build complete systems, systems level performance improvements, effective access to very large data sets, and exploitation of large heterogeneous collections of scientific and engineering data. Much of this work will be touched upon, but the focus will be on the work that is specifically driven by the need for a variety of users to exploit common collections of data from many diverse sources.

This work is very much in progress. There are results to show, but also problems not yet satisfactorily resolved.

December 15, 1998

Does the Large Scale Determine the Small Scale in Turbulent Flows?

Heinz-Otto Kreiss
kreiss@math.ucla.edu

University of California, Los Angeles

Abstract

In meteorological applications, one can only measure the large scale. It is believed, however, that the small scale (mesoscale) plays an important role for longer-time forecasts. Since the observational net does not resolve the small scale, we want to discuss, for a model problem (2D and 3D Navier–Stokes equations), whether we can reproduce the small scale by the time history of the large scale.

December 11, 1998

New Algorithms for Incomplete Factorization: Symmetric Reduction Redux

Alex Pothen

pothen@cs.odu.edu

Old Dominion University

Abstract

We describe new algorithms for computing incomplete factor preconditioners for solving systems of linear equations by iterative methods. These algorithms rely on a newly developed structure theory for identifying “fill” in incomplete factorization. We use a technique called symmetric reduction to reduce the time needed to compute the preconditioners when high levels of fill are permitted. The new algorithms are provably faster than current implementations, and are also inherently parallel. Our implementations show that the new algorithms are faster than ILUT preconditioners that rely on a numerical threshold to compute the incomplete factors.

December 10, 1998

Some New Parallel Adaptive Finite Element Methods

Michael J. Holst
mholst@math.ucsd.edu

University of California, San Diego

Abstract

We discuss a new approach to using parallel computers with adaptive multilevel finite element methods. We propose an algorithm, for general elliptic equations on arbitrary polyhedral domains, which has the following interesting features.

- The algorithm begins with one processor solving a special small “coarse” problem, and then broadcasting the coarse problem and its solution to the remaining processors.
- The computation then proceeds independently on each processor, without communication (i.e., no boundary exchange communications are required).
- The final adapted mesh on each processor has nearly the same number of elements, so that load balancing is not necessary.
- The resulting solution is (provably) “as good” as the solution produced by the usual communication-intensive boundary exchange methods coupled with dynamic load-balancing strategies.

The parallel algorithm will be described in detail. To illustrate that the parallel algorithm can be used in conjunction with any sequential adaptive finite element method, we will give some examples using two different adaptive codes: the well-known 2D package PLTMG, and a new similar package called MC. MC is a dimension-independent, simplex-based, ANSI-C finite element code for covariant problems on d -manifolds ($d=2,3,\dots$). MC implements many features of the PLTMG, including adaptive error control, unstructured algebraic multilevel methods, global Newton methods, and continuation.

To illustrate the generality of the approach, the examples will include scalar linear elliptic equations in 2D, nonlinear elasticity in 3D, and the Hamiltonian and momentum constraints in the Einstein equations (a coupled four-component nonlinear elliptic system in 3D). We finish the talk by outlining the some local a priori and a posteriori error estimates of Xu and Zhou, which provide a theoretical justification for the new algorithm.

November 24, 1998

Advising the President on Information Technology

David M. Cooper
cooper31@llnl.gov

Lawrence Livermore National Laboratory

Abstract

This presentation focuses on the recommendations of the President's Information Technology Advisory Committee in the Interim report to the president submitted in August 1998. The major findings of the committee are that the federal investment in Information Technology (IT) R&D is inadequate and that the federal IT investment is too heavily focused on near-term problems. To offset these trends, the committee recommends that the United States government fund a strategic initiative in fundamental IT research with specific investments in software, high-end computing, scalable information infrastructure, and socio-economic and workforce activities. Specific details of each of these elements are presented.

