Institute for Scientific Computing Research

Annual Report

Fiscal Year 2002

Lawrence Livermore National Laboratory, P.O. Box 808, L-561, Livermore, CA 94551

UNIVERSITY RELATIONS PROGRAM

The University Relations Program (URP) encourages collaborative research between Lawrence Livermore National Laboratory (LLNL) and the University of California campuses. The Institute for Scientific Computing Research (ISCR) actively participates in such collaborative research, and this report details the Fiscal Year 2002 projects jointly served by URP and ISCR. For a full discussion of all URP projects in FY 2002, please request a copy of the URP FY 2002 Annual Report by contacting

Lawrence Livermore National Laboratory
Edie Rock, University Relations Program
P. O. Box 808, L-413
Livermore, CA 94551

UCRL-LR-133866-02

Work performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

Contents

The Mission of the ISCR

The Institute for Scientific Computing Research (ISCR) at Lawrence Livermore National Laboratory is jointly administered by the Computing Applications and Research Department (CAR) and the University Relations Program (URP), and this joint relationship expresses its mission. An extensively externally networked ISCR cost-effectively expands the level and scope of national computational science expertise available to the Laboratory through CAR. The URP, with its infrastructure for managing six institutes and numerous educational programs at LLNL, assumes much of the logistical burden that is unavoidable in bridging the Laboratory's internal computational research environment with that of the academic community.

As large-scale simulations on the parallel platforms of DOE's Advanced Simulation and Computing (ASCI) become increasingly important to the overall mission of LLNL, the role of the ISCR expands in importance, accordingly.

Relying primarily on non-permanent staffing, the ISCR complements Laboratory research in areas of the computer and information sciences that are needed at the frontier of Laboratory missions. The ISCR strives to be the “eyes and ears” of the Laboratory in the computer and information sciences, in keeping the Laboratory aware of and connected to important external advances. It also attempts to be “feet and hands,” in carrying those advances into the Laboratory and incorporating them into practice. In addition to conducting research, the ISCR provides continuing education opportunities to Laboratory personnel, in the form of on-site workshops taught by experts on novel software or hardware technologies.

The ISCR also seeks to influence the research community external to the Laboratory to pursue Laboratory-related interests and to train the workforce that will be required by the Laboratory. Part of the performance of this function is interpreting to the external community appropriate (unclassified) aspects of the Laboratory's own contributions to the computer and information sciences—contributions that its unique mission and unique resources give it a unique opportunity and responsibility to make.

Of the three principal means of packaging scientific ideas for transfer—people, papers, and software—experience suggests that the most effective means is people. The programs of the ISCR are therefore people-intensive.

Finally, the ISCR, together with CAR, confers an organizational identity on the burgeoning computer and information sciences research activity at LLNL and serves as a point of contact within the Laboratory for computer and information scientists from outside.

Institute for Scientific Computing Research Fiscal Year 2002 Director's Report

Large-scale scientific computation, and all of the disciplines that support it and help to validate it, have been placed at the focus of Lawrence Livermore National Laboratory by the Advanced Simulation and Computing (ASCI) program and more recently by DOE's Scientific Discovery through Advanced Computing (SciDAC) initiative. The Laboratory operates computers with among the highest peaks of performance in the world and has undertaken some of the largest and most compute-intensive simulations ever performed. Energy Secretary Spencer Abraham announced in November 2002 the awarding to the Laboratory of two future machines, each of which is expected to be at the time of their delivery the world's most capable. However, computers at architectural extremes are notoriously difficult to use efficiently. Furthermore, each successful terascale simulation only points out the need for much better ways of interacting with the resulting data.

ISCR Visitor Program

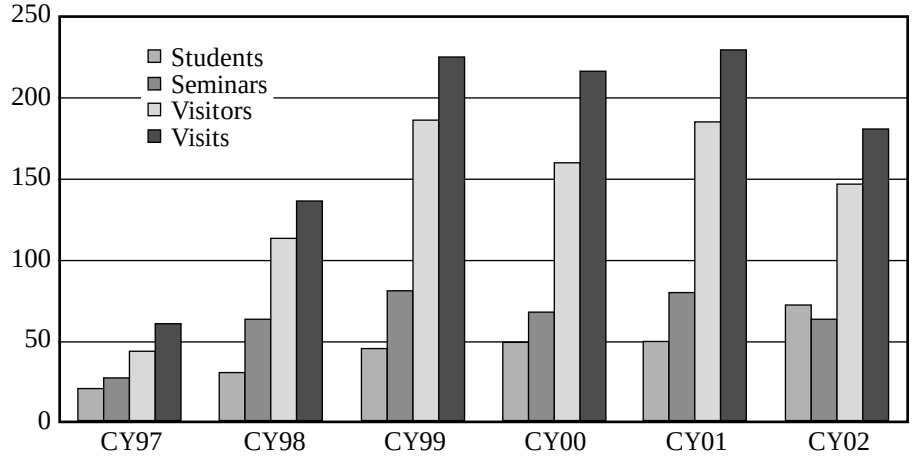

Advances in scientific computing research have therefore never been more vital to the core missions of the Laboratory than at present. Computational science is evolving so rapidly along every one of its research fronts that to remain on the leading edge the Laboratory must engage researchers at many academic centers of excellence. In FY 2002, the Institute for Scientific Computing Research (ISCR) has served as one of the Laboratory's main bridges to the academic community in the form of collaborative subcontracts, visiting faculty, student internships, workshops, and an active seminar series.

ISCR research participants are integrated into the Laboratory's Computing and Applied Research (CAR) department, especially into its Center for Applied Scientific Computing (CASC). These organizations, in turn, address computational challenges arising throughout the Laboratory. Administratively, the ISCR flourishes under the Laboratory's University Relations Program (URP). Together with the other Institutes of the URP, it navigates a course that allows the Laboratory to benefit from academic exchanges while preserving national security. While it is difficult to operate an academic-like research enterprise within the context of a national security laboratory, the results declare the challenges well met and worth the continued effort.

Fiscal year 2002 was the third full year under Acting Director David Keyes. Keyes, the Richard F. Barry Professor of Mathematics & Statistics at Old Dominion University and an ISCR faculty participant since October 1997, dedicated one-third of his time to the technical program of the ISCR. Dr. James McGraw assumed a critical role as the Deputy Director of the ISCR, just in time for a major expansion of its programs and accountabilities stemming from new support responsibilities within CAR. Linda Bodtker came aboard as the full-time Institute Administrator. Emma Horcabas and Leslie Bills also assisted with the large visitor and summer programs.

In May 2002, the ISCR relocated from its offices in the Building 451 complex, where staff and visitors were interspersed with their administrative sponsors and research collaborators, to newly renovated space in the east wing of Building 219. There, it joined the other institutes of the University Relations Program. While in some ways this move carried disadvantages that require more effort on the part of everyone involved in university collaborations to overcome – especially in terms of spontaneous interaction between sponsors and visitors and in terms of convenience in attending seminars of joint interest – it also created more camaraderie and opportunities for shared experiences between visitors, and increased community with the URP. The new level of interaction with the sister Materials

Research Institute has been especially fruitful, as is in evidence, for instance, in the new Quantum Computing seminar series.

In June, with the advent of our large student summer program and sponsorship from the Defense Programs office of DOE HQ, we ramped up our third annual Internships in Terascale Simulation Technology tutorial series. The tutors included Erick Cantu-Paz, Terence Critchlow, Alex Garcia, Jeff Hittinger, Tanya Kostova-Vassilevska, Gary Kumfert, Carol Woodward, CASC's Director Pete Eltgroth, and the ISCR Director. Though intended for students, permanent CASC researchers attended an occasional subseries of the lectures.

Throughout FY 2002, the ISCR brought to the Laboratory a vigorous contingent of post-docs, faculty visitors, and students. Twenty-four faculty visitors were in residence for more than just a seminar visit – for a week to a semester. Nine post-docs made the ISCR their home this past year. We also had 72 students in residence, mostly for 8–10 weeks of the summer, but several of them for a semester or a full year. Each of these students was in a research relationship with a full-time technical staff member.

The pages of this report summarize the activities of the faculty members, post-doctoral researchers, students, and guests from industry and other laboratories who participated in LLNL's computational mission under the auspices of the ISCR during FY 2002. Altogether, the ISCR hosted 180 visits from 147 different visitors, who gave a total of 66 seminars on site, an average of a little more than one per week. The vast majority of the visitors were from academia, with 8% from industry and 12% from other laboratories. Visitors from outside of the United States made up 25% of the total. The histogram left charts the numbers of visitors and seminars over the past six years. Students in residence were sharply up due to the expansion of the ISCR's responsibility in the larger CAR organization. The numbers of visitors and seminars hosted by the ISCR was down in part due to budget uncertainties and in part to participation in lab-wide seminar series and special events surrounding the fiftieth anniversary of LLNL.

Most of the material of this annual report comes directly from the visitors and principal investigators of the projects being reported, who selected formats convenient for their purposes. We thank Whitney Lacy for her editorial work and Dan Moore for his graphic artistry in producing an easily navigated and visually pleasing document.

We hope that you enjoy examining this report on the ISCR's diverse activities in FY 2002. For further information about the Institute, please contact us at the address below. Inquiries about how you might enhance the on-going FY 2003 program at the ISCR, or beyond, are welcome

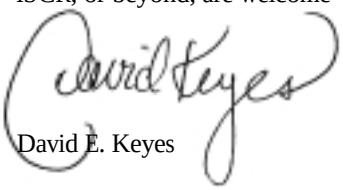

David E. Keyes

Institute for Scientific Computing Research, Lawrence Livermore National Laboratory, P. O. Box 808, L-561, Livermore, CA 94551
http://www.llnl.gov/casc/iscr/

David E. Keyes
Acting Director
http://www.llnl.gov/casc/people/keyes/
dekeyes@llnl.gov
925-422-1325

James McGraw
Deputy Director
http://www.llnl.gov/casc/people/mcgraw/
jmcgraw@llnl.gov
925-422-0541

Linda Bodtker
Administrator
http://www.llnl.gov/casc/people/bodtker/
bodtker2@llnl.gov
925-422-0421

Special Report on Employment Placement of ISCR Summer Student Alumni

In addition to supporting the Laboratory's computational science research programs, the ISCR serves as a valuable recruiting vehicle. The purpose of this special report is to summarize the success of the ISCR as a means of attracting outstanding young computational scientists from its summer program to the DOE national laboratories.

We emphasize, however, that the goal of the ISCR summer program is not direct DOE job placement, but rather making progress on LLNL scientific missions, educating the academic community about those Laboratory missions through immersion of personnel, and providing the technical resources, training, and direction necessary to allow them to help us achieve those missions. Nevertheless, much of the support and good will that the ISCR has enjoyed in recent years—from the University Relations Program, the Computation Directorate, the Director's Office, Laboratory Programs, and NNSA headquarters—is related to its potential for attracting the nation's most capable future scientists to career laboratory work. Therefore, the five-year snapshot of this report is gratifying.

Over the most recent five-year window (1998-2002, inclusive), the ISCR hosted 178 distinct students for a total of 260 student summers, with some students repeating for a second summer or even more. Many of these were undergraduates who were graduate school-bound when they interned, so even with a five-year window most of these are still in the pipeline and do not yet contribute to initial career placement statistics. Since a typical number of years in graduate school is five or six in the sciences and engineering, many of the students who first came to LLNL as graduate students are also still in the pipeline. Moreover, our earlier summer programs were small in comparison with our most recent summers, so most of our interns are still students. To be precise, 127 of our summer interns are known to be students still, and though we are missing data on a few of the remainder, we conclude that 51 have completed their studies. Of these 51 students who left school, we know of 23 ISCR summer students that have gone on to full-time DOE lab employment. Therefore at least 45% of ISCR summer interns have ended up at a DOE laboratory upon completion of their studies!

The 23 of whose lab placement we know break down as follows:

15	LLNL
3	Sandia
3	Los Alamos
1	Argonne
1	Brookhaven

Two categories of summer students, those enrolled year-round in the UC Davis Division of Applied Science (DAS) doctoral program and those recruited under the Institute for Terascale Simulation Technologies (ITST), have the greatest likelihood, statistically, of becoming laboratory employees upon graduation.

Is 45% good overall? We think it is very reasonable. Certainly 100% would not be good—one of the opportunities of a summer internship program is to filter out students who, while talented in laboratory mission areas, do not have the team spirit for career laboratory employment. It is certainly better to identify such students through a summer program than after consummating secure employment. In addition, we want a certain number of students coming through the summer program who intend to make a career of the professoriate — training the next generation.

In fact, we have returned over 20% of our seed corn to the field of academe. Eleven of our estimated 51 graduated former summer interns have taken positions in academia, a few initially as post-docs and most as assistant professors. They have ended up at: Duke University, Stanford University, UC Irvine, University of Minnesota, University of Montana, University of North Carolina, University of Pennsylvania, University of Utah, and University of Washington, for those holding domestic appointments. Abroad, we have ISCR intern alumni at the prestigious Swiss Federal Institute of Technology (ETH) and the University of Toronto. These alumni in highly ranked doctoral research universities are likely to be highly useful to us through the years. They have begun to contact us proactively about sending their own best students for an ISCR summer in 2003. We now, in effect, have several scouts and agents in academia to identify talented students and help them to overcome any negative stereotypes of NNSA lab work.

We are “completing the circle” on the Laboratory side in addition to on the academic side: former summer research mentorees now employed at LLNL are actively taking students as summer mentors. They probably make more effective mentors than LLNL career scientists at random, because they can relate to the difficulty of getting useful work done in just 8–12 weeks. There is a “science” to summer internships that is evolved by doing. We know that there are considerable improvements yet to be made to the program, logistically. At the same time, we believe that there are new scientific areas and new student demographics that can be opened up with similar success.

In both the research and the job placement, we can take considerable satisfaction over the five years during which the ISCR has been closely allied with the Center for Applied Scientific Computing. We have similar expectations for the ISCR's growing role in hosting computational summer students for other Laboratory organizations.

ISCR Fiscal Year 2002 in Review
FY2002 Seminar Series, in reverse chronological order

Eric Jones, Enthought, Inc. ..... September 25-26, 2002
Allen Malony, University of Oregon ..... September 9-20, 2002
Bodo Parady, LLNL ..... September 17, 2002
Frederick Wong, University of California, Berkeley ..... September 17, 2002
Andrew Finney, California Institute of Technology ..... September 16, 2002
Ira Baxter, Semantic Designs, Inc. ..... September 6, 2002
James Davis, Stanford University ..... September 4, 2002
Ian Buck, Stanford University ..... August 30, 2002
Roger Davis, University of California, Davis ..... August 23, 2002
Branden Fitelson, San Jose State University ..... August 7, 2002
Fabio Milner, Purdue University ..... July 30-August 2, 2002
Zhiqiang Cai, Purdue University ..... July 29-August 2, 2002
Frank Mueller, North Carolina State University ..... July 22-August 2, 2002
Gene Golub, Stanford University ..... July 25, 2002
Vance Faber, Mapping Science, Inc. ..... July 1, 2002
Hoanh Vu, University of California, San Diego ..... June 28, 2002
Greg Pope, LLNL ..... June 26, 2002
Kim Yates, LLNL ..... June 25, 2002
Tim Matson, Intel Corporation ..... June 21, 2002
Linda Petzold, University of California, Santa Barbara ..... June 10-21, 2002
Serge Belongie, University of California, San Diego ..... June 20, 2002
Andrew Strelzoff, University of California, Santa Barbara ..... June 19-20, 2002
Michael Minion, University of North Carolina ..... June 12, 2002
Ralf Hiptmair, Universitaet Tuebingen ..... June 10, 2002
Erin Parker, University of North Carolina, Chapel Hill ..... May 28-31, 2002
Lars Arge, Duke University ..... May 22-25, 2002
Shivkumar Chandresakaran, University of California, Santa Barbara ..... May 24, 2002
Ming Gu, University of California, Berkeley ..... May 24, 2002
Srinivasan Parthasarathy, Ohio State University ..... May 17, 2002
Lori Freitag, Argonne National Laboratory ..... May 9-10, 2002
Raymond Loy, Argonne National Laboratory ..... May 9-10, 2002
Brent Gorda, Lawrence Berkeley National Laboratory ..... April 23, 2002
John Harer, Duke University ..... April 15-20, 2002
Patrick Roache, Ecodynamics Research Associates, Inc. ..... April 18, 2002
Long Lee, University of Washington ..... April 11-12, 2002
Florian Potra, University of Maryland, Baltimore County ..... April 3-5, 2002
Gabriel Silberman, IBM TJ Watson Laboratory ..... March 26, 2002
Phil Roth, University of Wisconsin-Madison ..... March 18-21, 2002
Sally McKee, University of Utah ..... March 3-14, 2002

Doug Enright, Stanford University ..... March 8, 2002
Gregory Pope, LLNL ..... February 28, 2002
Rolfe Schmidt, University of Southern California ..... February 27-28, 2002
Harold Trease, Pacific Northwest National Laboratory ..... February 27-28, 2002
Mathew Colgrove, Portland State University ..... February 21-22, 2002
Karen Karavanic, Portland State University ..... February 21-22, 2002
Thomas Hagstrom, University of New Mexico ..... February 7-8, 2002
Minnie Kerr, North Carolina State University ..... February 1, 2002
Luiz De Rose, IBM TJ Watson Research Center ..... January 29-30, 2002
David Stevens, LLNL ..... January 14, 2002
Terence Critchlow, LLNL ..... December 20, 2001
David Butler, Limit Point Systems ..... December 14, 2001
Giovanni Lapenta, Los Alamos National Laboratory ..... December 12, 2001
Achi Brandt, Weizmann Institute of Science ..... December 10-12, 2001
Linda Harden, LLNL ..... December 6, 2001
Annette Molinaro, University of California, Berkeley ..... December 6, 2001
Bernardo Cockburn, University of Minnesota ..... December 3, 2001
Bryan Buck, University of Maryland ..... November 30, 2001
William Dally, Stanford University ..... November 29, 2001
Patrick Hanrahan, Stanford University ..... November 29, 2001
Andries van Dam, Brown University ..... November 27, 2001
Ariel Shamir, The Interdisciplinary Center ..... November 19-21, 2001
Jim Douglas, Purdue University ..... November 16, 2001
Brandon Whitcher, National Center for Atmospheric Research ..... November 16, 2001
Herbert Edelsbrunner, Duke University ..... November 6, 2001
Walid Aref, Purdue University ..... October 29, 2001
Daniel Meiron, California Institute of Technology ..... October 25, 2001
Richard Braun, University of Delaware ..... October 23-25, 2001
Rob Ross, Argonne National Laboratory ..... October 15-19, 2001
Zhiqiang Cai, Purdue University ..... October 18, 2001
Jelena Tesic, University of California, Santa Barbara ..... October 15, 2001
Mikhail Shashkov, Los Alamos National Laboratory ..... October 12, 2001
Claudio Silva, AT&T Labs ..... October 8-12, 2001
Mikhail Shashkov, Los Alamos National Laboratory ..... October 11, 2001
Robert Bosch Jr., Stanford University ..... October 3, 2001
Zachary Peterson, University of California, Santa Cruz ..... October 1, 2001

FY2002 Institute for Terascale Simulation Lecture

Andries van Dam, Brown University

Visiting Faculty, Guests, Consultants, and Researchers

Visiting and Collaborating Researchers

Carlo Bottasso, Politecnico di Milano, Italy
Xiao-Chuan Cai, University of Colorado, Boulder
Zhiqiang Cai, Purdue University
Umit Catalyurek, Ohio State University
Tim Chartier, University of Washington
Herbert Edelsbrunner, Duke University
Alejandro Garcia, San Jose State University
Rod Fatoohi, San Jose State University
Nicolas Hadjiconstantinou, Massachusetts Institute of Technology
John Harer, Duke University
Ralf Hiptmair, Universitaet Tuebingen
Andrew Knyazev, University of Colorado, Denver
Kenneth Joy, University of California, Davis
Raytcho Lazarov, Texas A&M University
Sally McKee, University of Utah
Michael Minion, University of North Carolina
Frank Mueller, North Carolina State University
Joe Pasciak, Texas A&M University
Michael Pernice, Los Alamos National Laboratory
John Ruge, Front Range Scientific Computations, Inc.
Don Schwendeman, Rensselaer Polytechnic Institute
Claudio Silva, AT&T Labs
Lee Taylor, TeraScale LLC
Jacob Ystrom, Royal Institute of Technology, Stockholm