November 6, 1998

Applications of ART Neural Networks

Gail Carpenter

gail@cns.bu.edu

Boston University

Abstract

ART (Adaptive Resonance Theory) neural networks for fast, stable learning and prediction have been applied in a variety of areas. Applications include airplane design and manufacturing, automatic target recognition, financial forecasting, machine tool monitoring, digital circuit design, chemical analysis, and robot vision. Supervised ART architectures, called ARTMAP systems, feature internal control mechanisms that create stable recognition categories of optimal size by maximizing code compression while minimizing predictive error in an on-line setting. Special-purpose requirements of various application domains have led to a number of ARTMAP variants, including fuzzy ARTMAP, ART-EMAP, ARTMAP-IC, Gaussian ARTMAP, and distributed ARTMAP. This talk will discuss some ARTMAP applications, including the following.

Multi-sensor fusion, with application to sonar target recognition

A new ARTMAP variant, called ARTMAP-FTR (fusion target recognition), has been developed for the problem of multi-ping sonar target classification (Carpenter and Streilein, 1998). The development data set, which lists sonar returns from underwater objects, was provided by the Naval Surface Warfare Center (NSWC) Coastal Systems Station (CSS). The ARTMAP-FTR network has proven to be an effective tool for classifying objects from sonar returns. The system also provides a procedure for solving more general sensor fusion problems.

Geospatial mapping from satellite remote-sensing data

A remote sensing testbed allows performance comparisons between neural network systems and state-of-the-art image processing and recognition techniques, in a collaborative project between researchers at the Center for Adaptive Systems and the Boston University Center for Remote Sensing (Carpenter, Gopal, Martens, and Woodcock, 1997). Network hierarchies that take advantage of the database size and structure have been developed. These systems provide an efficient method for producing accurate geospatial maps from high-dimensional satellite and terrain data. A new ARTMAP network is being considered as the candidate algorithm for the Moderate Resolution Imaging Spectrometer (MODIS) land cover product of the NASA Earth Observing System (EOS).

Self-organizing expert systems and computer-assisted medical diagnosis

Medical databases present many of the challenges found in general information management settings where speed, efficiency, ease of use, and accuracy are at a premium. A direct goal of improved computer-assisted medicine is to help deliver quality emergency care in situations that may be less than ideal. Working with these problems has stimulated a number of ART architecture developments in recent years, including ARTMAP-IC (Carpenter and Markuzon, 1998). A new collaborative effort brings together medical statisticians and clinicians at the New England Medical Center with researchers developing expert system and neural network learning systems.

November 5, 1998

Data Mining in Very Large Dimensional Data Sets

George Karypis
karypis@cs.umn.edu

University of Minnesota

Abstract

Data sets with high dimensionality pose major challenges for conventional data mining algorithms. For example, traditional clustering algorithms such as K-means or AutoClass fail to produce good clusters in large dimensional data sets even when they are used along with well known dimensionality reduction techniques such as Principal Component Analysis. Similarly, traditional classification algorithms such as C4.5 perform poorly on large dimensional data sets.

This talk presents a novel method for clustering related data items in large high-dimensional data sets. Relations among data items are captured using a graph or a hypergraph, and an efficient multi-level graph partitioning algorithm is used to find clusters of highly related items. We present results of experiments on several data sets

including S&P500 stock data for the period of 1994-1996, protein coding data, and document data sets from a variety of domains. These experiments demonstrate that our approach is applicable and effective in a wide range of domains, and outperforms techniques such as K-Means, even when they are used in conjunction with dimensionality reduction methods such as principal component analysis or latent semantic indexing scheme.

This talk also presents a graph-based nearest-neighbor classification scheme in which the importance of discriminating variables is learned using mutual information and weight adjustment techniques. Empirical evaluations on many real world documents demonstrate that this scheme outperforms state of the art classification algorithms such as C4.5, Ripper, Naive-Bayesian, and PEBLS.

November 4, 1998

Preconditioning Operators for Elliptic Problems with Bad Parameters

Sergei Nepomnyaschikh

Russian Academy of Sciences

Abstract

The construction of preconditioning operators for the iterative solution of systems of grid equations approximating elliptic boundary value problems with “bad” parameters is considered. These bad parameters characterize jumps in the coefficients, shape of the subdomains with composite materials, etc.