Participating Guests

Marsha Berger, New York Univerisity
Marian Brezina, University of Colorado
Alok Choudhary, Northwestern University
Richard Cook, University of California, Davis
Eric de Sturler, University of Illinois, Champaign-Urbana
Branden E. Fitelson, Argonne National Laboratory
John Fitzgerald, Lawrence Livermore National Laboratoy (Retired)
Sharon Frazier, Lawrence Livermore National Laboratory (Retired)
Alejandro Garcia, San Jose State University
Michael Griebel, Bonn University
Amarnath Gupta, San Diego Supercomputer Center
Bernd Hamann, University of California, Davis
Alan Hindmarsh, Lawrence Livermore National Laboratory (retired)

Kenneth Joy, University of California, Davis
Andrew Knyazev, University of Colorado, Denver
Johannes Kraus, University of Leoben, Austria
Falko Kuester, University of California, Irvine
Raytcho Lazarov, Texas A&M University
Byung Lee, University of Vermont
Lars Linsen, University of California, Davis
Ida Lozares, Lawrence Livermore National Laboratory (retired)
Bertram Ludaescher, San Diego Supercomputer Center
Kwan-Liu Ma, University of California, Davis
Jennifer Mariani, University of California, Davis
Sally McKee, Cornell University
Michael Minion, University of North Carolina
Frank Mueller, North Carolina State University
Christof Nuber, University of California, Davis
Beth Ong, Lawrence Livermore National Laboratory
Joseph Pasciak, Texas A&M University
Calton Pu, Georgia Institute of Technology
Elbridge Gerry Puckett, University of California, Davis
Ulrich Ruede, University of Erlangen
Paul E. Saylor, University of Illinois, Champaign-Urbana
Dan Schikore, CEI
Gregory L. Schussman, University of California, Davis
Rob van der Wijngaart, University of Bonn
Mladen Vouk, North Carolina State University
Gabriel Wittum, University of Heidelberg
Donald Wolitizer, California State University, Hayward
Jacob Ystrom, Royal Institute of Technology, Stockholm
Ludmil Zikatanov, Penn State University

Consultants

Bernie Alder, University of California (Professor Emeritus)
Randolph Bank, University of California, San Diego
Leo Breiman, University of California, Berkeley
Nancy Collins, University of Colorado, Boulder
Gene Golub, Stanford University
Anne Greenbaum, University of Washington
Charles Hansen, University of Utah
Michael Holst, University of California, San Diego
David Keyes, Old Dominion University
Heinz-Otto Kriess, University of California, Los Angeles
Luc Machiels, Swiss Federal Institute of Technology
Thomas Manteuffel, University of Colorado, Boulder

Stephen McCormick, University of Colorado, Boulder
Gregory Miller, University of California, Davis
Linda Petzold, University of California, Santa Barbara
Steve Schaffer, New Mexico Tech
Homer Walker, Worcester Polytechnic Institute

Department of Applied Science Faculty

Nelson Max
Garry Rodrigue
Rao Vemuri

Postdoctoral Researchers

Robert Anderson
Paul Castillo
Leonardo Colletti
Miguel Dumett
Jeff Hittinger
David Hysom
Bobby Philip
Markus Schordan
Leonid Tsap

University Collaborative Research Program Faculty and Students

Randolph Bank and Kathy Lu, University of California, San Diego
Padhraic Smyth and Scott Gaffney, University of California, Irvine
Mark van der Laan and Annette Molinaro-Clark, University of California, Berkeley
Berni Alder and Yihao Zheng, University of California, Davis
B.S. Manjunath and Jelena Tesic, University of California, Santa Barbara
Sutanu Sarkar, David Lopez, and Carlos Pantano, University of California, San Diego

LDRD Project Investigators

Terence Critchlow, LLNL, Center for Applied Scientific Computing
Bronis de Supinski, LLNL, Center for Applied Scientific Computing
Valerio Pascucci, LLNL, Center for Applied Scientific Computing

Students

Student Guests

Merico Argentati, University of Colorado, Denver
Peer-Timo Bremer, University of California, Davis
Paul Covello, University of California, Davis

Alessandro Croce, Politecnico di Milano, Italy
Davide Detomi, Politecnico di Milano, Italy
Ilja Friedel, California Institute of Technology
Boyce Griffith, New York University
Aglika Gyaourova, University of Nevada, Reno
Indrajeet Kumar, University of Utah
John Lai, University of California, Davis
Daniel Laney, University of California, Davis
Tushar Mohan, University of Utah
Vijay Natarajan, Duke University
Elijah Newren, University of Utah
Erin Parker, University of North Carolina
Jonathan Rochez, University of California, Davis
Jay Thomas, University of California, Davis
Yihao Zheng, University of California, Davis

Department of Applied Science Students

Aaron Fisher
Ben Gregorski
Ana Iontcheva
Joseph Koning
Robert Rieben
Josh Senecal

ISCR Students

Lucas Ackerman, Worcester Polytechnic Institute
Lisa Alano, University of San Francisco
Dave Alber, University of Illinois, Urbana-Champaign
John Anderson, University of the Pacific
Cheryl Barkauskas, Washington University, St. Louis
Bridget Benson, Cal Poly, San Luis Obispo
Rita Borgo, University of Pisa, Italy
Sadik (Han) Caglar, University of San Francisco
Timothy Campbell, University of Arizona
Karl Chen, University of California, Berkeley
John Clark, Northern Arizona University
Hillary Davis, Sierra High School
Paul Dostert, Texas A&M University
Roger Elion, Purdue University
Jason Estrada, Baylor University
Craig Falls, Hendrix College
Aaron Fisher, University of California, Davis
Jessica Fisher, Harvey Mudd College
Ilja Friedel, California Institute of Technology

Karen Glocer, University of California, Santa Cruz
Aglika Gyaourova, University of Nevada Reno
Matt Haddox, University of the Pacific
Chaz Hales, Brigham Young University
Keith Henderson, Purdue University
Amy Henning, University of California, Santa Cruz
Taylor Holliday, University of California, Davis
Bret Hull, University of California, Berkeley
Bryan Hunter, Allegheny College
Lorenzo Ibarria, Georgia Institute of Technology
David Jiambalvo, Rochester Institute of Technology
Ming Jiang, Ohio State
Kristaps Johnson, University of Rochester
Tzanio Kolev, Texas A&M University
Markus Kowarschik, University of Erlangen
Dedaimia Kozlovsky, University of Wisconsin-Madison
Ajith Mascarenhas, University of North Carolina
Deanna Midtaune, San Jose State University
Mohammed Mokbel, Purdue University
Evan Moran-Bernard, Carnegie Mellon University
Arne Naegel, Universitaet Heidelberg, Germany
James Newsome, University of Michigan
Luke Olson, University of Colorado
Susan Overstreet, Purdue University
Vera Pavel, Las Positas College
Ricardo Portillo, University of Texas at El Paso
Serban Porumbescu, University of California, Davis
Rachel Post, San Jose State University
Dan Rocco, Georgia Institute of Technology
Greg Scharlemann, Loyola Marymount University
Sunjeev Sikand, University of California, San Diego
Jonathan Strasser, University of California, Davis
Mark Stuppy, University of Missouri
Ryuta Suzuki, University of Minnesota
Erika Tarte, University of California, Berkeley
Charles Taylor, Brigham Young University, Idaho
Jeremy Thornock, University of Utah
Brian Truitt, New Mexico Institute of Mining and Technology
Ian Webb, Colorado State University
Aaron Wegner, Baylor University
Sanith Wijesinghe, Massachusetts Institute of Technology
Candita Woodis-Rucker, Diné College
I-Hsuan (Sandy) Wu, University of California, Davis

National Physical Science Consortium (NPSC) Student

Rachel Karchin, University of California, Santa Cruz

Workshops and Conferences

Copper Mountain Conference, Copper Mountain, CO ..... March 2002
Nonlinear Solvers and Differential Equations Workshop, Livermore, CA ..... March 2002
KAI ASCI Pathforward Workshop, Livermore, CA ..... March 2002
Conference on High Speed Computing, Gleneden Beach, OR ..... April 2002
XVth Householder Symposium on Numerical Linear Algebra, Peebles, Scotland ..... June 2002
BlueGene/L Workshop 2002, Tahoe, CA ..... August 2002
Algebraic Multigrid Summit Summit, Lake City, CO ..... September 2002

Institute for Scientific Computing Research

Seminar Series Abstracts

(in reverse cronological order)

September 25, 2002

Open Source Python Development in SciPy and the Chaco Plotting Package

Eric Jones
Enthought, Inc.

Email: eric@enthought.com

Abstract:

Enthought, Inc., is developing the next generation of open source scientific computing tools for Python. These packages inlcude SciPy (www.scipy.org), Chaco, and Weave. SciPy provides a comprehensive toolbox of algorithms that are indespensible to scientist and engineers. It covers a wide range of areas including linear algebra, optimization, integration, special functions, signal and image processing, genetic algorithms, ODE solvers, parallel programming, and others. Plotting and 2D data visualization are provided by Chaco, a cross platform interactive plotting package. Based on DisplayPDF, Chaco is designed to work with Python, TkInter, OpenGL, and other windowing systems as well as produce high quality PDF output. Weave solves problems where pure Python just isn't fast enough. Its allow users to seamlessly integrate C/C++ code directly within Python code. This provides maximum speed with minimum fuss. Enthought is collaborating with Patrick Miller of LLNL to make this even simpler by compiling Python code to efficient C code on the fly. This talk will provide a general overview of SciPy, Chaco, and Weave as well as a discussion of future directions for the tools.

Research web page: http://www.scipy.org/

Institution web page: http://www.enthought.com/

September 17, 2002

FPGA-based Optimizations of MPI Implementations

Fred Wong
University of California, Berkeley

Email: fredwong@fredwong.com

Abstract:

Message passing libraries like MPI provide high portability across different parallel architectures, but performance suffers. Previous research has shown that performance of MPI collective operations are very sensitive to the patterns and protocols they use for communication. Therefore, researchers have put much effort into rewriting their parallel applications to avoid using high-latency collective operations like MPI reduction. Many optimized MPI implementations utilized pre-configured hardware for fast manipulation of distributed data. However, these schemes are limited fundamentally by the hardware design.

Field Programmable Gate Arrays (FPGAs) have been used traditionally as a tool for developing new architectures before mass production. Today, many systems are pre-configured with FPGA units for the sole propose of fast vector computations. This presentation gives an overview of implementing an MPI library that utilizes FPGAs versatility for collective operations with data redistribution and manipulation, as well as process synchronization. The target platform is a cluster of SMPs. Each SMP consists of an FPGA-based memory bridge connecting the host with external memory modules that can be shared with other SMP nodes. Moreover, external FPGA-based processing boards are connected to the memory modules for direct access to the shared memory.

Speaker's web page: http://www.cs.berkeley.edu/~fredwong/

Institution web page: http://www.berkeley.edu/

The Systems Biology Workbench (SBW) and the Systems Biology Markup Language (SBML)

Andrew Finney
California Institute of Technology

Email: afinney@cds.caltech.edu

Abstract:

The tremendous amounts of data and research now emerging from molecular biotechnology have fueled an explosion in the development of software tools. Regrettably, developers often end up duplicating each other's efforts when writing different packages. In an effort to make it more attractive for developers to share rather than reimplement resources, we have implemented the Systems Biology Workbench (SBW), a free, open-source, application integration environment.

SBW uses a portable architecture that enables applications (potentially running on separate computers via SSH) to learn about and communicate with each other. SBW's communications facilities allow heterogeneous packages to be connected together using a remote procedure call mechanism; this mechanism uses a simple message-passing network protocol. The interfaces to the system are encapsulated in client libraries for different programming languages (currently C, C++, Delphi, Java, Perl, and Python). SBW is portable to Windows, Linux, and MacOS X.

In this seminar, I will discuss the SBW framework and the facilities it provides for software developers. I will demonstrate some of the many SBW-enabled modules now available for users. These include a high level parser and validator for SBML; a translator from SBML to MATLAB Simulink and ODE forms; a stochastic simulator; Jarnac, an ODE-based simulator; and JDesigner, a visual tool for drawing biochemical network models. With the release of SBW version 1.0 in June 2002, development has shifted toward more module development. For example, we are developing simulation tools with support for two- and three-dimensional spatial models.

Molecular biotechnology now makes it possible to build elaborate systems models, but the systems biology community needs agreed-upon information standards if models are to be shared, evaluated and developed cooperatively. I will discuss the Systems Biology Markup Language (SBML), an open XMLbased format developed to facilitate the exchange of models of biochemical reaction networks between software packages. SBML is being developed in collaboration with the groups responsible for BioSpice, Cellerator, CellML, DBsolve, E-CELL, Gepasi, Jarnac, ProMoT/DIVA, Virtual Cell, and StochSim. Specifically, I will discuss SBML Level 2, recently developed to provide a foundation for future development of SBML. Level 2 is being developed through the close collaboration between the SBML developers forum and the DARPA BioSPICE consortium.

The SBW version 1.0 package and documentation on SBW and SBML are now available from our project web site at http://www.sbw-sbml.org/.

Research web page: http://www.sbw-sbml.org/index.html

Institution web page: http://www.caltech.edu/

September 9, 2002

Recent Advances in the TAU Performance System

Allen Malony
University of Oregon

Email: malony@cs.uoregon.edu
Sameer Shende

Abstract:

Increasing complexity in parallel and distributed systems and software necessitates advances in performance technology towards robust tools and broad implementation. However, system and software complexity introduces challenges for performance instrumentation, measurement, analysis, and visualization. This talk presents recent advances in the TAU performance system in several areas:

- selective instrumentation control
- performance mapping and dynamic callpath profiling
- online performance analysis and visualization
- performance analysis and component software
- performance databases

In the area of instrumentation control, we are concerned with the removal of instrumentation in cases of high measurement overhead using rule-based analysis. Performance mapping focuses on measuring performance with respect to dynamic calling paths in which the static callgraph cannot be determined prior to execution. A prototype online performance data access, analysis, and visualization tool is reported as an example of performance interaction and steering system. We then describe our preliminary work on performance analysis of component software based on performance technology integration in component systems.

Lastly, we introduce the TAU Performance DataBase Framework (PerfDBF). Empirical performance evaluation of parallel and distributed systems often generates significant amounts of performance data and analysis results from multiple experiments as performance is being investigated and problems diagnosed. To better manage performance information, there is a strong motivation to develop performance database technology that can provide a common foundation for performance data storage and access. Such technology could offer standard solutions for how to represent the performance data, how to store them in a manageable way, how to interface with the database in a portable manner, and how to provide performance information services to a broad set of analysis tools and users. The talk reviews a prototype PerfDBF for TAU which stores parallel performance profiles in the database and allows analysis tools to perform intra-trial, inter-trial, and cross-experiment query and analysis.

Speaker's web page: http://www.cs.uoregon.edu/~malony/

Research web page: http://www.cs.uoregon.edu/research/paracomp/tau/

Institution web page: http://www.uoregon.edu/

September 6, 2002

DMS: Software Quality Enhancement via Automated Software

Ira Baxter
Semantic Designs, Inc.

Email: idbaxter@semdesigns.com

Abstract:

Remarkably, much of software engineering today is still carried out by manual methods. Significant productivity enhancements require automation, which in turn require tools that deeply understand programs. Generative programming is a class of tool technology that captures knowledge about how to generate code, enabling automation. While generative programming is normally considered to enhance forward engineering of software, the bulk of software engineering goes into software enhancement and maintenance.

This talk will show a practical, commercial "DMS" that uses program transformations on arbitrary source languages, a kind of generative programming, to carry out a surprisingly wide variety of software enhancement tasks, including installation of test probes, duplicate code detection and removal, and automated translation from one programming language to another. Such a tool could also be used to build custom program optimizers for scientific codes.

Speaker's web page: http://www.semdesigns.com/Company/People/idbaxter/index.html

Institution web page: http://www.semdesigns.com/index.html

September 4, 2002

Filling Holes in Complex Surfaces Using Volumetric Diffusion

James Davis
Stanford University

Email: jedavis@graphics.stanford.edu

Abstract:

We address the problem of building watertight 3D models from surfaces that contain holes, for example, sets of range scans that observe most but not all of a surface. We specifically address situations in which the holes are too geometrically and topologically complex to fill using triangulation algorithms. Our solution begins by constructing a signed distance function, the zero set of which defines the surface. Initially, this function is defined only in the vicinity of observed surfaces. We then apply a diffusion process to extend this function through the volume until its zero set bridges whatever holes may be present. If additional information is available, such as known-empty regions of space inferred from the lines of sight to a 3D scanner, it can be incorporated into the diffusion process. Our algorithm is simple to implement, is guaranteed to produce manifold non-interpenetrating surfaces, and is efficient to run on large datasets because computation is limited to areas near holes.

Speaker's web page: http://graphics.stanford.edu/~jedavis/

Research web page: http://graphics.stanford.edu/~jedavis/projects/index.html

Institution web page: http://www.stanford.edu/

August 30, 2002

Programmable Graphics Hardware

Ian Buck
Stanford University

Email: ianbuck@graphics.stanford.edu

Abstract:

A recent breakthrough has occurred in graphics hardware: fixed function pipelines have been replaced with programmable vertex and fragment processors. Instead of configuring the traditional OpenGL and DirectX pipeline with awkward API calls, today's hardware allows the execution of a programmer-specified assembly level program at every fragment or vertex. This flexibility opens endless possibilities for new rendering effects and realism.

In this talk, I will outline the brief history of programmable graphics hardware as well as describe in detail how to program these new architectures. Topics covered include NVIDIA vertex programs, DirectX pixel shaders, and the CG language as well as the Radeon 9700 and the NV30 architectures. I will also present material on trends and future directions in programmability.

Speaker's web page: http://graphics.stanford.edu/~ianbuck/

Institution web page: http://www.stanford.edu/

August 23, 2002

A Modular, Parallel Grid-Embedding Adaptation Scheme for General Applications

Roger Davis
University of California, Davis

Email: davisrl@ucdavis.edu

Abstract:

The push toward high-performance aerospace designs requires the development of new advanced, fluid dynamic simulation techniques that can provide rapid and accurate predictions. Many simulation procedures used in government, industry, and academia today are based on multi-block, structured-grid systems for reasons of computational efficiency and handling of complex geometries. With these fixed-grid numerical procedures, however, an increase in solution accuracy coincides with an increase in global computational grid density and, therefore, solution cost and time. This increase in solution time becomes unacceptable as the problem size and flow complexity increases. Also, if the solution evolves in such a way that high-gradient features, such as shocks and/or high-shear viscous flows, develop in regions of coarse grid spacing, the error in the solution becomes spatially non-uniform. As a result, the solution as a whole has the accuracy associated with the region of highest numerical error. Applications, such as sonic boom or dynamic stall control that fall into this category require intelligent numerical procedures that can recognize regions of high local error and enrich the local computational region in these regions to provide the desired uniform accuracy at minimal cost.

A numerical procedure has been developed that can provide the desired level of accuracy without the penalty of large solution time through the use of grid-embedding adaptation. In this numerical procedure, computational grid cells are directionally sub-divided only in the vicinity of high numerical error. As a result, the solution evolves to one of high global accuracy and low, uniform numerical error. Since computational points are only added where they are required, the solution is obtained with minimal computational cost and time. Grid embedding is performed by sub-dividing some portion of the cells in a block depending on the desired granularity of adaptation and the areas where adaptation is desired. Sub-division of cells may be performed multiple times to create a multi-layered region of grid embedding. Sub-division of the cells may also be performed "directionally" to improve the resolution of flow features or regions where the error is proportional to the grid spacing in the direction(s) of high gradients. Adaptation granularity (i.e., the size of the region to be sub-divided) can be made arbitrary and may be controlled through the use of domain decomposition of the original block(s) into mini-blocks. By performing this additional decomposition, the overhead of the grid adaptation scheme can be kept to a minimum. Solution costs are also reduced through the use of scalable distributed memory parallel computing across groups of mini-blocks along with shared memory parallel computing within the mini-blocks themselves. This strategy is consistent with new evolving computer architectures such as the "streaming" computer.

The application of this numerical procedure is demonstrated on a number of internal and external flow configurations. In addition, details of the existing numerical algorithms, computational strategy, and data processing will be described. Finally, the extension of this adaptive-grid system to create a library of API-callable "add on" routines to allow any existing multi-block structured-grid procedure to take advantage of grid-embedding adaptation is proposed and discussed.