The suggested technique is based on the domain decomposition, i.e., the original problem is decomposed into subproblems in which the coefficients of equations inessentially change. We consider two variants of the domain decomposition technique: splitting into nonoverlapping subdomains and splitting into overlapping subdomains “without overlapping in the coefficients.” For the first domain decomposition method constructing of preconditioning operators involves constructing easy invertible norms for grid functions on the boundaries of the subdomains and norm-preserving explicit extension operators of func-

tions from the boundaries into inside subdomains. These norms should be equivalent to the norms which are generated by original problems on these boundaries (Schur Complements).

The second domain decomposition method is based on the theorem of extension of functions from subdomains with “strong” coefficients into subdomains with “weak” coefficients. Multilevel preconditioning operators for the subproblems on unstructured grids with the smooth coefficients are suggested and used in both methods. Design of these preconditioning operators is based on the fictitious space method. This method reduces the problem from the unstructured grid to a structured grid (at the first step) and then (at the second step) reduces the problem to a structured hierarchical grid where BPX-like preconditioners can be used. The convergence rate of the preconditioned iterative process is independent of the grid size and the bad parameters.

October 29, 1998

Parallel Particle-in-Cell Modeling of Semiclassical Quantum Models

Roy Hemker

University of California, Los Angeles

Abstract

We are modeling many-particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas.

October 22, 1998

Clustering of Large Data Sets: The Possibilities Are Endless!

Wray Buntine

wray@ultimode.com

Ultimode Systems, LLC

and University of California, Berkeley

Abstract

Clustering or unsupervised learning is a basic form of pattern recognition developed to a fairly advanced state in the Sixties. Techniques for automatically selecting the “right” number of clusters, determining irrelevant features, and clustering of more sophisticated models have been in successful use for many years (despite claims by some that these are open research questions). Following well-publicised efforts at JPL some years back, the data-mining community is currently seeing a flurry of activity in scaling up clustering to large data sets.

First, I review potential methods available in the scientific literature, many of which are developed outside of the data-mining community and not well explored. Included are bump-tree and k-d tree techniques used for speeding up nearest neighbor algorithms, “clustering” methods from optimization in CAD, numerical analysis and high-performance computing (the term clustering here is used differ-

ently), automated software engineering methods for tuning on parallel processes and code-optimization, iterative, “hard” and deterministic-annealed versions of the underlying algorithms, and hierarchical clustering. Combinatorially, there are an enormous range of possibilities here and data mining researchers have barely scratched the surface. Thus I view the scaling up of clustering as largely a sophisticated software and engineering issue.

Second, I discuss some of our practical experiences in applying clustering for both commercial and government clients: We firmly believe intelligent preprocessing and post-visualization needs to be done by an innovative professional, and that the core software used in clustering sometimes need adaptation to individual problems, simple clustering models used for instance in SAS and standard earth science packages are too inflexible and the software itself may need to be customized.

October 20, 1998

Computational Fluid Dynamic Studies of Arterial Flow Disturbance Induced by Intravascular Stents

E. Tina Cheng
etcheng@ucdavis.edu

University of California, Davis

Abstract

Atherosclerosis is an arterial disease whose pathological complications, namely heart disease and stroke, are the leading causes of mortality in the industrialized world. In its advanced form, atherosclerosis leads to plaques which protrude into arterial lumens and form stenoses, or even complete vessel occlusions that obstruct blood flow and give rise to the pathological events. One common interventional procedure involves the placement of an intravascular stent, an expandable wire mesh structure that is introduced into the diseased artery in a compressed state and inflated at the stenosis or occlusion site to both restore blood flow and provide structural stability to the arterial wall. The major limitation to the success of this procedure, however, is restenosis, a complex and incompletely understood process by which plaques reform and re-protrude into the vessel lumen within a period of a few months.

The placement of a stent in an artery mechanically damages the endothelium, the monolayer of cells lining the inner surfaces of all blood vessels. In vitro data indicate that the rate of endothelial repair after injury may be significantly slower in regions in which endothelial cells are exposed to relatively large fluid mechanical shear stress gradients, as occurs at the end points of flow separation

zones. Therefore, flow separation in the vicinity of a stent may contribute to restenosis. Furthermore, clinical data suggest that the incidence of restenosis is higher for thicker-wire stents. We hypothesized that the occurrence of flow separation depends on hemodynamic matching between the stent wire thickness and the flow and geometric properties of the arterial segment in which the stent is positioned. To test this hypothesis, we have been studying the impact of stent wire thickness on the occurrence of flow separation under various arterial geometric and flow parameters using computational fluid dynamic techniques.