Institution web page: http://www.ucdavis.edu/

Some Philosophical Issues in the Foundations of Statistics

Branden Fitelson
San Jose State University

Email: branden@fitelson.org

Abstract:

What are the ends of statistical (and/or scientific) inference? What are the most effective means toward these ends? Several answers to these questions (along with further questions, of course!) will be discussed. Several well-known controversies in the foundations of statistical and scientific inference will be addressed along the way.

Speaker's web page: http://fitelson.org/

Institution web page: http://www.sjsu.edu/

July 31, 2002

Structured Population Models

Fabio Milner
Purdue University

Email: milner@math.purdue.edu

Abstract:

Some basic population models will be briefly described, both structured and unstructured. Different areas of their applicability will be discussed, including demography, mathematical epidemiology, and parasitology.

Some numerical difficulties associated with their use will be discussed, including the size of the problem, as well as that of some of the numbers involved in simulations.

Speaker's web page: http://www.math.purdue.edu/~milner/

Institution web page: http://www.purdue.edu/

July 30, 2002

Numerical Methods for Elasticity: Stress-Displacement Formulation

Zhiqiang Cai
Purdue University

Email: zcai@math.purdue.edu

Abstract:

There have been many efforts during the past four decades to develop stable mixed-finite element methods for linear elasticity. This approach is based on the Hellinger–Reissner variational principle and approximate both the displacement and the stress tensor simultaneously. Mixed methods are preferable to standard displacement methods for some important practical problems, e.g., the modelling of nearly incompressible or incompressible materials, and the modelling of plastic materials where the elimination of the stress tensor is difficult. Stable stress-displacement finite elements are extremely difficult to construct. This is caused by the symmetry constraint of the stress tensor. In this talk, I will introduce a new stable stress-displacement pair that has the least degrees of freedom and that preserves exact equilibrium for constant body force and exact symmetry of the stress tensor.

The second part of the talk is on least-squares approaches. The principal attractions of the method include freedom in the choice of finite element spaces, fast multigrid solver for the resulting algebraic equations, and free practical and sharp *a posteriori* error measure for adaptive mesh refinements. The first approach is based on the stress-displacement formulation. The second one is based on the stress-displacement-pressure formulation. The introduction of the pressure variable enables us to treat incompressible materials and to use simple standard finite element spaces approximating stress, displacement, and pressure.

Speaker's web page: http://www.math.purdue.edu/~zcai/

Institution web page: http://www.purdue.edu/

July 29, 2002

SPAN: Shared-Memory Performance Analysis

Frank Mueller
North Carolina State University

Email: mueller@cs.ncsu.edu

Abstract:

For contemporary high-performance clusters of SMPs, it has been found that a number of scientific applications utilizing a mixed mode of MPI+OpenMP are performing worse than when relying on MPI only. The objective of our work is to determine the sources of inefficiencies in utilizing memory hierarchies for threaded programs vs. parallel processes and to assist the programmer in alleviating these problems.

We are developing a technique for generating partial data traces through binary instrumentation. We employ a portable method for extracting precise data traces for partial executions of arbitrary applications and a set of hierarchical structures for compactly representing these accesses. We utilize incremental cache analysis based on partial traces to detect memory bottlenecks by correlating cache conflicts to their sources and by distinguishing intra-processor conflicts vs. inter-processor coherence misses. The conflicts will be depicted with a reference to the source program, thereby guiding the programmer to hot spots of cache misses. This enables the programmer to identify opportunities for changing the data layout and for performing algorithmic transformations to improve memory performance. Initial results on benchmarks show the potential to find memory bottlenecks and the ability of guiding the programmer to the causes of problems. These results also indicate that there is a potential to automate some optimizations in response on a running application via dynamic binary rewriting.

Speaker's web page: http://moss.csc.ncsu.edu/~mueller/

Institution web page: http://www.ncsu.edu/

July 25, 2002

Solution of Non-Symmetric, Real Positive Linear Systems

Gene Golub
Stanford University

Email: golub@sccm.stanford.edu

Abstract:

The methods we discuss use a Hermitian/skew-Hermitian splitting (HSS) iteration and its inexact variant, the inexact Hermitian/skew-Hermitian splitting (IHSS) iteration, which employs inner iteration processes at each step of the outer HSS iteration. Theoretical analyses show that the HSS method converges unconditionally to the unique solution of the system of linear equations. Moreover, we derive an upper bound of the contraction factor of the HSS iteration, which is dependent solely on the spectrum of the Hermitian part. Numerical examples are presented to illustrate the effectiveness of both HSS and IHSS iterations. In addition, a model problem of three-dimensional convection-diffusion equation is used to illustrate the advantages of our methods.

Speaker's web page: http://www-sccm.stanford.edu/Faculty/Golub.html

Institution web page: http://www.stanford.edu/

July 1, 2002

Convergence and Non-convergence in Algorithms for Prediction

Vance Faber
Mapping Science, Inc.

Email: vance@mappingscience.com

Abstract:

Many popular methods for pattern recognition and prediction involve iteration. These algorithms have been tested by using them on a wide range of problems, but no mathematical proof of convergence (in any useful sense of the word) exists. In other cases, mathematical proofs of convergence do not include information on rates of convergence.

I illustrate this situation by examining three fairly well-known algorithms: (A) Lloyd's method for clustering (a variant of k-means); (B) Breiman's algorithm for approximation of data by a hinge function; and (C) a prediction algorithm of mine which is a variant of the training stage for radial basis function neural networks. Algorithm A is known to converge sublinearly, but I construct an example to show that it is numerically non-convergent (on a finite precision machine it stalls significantly short of the solution). I analyze algorithm B and show how to produce examples where not only does it fail to converge but it actually diverges from any starting guess. Algorithm C not only converges but it can be proven that the sum of the squares of the errors that it makes will be forever bounded.

Institution web page: http://www.mappingscience.com/

OSCAR: Taking Clusters Into the Mainstream

Tim Mattson
Intel Corporation

Email: timothy.g.mattson@intel.com

Abstract:

The academic and national laboratory research communities love clusters. I doubt that there is a national lab or research university in the United States that doesn't have several of them. If you look at the corporate technical computing world, however, clusters are rare. This is slowly changing, but with few exceptions, the corporate world with its expensive labor costs can't deal with clusters today.

I believe it is in all of our interests to change this situation. We created the cluster-computing craze, and it's our job to take it into the mainstream. In this talk, I will discus one long-range effort to accomplish this; the Open Cluster Group with our first "product" OSCAR. I will describe the motivation for creating OSCAR, its current status, and our plans for the future.

Research web page: http://www.openclustergroup.org/

Institution web page: http://www.intel.com/

June 21, 2002

OpenMP: a Status Report and Look to the Future

Tim Mattson
Intel Corporation

Email: timothy.g.mattson@intel.com

Abstract:

Since its introduction in 1997, OpenMP has rapidly grown to become the standard API for writing multi-threaded applications. Its rapid transformation into a de-facto standard wasn't due to brilliance by its creators. Rather, OpenMP was a standardization of accepted practice with a narrow scope restricted to SMP machines. Coupled with support from all the standard SMP vendors, OpenMP was practically guaranteed to become a successful API.

HPC architecture marches forward and so too must the APIs that programmers use. If OpenMP wants to remain relevant, it must adapt to the increasing dominance of non-SMP systems. In addition, OpenMP needs to add constructs to support a wider range of applications. While we serve regular data driven applications quite well, we need to better handle applications that are driven by traversal of linked lists and algorithms with highly variable and unpredictable data access patterns.

In this talk, I will provide an overview of where we are today with OpenMP. I will then discuss some of the innovations we are considering as we move into the future. Finally, I plan to sit back and take notes and learn as the audience tells me what they want in future releases of OpenMP.

Research web page: http://www.openmp.org/

Institution web page: http://www.intel.com/

Shape Matching and Image Segmentation

Serge Belongie
University of California, San Diego

Email: sjb@cs.ucsd.edu

Abstract:

In this talk I will review some recent work I have done with colleagues at UC Berkeley and UCSD on shape matching and image segmentation. In the first half of the talk, I will present an approach to measuring similarity between shapes based on a novel descriptor called the "shape context." In this framework, the measurement of similarity is preceded by (1) solving for correspondences between points on the two shapes, and (2) using the correspondences to estimate an aligning transform. Object recognition results are presented for silhouettes, trademarks, handwritten digits, and 3D objects.

The preceding shape matching method assumes the objects to be recognized have been pre-segmented from the background, or that segmentation is not necessary, e.g., due to a uniform background. In the second half of the talk, I will discuss recent advances we have made on the problem of image segmentation using spectral graph theoretic methods.

Speaker's web page: http://www-cse.ucsd.edu/~sjb/

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

June 19, 2002

JMPL: Java Math Package Launcher Automating GUI Development and Maintenance For Science and Engineering

Andrew Strelzoff
University of California, Santa Barbara

Email: strelz@engineering.ucsb.edu

Abstract:

Developers of scientific applications typically lack the time and expertise necessary to produce and maintain high quality Graphic User Interfaces (GUIs). In this presentation we introduce JMPL: Java Math Package Launcher, an environment for developing and maintaining GUI front ends for scientific computing. JMPL assists the development process by automatically generating matching GUI skeletons from scientific applications. These skeletons may be further developed and customized in the JMPL editor before distribution to the community of users. The JMPL revision manager assists in the process of updating the GUI as the underlying program changes. The result is an environment for continuing publication and improvement.

In this talk we will focus on two critical areas of the JMPL project. The first is the fundamental task of discovering interface requirements from a program and the translation of these requirements into an intermediate format, XML/UML. This intermediate format is then parsed by a browser which dynamically constructs the GUI. The second area covered will be the development of an algorithm to estimate which parts of an underlying application have changed. This is the input to the revision manager which will allow it to assist in keeping the GUI up to date. Preliminary results will be presented, and the future development and direction of the project including extension to object-oriented languages will be discussed.

Speaker's web page: http://www.cs.ucsb.edu/~strelz/

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

June 13, 2002

Toward the Multiscale Numerical Solution of Chemically Reacting Systems

Linda Petzold
University of California, Santa Barbara

Email: petzold@engineering.ucsb.edu

Abstract:

Stochastic simulation enables the numerical simulation of the time evolution of a well-stirred chemically reacting system in a way that takes proper account of the randomness that is inherent in such a system. However, the computer times required to simulate over reasonable time periods are prohibitively long if the molecular populations of any of the reactant species are large. In cellular systems, for example, this is nearly always the case. At the same time, the deterministic reaction rate equations are well-suited for modeling of systems with large molecular populations, but cannot accurately capture the evolution of species whose populations are small.

In this talk we will examine a range of models from stochastic simulation to the deterministic reaction rate equations. Since the models naturally segue to each other, there is some reason to hope that a multiscale numerical method that treats each species and reaction at the most appropriate scale could be developed. We will describe our progress on the first phase of this effort: the development of efficient and robust accelerated stochastic simulation algorithms inspired by the recently-developed tau-leaping method of Gillespie.

This is joint work with Dan Gillespie and Muruhan Rathinam.

Speaker's web page: http://www.me.ucsb.edu/dept_site/people/new_faculty_pages/petzold_page.html

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

June 12, 2002

Multi-Implicit Methods for Advection-Diffusion-Reaction and Related Problems

Michael Minion
University of North Carolina, Chapel Hill

Email: minion@amath.unc.edu

Abstract:

A new method for the temporal integration of partial differential equations possessing multiple time scales will be presented. The method is based on a multi-implicit variation of spectral deferred corrections (MISDC) in which two or more components of the solution are treated implicitly. As in traditional operator splitting methods, when applied to Advection-Diffusion-Reaction (A-D-R) equations, MISDC methods allow both the diffusion and the reaction terms to be treated in an implicit but uncoupled manner. Furthermore, the size of the time steps used for the the diffusion and reaction terms can be chosen independently. Higher-order temporal accuracy is achieved by reducing both the splitting and truncation errors during the deferred correction iterations. Examples of MISDC method of order 3, 4, and 5, applied to model A-D-R problems will be presented. The difficulties in imposing accurate boundary conditions for time-dependent boundary conditions will also be discussed.

Speaker's web page: http://www.amath.unc.edu/Faculty/minion/

Institution web page: http://www.unc.edu/

June 11, 2002

Multigrid for Eddy Current Computation

Ralf Hiptmair
Universitaet Tuebingen

Email: hiptmair@na.uni-tuebingen.de

Abstract:

I set out from the quasi-state eddy current model in time domain. Because of stability, implicit time-stepping in required. In a finite element setting, each timestep involves the solution of a discrete variational problem for a bilinear form (•, •) + (curl •, curl •) posed over H(curl). I rely on Nedelec's H(curl)-conforming finite elements (edge elements), which yield a viable discretization for Maxwell's equations. Ultimately, we are faced with a large sparse linear system of equations to be solved in each timestep, for which we aim to construct a fast iterative solver.

A naïve multigrid approach is doomed, since the problem is hardly elliptic on the large kernel of the curl-operator. Hence, plain nodal relaxations are only effective for high-frequency error components that are sufficiently orthogonal to this kernel. Yet, the exceptional properties of edge elements provide discrete potentials in the space of piecewise polynomial Lagrangian finite element functions. Thus, additional smoothing in the kernel becomes computationally feasible and, ultimately, leads to a multigrid method, which achieves “textbook multigrid efficiency.”

Speaker's web page: http://na.uni-tuebingen.de/~hiptmair/

Institution web page: http://www.uni-tuebingen.de/

June 10, 2002

Higher Order Whitney Forms

Ralf Hiptmair
Universitaet Tuebingen

Email: hiptmair@na.uni-tuebingen.de

Abstract:

The calculus of differential forms can be used to devise a unified description of discrete differential forms of any order and polynomial degree on simplicial meshes in any spatial dimension. A general formula for suitable degrees of freedom is also available. Fundamental properties of nodal interpolation can be established easily. It turns out that higher order spaces, including variants with locally varying polynomial order, emerge from the usual Whitney-forms by local augmentation. This paves the way for an adaptive p-version approach to discrete differential forms.

Speaker's web page: http://na.uni-tuebingen.de/~hiptmair/

Institution web page: http://www.uni-tuebingen.de/

May 29, 2002

Compile-time Analysis of Loop Nest Behavior in Cache Memory

Erin Parker
University of North Carolina, Chapel Hill

Email: parker@cs.unc.edu

Abstract:

It is difficult to talk about high-speed computing without considering the bottleneck of accessing memory. For large applications, achieving good performance means taking full advantage of fast, small cache memories. Applications exhibiting good locality realize the benefits of caches; however, locality is often elusive and volatile. As an application evolves, modifications to its code and platform lead to changes in its locality. The need for efficient, analytical models of the memory behavior of programs is evident.

We provide an exact model of the behavior of loop nests executing in a memory hierarchy. Our model is flexible enough to describe whether a memory access hits or misses in the cache, whether a memory block experiences reuse while in the cache, and at what point a memory block experiences no more reuse before being evicted from the cache. We use Presburger arithmetic to express the scenarios above and a combination of tools to simplify and identify such memory accesses. We can then use this information to predict the number of cache misses incurred by loop nests, to avoid polluting the cache with memory blocks that are not reused, to choose when to evict memory blocks from cache in order to reduce bus congestion, and even to turn off cache lines holding data not reused to reduce cache leakage power.

Speaker's web page: http://www.cs.unc.edu/~parker/

Institution web page: http://www.unc.edu/

May 24, 2002

Fast and Stable Direct Solvers for Matrices from PDEs and Integral Equations

Shivkumar Chandresakaran
University of California, Santa Barbara

Email: shiv@ece.ucsb.edu

Abstract:

We will present the "sequentially semi-separable (SSS)" class of dense structured matrices. This class includes dense matrices, semi-separable matrices, the inverse of both dense and semi-separable matrices, and many others. We will present fast and stable algorithms for the solution of linear systems of equations with SSS matrices. We will also present fast and memory efficient techniques for constructing the SSS representation. Finally we will show how two spectral methods, one for two-point boundary value problems and one for exterior two-dimensional scattering problems, can be solved in linear time using the SSS representation. Experimental results on the efficiency of the fast algorithms will also be presented.

This is joint work by S.Chandrasekaran, M.Gu, P.Dewilde, T.Pals, A.-J.van der Veen.

Speaker's web page: http://www.ece.ucsb.edu/Faculty/Chandrasekaran/default.html

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

May 23, 2002

Cache-Oblivious Priority Queue and Graph Algorithm Applications

Lars Arge
Duke University

Email: large@cs.duke.edu

Abstract:

In recent years algorithms designed to take advantage of the hierarchical memory system of modern machines have been extensively researched. Most research on such algorithms has been done in two-level (or external memory) memory models. The advantage of such models is that they are simple, while the disadvantage is that they focus efforts on one level of the multilevel hierarchy. Recently a new model that combines the simplicity of two-level models with the realism of more complicated and realistic hierarchical models was introduced. The idea in this "cache oblivious" model is to design algorithms in a two-level model but avoid using the parameters describing the two-level hierarchy in the algorithm description. This way an algorithm that is efficient in the two-level model is efficient on any level of an arbitrary multilevel memory hierarchy.

In this talk we first discuss fundamental sorting and searching results in the cache-oblivious model. Then we describe an optimal cache-oblivious priority queue data structure, supporting insertion, deletion, and deletemin operations in O((1/B) log(M/B) (N/B)) amortized memory transfers, where M and B are the memory and block transfer sizes of any two consecutive levels of a multilevel memory hierarchy. These bounds match the bounds of several previously developed external-memory (cache-aware) priority queue data structures, which all rely crucially on knowledge about M and B. Priority queues are a critical component in many of the best known external-memory graph algorithms, and the cache-oblivious priority queue can, for example, be used to develop several cache-oblivious graph algorithms. In most cases the memory-transfer bounds of these algorithms match the bounds of the best known external-memory algorithms.

Joint work with Michael Bender, Erik Demaine, Bryan Holland-Minkley and Ian Munro presented at STOC'02.

Speaker's web page: http://www.cs.duke.edu/~large/

Institution web page: http://www.duke.edu/

May 17, 2002

Incrementally Mining Frequent Patterns and Mining Bio-medical Datasets

Srinivasan Parthasarathy
Ohio State University

Email: srini@cis.ohio-state.edu

Abstract:

In the first part of the talk I will describe our work in incrementally mining frequent patterns in evolving datasets. Changes to a dataset can invalidate existing patterns or introduce new ones. Simply re-executing algorithms from scratch on a database update can result in an explosion in the computational and I/O resources required. What is needed is a way to incrementally process the data and update the information that is gleaned while being cognizant of the interactive requirements of the knowledge discovery process.

During the second part of the talk I will describe some recent work in mining biomedical datasets. A common theme across the proposed techniques here is that they are cognizant of the structure/shape-function relationship that pervades a lot of applications in these domains. I will describe techniques that are cognizant of features relating to the shape and structure of biological elements (corneal shapes, protein structures) and how these techniques can be used in conjunction with classical data mining techniques to mine such datasets.

Acknowledgements: The first part of the talk is joint work with researchers from RPI and UFMRG (Brazil). The second part of the talk is joint work with graduate students at OSU.

Speaker's web page: http://www.cis.ohio-state.edu/~srini/

Research web page: http://www.cis.ohio-state.edu/~srini/

Institution web page: http://www.ohio-state.edu/index.php/

May 10, 2002

Data Reduction Techniques using Subsampling

Lori Freitag and Raymond Loy

Argonne National Laboratory

Email: freitag@mcs.anl.gov

Abstract:

The interactive visualization and exploration of large scientific data sets is a challenging and difficult task; their size often far exceeds the performance and memory capacity of even the most powerful graphics workstations. To address this problem, we have created a technique that combines hierarchical data reduction methods with parallel computing to allow interactive exploration of large data sets while retaining full-resolution capability. Our technique is based on subsampling the original data set using a distributed octree, and we describe the software architecture of the system in some detail. We develop performance models for this and uniform subsampling strategies to investigate the tradeoffs between multiresolution and uniform grid methods in terms of performance and approximation errors.

Speaker's web page: http://www-unix.mcs.anl.gov/~freitag/

Research web page: http://www-unix.mcs.anl.gov/~freitag/

Institution web page: http://www.anl.gov/

April 23, 2002

Future Technologies Activities at LBL

Brent Gorda
Lawrence Berkeley National Laboratory

Email: bgorda@lbl.gov

Abstract:

Brent Gorda, newly appointed group lead for the Future Technologies Group at the Lawrence Berkeley National Laboratory, will talk about the activities of the group in the area of High Performance Computing.