Our results to date have demonstrated that for straight arterial segments in steady flow, flow separation is more likely to occur in the case of thicker-wire stents and for higher flow Reynolds numbers. Flow pulsatility leads to periodic appearance and disappearance of flow recirculation zones thereby introducing significant temporal (in addition to the spatial) shear stress gradients. Finally, vessel curvature leads to additional flow recirculation zones not present in straight arterial segments. These results may provide insight into the mechanisms governing the clinical observation of increased incidence of restenosis in thicker-wire stents and may guide strategies for targeting particular stent structures for specific vascular sites.

October 19, 1998

Wavefunction Engineering of Quantum Devices Computational Issues in Modeling Nanostructures

Ramdas Ram-Mohan

lrram@wpi.edu

Worcester Polytechnic Institute

Abstract

Quantum semiconductor heterostructures provide new opportunities to investigate fundamental quantum mechanical effects while holding forth the promise of new optoelectronic devices. The talk will focus on the theoretical and computational issues that have to be addressed in understanding the optical properties of such structures. The efficacy of the finite element method developed by the speaker for the calculation of energy levels in quantum semiconductors is demonstrated, and their application to structures such as quantum wells, superlattices, quantum wires, and checker-board superlattices, and their optical properties are presented. Computational challenges in this area of quantum applications are as demanding as, for example, in fluid mechanics.

The new developments in computation have led to the concept of band-gap engineering of semiconducting materials being replaced by the new paradigm of wavefunction engineering for quantum semiconductor structures. The computational problems faced in solving sparse banded matrix equations and in determining eigenvalues of large sparse complex banded matrices will be considered. The need for parallel computing, the development of faster algorithms for parallel computers for matrix analysis, and the visualization of wavefunction calculations are some of the computational issues and future directions in quantum device modeling. This talk includes elementary conceptual introductions to the topics.

October 16, 1998

Stochastic Programming: A Model for Decision-Making under Uncertainty

Roger Wets

rjbow@math.ucdavis.edu

University of California, Davis

Abstract

Almost all (important) decision problems involve some level of uncertainty, either about data measurements, the values to assign to parameters describing future evolution or even about the environment in which one has to operate. A few typical examples are: trajectory selection with data uncertainty about the initial conditions; management of energy resources with uncertainty about hydro-power supply and prices; and groundwater remediation with uncertainty about the (nonhomogeneous) soil composition.

The main objective is to provide a brief introduction to stochastic programming models and solution techniques. Because the choice of an optimal (or a good) decision must take into account a huge number of possible realizations of the uncertain (stochastic) parameters of the problem, the development of algorithmic procedures must take advantage of the available architectures that are now becoming available.

October 13, 1998

Parallel, High-Resolution Finite Element Analysis of Trabecular Bone Biomechanics

Glen Niebur

gln@biomech2.ME.Berkeley.EDU

University of California, Berkeley

Abstract

We will describe ongoing research at the Orthopaedic Biomechanics Laboratory at UC Berkeley on development of spatially high-resolution (10–20 micron) finite element models of trabecular bone, the spongy bone tissue found in the spine and at all articulating joints such as the hip and knee. Trabecular bone is most involved in diseases such as osteoporosis and in prosthesis fixation.

The overall long-term goal of this work is to develop computational modeling techniques that can be used to investigate the effects of aging and disease on the mechanical behavior of this material. Over the past three years with funding from LLNL, we have successfully developed and optimized the parallel finite element codes to solve this class of problems, with models having up to 5 million elements. We now describe our work on the numerical convergence behavior of this technique, and applications to prediction of lower level material problems of this cellular network, as well as simulation of mechanical testing techniques, and investigation of failure behavior and mechanisms.