Research web page: http://www.nersc.gov/research/ftg/

Institution web page: http://www.lbl.gov/

April 18, 2002

An Overview Of Verification and Validation for Computational Science and Engineering

Patrick Roache
Ecodynamics Research Associates, Inc.

Email: hermosa@swcp.com

Abstract:

This seminar presents an overview of Verification and Validation for Computational Science and Engineering and related areas. Topics include: semantics and definitions; Verification of Codes by the Method of Manufactured Solutions (MMS); Partitioning the Option Matrix; Verification of Calculations using the Grid Convergence Index (GCI); Single-Grid Error Estimators; validation experiments and metrics; "nearby" problems; scale of unsteadiness; scaling-up; prevalence of errors in scientific software; use of static analyzers in QA systems; and making QA your friend.

Speaker's web page: http://kumo.swcp.com/hermosa/html/pjr.html

April 12, 2002

An Immersed Interface Method for Incompressible Navier–Stokes Equations

Long Lee
University of Washington

Email: longlee@amath.washington.edu

Abstract:

Peskin's immersed boundary (IB) method is one of the most robust numerical methods for solving fluid problems with immersed interfaces. In this talk, I will present a method motivated by the immersed boundary method, which allows one to model the motion of flexible membranes or other structures immersed in viscous incompressible fluid using a fluid solver on a fixed Cartesian grid. The IB method uses a set of discrete delta functions to spread the entire singular force exerted by the immersed boundary to the nearby fluid grid points. Our method instead incorporates part of this force into jump conditions for the pressure, avoiding discrete dipole terms that adversely affect the accuracy near the immersed boundary. This has been implemented for the two-dimensional incompressible Navier–Stokes equations using a high-resolution finite-volume method for the advective terms and a projection method to enforce incompressibility. In the projection step, the correct jump in pressure is imposed in the course of solving the Poisson problem. This gives sharp resolution of the pressure across the interface and also gives better volume conservation than the traditional IB method. Comparisons between this method and the IB method are presented for several test problems. Numerical studies of the convergence and order of accuracy are included. This work is joint with Randy LeVeque.

Speaker's web page: http://amath.washington.edu/~longlee/

Institution web page: http://www.washington.edu/

April 3, 2002

Numerical Methods for Air Pollution Modeling and Simulation

Florian Potra

University of Maryland, Baltimore County

Email: potra@math.umbc.edu

Abstract:

A detailed understanding of the relationships between the emissions and the resulting distribution of primary and secondary species in the atmosphere is a requisite to designing actions for the maintenance of a healthy environment. To better predict the transport and fate of trace gases and pollutants in the atmosphere, comprehensive atmospheric transport-chemistry models have been developed. These models treat transport, chemical transformations, emissions and deposition processes in an integrated framework. They can be used to study the response of the pollutant distributions to system perturbations, and to link pollutant distributions to environmental effects. However, these comprehensive atmospheric chemistry models are computationally intensive because the governing equations are nonlinear, highly coupled, and extremely stiff. Our ability to fully utilize these models remains severely limited by today's computer technology and algorithmic development.

The talk will focus on improved numerical methods for solving the stiff differential equations associated with the chemical reactions and on new software tools for generating the computer code.

Speaker's web page: http://www.math.umbc.edu/~potra/

Institution web page: http://www.umbc.edu/

March 26, 2002

Cooperative Research at IBM's Centre for Advanced Studies (CAS)

Gabriel Silberman
IBM, TJ Watson Laboratory

Email: gabbys@us.ibm.com

Abstract:

The Centre for Advanced Studies (CAS) was established by the IBM Toronto Laboratory in 1991 to strengthen the links between academia and industry. Since then, this research group has been helping bridge the gap between industry requirements for new technology and academia.

CAS places an emphasis on solving real problems encountered by developers. These lead to long-term projects from which prototypes and publication-quality research can be delivered, all in collaboration with other members of the research community.

Speaker's web page: http://www.cas.ibm.ca/director/

Research web page: http://www.cas.ibm.ca/index.shtml

Institution web page: http://www.watson.ibm.com/

March 18, 2002

Scalable On-line Automated Performance Diagnosis

Philip Roth
University of Wisconsin–Madison

Email: pcroth@cs.wisc.edu

Abstract:

Automated search is an effective strategy for finding application performance problems. Search-based automated tools are especially important for performance diagnosis of large-scale applications. Unfortunately, existing automated performance diagnosis tools do not scale well.

In the first part of this talk I present Deep Start, a new algorithm that attacks the scalability problem of automated performance diagnosis for applications with a large number of functions. Deep Start uses stack sampling to augment a simple search-based automated performance diagnosis strategy. Our hybrid approach locates performance problems more quickly and finds problems hidden from a more straightforward search strategy. Deep Start uses stack samples collected as a by-product of normal search instrumentation to find deep starters, functions that are likely to be application bottlenecks. Deep Starters are examined early during a search to improve the likelihood of finding performance problems quickly. We implemented the Deep Start algorithm in the Performance Consultant, the automated bottleneck detection component of the Paradyn performance tool. Deep Start found half of our test applications' known bottlenecks 32% to 59% faster than the Performance Consultant's current call graph-based search strategy, and finished finding bottlenecks 10% to 61% faster than the call graph strategy. In addition to improving search time, Deep Start often found more bottlenecks than the call graph search strategy.

In the second part of the talk, I will outline our plans for improving the scalability of automated performance diagnosis tools for applications that create a large number of processes. I will discuss our strategy for distributing the Performance Consultant, including our design for a general-purpose software-based multicast/reduction network. I will also present the Sub-graph Folding Algorithm, our approach for improving the scalability of the Performance Consultant's results display.

Speaker's web page: http://www.cs.wisc.edu/~pcroth/

Research web page: http://www.paradyn.org/

Institution web page: http://www.wisc.edu/

March 8, 2002

A Hybrid Particle Level Set Method for Improved Interface Capturing

Doug Enright
Stanford University

Email: enright@stanford.edu

Abstract:

In this talk, a new numerical method for improving the area (volume in three spatial dimensions) conservation properties of the level set method when the interface is passively advected in a flow field is presented. This method uses Lagrangian marker particles to rebuild the level set in regions that are under-resolved as is often the case for flows undergoing stretching and/or tearing. The overall method maintains a smooth geometrical description of the interface and the implementation simplicity characteristic of the level set method. The method compares favorably with volume of fluid methods in the conservation of mass and purely Lagrangian schemes for interface resolution in highly vortical flows. Numerical examples in both two and three spatial dimensions are shown along with preliminary results of the use of this method in free-surface calculations.

Speaker's web page: http://www-sccm.stanford.edu/~enright/

Research web page: http://cits.stanford.edu/

Institution web page: http://www.stanford.edu/

March 5, 2002

Smarter Memory Controllers: Improving Memory System Performance from the Bottom Up

Sally McKee
University of Utah

Email: sam@cs.utah.edu

Abstract:

Microprocessor speed is increasing much faster than memory system speed: the speed-growth curves are diverging exponentials. The traditional approach to attacking the memory system bottleneck has been to build deeper and more complex cache hierarchies.Although caching may work well for parts of programs that exhibit high locality, many important commercial and scientific workloads lack the locality of reference that makes caching effective. Studies have shown that memory bus and DRAM latencies cause an 8X slowdown from peak performance to actual performance on commercial database workloads, that the efficiency of current caching techniques is generally less than 20% of an optimal cache's, and that cache sizes are up to 2,000 times larger than an optimal cache would be. The evidence is clear: no matter how hard we push it, traditional caching cannot bridge the growing processor-memory performance gap.

This talk presents research that attacks the memory problem at a different level—the memory controller. We describe the Impulse Adaptable Memory System being built at the University of Utah. This general-purpose, uniprocessor system that improves performance within the cache hierarchy and the memory back end for both regular and irregular computations. It does this in three ways: by optimizing the use of DRAM resources in the memory controller back end, by prefetching data within the memory controller and delivering it to the processor only when requested, and by remapping previously unused physical addresses within the memory controller. Extending the virtual memory hierarchy in this way allows optimizations that improve the efficiency of the system bus and the performance of the CPU's cache hierarchy and translation lookaside buffer (TLB). Impulse represents a combined hardware/software solution in that the compiler, OS, or application writer supplies the access pattern information that the memory controller exploits.

We provide an overview of Impulse functionality, and then present details of the Parallel Vector Access (PVA) mechanism that optimizes the use of DRAM resources by gathering data in parallel within the memory controller. The PVA performs cache-line fills as efficiently as a normal, serial controller, and performs strided vector accesses from three to 33 times faster. The scalable design is two-five times faster than other-gathering mechanisms with similar goals, at the cost of only a slight increase in hardware complexity.

Speaker's web page: http://www.cs.utah.edu/~sam/

Institution web page: http://www.utah.edu/

February 28, 2002

Fast Wavelet-based Evaluation of Polynomial Range-Sum Queries

Rolfe Schmidt
University of Southern California

Email: rrs@pollux.usc.edu

Abstract:

Many range aggregate queries, including COUNT, SUM, and COVARIANCE, can be efficiently derived from a fundamental class of queries: the polynomial range-sums. This talk will describe how range-sum queries can be thought of as linear operators on the data density function, and evaluated exactly in the wavelet domain. With this intuition, I will introduce ProPolyne, a novel technique to evaluate arbitrary polynomial range-sums progressively. At each step of the computation, ProPolyne makes the best possible wavelet approximation of the submitted query. When the filters are chosen appropriately, polynomial range-sum queries have very sparse wavelet representations. The result is an excellent exact-query answering technique that provides highly accurate data-independent approximate results long before query evaluation is finished. I will conclude by discussing a few of the challenges that must be addressed to fully exploit this technique, and our progress in overcoming them.

Research web page: http://dimlab.usc.edu/research.html

Institution web page: http://www.usc.edu/

February 27, 2002

Mesh-based Modeling Methods for Computational Biology Applications

Harold Trease
Pacific Northwest National Laboratory

Email: het@pnl.gov

Abstract:

Biological systems are inherently multi-dimensional because of their geometry, heterogeneity, and time-dependent behavior. One approach to capturing these aspects is to use mesh-based modeling methods as the basis for a computational simulation framework that explicitly represents complex geometry, non-uniform heterogeneous distributions of concentrations, time-dependence, and complex boundary/initial conditions. To build such simulation models there are several problems that must be solved to integrate biology, biophysics, mathematics, and computer science into a useable model.

This presentation demonstrates the approaches that we at the Pacific Northwest National Laboratory are taking toward creating systems biology models from everything from computational cell biology to complex organ models. Specifically, we will show the virtual microbial cell and virtual mammalian nose/lung models that we are constructing. This includes using various forms of digital image reconstruction and feature extraction for NMR/MRI, confocal, and ultrasound to obtain computation geometry models, from which our mesh models are generated. We also show the mathematical models that we are using to simulate the combined effects of fluid-dynamics, tissue mechanics, reaction/diffusion, and discrete particle transport for these multi-component biological systems. The simulation framework presented is part of the DOE SciDAC Terascale Simulation Tools and Technology (TSTT) Center. Both LLNL and PNNL are contributing members of this center.

Institution web page: http://www.pnl.gov/

February 21, 2002

Parallel Performance Diagnosis with PperfDB

Karen Karavanic
Portland State University
Mathew Colgrove
Email: karavan@cs.pdx.edu

Abstract:

The PPerfDB research project is developing methods for diagnosing large-scale parallel applications using data from more than one program execution. This multi-execution view enables comparison of different code versions, different communication libraries, different platforms, and even different monitoring tools. The motivating vision driving the PPerfDB project is a tool that can diagnose the performance of large-scale parallel applications automatically, without user involvement.

PPerfDB can launch new performance experiments using IBM's Dynamic Probe Class Library (DPCL). PPerfXchange, a tool component under development, allows geographically dispersed data stores to be included as sources of data for PPerfDB. PPerfXchange models each site's performance data as XML documents, based upon a global schema, allowing client applications to retrieve views on performance data using the XQuery Language.

In this talk we will describe PPerfDB and report results that demonstrate our prototype's ability to combine performance data from traditional tracing tools and dynamic instrutmention-based tools for comparative performance profiling.

Speaker's web page: http://www.cs.pdx.edu/~karavan/research.html

Institution web page: http://www.pdx.edu/

February 7, 2002

Whats New with the Wave Equation

Thomas Hagstrom
University of New Mexico

Email: hagstrom@math.unm.edu

Abstract:

In recent years there have been a number of subtantial advances in numerical techniques for solving the wave equation. These include:

(i) Efficient methods for truncating unbounded domains which can provide arbitrary accuracy.

(ii) Fast methods for evaluating the standard integral solution formulas, akin to the fast multipole method in the frequency domain.

(iii) New discretization strategies including stabilized one-sided differencing near boundaries and unconditionally stable, explicit 2-step marching methods based on new integral formulas for the solution.

I will review these new methods, discuss their applicability to more general hyperbolic initial-boundary value problems, and speculate on future developments.

Institution web page: http://www.unm.edu/

February 1, 2002

New Multivalue Methods for Differential Algebraic Equations

Minnie Kerr
North Carolina State University

Email: mykerr@unity.ncsu.edu

Abstract:

Multivalue methods are slightly different from the general linear methods John Butcher proposed over 30 years ago. Multivalue methods capable of solving differential algebraic equations (DAEs) have not been developed. In this paper, we have constructed three new multivalue methods for solving DAEs of index 1, 2, or 3, which include multistep methods and multistage methods as special cases. The concept of stiff accuracy will be introduced, and convergence results will be given based on the stage order of the methods. These new methods have the diagonal implicit property and thus are cheap to implement and will have order 2 or more for both the differential and algebraic components. We have implemented these methods with fixed step size, and they are shown to be very successful on a variety of problems. Some numerical experiments with these methods are presented.

Institution web page: http://www.ncsu.edu/

January 29, 2002

SiGMA: The Simulator Guided Memory Analyzer

Luiz DeRose
ACTC-IBM Research

Email: laderose@us.ibm.com

Abstract:

Application developers have been facing new and more complex performance tuning and optimization problems as parallel architectures become more complex. Users must manage the intricate interactions of many hardware and software components in order to achieve high performance. The efficient utilization of the memory subsystem is one of the critical factors to obtain performance on current and future architectures. However, the lack of performance tools that could help users to understand the behavior of the memory hierarchy makes it even more difficult for programmers to tune and optimize their codes.

In this talk I will present the Simulator Guided Memory Analyzer (SiGMA), an application performance tool under development at the Advanced Computing Technology Center, at IBM Research, for identification of problems, bottlenecks, and inefficiencies in a program caused by the mapping of data structures into the memory hierarchy. SiGMA is being designed to become an autonomic tool that will provide information on how well a code uses the memory system and guide application changes to improve the memory system performance. In addition, SiGMA is being designed to help performance prediction in current and future architectures.

Institution web page: http://www.research.ibm.com/actc/

January 14, 2002

Terascale Simulation of Trade Cumulus Convection

David Stevens

Lawrence Livermore National Laboratory

Email: dstevens@llnl.gov

Abstract:

This talk presents three-dimensional numerical simulations of oceanic trade cumulus clouds underlying stratocumulus clouds. The simulations are based upon a case studied in a Global Energy and Water Experiment Cloud System Study (GCSS) model intercomparison that is loosely based on observed conditions during the Atlantic Trade Cumulus Experiment (ATEX). It is motivated by the importance of this cloud type to global cloud radiative forcing, and their role as a feeder system for deep convection in the tropics. This study focuses on the sensitivity of the modeled cloud field to the domain size and the grid spacing. Domain widths from 6.5 to 20 kilometers and horizontal grid spacings ranging from 10 to 80 meters, with corresponding vertical grid spacing ranging from 5 to 40 meters, are studied, involving massively parallel computations on up to 2.5 billion grid cells. The combination of large domain size and small grid resolution provides an unprecedented perspective on this type of convection.

The mean stratocumulus cloud fraction, optical depth, and vertical fluxes of heat, moisture, and momentum are found to be quite sensitive to both the domain size and the resolution. The sensitivities are associated with a strong feedback between cloud fraction, cloud-top radiative cooling, and entrainment. The properties of individual cumulus clouds rising into the stratocumulus are less affected than the stratocumulus clouds. The simulations with 80 m horizontal/40 m vertical resolution are clearly underresolved, with distinctly different distributions of liquid water within the clouds. Increasing the resolution to finer than 40 m horizontal/20 m vertical affects the inversion structure and entrainment processes somewhat, but has less impact on the structure of individual clouds. Large-domain simulations exhibit mesoscale structure in the cloud organization and liquid water path. This mesoscale variability feeds back on the domain-mean properties through the cloud-radiative feedback. These simulations suggest that very large computations are required to obtain meaningful cloud statistics for this case.

Speaker's web page: http://www.llnl.gov/CASC/people/stevens/

Institution web page: http://www.llnl.gov/

December 20, 2001

An Overview of the DataFoundry Project

Terence Critchlow

Lawrence Livermore National Laboratory

Email: critchlow1@llnl.gov

Abstract:

The DataFoundry project is an ongoing effort to help scientists better interact with their data. As a result of this broad goal, the project has evolved over time to focus on different problems being faced by Lab scientists. This talk will outline, in an approximately chronological order, the various research areas we have been involved with. I will begin with an overview the project's original focus of data warehousing and integration, then I will briefly discuss each of the three efforts we are currently working on: large-scale data access, query pattern analysis, and ad hoc queries over mesh data.

Speaker's web page: http://www.llnl.gov/CASC/people/critchlow/

Institution web page: http://www.llnl.gov/

December 14, 2001

Introduction to the Sheaf Data Management System

David Butler
Limit Point Systems Inc.

Email: dmbutler@limitpt.com

Abstract:

The sheaf system is under development for the ASCI project to create both a theoretical data model and a concrete representation for scientific data that addresses the complexity of ASCI simulation data structures. The sheaf data model is based on mathematical sheaf theory, which incorporates topology, lattice theory, and graph theory. The sheaf data management system implements this model in C++, providing both a concrete representation for ASCI data and a wide range of powerful operations for manipulating data.

We begin by describing at a high level the fundamental functionality of the sheaf data model and architecture of the sheaf software system. We briefly review its history and development. We introduce a table + graph metaphor to describe the mathematical basis of the system. We then use the metaphor to discuss as many examples as time allows. Examples include conventional mesh-related applications, such as unstructured meshes, mesh queries, adaptive meshes, domain decomposition and parallelism, as well as less conventional, graph and order-related applications, such as inheritance hierarchies, version control, and project plans.

Institution web page: http://www.limitpt.com/index2.html

December 12, 2001

Variational Grid Adaptation Based on the Modified Equation Error Estimator

Giovanni Lapenta
Los Alamos National Laboratory

Email: lapenta@lanl.gov

Abstract:

A new method for grid adaptation will be presented. The method uses a variational approach based on a rigorous definition of the error.

In variational grid adaptation, appropriate functionals are minimized using the Euler–Lagrange approach. Typically the functionals are based on a heuristic understanding of the physics of the problem. A classic example is the Winslow approach where a merit function is used to adapt the grid to regions of strong gradients.

Recently, considerable attention has been devoted to using more mathematically rigorous definitions of the numerical error to construct automatic grid adaptation methods.

In the present seminar, I will show that a new approach can be proposed to combine the best and most reliable features of classic variational methods, such as the Winslow method, with more modern error estimators.

The seminar will be organized in a theoretical part and a series of examples.

The theoretical section will show how a rigorous definition of the error based on the modified equations approach can be used to derive new functionals for variational adaptation. The theoretical derivation will then be shown to lend itself to a simple practical implementation that can be used directly in pre-existing variational adaptation tools with minimal changes.

Examples will be shown to highlight the main features of the new approach. Elliptic problems and systems of Hyperbolic equations (gas dynamics and plasma physics) will be solved using adaptive grids. Particularly, I will compare the classic Winslow approach with the new method showing the remarkable improvements that can be obtained at no extra cost.

Institution web page: http://www.lanl.gov/

December 11, 2001

Multiscale Molecular Computation

Achi Brandt
Weizmann Institute of Science

Email: achi@wisdom.weizmann.ac.il

Abstract:

A general approach will be presented for describing macromolecules and fluids at increasingly coarser scales, and for deriving corresponding coarse-level Hamiltonians, based either on a given fine-scale Hamiltonian or, more fundamentally, on quantum chemistry. The coarse levels can be used for orders-of-magnitude acceleration of equlibrium (and also dynamics) simulations, avoiding slow-downs and attraction-basin traps typical to single-level simulations. They may further lead to "macroscopic equations," enabling simulations of indefinitely large systems.

Development of highly efficient multigrid methods for the level of electronic structure calculations will also be briefly surveyed.

Speaker's web page: http://www.wisdom.weizmann.ac.il/

Institution web page: http://www.weizmann.ac.il/

December 6, 2001

Subsetting and Clustering of array Comparative Genomic Hybridization Data

Annette Molinaro
UC Berkeley

Mark van der Laan, Dan H. Moore II

Email: molinaro@cc.ucsf.edu

Abstract:

Microarray-based comparative genomic hybridization is a unique way to measure genomic DNA aberrations and map them directly to the genome sequence. We will discuss the methodology behind this process and a statistical approach to the data. In this approach we will compare various ways of minimizing the number of variables by subsetting and discuss the results of clustering to predict tumor type.

Speaker's web page: http://oz.Berkeley.EDU/~laan/Students/Students_subpages/Annette/annette.html

Research web page: http://oz.Berkeley.EDU/~laan/Research/research.html

December 3, 2001

Enhanced Accuracy by Post-processing for Hyperbolic Problems

Bernardo Cockburn
University of Minnesota

Email: cockburn@math.umn.edu

Abstract:

We consider the enhancement of accuracy, by means of a simple post-processing technique, for finite element approximations to transient hyperbolic equations. The post-processing is a convolution with a kernel whose support has measure of order one in the case of arbitrary unstructured meshes; if the mesh is locally translation invariant, the support of the kernel is a cube whose edges are of size of the order of Delta x only. For example, when polynomials of degree k are used in the discontinuous Galerkin (DG) method, and the exact solution is globally smooth, the DG method is of order k+(1/2) in the L2 norm, whereas the post-processed approximation is of order 2k+1; if the exact solution is in L2 only, in which case no order of convergence is available for the DG method, the post-processed approximation converges with order k+(1/2) in L2(Omega_0) where Omega_0 is a subdomain over which the exact solution is smooth. Numerical results displaying the sharpness of the estimates are presented.

Speaker's web page: http://www.math.umn.edu/~cockburn/

Research web page: http://www.math.umn.edu/~cockburn/Research.html

Institution web page: http://www.umn.edu/

November 30, 2001

Using Hardware Monitors to Measure the Cache Eviction Behavior of Application Data Structures

Bryan Buck
University of Maryland

Email: buck@cs.umd.edu

Abstract:

This talk will discuss a proposed hardware performance monitor that provides support not only for measuring cache misses and the addresses associated with them, but also for discovering what data is evicted from the cache at the time of a miss. We will describe how to use this hardware support to efficiently measure the cache eviction behavior of application data structures at the source code level. We will then present the results of a study in which we implemented this technique in simulation and examined the overhead, perturbation of results, and usefulness of collecting this information. This talk will also discuss current developments on the Dyninst API project at the University of Maryland, which provides a library for modifying the code of running applications by inserting or removing instrumentation, and how it could be used in work like that presented here.

Speaker's web page: http://www.cs.umd.edu/users/buck/

Research web page: http://www.dyninst.org/

Institution web page: http://umd.edu/

November 29, 2001

A Streaming Supercomputer

William Dally and
Patrick Hanrahan
Stanford University

Email: billd@csl.stanford.edu

Abstract:

Modern VLSI technology enables 100GFLOPS chips (e.g., NVidea), 20 cent/MByte memory, and chips with 1Tbit/s of pin bandwidth, all of the ingredients for powerful and cost-effective scientific computing. Building supercomputers by clustering workstations and SMPs, however, doesn't realize the potential of the technology, instead giving machines that cost more per GFLOPS, GByte, and GUPS than their low-end counterparts.

Stream architectures have the potential to improve the performance (GFLOPS and GUPS) per unit cost by two orders of magnitude. Casting an application as a stream program makes all communication explicit, allowing much of it to be kept local (on chip), and hides the latency of global communication when it is needed. This in turn enables architectures with very high arithmetic intensity, e.g., 100GFLOPS chips with 20GBytes/s of memory bandwidth.

Realizing the performance and cost advantages of a streaming supercomputer, of course, requires recoding applications in a streaming style. We envision a layered programming system that simplifies the process of coding applications in this style and makes the resulting stream applications easily portable across many platforms.

hanrahan@Graphics.Stanford.edu

Speakers' web pages: http://csl.Stanford.EDU/~billd/

http://www-graphics.stanford.edu/~hanrahan/

Research web pages: http://www.graphics.stanford.edu

http://cits.stanford.edu

Institution web page: http://www.stanford.edu/

November 27, 2001

Immersive Virtual Reality for Scientific Visualization: A Progress Report

Andries van Dam

Thomas J. Watson, Jr. Professor,
Brown University

Email: avd@cs.brown.edu

Abstract:

Immersive virtual reality (IVR) has the potential to be a powerful tool for the visualization of burgeoning scientific datasets and models. While IVR has been available for well over a decade, its use in scientific visualization is relatively new, and many challenges remain before IVR can become a standard tool for the working scientist. In this presentation we provide a progress report and sketch a research agenda for the technology underlying IVR for scientific visualization. Among the interesting problem areas are how to do computational steering for exploration, how to use art-inspired visualization techniques for multi-valued data, and how to construct interaction techniques and metaphors for pleasant and efficient control of the environment. To illustrate our approaches to these issues, we will present specific examples of work from our lab, including immersive visualizations of arterial blood flow and medical imaging.

Speaker's web page: http://www.cs.brown.edu/people/avd/

Research web page: http://www.cs.brown.edu/research/

Institution web page: http://www.cs.brown.edu/

November 19, 2001

Temporal and Spatial Level of Details for Dynamic Meshes

Ariel Shamir
The Interdisciplinary Center

Email: arik@idc.ac.il

Abstract:

Multiresolution techniques enhance the abilities of graphics and visual systems to overcome limitations in time, space, and transmission costs. Numerous techniques have been presented that concentrate on creating level of detail models for static meshes. Time-dependent deformable meshes impose even greater difficulties on such systems. In this talk we describe a solution for using level of details for time-dependent meshes. Our solution allows for both temporal and spatial level of details to be combined in an efficient manner. By separating low- and high-frequency temporal information, we gain the ability to create very fast coarse updates in the temporal dimension, which can be adaptively refined for greater details.

Institution web page: http://www.idc.al.il/

November 16, 2001

Stochastic Multiresolution Models for Turbulence

Brandon Whitcher
National Center for Atmospheric Research

Email: whitcher@cgd.ucar.edu

Abstract:

The efficient and accurate representation of two-dimensional turbulent fields is of interest in the geosciences because the fundamental equations that describe turbulence are difficult to handle directly. Rather than extract the coherent portion of the image from the background variation, as in the classical signal plus noise model, we present a statistical model for individual vortices using the non-decimated discrete wavelet transform (MODWT). Unlike the orthonormal two-dimensional discrete wavelet transform (2D DWT), the 2D MODWT produces a redundant non-orthogonal transform. The main reason for this discrepancy is that the 2D MODWT does not downsample after convolving the wavelet filters with the input field. Hence, each wavelet field will have the same dimension as the original field.

A template image, supplied by the user, provides the features we want to extract from the observed field. By transforming the vortex template into the wavelet domain specific characteristics present in the template, such as size and symmetry, are broken down into components associated with spatial frequencies. Multivariate multiple linear regression is used to fit the vortex template to the observed vorticity field in the wavelet domain. Framing the regression model in the wavelet domain, through the multiresolution analysis, a template function that is idealized and simplistic in nature provides the foundation of a semi-parametric fit.

Application to a vortex census algorithm that tracks quantities of interest (such as size, peak amplitude, circulation, etc.) as the vorticity field evolves is given. Extensions to three dimensions are also possible.

Speaker's web page: http://www.cgd.ucar.edu/~whitcher/

Research web page: http://www.cgd.ucar.edu/stats/

Institution web page: http://www.cgd.ucar.edu/

November 16, 2001

Modelling Nuclear Contamination In Fractally Porous Media

Jim Douglas
Purdue University

Email: douglas@math.purdue.edu

Abstract:

The objects of this paper are to formulate a model for the transport of a chain of radioactive waste products in a fractured porous medium, to devise an effective and efficient numerical method for approximating the solution of the model, to implement the numerical method, and to present the results of some experimental calculations.

The derivation of the model begins with the single-porosity model considered in the thesis of Spagnuolo and first extends the model to a double-porosity model. Then, taking the time scale of the problem into account, this model is reduced to an analogue of the "limit model" for immiscible flow in a porous medium consdiered by Douglas, Paes, Leme, and Hensley. The well-posedness of the model follows from the theoretical results. A locally conservative Eulerian–Lagrangian numerical method is applied to the resulting system of equations; the experimental results cover both smooth and fractal distributions of permeability in the fractured reservoir.

Speaker's web page: http://www.math.purdue.edu/~douglas/

Institution web page: http://www.purdue.edu/

November 6, 2001

Morse Complexes and Topological Persistence

Herbert Edelsbrunner
Duke University

Email: edels@cs.duke.edu

Abstract:

We consider Morse complexes decomposing a manifold with a smooth height function into regions that have the same gradient flow pattern. We use a combinatorial algorithm with numerical components to construct such a complex via handle slides. A hierarchy of progressively simpler Morse complexes is then constructed by cancelling critical points in pairs. These cancellations are performed in the order of increasing importance, or persistence of critical points.

Speaker's web page: http://www.cs.duke.edu/~edels/

Institution web page: http://www.duke.edu/

October 29, 2001

Advances in Database Technologies for Emerging Applications

Walid Aref
Purdue University

Email: aref@cs.purdue.edu

Abstract:

The wide-spread use of the Internet and the explosion of modern applications demand new database engine technologies that accommodate the fluctuation, unpredictability,and new data types in data-intensive distributed information sources. The target is to develop and implement a broad class of innovative database technologies to address the demands of these emerging applications. In this talk, I will focus on data-intensive applications that handle distributed, spatial and multimedia information sources. I will present a collection of extensible database tools that address the needs of these applications. These include: high-dimensional space-filling curves and their applications in multimedia data scheduling, space-partitioning trees and their applications in indexing non-traditional data types in commercially used database systems, and pipelinable rank join algorithms for answering multi-feature nearest-neighbor queries that are common in multimedia database systems.

Speaker's web page: http://www.cs.purdue.edu/people/aref/

Research web page: http://www.cs.purdue.edu/homes/seh/PLACE.html

Institution web page: http://www.purdue.edu/

Coupled Eulerian-Lagrangian Simulations of Detonation Induced Shock Response in Tantalum

Daniel Meiron
California Institute of Technology

Email: dim@its.caltech.edu

Abstract:

The ASCI Alliance Center at Caltech is constructing a virtual shock physics facility with the aim of facilitating fully three dimensional simulations of the interaction of strong shock waves initiated by the detonation of high explosives (HE) with solid targets. In this talk we will present an overview of the current capabilities of the software environment under development at the center. A parallel implementation of an algorithm to dynamically couple Eulerian CFD (used for simulation of HE detonation) with Lagrangian solid mechanics will be presented along with computations on the ASCI terascale platforms utilizing this approach. We will also present results from a multi-scale modeling effort being pursued at the center to provide improved sub-grid scale models for those physical length and time scales that cannot be captured directly by simulation. We discuss the application of this approach to the development of a reduced reaction network for nitramine explosives and the development of models for plastic response of metals such as Tantalum.

Speaker's web page: http://www.acm.caltech.edu/people/faculty/meiron-d.html

Institution web page: http://www.caltech.edu/

October 24, 2001

Models for a Vertical Draining Film with an Insoluble Surfactant

Richard Braun
University of Delaware

Email: braun@math.udel.edu

Abstract:

The drainage of a thin Newtonian film with an insoluble surfactant is studied theoretically. Lubrication theory is applied to the thin film, which is suspended vertically from a frame and drains into a bath. In 1+1 dimensions, three nonlinear PDEs govern the evolution of the film shape, surface velocity and surfactant concentration. The surface viscosity and the surface tension of the films may be nonlinear functions of the surface concentration; we will consider linearized models primarily in this talk. Spatial discretization of the PDEs gives a large system of differential algebraic equations (DAEs) that may solved using perturbation methods, but results are typically obtained computationally. Slow and fast draining limits for the film drainage can be reached and a Marangoni-driven wave that is localized may be observed.

A model in 2+1 dimensions will also be presented that consists of 4 nonlinear PDEs and also requires the solution of a large set of DAEs. The results reveal an instability that is caused by a competition between gravity and Marangoni forces. The results obtained so far will be compared with experiment.

This work is in collaboration with S. Naire (now at U of Nottingham, UK) and S.A. Snow (Dow Corning) and has been supported by NSF DMS-9631287, DMS-9722854 and Dow Corning.

Speaker's web page: http://www.math.udel.edu/~braun/

Institution web page: http://www.udel.edu/

October 18, 2001

The Rate of Corrections and its Application in Scientific Computing

Zhiqiang Cai
Purdue University

Email: cai3@llnl.gov

Abstract:

One of the major issues in numerical analysis/scientific computing is the accuracy of the underlying numerical methods. It is a common knowledge that the accuracy of all current numerical methods for differentiation, integration, ordinary and partial differential equations, etc., are limited by the highest derivative of the underlying approximated function. Since solutions of many partial differential equations are not smooth, low order methods with adaptive mesh refinements seem to be the method of choice for such problems without the prior knowledge of singular behavior. But they are expensive. Obviously, it is desirable to develop higher order accurate methods for non-smooth problems. This seems to be a paradox.

This talk presents an approach for computing higher order accurate approximations for differentiation, integration, ODEs, and PDEs when the underlying approximated functions are not smooth. The key idea of this work is the introduction of the rate of corrections that is of universality, that quantifies the accuracy of the numerical method used, and that is computationally feasible. The rate of corrections is then used to increase the accuracy of the approximation. Also, this approach can be applied to problems without continuum background, such as the sequence of period-doubling bifurcations in chaos.

Speaker's web page: http://www.math.purdue.edu/~zcai/

Institution web page: http://www.purdue.edu/

October 17, 2001

Open Topics in Parallel I/O

Rob Ross
Argonne National Laboratory

Email: rross@mcs.anl.gov

Abstract:

Topics will include: implementing parallel i/o systems; design or choice of request languages; is atomicity always a requirement; and metadata issues in file system performance. I'm hoping that this talk will spark discussions during the rest of the week of my visit to LLNL.

Speaker's web page: http://www-unix.mcs.anl.gov/~rross/resume.htm

Research web page: http://www.parl.clemson.edu/pvfs/

Institution web page: http://www.anl.gov/

October 15, 2001

Mining Large Image Datasets

Jelena Tesic
UC Santa Barbara

Bangalore Manjunath

Email: jelena@ece.ucsb.edu

Abstract:

Mining large image datasets places a number of challenging requirements on the analysis framework. Some initial success has been achieved with systems that represent images as an organized collection of summarized information obtained from the feature descriptors and spatial constraints. However, the high dimensionality of the feature spaces and the size of the image datasets make meaningful summarization a challenging problem. A visual thesaurus based on low-level image descriptors provides a scalable conceptual framework for analyzing perceptual events. The heart of this method is a learning system that gathers information by interacting with database users.

Our main objective is to find clusters that represent similar feature points located in a small subset of a feature space. High-dimensional spaces represent challenges for clustering, due to the sparseness of the space. However, clusters may be formed from a couple of visually different elements that inhabit a large part of a high-dimensional space. Co-occurrence of clusters in an image helps us distinguish visually meaningful representatives. We are currently conducting experiments on texture feature sets to determine the dependency of the clusters in the texture feature space on feature vectors and spatial image layout. The objective of the visual thesaurus is to classify the image regions into perceptually similar categories. Spatial Event Cubes (SECs), are used represent and analyze the spatial relationships. SECs are computed with respect to particular spatial relationships. Detailed analysis shows that SECs can be used for visualization, discovery of latent spatial configurations, and for constructing indices for efficient and meaningful data access.

Speaker's web page: http://vision.ece.ucsb.edu/~jelena/

Research web page: http://www-iplab.ece.ucsb.edu/

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

October 12, 2001

Optimization-Based Reference-Matrix Rezone Strategies for Arbitrary Lagrangian-Eulerian Methods on Unstructured Grids

Mikhail Shashkov
Los Alamos National Laboratory

Email: misha@t7.lanl.gov

Abstract:

The philosophy of the Arbitrary Lagrangian-Eulerian methodology (ALE; [1]) for solving multidimensional fluid flow problems is to move the computational grid, using the flow as a guide, to improve the accuracy and efficiency of the simulation. A principal element is the rezone phase in which a "rezoned" grid is created that is adapted to the fluid motion. Here we describe a general rezone strategy that ensures the continuing geometric quality of the computational grid, while keeping "rezoned" grid at each time step as close as possible to the Lagrangian grid. Although the methodology can be applied to more general grid types, we restrict ourselves here to unstructured triangular grids with fixed connectivity.

The unstructured grid is defined by positions of the nodes and connectivity. Between the node and each of its neighbors is an "edge-vector." Connectivity determines which edge-vectors form the columns of matrices (Jacobians) used to form the objective function minimized in the optimization.

The rezoning procedure consists of two phases: a sequence of local optimizations followed by a single global optimization. The local optimization leads to local "reference" Jacobians to be used in the global optimization. At each node we form a local patch from the Lagrangian grid and construct a local analog of a condition number functional [2]. Minimization of this functional with respect to position of central node defines its 'virtual' location (the node is not actually moved). By connecting this virtually-moved node to its (stationary) neighbors, we define "reference" edge-vectors and Jacobians that represent the best locally achievable geometric grid quality.

The "Rezoned" grid results from minimizing a global objective function, which measures the distance (in a least-squares sense) between the Jacobian of the "rezoned" grid and the "reference" Jacobian. This objective function includes a "barrier" that penalizes grids whose cells are close to being inverted. The global objective function is minimized by direct optimization. We provide numerical examples to demonstrate the robustness and the effectiveness of our methodology on model examples as well as for ALE calculations of hydrodynamics problems.

1. L. Margolin, Introduction to "Arbitrary Lagrangian-Eulerian Computing Method for All Flow Speeds", J. Comp. Phys., 135 (1997), pp. 198-202.

2. P. Knupp, Matrix Norms and the Condition Number - A General Framework to Improve Mesh Quality via Node-Movement, Proc. of the 8th Meshing Roundtable, S. Lake Tahoe CA, 1999, pp. 13-22.

Speaker's web page: http://cnls.lanl.gov/~shashkov

Institution web page: http://www.lanl.gov/

The ZSWEEP Algorithm for Rendering Irregular Grids

Claudio Silva
AT&T Labs-Research

Email: csilva@research.att.com

Abstract:

The need to visualize unstructured volumetric data arises in a broad spectrum of applications including structural dynamics, structural mechanics, thermodynamics, fluid mechanics, and shock physics. In this talk, I will focus on direct volume rendering, a term used to define a particular set of rendering techniques which avoids generating intermediary (surface) representations of the volume data. Instead, the scalar field is generally modeled as a cloud-like material, and rendered by computing a set of lighting equations.

I will focus on a particular technique called the ZSWEEP algorithm. The main idea of the ZSWEEP algorithm is very simple; it is based on sweeping the data with a plane parallel to the viewing plane in order of increasing depth (or "z"), "projecting" the faces of cells that are incident to vertices as they are encountered by the sweep plane. During ZSWEEP's face projection, we simply compute the intersection of the ray emanating from each pixel, and store their z-value, and other auxiliary information, in "sorted" order in a list of intersections for the given pixel, similar to the way an A-buffer operates. Compositing is performed as the "target Z" plane is reached. The efficiency arises from: (1) the fact that the algorithm exploits the implicit (approximate) global ordering that the z-ordering of the vertices induces on the cells that are incident on them, thus leading to only a very small number of ray intersection are done out of order; and (2) the use of early compositing which makes the memory footprint of the algorithm quite small.

In this talk, I will describe several features and (parallel, distributed, and out-of-core) extensions of the ZSWEEP algorithm. I will also discuss a technique for efficiently implementing a variant of ZSWEEP in hardware.

This work was done in collaboration with R. Farias (Mississippi State University), and J. Mitchell (State University of New York at Stony Brook).