October 12, 1998

Algorithms for Rapid IsoSurface Extraction

Charles Hansen

hansen@cs.utah.edu

University of Utah

Abstract

Exploratory scientific visualization is a valuable paradigm for understanding complex physical phenomena. When these phenomena have associated volumetric (3D) scalar fields, isosurface extraction is a critical tool. An isosurface is the set of points where the scalar field has a particular value, the "isovalue." The position of an isosurface, as well as its relation to other neighboring isosurfaces, can yield clues to the underlying structure of the scalar field. The ability to interactively change the isovalue provides insight into structure and underlying physical meaning of the scalar field. This talk will focus on the work being done at Utah for rapid isosurface extraction algorithms that enable exploration of large datasets for both local and remote visualization. These algorithms provide fast localization of an isosurface within very large datasets, as well as giving scientists prompt feedback during exploratory sessions. This work was the Ph.D. Dissertation topic for Yarden Livnat, who defended at Utah in 1998.

October 6, 1998

Parallel Multigrid Solvers for Finite Element Problems in Solid Mechanics

Mark Adams

madams@CS.Berkeley.EDU

University of California, Berkeley

Abstract

Prometheus is our parallel linear equation solver for 3D finite element problems on unstructured meshes that uses a classical (geometric) multigrid method with an algebraic architecture—that is, the user need only provide the fine grid. This talk discusses our methods of automatically constructing good coarse grids for solid mechanics problems, as well as methods for optimizing parallel performance of parallel multigrid codes and PETSc, useful on machines with relatively slow communication (e.g., a large machine with a good network or a small machine with high latency).

Some general issues of multigrid solver behavior on problems with large jumps in material coefficients and high Poisson ratio (incompressible materials) are also discussed. We will show numerical results for linear and non-linear problems in elasticity and plasticity, of up to 9,600,000 degrees of freedom (and perhaps 16.5 M degrees of freedom, with any luck), on a Cray T3E and an IBM PowerPC cluster.

October 5, 1998

Aspects of Robust Scalable Multigrid: Interpolation, Smoothing & Coarsening

Justin Wan

wlwan@math.ucla.edu

University of California, Los Angeles

Abstract

We discuss three major integrated components of multigrid to design robust multigrid methods. First, we describe a new approach to constructing robust interpolation operators, which can handle problems with problematic (discontinuous, oscillatory) coefficients in a unified fashion for both regular or irregular grids. The basic idea derives from recent domain decomposition theory, and is based on defining coarse basis functions that are stable (minimizing the total energy) and have good approximation property (preserving constants).

Second, we discuss the use of sparse approximate inverses as parallel smoothers for multigrid. They have been recently demonstrated to be useful technique for preconditioning in the parallel environments. We prove further that they are effective for eliminating local high-

frequency errors; hence they can be used as smoothers for multigrid. Furthermore, a distinctive feature over relaxation smoothers is that we may improve their quality by adjusting the nonzero patterns of the sparse approximate inverses, which is particularly useful for anisotropic problems.

Third, we describe a new geometric technique to construct coarse grids/ subspaces for discontinuous coefficient problems. The interface preserving coarsening selects coarse grid points so that all the coarse grids are aligned with the interfaces for regular interface problems on structured grids, and that the interfaces are resolved as much as possible for irregular interface problems. As a result, multigrid with simple linear interpolation is sufficient to obtain fast convergence.

October 1, 1998

Neutron Transport Computation with Variational Nodal Methods

Elmer Lewis
e-lewis@hwu.edu

Northwestern University

Abstract

A unified variational formulation of the even-parity form of the Boltzmann equation for neutrons is presented. Spherical harmonics (P_n), simplified spherical harmonics (SP_n) and discrete ordinates (S_n) approximations in angle reduce the transport equation to a set of coupled second-order partial differential equations, where the number of equations increases with the order n . These equations are formulated variationally, and a hybrid finite element method is applied to obtain spatial discretization which imposes particle balance over a coarse nodal mesh.

Two classes of trial functions are employed within the nodes: high-order orthogonal polynomials, and piecewise bilinear functions over finite subelements. In both cases changes of variables result in nodal response matrix equations. These are solved using red-black or multicolored iterative algorithms. Implementations of the P_n and SP_n approximations in the Argonne National Laboratory code VARIANT are discussed. New results are presented for piecewise bilinear trial functions; these allow each response matrix to include highly heterogeneous regions. The suitability of the methods for parallel computation is examined.