Related papers can be obtained from the author's web page:

http://www.research.att.com/~csilva/papers

Speaker's web page: http://www.research.att.com/~csilva/

Institution web page: http://www.research.att.com/

October 11, 2001

Mimetic Finite Difference Methods for Diffusion Equations

Mikhail Shashkov
Los Alamos National laboratory

Email: misha@t7.lanl.gov

Abstract:

In this talk we review our ideas and present some new results related to construction of mimetic FDMs for solution of diffusion problems in strongly heterogeneous non-isotropic materials on nonorthogonal, nonsmooth, structured and unstructured computational grids. Development of high-quality finite-difference methods (FDMs) for diffusion equation is a part of a bigger effort to create a discrete analog of vector and tensor calculus, that can be used to accurately approximate continuum models for a wide range of physical processes. These FDMs preserve fundamental properties of the original continuum differential operators and allow the discrete approximations of partial differential equations (PDEs) to mimic critical properties including conservation laws and symmetries in the solution of the underlying physical problem.The discrete analogs of differential operators satisfy the identities and theorems of vector and tensor calculus and provide new reliable algorithms for a wide class of PDEs.

We describe discrete analogs of divergence and flux operators and discuss their properties. First, we describe basic (global) FDM, which uses these operators. Basic method uses cell-centered discretization for scalar unknowun and normal component of flux on each face. Next, we describe local modification of our method where we have addtional unknowns on the faces of the cell. Then, we present possible generalizations,which includes more general boundary conditions, unstructured grids, case of three dimensions, cylindrical coordinates and so on. Finally, we present some numerical examples.

Speaker's web page: http://cnls.lanl.gov/~shashkov

Institution web page: http://www.lanl.gov/

October 3, 2001

Using Visualization to Understand the Behavior of Computer Systems

Robert Bosch Jr.
Stanford University

Email: bosch@cs.stanford.edu

Abstract:

This talk introduces Rivet, a general-purpose environment for the development of computer systems visualizations. Rivet can be used for both real-time and post-mortem analyses of data from a wide variety of sources. The modular architecture of Rivet enables sophisticated visualizations to be assembled using simple building blocks representing the data, the visual representations, and the mappings between them. The implementation of Rivet enables the rapid prototyping of visualizations through a scripting language interface while still providing high-performance graphics and data management.

Rivet has been used for a wide variety of studies, ranging from processors and memory systems to multiprocessors and clusters to wide-area networks. The talk includes two examples. The first, a visualization displaying the behavior of superscalar processor pipelines, demonstrates the value of coordinated multi-view visualizations for understanding computer systems components. The second, a detailed performance analysis of the Argus parallel graphics library, shows how Rivet can be used in conjunction with comprehensive data sources such as the SimOS complete machine simulator to provide a powerful iterative analysis framework for understanding computer systems as a whole.

Speaker's web page: http://www-flash.stanford.edu/~bosch/

Research web page: http://graphics.stanford.edu/projects/rivet/

Institution web page: http://www.stanford.edu/

October 1, 2001

Application of Probe-Based Storage to High Performance Computing

Zachary Petersen
UC Santa Cruz

Darrell Long

Email: zachary@cse.ucsc.edu

Abstract:

Although there is continual improvement in speed and capacity of storage systems, traditional longitudinal magnetic recording is approaching a hard physical limit. A performance gap between the speed and capacity of RAM and disk is increasing at a rate of 50% a year. This performance gap is especially visible in high-performance workloads such as those employed in NNSA's Advanced Simulation and Computing Program (ASCI). We in the Computer Systems Research Group at the University of California, Santa Cruz are investigating two lines of new and innovative high-performance storage research.

The first involves an exciting new storage technology based on MicroElectroMechanical Systems (MEMS) which promises a significant increase in performance, capacity, and reliability relative to modern storage devices. The second, approaches the performance problems of current storage devices, applying novel ideas of grouping like data to improve read performance, while eliminating the amount of rewriting needed to be done to keep these data contiguous. Using sanitized workload traces from LLNL and support from the ISCR program, we have shown both lines of research successful in addressing the I/O problems posed by ASCI applications.

Speaker's web page: http://www.cse.ucsc.edu/~zachary/

Research web page: http://csl.cse.ucsc.edu/

Institution web page: http://www.ucsc.edu/public/

Institute for Scientific Computing Research

University Collaborative Research Program Subcontract Research Summaries

Multiscale Simulation Combining Direct Simulation Monte Carlo and Navier–Stokes Solvers

Berni J. Alder, PI
UC Davis

Yihao Zheng, RA
UC Davis

Alejandro L. Garcia, co-PI
San Jose State University

Andrew Wissink
LLNL, Collaborator

Summary

Numerical modeling of fluids is challenging when the physics of interest spans length scales differing by orders of magnitude. In the computation of hydrodynamic flows, structured adaptive mesh refinement (AMR) is used to efficiently perform calculations whose length scales span several orders of magnitude. However, when the finest mesh approaches the molecular scale, the macroscopic partial differential equations of fluid mechanics themselves are no longer valid. Adaptive mesh and algorithm refinement (AMAR) has been developed to solve these types of multiscale/multiphysics problems. In an AMAR calculation, at the particle scale a molecular algorithm such as the direct simulation Monte Carlo method (DSMC) is used. The objective of the current period is to develop for the continuum method a Navier–Stokes solver for integration with the DSMC method in the AMAR context.

An AMAR code has been developed at LLNL using Structured Adaptive Mesh Refinement Application Infrastructure (SAMRAI). The present code solves the inviscid Euler equations of gas dynamics at multiple (macroscopic) levels of refinement with the DSMC particle method.

We have replaced an existing SAMRAI inviscid Euler solver with a viscid Navier–Stokes solver and integrated the new solver into the existing SAMRAI scheme. The Navier–Stokes solver is a second-order Godunov-type explicit solver and has to date been tested under full multiregion, multilevel conditions for the following cases:

- one-dimensional acceleration-driven Poiseuille flow and
- two-dimensional pressure-driven Poiseuille flow.

The results agree well with a comparably refined stand-alone one-region one-level Navier–Stokes solver.

Our future plan is to apply the hybrid code combining the DSMC and the SAMRAI Navier–Stokes solver to study various classical hydrodynamic instability problems, such as flow in a pipe, with this refined computational tool.

Scalable Algebraic Domain Decomposition Preconditioners

Randolph E. Bank, PI
UC San Diego

Kathy Lu, RA
UC San Diego

Charles H. Tong and Panayot S. Vassilevski
LLNL Collaborators

Summary

We study the domain decomposition approach to the parallel solutions of partial differential equations with the novel feature that the subproblem residing on each processor is defined over the entire domain, although coarsely refined outside of the processor associated with the subdomain. This feature ensures that a global coarse description of the problem is contained within each of the subproblems. The advantages of this approach are that interprocessor communications are minimized while optimal order of convergence rates is preserved, and the convergence rate of the local subdomain solves can be optimized using the best existing sequential algebraic solvers.

This procedure is similar in philosophy to the parallel adaptive mesh refinement paradigm introduced by Bank and Holst, except that the present project deals with an algebraic version of the Bank-Holst paradigm in the sense that, instead of mesh refinement here we coarsen the degrees of freedom (or matrix) outside the prescribed subdomain. Thus, the domain decomposition method applied in each processor involves the local subdomain plus a small coarse space defined on the whole domain outside the processor, and each solve utilizes only a single interprocessor communication to retrieve the global vector. This approach can be applied to general sparse matrices, although matrices arising from discretization of partial differential equations are the principal target.

Besides the original global matrix representing the problem on the fine grid, the algorithm introduces a set of rectangular matrices (prolongation operators) and a set of coarse matrices that are constructed and stored in sparse parallel matrix format. A unique feature in our coarsening algorithm is that, even though coarsening is performed uniformly on the entire domain in parallel, to each processor coarsening occurs exclusively on the parts of the global linear system associated with other processors, and the matrix and right hand side for the local subdomain are preserved in the local subproblem construction.

We have focused mainly on developing and implementing a specific algebraic coarsening procedure that plays a central role in our overall algorithm. The software package, called FocusDD, consists of a parallel preconditioning algorithm built on top of the *hypre* library, developed at LLNL. As a by product of the Focus solver, an enriched set of functions to handle the *hypre* ParCSRMatrix and CSRMatrix data classes have been written and debugged. We extensively reuse existing subroutines in *hypre* for coarsening. For solving local subproblems, we use the SuperLU direct solver library. A detailed description of FocusDD is now available in reference manual form.

Since we now have an overall working code, our attention has turned to improving specific parts of the algorithm. In particular, we have been studying the linear systems arising from 5-point finite difference discretization of the Poisson equation on a rectangular domain. For this problem, we have made a matlab implementation of our

Randolph E. Bank, PI
UC San Diego

Kathy Lu, RA
UC San Diego

Charles H. Tong and
Panayot S. Vassilevski
LLNL Collaborators

Summary (continued)

algorithm, substituting a non-algebraic procedure for computing the prolongation matrix. Since this a very structured situation, we are able to easily construct a simple, structured prolongation operator. In the matlab code, we allow only one level of coarsening. We ran some numerical experiments and verified that the rate of convergence was less than 0.1, independent of both the size of the problem and the number of subdomains. We do not know if this prolongation operator is optimal in any sense, but the main point of the exercise was to demonstrate "proof of concept" in a simple setting. The FocusDD solver presently does not exhibit such ideal behavior for this model. Thus, we are presently examining the parallel coarsening procedure to determine what details in our implementation are responsible, and hopefully improve the overall performance.

In the longer term, our goal is to implement a more sophisticated coarsening procedure. In the current implementation, parts of the matrix on each processor are either fine or uniformly coarse. We have concluded from our initial analysis that it is useful to have graded coarsening. In the geometric sense, it means that near the boundary of the fine region, the coarsening should be modest. As one moves further from the fine region, the coarsening can be more generous. Thus, a geometric reduction of number of unknowns with distance from the subdomain seems appropriate. The overall order of the locally coarsened matrix A could be unchanged; it is its structure that will be modified. Our initial approach to achieve graded coarsening involves two different strategies. First, we will utilize our present parallel coarsening strategy (and software) to achieve a modest uniform coarsening. Each processor will then independently compute further coarsening to achieve the appropriate grading. While the algorithms used to achieve this grading will be sequential on each processor, all processors can do this step independently in parallel. The computation will be organized such that existing sequential graph coarsening software can be utilized, and less inter-processor communication will be required compared to parallel AMG in *hypre* library.

Our ultimate goal is to fine-tune our parallel code and benchmark its performance on selected applications of interest to LLNL. Our target computing platforms are massively parallel machines, including those at both UCSD and LLNL.

Similarity Query Processing for Quadratic Queries and Relevance Feedback

B. S. Manjunath and
S. Chandrasekaran, co-PIs
UC Santa Barbara

Helena Tesic
RA, UC Santa Barbara

Imola K. Fodor
LLNL Collaborator

Summary

The main objective of this research is to investigate user interactions with large multimedia databases in the context of direct access to and mining of the data. This work supports development of scalable algorithms for the interactive exploration of large, complex, multi-dimensional scientific data. The exponential growth of such multi-media and multi-sensor data in consumer as well as scientific applications pose many interesting and task critical challenges. There are several inter-related issues in the management of such data, including feature extraction, similarity-based search, high dimensional indexing, scalability to large data sets, and personalizing search and retrieval.

As technology advances and more visual data are available, we need more effective systems to handle the image data processing and user interaction. The framework must efficiently summarize information contained in the image data; it must provide scalability with respect to the nature, size and dimension of a dataset; and it must offer simple representations of the results and relationships discovered in the dataset. In particular, the database support to carry on the learning process in large image databases is an important issue that has been largely ignored. Our research focuses on similarity search and indexing of a real image/video databases that use relevance feedback mechanisms to improve retrieval results. We also introduce an efficient scheme that preserves the important data characteristics for data summarization.

Adaptive Nearest Neighbor Search in Relevance Feedback. We introduce the problem of repetitive nearest neighbor search in relevance feedback and propose an efficient search scheme for high dimensional feature spaces. Relevance feedback learning is a popular scheme used in content-based image and video retrieval to support high-level concept queries. This work addresses those scenarios in which a similarity or distance matrix is updated during each iteration of the relevance feedback search and a new set of nearest neighbors are computed. This repetitive nearest neighbor computation in high dimensional feature spaces is expensive, particularly when the number of items in the data set is large. We propose a search algorithm that supports relevance feedback for the general quadratic distance metric. The scheme exploits correlations between two consecutive nearest neighbor sets, thus significantly reducing the overall search complexity.

Dimensionality Reduction in Gabor Texture Features. Texture has been recognized as an important visual primitive in image analysis. A widely used texture descriptor is that computed us-ing multiscale Gabor filters. The high dimensionality and computational complexity of this descriptor adversely affect the efficiency of content-based retrieval systems. Our work shows how the dimensionality and complexity of the descriptor can be significantly reduced, while retaining comparable performance. This gain is based on a claim that the absolute values of the filter outputs follow a Rayleigh distribution. Ex-perimental results show that the dimensionality can be reduced by almost 50%,

B. S. Manjunath and
S. Chandrasekaran, co-PIs
UC Santa Barbara

Helena Tesic
RA, UC Santa Barbara

Imola K. Fodor
LLNL Collaborator

Summary (continued)

with a tradeoff of less than 2% on the error rate. We also propose a new normalization method that improves similarity retrieval and reduces indexing overhead. Details are discussed.

Data summarization and indexing for efficient similarity retrieval. This work introduces a conceptual representation for complex spatial arrangements of image features in large multimedia datasets. A data structure, termed the Spatial Event Cube (SEC), is formed from the co-occurrence matrices of perceptually classified features with respect to specific spatial relationships. A visual thesaurus constructed using supervised and unsupervised learning techniques is used to label the image features. SECs can be used to not only visualize the dominant spatial arrangements of feature classes but also to discover non-obvious configurations. Given a large collection of images with similar content—over space or time— associations are found among the image regions that visually characterize the dataset. These associations subsequently aid in the discovery of connections between image structure and semantic events, such as spatial rules. The emphasis is to mine the image content using little or no domain knowledge. An inclusive association rule algorithm is proposed as an extension of traditional association rules for application to image datasets. The definitions of standard transaction and association mining rules are extended for image datasets. Experiments show the approach has a wide range of application in image and video datasets.

This research has been presented at CASC seminars and at the Scientific Data Mining workshop held at IPAM in 2002 and has been documented in numerous refereed papers.

Simulation of Compressible Reacting Flows

Sutanu Sarkar, PI
UC San Diego

David Garrido Lopez, RA
UC San Diego

Andrew W. Cook
LLNL Collaborator

Summary

A fundamental problem of interest to NIF capsule design is the behavior of a burn front when it propagates through a deuterium/tritium (DT) mix contaminated by inert shell material introduced by Rayleigh-Taylor instabilities. In order to accurately describe burn front propagation, we have developed a scheme that is high-order on arbitrarily nonuniform grids. Nonuniform stretched grids are needed because the burn front, being much thinner than other flow features, needs a much smaller grid step for resolution of the transport and reactive processes that determine the burn velocity. We have simulated two-dimensional burn propagations into a contaminated mix.

Simulations on a nonuniform grid in physical space are often performed by using a computational grid that is uniform, and incorporating the Jacobian of the transformation between physical and computational grids into the governing equations. We refer to such an approach as Method 1. Although simple in principle, Method 1 does not ensure that the formal accuracy of the derivative in the case of uniform grid step is maintained on a nonuniform grid. The actual accuracy can be as low as first-order depending on the grid stretching as well as the value of the local derivative. In turbulent combustion, where sharp changes of the field variables occur locally, it is of interest to use a numerical scheme that has a defined high order of accuracy of the spatial derivative independent of the solution. One such approach, called Method 2, is to write an expression for the spatial derivative that, by construction, has the required spatial accuracy on an arbitrary nonuniform grid. We have demonstrated the advantage of Method 2 by evaluating the derivatives of known smoothly differentiable functions under both methods. Versions of Method 2 that are second, fourth and sixth-order accurate on arbitrary grids show the expected order of accuracy as resolution increases until roundoff error associated with very small grid spacing begins to be important. On the other hand, applications of Method 1 with derivatives that are fourth and sixth-order accurate on a uniform grid do not maintain high-order accuracy on the nonuniform grid. Therefore, Method 2 is used for the evaluation of spatial derivatives in simulations of the burn front.

We simulate burn front propagation in a reactant mix contaminated by a multi-scale field of inert fluctuations. Both Arrhenius chemistry and DT chemistry have been explored. The flame zone is thicker in the latter case; however, the chief results regarding the burn propagation are similar. The unsteady, three-dimensional form of the compressible equations for a reactive mix-ture of fluids is solved. A nonuniform grid is used in the direction of flame propagation, with clustering of points in the burn region. Derivatives in this direction are computed using a 6th-order compact scheme valid for arbitrary nonuniform grids. The code is parallelized using MPI and runs on the unclassified IBM SP machine at LLNL using up to 896×1024 grid points.

Sutanu Sarkar, PI
UC San Diego

David Garrido Lopez, RA
UC San Diego

Andrew W. Cook
LLNL Collaborator

Summary (continued)

In the zones contaminated by inert, the temperature rise due to the burn energy release is smaller than average leading to a lower reaction rate and a local burn velocity that is smaller than average. Thus, the nonuniform inert introduced into the reactant mix by flow instabilities leads to distortions of the burn front. An additional mechanism for distorting the burn front that is operative even in an uncontaminated mix is the so-called Darrieus-Landau (DL) instability. Linear stability analysis gives the growth rate of infinitesimal perturbations of an initially planar burn front. The value of an associated with the basic hydrodynamic instability increases with wavenumber but, due to the stabilizing thermal effect associated with heat diffusion, the growth rate eventually decreases at large wavenumbers. Thus, there is a preferred wavelength of maximal growth of the DL instability. At any given time, an 'average' burn velocity based on the overall reactant consumption rate can be defined.

The research was presented in the annual meeting of the APS Division of Fluid Dynamics and at a seminar in ISCR.

Probabilistic Clustering of Dynamic Trajectories for Scientific Data Mining

Padhraic Smyth
PI, UC Irvine

Scott Gaffney, RA
UC Irvine

Chandrika Kamath
LLNL Collaborator

Summary

Data-driven exploration of massive spatio-temporal data sets is an area where there is particular need of new data mining and data analysis techniques. Analysis of spatio-temporal data is inherently challenging, yet most current research in data mining is focused on algorithms based on more traditional feature-vector data representations. In this project we have developed a set of flexible and robust algorithms as well as software for tracking and clustering time-trajectories of coherent structures in spatio-temporal grid data. These algorithms and software provide a basic set of data analysis tools for exploration and modeling of dynamic objects, in a manner analogous to the much more widely-available techniques for clustering of multivariate vector data (e.g., k-means, Gaussian mixtures, hierarchical clustering, etc).

In our second year we have focused on extending our earlier probabilistic mixture-based modeling of trajectories of objects to a more general scheme known as "random effects" models. We have described a variety of algorithms for "trajectory clustering," including methods based on polynomial curve models and splines. The use of random effects models (also known as Bayesian hierarchical models) is shown to increase the predictive power of these clustering algorithms, by allowing each trajectory to have its own parameters (influenced by a population prior), which in effect leads to clustering in parameter space.

Clustering is typically used as a tool for understanding and exploring large data sets. Most clustering algorithms operate on so-called feature vectors of fixed dimension. In contrast to this we address the problem of clustering sets of variable-length curve or trajectory data generated by groups of objects or individuals. The curves are sequences of observations measured over time (or some notion of time), functionally dependent on an independent variable or set of variables. Typically, the independent variable is time, but in general it can be any number of variables measured over the same interval as the observations. It is assumed that the trajectory data is generated by groups of objects or individuals such as humans, animals, organizations, natural phenomena, etc. Unlike fixed-dimensional feature vectors, trajectories have variable lengths and can be observed at different measurement intervals as well as contain missing observations.

This type of data is quite common due in part to large-scale data collection in the scientific community. Such data cannot be clustered with any of the standard vector-based clus-tering algorithms absent some *ad hoc* preprocessing procedure to reduce the data to a set of fixed-dimensional feature vectors. Our focus is to model trajectory data directly using a set of model-based curve-clustering algorithms for which we use the all-encompassing name "mixture of regressions" or "regression mixtures."

For this project we have developed and extended several different types of regression mixture models to deal with curve data sets. All our algorithms explicitly model the dependence of trajectories on the independent set and all use an Expec-tation Maximization (EM) procedure to perform the clustering.

Padhraic Smyth
PI, UC Irvine

Scott Gaffney, RA
UC Irvine

Chandrika Kamath
LLNL Collaborator

Summary (continued)

We have applied of this approach to the problem of tracking and clustering of extra-tropical cyclones (ETCs) in the North Atlantic. Understanding ETC trajectories is scientifically important for understanding both the short-term and long-term dynamics of atmospheric processes. Atmospheric scientists are interested in the spatio-temporal patterns of evolution of ETCs for a number of reasons. For example, it is not well-understood how long-term climate changes (such as global warming) may influence ETC frequency, strength, occurrence and spatial distribution. Similarly, changes in ETC patterns may provide clues of long-term changes in the climatic processes that drive ETCs. The links between ETCs and local weather phenomena are also of interest: clearly ETCs have significant influence on local precipitation, and in this context, a better understanding of their dynamics could provide better forecasting techniques both on local and seasonal time-scales. Furthermore. an explicit model of ETC evolution can serve as an intermediate link between global atmospheric phenomena (e.g., geopotential height patterns and regimes) and local “weather-related” phenomena such as precipitation, addressing the long-standing problem in atmospheric science known as “downscaling” (i.e., how to link and model global-scale arid local-scale phenomena in a coherent manner). The data sets that we are working with are widely used simulation data sets known as general circulation model data. We have demonstrated that our clustering methodology, when applied to ETC data, reveals useful and interpretable information about ETC patterns and behavior.

Several software products have been produced during this project. Initially, all our code was prototyped in Matlab and later moved to a more efficient language. In more recent work we have designed, implemented, and tested a general-purpose curve clustering toolkit that takes any type of curve data and produces the desired clustering. The code is written in C and C++ and is highly optimized and integrated with the C-LAPACK mathematical library.

We have made three visits to LLNL since the start of thisp roject, and the doctoral student on the project has presented two seminars on this work to LLNL personnel. Several conference and journal papers document new methodological and scientific results under this project.

Statistical Inference from Microarray Data

Mark van der Laan, PI
UC Berkeley

Annette Molinaro-Clark, RA
UC Berkeley

Dan Moore
LLNL Collaborator

Summary

The overall goal of this collaboration is to link gene microarrays to disease progression for eventual applications in treatment decisions and disease prevention. This goal includes the following components: to assess the reproducibility and reli-ability of numerous summary measures on gene microarrays; to determine reliable methods for classifying tumors based on their genetic profile; to link the classified tumors to covariates and clinical outcomes; to link the gene microarrays to covariates; and to link these gene microarravs and covariates to clinical outcomes.

This year was comprised of work on three projects: utilizing genetic profiles to classify tumors; model selection when linking covariates to clinical outcomes; and prediction of survival when linking covariates to clinical outcomes.

The first project included a dataset containing 41 renal tumors measured on CGH arrays of 90 BACs with three sub-types of cancer and two benign groupings. The goal of this study was to determine whether loci showing changes can be used to discriminate between the sub-types of cancer and between the benign and non-benign groupings. This dataset allowed us to explore: methods of simultaneous inference; prediction of sub-type by gain and loss of chromosomal region; identification of homogeneous groups of tumors by using hierarchical clustering routines with varying distant metrics; clustering BACs; and clustering tumors within clusters of BACS.

The second ongoing project involves model selection when linking gene microar-rays and covariates to clinical outcomes. In this project we are addressing how to choose bivariate and trivariate models given hundreds or thousands of independent predictors. We are using simulated data and a clinical dataset for this project. For both, we observe data on time to recurrence, follow-up time, and covariates. Due to the possibility of multiple outcomes of interest we are implementing counting processes and mul-tiplicative intensity models. A special case of this is a counting process that is equivalent to the Cox Proportional Hazards model.

Our approach entails bootstrapping the datasets m times and subsequently split-ting each bootstrap into a test and validation set. The training set is used to estimate the coefficients from the multiplicative intensity model. In the bivariate models this is done four times for each combination of covariates. Once the coefficients are attained, the test set is used to estimate the partial log-likelihood. The mean of the m partial log-likelihoods is calculated for each model within each combination of variables. In order to choose those combinations that are most related to the clinical outcomes we are presently investigating different criteria to use on the ranked means.

The third project involves prediction of survival, or other clinical outcomes, by hundreds or thousands of covariates. We are interested in nonparametric regression methods including classification and regression trees, CART, which are conducive both to

Mark van der Laan, PI
UC Berkeley

Annette Molinaro-Clark, RA
UC Berkeley

Dan Moore
LLNL Collaborator

Summary (continued)

censored and complete (i.e., uncensored) data. In the past, numerous ap-proaches to adapt CART specifically to censored survival data have been suggested. However, these methods are strictly restricted to censored data and do not propose performance assessments of predictors given a user supplied loss function. In this project we are building predictors with CART, functional in both the censored and uncensored setting, and assessing their performance.

The algorithm for CART relies on a split criterion for splitting a node into two nodes for continuous (uncensored) outcomes. We have modified this criterion with an inverse probability of treatment weighted (IPTW) estimator, which allows for informative censoring and a gain in efficiency. After building the candidate pre-dictor using CART and the modified split criterion on a training set, we evaluate the performance of this predictor by estimating its conditional risk with the IPTW estimator of the risk based on a validation set. Our goal is to select a predictor, which has a risk approximating the optimal risk, i.e., the truth.

Institute for Scientific Computing Research

ISCR Subcontract Research Summaries

Self-correcting Algebraic Multigrid

Timothy P. Chartier,
University of Washington

Summary

We report on the research and development of spectral AMGe (sAMGe), including:

1) enhancing the operator dependent decision–making process in the algebraic agglomeration routine developed by LLNL's Jim Jones and Panayot Vassilevski,

2) researching less computationally costly algorithms based on sAMGe that do not demand the use of spectral decomposition to determine smoothness, and

3) generating code to facilitate numerical experiments that would compare the efficacy of the sAMGe algorithms.

We are interested in element-free AMGe which does not demand access to the element matrices. Instead, the algorithm requires spatial coordinates for each degree of freedom. Then, polynomials would serve as initial guesses to globally smooth vectors. After a few relaxation steps, such polynomials would satisfy boundary conditions and become estimates for globally smooth error in sAMGe. Though requiring certain assumptions on smoothness, such an algorithm could harness sAMGe's ability to choose coarse grids and interpolation simultaneously while at the same time considerably reducing the computational cost of the algorithm.

The method produced impressive 2-level multigrid results for scalar elliptic problems. However, the multilevel tests were much less satisfying. Chow and Chartier continued discussions on this area of research throughout the summer and results of this work led to advancements in self-correcting methods detailed below. Still, further investigation is needed in order to determine if changes in the algorithm might produce multilevel results similar to those in the 2-level tests.

We began researching algorithmic development of self-correcting algebraic multigrid methods that use the sAMGe framework. Such methods automatically and adaptively correct their multigrid components to accelerate convergence.

As with sAMGe, such algorithms generalize the sense of smoothness. Initially, the work concentrated on adaptively correcting the columns of interpolation needed for coarse-grid correction in multigrid cycling.

Chow's research led to an alternate self-correcting algebraic multigrid scheme that creates additional V-cycles to dampen unresolved error components. In this way, convergence is accelerated adaptively within the method. We implemented this idea within the sAMGe framework and obtained excellent 2-level results possessing invariance of

Timothy P. Chartier,
University of Washington

Summary (continued)

convergence factors to diagonal scaling of the discrete operator – an important feature of self-correcting methods that is not found in most AMG methods. We also completed an initial multilevel version of the algorithm. The preliminary results look promising as to the efficacy of the method. Further enhancement of the algorithm should further improve these results. Testing the self-correcting ideas on systems of partial differential operators is an important step in the research and should be the next focus of the work.

On the matter of conducting numerical experiments to compare the efficacy of the sAMGe algorithms, Chartier completed code that reformats input data for the sAMGe code developed for his dissertation into a format suitable for the sAMGe code currently utilized at LLNL. This work will facilitate needed testing between the two sAMGe methods on a variety of problems.

sAMGe is a notable development in algebraic multigrid research. The promising results from the work of this subcontract bode well for the ability of sAMGe research to widen the problems efficiently solved by algebraic multigrid methods. Continued research is needed; yet important advances have been made as a part of this research.

Morse Complexes and Visualization of Complex Topology Phenomena

Herbert Edelsbrunner
Duke University

Summary

My LLNL sabbatical was divided between pursuing research on Morse complexes and learning about specific application problems at the laboratory. The research on Morse complexes and related topics progressed nicely. A non-exhaustive list of topics researched is

- The loop structure in Reeb graphs
- A simplification hierarchy of Reeb graphs
- An algorithm for hierarchical 3D Morse-Smale complexes
- The combination of topological and geometric simplification
- Jacobian sub manifolds of multiple Morse functions
- A definition of the similarity between continuous functions

We worked on all these problems as a group, always including Valerio Pascucci, Kree Cole-McLaughlin, Vijay Natarajan, and myself. For a week during my sabbatical, we regularly met with John Harer, mathematician (and vice-provost) at Duke University, working on the same set of problems. Topological and geometric simplification work also involved Timo Bremer, student at UC Davis. To different degrees, the projects are on going and require further mathematical work, the development of algorithms, the design of visualization tools, and implementation work.

I also spent time with other researchers at LLNL discussing application problems for which we feel that our theoretical work could lead to new insights and developments. A non-exhaustive list of applications in this category is

- The study of breaking copper wires on an atomic level
- The behavior of gas under extreme pressure
- The state transition problem in computational chemistry
- The simulation of the heat chamber in atomic fusion
- The mixing of two liquids after a passing shock wave

In each case, we observe rather complicated topological phenomena developing over time. We made steps towards developing some of our ideas into software and to use one or two application problems in the above list as a case study to validate the tool, but primarily to serve the application research.

In summary, my sabbatical at LLNL was extremely productive for my own research. The visit has got the collaboration between Valerio Pascucci's and my own group off the ground, and I look forward to a continuation of that collaboration in the near and further future.

Enabling Communication between Components for Scientific Computing

Professor Rod Fatoohi
San Jose State University

Summary

There has been a growing interest within the scientific computing community in component technologies. These technologies have shown their usefulness in the business community with examples like Microsoft's Component Object Model (COM), Sun's Java Beans, and OMG's CORBA Component Model (CCM). The Component Project at LLNL has developed a software tool called Babel that enables language interoperability among a variety of scientific programming languages, including Fortran, C, C++, Java, and Python. Babel enables the creation and distribution of software libraries specified by a Scientific Interface Definition Language (SIDL), which is similar to CORBA and COM IDLs but it focuses on the needs of scientific computing. Babel generates stubs and skeleton code, as in CORBA and others, from SIDL interfaces, but no communication code. Here we are developing the communication mechanism to enable invocation across multiple memory address spaces. Several communication protocols are considered including: ONC RPC, BSD Sockets, CORBA IIOP, and SOAP.

The goal of this project is to enable Babel users to make remote calls seamlessly, regardless to the location of the requested component. We are developing an abstract sub-layer underneath the Babel generated code to interface with multiple communication protocols. This sub-layer, similar to the abstract protocol layer of the Inter-Language Unification (ILU) system by Xerox, defines message types and parameter types between Babel runtime system and the communication protocols. The application programming interface for these protocols will be integrated with the Babel tool so that calls to them will be automatically generated by Babel.

Different communication protocols are considered including ONC RPC, BSD Sockets, CORBA IIOP, and SOAP. We have already specified how Babel data types are mapped into the ONC RPC wire protocol: eXternal Data Representation (XDR). Our next step is to integrate our interface with the Intermediate Object Representation (IOR) of Babel. We are also leveraging from the research activity at Indiana University in this area. We are aiming at implementing all four protocols (listed above), and give the user the option to select any specific protocol based on several criteria including performance, reliability, flexibility, scalability, and maintenance cost.

Scalable Symmetric Eigenvalue Solver using *hypre*

Andrew Knyazev
University of Colorado, Denver

Summary

The goal of this project was to develop Scalable Symmetric Eigenvalue Solver algorithm sofware for the solution of partial eigenvalue problems for large, sparse symmetric matrices on massively parallel computers, taking advantage of advances in the Scalable Linear Solvers project, in particular in multigrid technology and in incomplete factorizations (ILU) developed under the *hypre* project at the LLNL. This implementation allows the utilization of *hypre* preconditioners for symmetric eigenvalue problems, thus significantly extending their capability. LLNL plans to include this capability in the next *hypre* release.

The base iterative algorithm that has been implemented is the PI's locally optimal block preconditioned conjugate gradient (LOBPCG) method. The LOBPCG solver finds one or more of the smallest eigenvalues and eigenvectors of a positive definite matrix. The code is written in C and uses *hypre* and LAPACK/BLAS libraries. The user interface to the solver is implemented using *hypre* style object-oriented function calls. The matrix-vector multiply and the preconditioned solve are done through user supplied functions, which provide significant flexibility. The implementation illustrates that this matrix-free algorithm can be implemented successfully and efficiently using parallel libraries.

An LOBPCG driver was developed as part of the package, which exercises the software, and provides an implementation example of how to use the LOBPCG hypre user interface. It can be easily run after compilation on various platforms to test different configurations. This driver is patterned after a hypre driver that tests IJ_linear_solvers. It optionally generates several Laplacians and optionally reads a matrix from a matrix market file. The following preconditioners are available and have been tested: AMG–PCG; DS–PCG; ParaSails–PCG; Schwarz–PCG (default); and Euclid–PCG. Partition of processors is determined by user input consisting of an array of parallel vectors. The testing was done on a variety of platforms using up to 60 processors.

LOBPCG has been mainly developed and tested on the Beowulf cluster at the Colorado University at Denver. This system includes 36 nodes, 2 processors per node, PIII 933MHz processors, 2GB memory per node running Linux Redhat, and a 7.2SCI Dolpin interconnect. LOBPCG has been compiled and tested on a subset of the OCF Production Systems at LLNL as well.

Source code and documentation were delivered to LLNL and a seminar on their use presentated.

Non-conforming Finite Elements, Mesh Generation, Adaptivity and Related Algebraic Multigrid and Domain Decomposition Methods in Massively Parallel Computing Environment

Raytcho Lazarov and Joseph Pasciak
Texas A & M University

Summary

Construction, analysis, and numerical testing of efficient solution techniques for solving elliptic PDEs that allow for parallel implementation have been the focus of the research. A number of discretization and solution methods for solving second order elliptic problems that include mortar and penalty approximations and domain decomposition methods for finite elements and finite volumes have been investigated and analyzed. Techniques for parallel domain decomposition algorithms in the framework of PETSc and HYPRE have been studied and tested. Hierarchical parallel grid refinement and adaptive solution methods have been implemented and tested on various model problems. A parallel code implementing the mortar method with algebraically constructed multiplier spaces was developed.

The discretization methods of PDEs on non-matching grids provides greater flexibility in the grid generation process, increases the portability of various approximation methods and computer implementations, enhances the capabilities of coarsening strategy in parallel algebraic multigrid methods, and provides a natural and practical way for parallel domain decomposition methods and parallel adaptive methods based on a posteriori error analysis. Our research focused on the analysis, implementation, and testing of some discretization methods of PDEs on non-matching grids. This approach seems to lead to competitive algorithms that can be used in various codes for complex applications in physics and engineering.

In the area of domain decomposition methods for non-mortar approximations on non-matching grids, the interior penalty method aims at eliminating the need for additional (Lagrange multiplier or mortar) spaces and imposes (only approximately) the required continuity across the interfaces by an appropriate penalty term. In our approach, the jumps in the values of the functions along these interfaces are penalized in the variational formulation. For smooth solutions we lose the optimal accuracy due to lower approximation at the interface, but on the other hand we produce symmetric and positive definite discrete problems, which have optimal condition number.

We also address the issue of constructing preconditioners for solving the system on the composite non–matching grids. We proposed and investigated an interface domain decomposition type preconditioner that is spectrally equivalent to the reduced (on the interface) algebraic problem. We have tested both the accuracy of the method and the preconditioning technique on a series of model problems. We were able to prove an almost optimal error estimate for the interior penalty approximation of the original problem based on the partition of the domain into a finite number of subdomains. Further, an improved error analysis for the finite element approximation of the penalty formulation was derived. Numerical experiments on a series of model second order problems were performed in order to test and verify computationally our theoretical findings.

Raytcho Lazarov and
Joseph Pasciak
Texas A & M University

Summary (continued)

In mortar approximations of finite element and finite volume methods on non-matching grids, we considered an algebraic extension of the local construction for the mortar multipliers based on the general 3–d dual finite element basis. The purely algebraic construction of the mortar interpolation in the present case requires the inversion of local mass matrices on the non-mortar interface. Since it is purely algebraic, it can be applied to the generalized objects (elements, faces and degrees of freedom) produced by the subdomain AMGe. A parallel code implementing the mortar method with algebraically constructed multiplier spaces was developed. The target application of this code is in the construction of parallel (including multigrid) preconditioners using element based (AMGe) coarsening in each subdomain. A general code was implemented to illustrate the behavior of the proposed method. It requires input data for each subdomain that includes the element topology, the local subdomain stiffness matrices as well as the mass matrices on the interfaces. This information is independent of the dimension and structure of the problem and is regenerated after an AMGe coarsening.

In multilevel adaptive grid refinement and error control, we worked on a problem of multilevel grid refinement and error control for both finite volume approximations and penalty domain decomposition methods. This work was in direct connection with research in CASC on developing and testing of parallel algorithms. We have developed 2-D and 3-D codes for parallel adaptive grid refinement that produces nested (and matching) grids. The new software was connected to the Hypre Preconditioner Library. A parallel mesh generation tool, called ParaGrid, was further developed. The development was a continuation of a 2-D project that was started last summer in CASC. ParaGrid is software that takes as input a coarse tetrahedral mesh, which describes suitably the domain, splits it using METIS, distributes the partitioning among the available processors and generates, in parallel, a sequence of meshes. It has internal solvers and is able to generate various Finite Element/Volume discretizations. The data structures allow ParaGrid to be easily connected to (or used to provide data to) external parallel finite element/volume solvers based on domain decomposition. Generation and solution routines for elasticity problems were added to the code. HYPRE preconditioners and solvers can be used. The connection is done through FEI 3.0. It has been successfully used by several researchers in CASC for algorithm testing purposes.

The software was developed in close collaboration with our student S. Tomov from Texas A&M University and Dr. Charles Tong from CASC. It was used for testing various ideas and strategies in the a posteriori error analysis and error control for convection-diffusion-reaction problems in 3-D domains with complex structure.

Bioinformatics Website Discovery and Analysis

Ling Liu
Georgia Institute of Technology

Summary

This report documents: (1) wrapper technology aiming at incorporating more functionality to XWRAP toolkit to provide richer content extraction and high automation, (2) initial design of a dynamic web crawler, and (3) website discovery and analysis which is built on top of (1) and (2).

Automatically extracting data from Web sites is recognized as an important task in many Web services and the Semantic Web in general. Most advances in Web information extraction are semi-automated or manual approaches, which require hand-coding information extraction logic and semantics. In addition, most existing extraction tools and wrapper generation tools limit their abilities to extracting information from single pages. They cannot extract related information from multiple pages or extracting information by following links. One example application that requires multi-page information extraction and more sophisticated extraction logics is the SciDAC SDM data integration pilot scenario .

We have observed that wrappers generated by *wrapper generators*, such as XWRAP Original or XWRAP Elite, usually perform well when extracting information from individual documents (single web pages) but are poorly equipped to extract information from multiple linked Web documents. An obvious reason for the inefficiency is due to the lack of system-wide support in the wrapper generators. For example, a wrapper in XWRAP Elite or Original can only perform information extraction over one type of Web document, while typical bioinformatics sources use several types of pages to present their information, including HTML search forms and navigation pages, summary pages for search results, and strictly formatted pages for detailed search results.

Information extraction over multiple different pages imposes new challenges for wrapper generation systems due to the varying correlation of the pages involved. The correlation can be either horizontal when grouping data from homogeneous documents (such as multiple result pages from a single search) or vertical when joining data from heterogeneous but related documents (a series of pages containing information about a specific topic). Furthermore, the correlation can be extended into a graph of workflows. A multi-page wrapper not only enriches the capability of wrappers to extract information of interests but also increases the sophistication of wrapper code generation.

XWRAP Composer is a semi-automated wrapper generation system that generates wrappers capable of extracting information from multiple heterogeneous Web documents. The XWRAP Composer script usually contains three types of root commands, *document retrieval*, *data extraction* and *post processing*. The document retrieval commands construct a file request or an HTTP request and fetch the document. The data extraction commands extract information from the fetched document. The post processing commands allow adding semantic filters to make the output conform to the outerface specification.

Ling Liu
Georgia Institute of Technology

Summary (continued)

One design approach to Web site discovery and analysis is to use Crawler and Service Class Description. The Crawler and Service Class Description framework provides core service discovery, classification, and integration services. The heart of this system is the service class description, which outlines the relevant aspects of a service. The description provides a mechanism for encapsulating the important components of a particular service that are common to all members of the class and is the mechanism for hiding insignificant differences between individual services. We expect that service class descriptions will be crafted by scientists interested in utilizing automatic service discovery for their own application areas, so discrimination of important or significant service aspects is left to the description writer. Indeed, since Web services may fall into several classes, aspects of a service that are important to one service class may be irrelevant in another. The service class description framework has been designed with these considerations in mind and allows service class descriptions to be tailored to the particular needs of the application domain.

Service classes are described with an XML document consisting of three main components: the data types used by the service, example queries and output, and a control flow graph representing how types can interact. The type definitions describe the atomic and complex types needed by members of the service class. The service class type system is modeled after the XML Schema type system for compatibility and provides the mechanism for declaring the data types that will be needed when processing a site. For example, a simplified view of the DNA sequence BLAST service class includes a DNA sequence input type, a DNA BLAST result output type.

The second section of a service class description enumerates the expected navigational paths used by members of the service class. Nodes in the graph represent control states; these are typically pages encountered while interacting with the site, such as the query entry page; entry and exit points to the graph are stated explicitly and represent a service's interface. Edges are directed and represent possible execution paths between control states. A BLAST control flow graph will include states for the input, summary, and results as well as control features like indirection and delay pages.

The final component of the service class description is the example queries. The crawler uses these examples to test a site and see if it produces reasonable results when queried. An input example can be used to check if a site accepts input as required by the service class. The examples may also include negative or null queries that are expected to produce no results. Negative queries are useful for both validation purposes and for determining what parts of a result page are data.

The crawler provides the discovery and integration mechanics. Starting from a root set of relevant sources, the crawler searches the Web for sites providing a Web service---e.g. those sites that support Web forms. When it finds a site, the Web crawler uses the

Ling Liu
Georgia Institute of Technology

Summary (continued)

service class description to analyze the site; if it matches the description, we say that the site is determined to be a member of the service class described.

After classifying a site as a service class member, the crawler generates a wrapper description that maps the general concepts described in the service class description to the implementation details of the particular site. Once generated, the XWrap wrapper-generation system processes the wrapper description and produces a Java wrapper for the site. The XWrap system will also generate a Java interface definition from the service class description, which is used by the mediation and user interface subsystem and allows all members of a particular service class to be queried via a common interface.

Our crawler uses the three components of the service class description described above to both identify a site as a class member and to generate a description specific to the individual site. Once a site has been identified as a member of the service class, the system must produce a wrapper that can be used by applications interested in interacting with members of the service class. Such applications expect a single, uniform interface to all members of the class, so every generated wrapper must support the same interface signature despite differences between sites with respect to input parameter names, data values, supported features, process flow, etc. The wrapper must be able to reconcile the application's expected input parameters with the nuances of the wrapped site and, once the site returns an answer, the wrapper must reverse the process by transforming the result into the expected output format. Finally, the wrapper needs to handle exception conditions gracefully.

Some future plans for XWRAP Composer continues can be described from three aspects. First, we will finish the design of XWRAP Composer scripting language. Second, we will develop the first prototype of XWRAP Composer system. Particularly, we need to finish the script compiler that generates the wrapper code. Third, we plan to design an evaluation metrics and a set of experiments to provide experimental analysis of the performance of XWRAP Composer. Most of this work will be conducted under the SciDAC project.

Enhancing the TBEPI Electronic Structure Code

Professor Calvin Ribbens
Virginia Polytechnic Institute

Summary

We have collaborated with Patrice Turchi and Ben Torralva from the Materials Science and Technology Division of CMS at LLNL to support the development large-scale codes for simulating the electrical and transport properties of metallic alloys, by developing parallel algorithms and implementations for these codes, so that LLNL parallel computing resources can be used to tackle much larger problem sizes, yielding increased accuracy in the simulations.

The recent focus has been on enhancing the modeling capabilities and the performance of the *TBEPI* electronic structure code. The goal was to develop one code that will execute efficiently on a wide range of high-end machines, including ASCI Blue-Pacific, ASCI White, and the Compaq and Linux clusters. Significant progress was made in three areas:

Incorporation of new physics into the model. The so-called "off-diagonal disorder" case was included in the parallel simulation code. This required a nontrivial modification to the communication strategy used in a major phase of the computation.

Improvement of memory-scaling. Previous versions of the parallel code required each thread to have a workspace proportional to *nthreads*natoms*, where *nthreads* is the number of threads used in each process and *natoms* is the total number of atoms in the simulation. The implementation strategy was modified so that each thread only needs memory proportional to the number of atoms assigned to that thread. Previously, memory constraints became a problem on ASCI Blue for problem sizes larger than about 2000 atoms. With this modification we are now able to do simulations with over 5000 atoms.

Allowing variable number of threads per process. Initial parallel versions of the code were developed for the ASCI Blue-Pacific machine, which has four processors per compute node. Hence, the implementation was tuned to perform significantly better with four threads per process. The parallelization strategy has now been changed so that essentially any number of threads can be used effectively. This will be important for future simulation runs on ASCI White, for example, where each compute node has 16 processors.

The final version of *TBEPI* is now in production use by Turchi, Torralva, and colleagues in CMS. A journal article is in preparation, reporting on parallel scalability up to 732 processors on ASCI Blue.

Spectral AMGe, ALE3D, and the FOSPACK Codes

John W. Ruge
Front Range Scientific Computations, Inc.

Summary

In this project, we extended the FOSPACK code and applied it to problems of elasticity. Our first task was to extend the 3D code to allow for general trilinear functions over arbitrary hexahedral meshes. The internal mesh generation capability, while limited to logically rectangular hexahedral meshes, was enhanced. The capability for reading TrueGrid files and converting the data to the internal FOSPACK data structures was added. In addition, the capability was added to be able to easily modify grids, mark subdomains, and mark domain boundaries, whether the grids were generated internally or read from files. This latter part is important since FOSLS formulations require explicit addition of boundary conditions even on free boundaries where no conditions are needed in more standard formulations.

An option was added for "strong" handling of boundary conditions (imposing them on the finite element space) in addition to the existing "weak" handling, where they are incorporated as functional terms (similar to a penalty method). A related task was the treatment of boundaries as true curved surfaces where necessary, rather than as piecewise bilinear element surfaces. This was found to be important in 2D test problems. That is, the capability was added to find normal and tangential vectors to the true surface at a given node, rather than the normal or tangential vectors to the finite element surface.

The most extensive task was the modification of the code to allow for either full solution of the FOSLS problems by AMG, or a block-diagonal preconditioning for lowering storage requirements, extremely important in the solution of large 3D problems, since these requirements are much higher than standard formulations on the same mesh. This required partial matrix assembly, as opposed to the full assembly previously used. Previously, the full form of the matrix was determined before actual assembly (which unknowns connect to which ones and where, what form will it have after elimination of boundary conditions and slave nodes, etc.) This has been overhauled, and now the main connectivity (using only the desired connections from the internal equations is determined beforehand. Additional auxiliary storage is used for any desired entries that do not fit the predetermined form, and these connections are assembled into the matrix afterwards. This also allows for flexible handling of boundary conditions, and which other connections may be included in the preconditioner (for example, keeping connections between strongly coupled unknowns).

Such a preconditioning approach (either in a straightforward defect-correction approach or in conjunction with conjugate gradient acceleration) required addition of code for the computation of the full residual when the full matrix is not assembled. The needed routines were first implemented and tested in the 2D code, then extended to the 3D version.

John W. Ruge
Front Range Scientific Computations, Inc.

Summary (continued)

Some modifications to the AMG solver were introduced to reduce the storage and computational requirements, particularly for large 3D problems. These included the use of long range interpolation, which allows for faster coarsening.

For FOSLS approaches, especially in 3D, storage can be a concern, since 12 unknowns (9 when only the first stage is solved) are required. This, combined with increased complexity of the current AMG algorithm in 3D, can require much more storage than standard finite element methods. The use of partial matrix assembly as noted above can cut down storage requirements by a factor of 3-4 (with some loss of linear convergence) while modifications to the AMG coarsening can cut complexity (and thus storage requirements) by another factor of 2-3.

Many numerical tests were performed in both 2D and 3D, with comparisons made between FOSLS and standard Galerkin formulations. While standard Galerkin finite elements appear to be more efficient for “nice” problems, FOSLS formulations allow for solution of elasticity problems approaching the incompressible limit with no degradation of convergence, while linear solvers for the standard approaches degrade badly. FOSLL* formulations are also useful in the presence of singularities, while providing direct approximations to the stresses.

Numerical Methods and Studies of High-Speed Reactive and Non-reactive Flow

Donald W. Schwendeman
Rensselaer Polytechnic Institute

Summary

The work carried out under this subcontract involved the development and use of an adaptive numerical method for the accurate calculation of high-speed reactive flows on overlapping grids. The flow is modeled by the reactive Euler equations with an assumed equation of state and with various reaction rate models. A numerical method has been developed to solve the nonlinear hyperbolic partial differential equations in the model. The method uses an unsplit, shock-capturing scheme, and uses a Godunov-type scheme to compute fluxes and a Runge–Kutta error control scheme to compute the source term modeling the chemical reactions. An adaptive mesh refinement (AMR) scheme has been implemented in order to locally increase grid resolution. The numerical method uses composite overlapping grids to handle complex flow geometries. The code is part of the Overture–OverBlown framework of object-oriented codes, and the development has occurred in close collaboration with Bill Henshaw and David Brown, and other members of the Overture team within CASC.

During the period of this subcontract, a number of tasks were accomplished, including

- an extension of the numerical method to handle "ignition and grow" reaction models and a JWL equations of state
- an improvement in the efficiency of the AMR scheme and the error estimator
- an addition of a scheme of numerical dissipation designed to suppress numerical oscillations/instabilities near expanding detonations and along grid overlaps
- an exploration of the evolution to detonation in an annulus and of detonation failure in an expanding channel

Our method has been extended to handle an ignition and growth model of reactive flow. In this model, the explosive mixture consists of two components, a reactant and a product. Each component is assigned its own equation of state, typically of JWL form, from which is constructed a mixture equation of state under certain closure assumptions such as pressure and temperature equilibrium between the two components. A single variable measures the progress of the reaction. Unlike a simple one-step Arrhenius reaction rate that was implemented in the original code, the ignition and growth model considers a rate function that undergoes sudden changes in form as the progress variable crosses certain values (although the overall rate function remains continuous). The model involves many parameters that require experimental calibration, but our aim is the development of the numerical method to handle the model and the mathematical issues surrounding the model itself. An added difficulty in the implementation of the model is that the assumed equilibrium between states must be maintained throughout the calculation. This is done in a new separate section of code that is designed to handle mixture equations of state.

In a previous subcontract, an automatic mesh refinement (AMR) scheme was implemented in the code. The scheme is a patch-type refinement scheme similar in flavor to that developed originally by Berger. For a particular domain, a composite overlapping

Donald W. Schwendeman
Rensselaer Polytechnic Institute

Summary (continued)

grid is generated using Ogen and this grid becomes the base grid on which AMR grids may be built. As the solution evolves, AMR grids are created, modified, or removed based on an error estimate of the solution. The error estimate is based on first and second differences of the solution, as was done in the first implementation of the method, and on a measure of the error due to the source term that has been improved. The method handles the source modeling the chemical reactions using a Runge–Kutta error control scheme. (There are separate chemistry time steps in the method corresponding to an integration of a set of ordinary differential equations.) The error-control scheme estimates the truncation error and this is now used in the AMR error estimator in order to determine whether a grid cell should be tagged for refinement based on the fast scales due to the chemical reaction model.

In addition to improvements in the way the source term contributes to the error estimate for AMR, the AMR grids themselves are now handled more efficiently in the implementation of the method. These improvements involve, in particular, better memory management and more efficient interpolation. This has led to a significant improvement in the computational time spent managing the AMR grids in a typical reactive flow calculation.

A scheme of numerical dissipation is now included. This was done to suppress numerical oscillations/instabilities that have been observed in various reactive flow calculations. A typical case in which the oscillations were observed involves detonation diffraction. If the detonation is expanding along a wall, numerical oscillations in the state variables (in the density for example) occur in the direction normal to the wall. For problems of this type, there is very little numerical dissipation provided by the calculation of the numerical fluxes in the Godunov scheme so that an additional scheme of numerical dissipation was considered to be appropriate. A similar problem has been observed near shocks in non-reactive flows. The problem is worse for reactive flows where small numerical oscillation can lead to large errors in the solution due to the state-sensitivity of the solution. Numerical oscillations have also been observed for problems in which a detonation crosses a grid overlap, and our scheme of numerical dissipation helps in these situations as well.

Our code has been used to study a number of problems involving high-speed reactive flow, including evolution to detonation in an annulus and detonation diffraction and failure in expanding geometries.

Numerical Methods for Electromagnetic Radiation and Scattering Problems

Karl F. Warnick
Brigham Young University

Summary

The purpose of this project is to develop a multilevel fast multipole algorithm (MLFMA) implementation that can be integrated into LLNL's EIGER electromagnetic radiation and scattering solver, in order to increase the computational efficiency of numerical simulation of electromagnetics problems. For large and complex problems, hundreds of thousands or millions of unknowns can be required. Since the boundary element method (BEM) used by EIGER leads to dense matrices, filling and solving such large systems is not feasible. MLFMA reduces the computational complexity of BEM from at least O(N**2) to O(N log N), making it competitive with sparse-matrix methods and dramatically increasing the range of problems that can be solved.

Over the period of this project, an MLFMA test code was developed. A derived version of this code will be integrated into EIGER's parallel, object-oriented framework. From a research point of view, MLFMA is fairly well understood and has been implemented in commercial codes. The accuracy and efficiency of the algorithm are strongly sensitive to physics-dependent parameters, and many aspects of low-level implementation required to obtain a working algorithm have not been published. Part of the effort of this project was devoted to working through these implementation details so that a more complete, tutorial-level technical document on the method can be produced.

During the development of the MLFMA code, we addressed several method and algorithm issues. MLFMA requires the segmentation of a mesh into a heierarchical, multigrid-like group structure. We implemented the standard oct-tree approach and an irregular method that equalizes the number of scatterers in each group. The underlying multipole expansion on which MLFMA is based is a divergent series, so the algorithm is sensitive to parameter choices and numerical precision. We found, for example, that some commonly used special function packages are not sufficiently accurate for MLFMA. The algorithm also requires interpolation on a sphere, which is another source of numerical error. We compared simple bilinear interpolation with tensor product Lagrange interpolation, and found that the less accurate scheme fails. Special care was given in examining how the interpolation methods should handle points clustered at the poles of the sphere, where the coordinate system is singular.

Institute for Scientific Computing Research

Laboratory Directed Research and Development Project Research Summaries

Enabling Large-scale Data Access

Terence Critchlow
Center for Applied Scientific Computing

Summary

The web is the preferred method for distributing scientific data because it is easy to develop a customized web interface to data that provides scientists from around the world access to colleagues' results with the click of a mouse. This revolution has the potential to advance scientists' ability to perform research by enabling large-scale data exploration, improving communication among groups, increasing academic scrutiny of results, reducing duplication of effort, and encouraging collaborations. Unfortunately, these goals are floundering because scientists are overwhelmed by the hundreds of custom interfaces they must use to access the data.

This project is developing an architecture capable of identifying, categorizing, and wrapping (i.e., writing a class to interact with) these interfaces, allowing us to provide scientists a single interface to access hundreds of data sources. This interface will simplify scientists' interaction with the data and enable them to answer more complex questions than currently possible. Our proposed architecture uses service class descriptions to describe interfaces that are of interest to our scientists, a spider (i.e., a program that parses web pages and follows html links) capable of identifying instances of these interfaces when they are encountered, an interface description that is generated by the spider detailing how to interact with an interface of interest, and a wrapper generator to take the interface description and generate code capable of interacting with the interface. For the wrapper generator, we are expanding on the Xwrap Elite program developed by our collaborators at Georgia Institute of Technology.

This project will implement the architecture through a series of increasingly complex prototypes. Each successive prototype will improve access to real-world data sources, with the final one capable of categorizing several dozen complex, scientific data sources. The resulting wrappers will be accessed by our programmatic collaborators through a customized graphical query interface. This approach will allow us to evaluate alternative strategies in a realistic environment while quickly transferring the technology to our programmatic collaborators.

In FY02, we (1) wrote an initial specification of the interface description format, which will be used to describe an interface once spider has determined how to interact with it; (2) designed an extension to the existing XWrap Elite architecture to generate wrappers for complex interfaces; and (3) wrote a simple spider that takes a set of input pages as input, and goes through each page sequentially, identifying all of the links contained on the page and queuing them for traversal. While the current implementation of the spider is not yet capable of categorizing web pages, it is able to identify those pages that utilize web forms to provide data input.

For FY03, we will (1) write a white-paper specification of the service-class descriptions; (2) complete initial development of the XWrapComposer program; (3) create a version of the web spider program capable of automatically categorizing simple interfaces; and (4) provide our collaborators with access to new data sources through prototype inter

Terence Critchlow
Center for Applied Scientific Computing

Summary (continued)

face. This work supports all Laboratory missions by providing scientists direct and efficient access to more external data such as scientific publications, chemistry databases, material descriptions, production techniques, weather data, and urban planning information.

Overcoming the Memory Wall in SMP-Based Systems

Bronis R. de Supinski, Andy Yoo, Sally A. McKee, Frank Mueller, and Tushar Mohan

Center for Applied Scientific Computing

Summary

Both CPU and memory speeds are increasing at exponential rates, as expressed in Moore's Law. Unfortunately, memory hardware is slower than CPUs, and memory speeds are not increasing as rapidly as CPU speeds. For example, on snow, the ASCI White testbed system, an average of 87 floating point operations can be completed in the time required to load one operand from main memory. Even worse, CPU speeds are increasing faster than memory speeds; thus, the number of CPU cycles required to access memory is increasing. This divergence will exacerbate an existing problem for codes with large memory footprints, including the codes typically in use at LLNL: memory accesses dominate performance. Not only is the performance of many LLNL codes dominated by the cost of main memory accesses, but many current trends in computer architecture will lead to substantial degradation of the percentage of peak performance obtained by these codes. Many researchers anticipate a "Memory Wall" in which memory accesses imply an absolute performance limit, and improvements in CPU speed provide no performance benefit.

We are extending dynamic access optimizations (DAO), a promising set of techniques for overcoming the Memory Wall, to symmetric multiprocessors (SMPs); SMP-based systems are common at LLNL. DAO techniques have shown significant promise to overcome the Memory Wall without requiring complex source code changes. These techniques change the order or apparent locations of memory accesses from those generated by the issuing program to ones that use the memory system more effectively without changing the results. For example, altering the execution order can exploit memory hardware characteristics such as interleaved memory banks and hot dynamic random access memory (DRAM) rows, while techniques that alter the apparent location can significantly increase cache hit ratios. DAO mechanisms can reduce run times of memory intensive portions of programs by factors of two to an order of magnitude. Previous projects investigating DAO focus on uniprocessor systems and require special-purpose hardware. Although promising, DAO techniques for uniprocessors do not target the systems in use at LLNL. All major LLNL computing resources are clusters with SMP nodes. Thus, we need DAO techniques that support simultaneous access to the memory system by multiple processors.

Implementing DAO techniques for SMPs is more difficult than for uniprocessor systems, since multiple processors may access the same physical address through different apparent locations. This aliasing of actual physical address with other apparent physical addresses creates a remapping coherence problem, a variant of the cache coherence problem. SMP-aware DAO techniques must ensure that the results are consistent to those that occur without remapping. SMPs use hardware mechanisms, such as cache invalidations, to solve the cache coherence problem. However, these mechanisms are keyed on (apparent) physical addresses, and thus are not automatically invoked for all of the aliases created through remapping.

Bronis R. de Supinski, Andy Yoo, Sally A. McKee, Frank Mueller, and Tushar Mohan
Center for Applied Scientific Computing

Summary (continued)

We have designed three distinct mechanisms for solving the remapping coherence problem. Our most promising technique modifies the coherence controller to account for remapping. This solution requires that the coherence controller can access the alias translation mechanisms that support uniprocessor DAO techniques. Our hardware-based solution guarantees that coherence operations are invoked on all of the aliases of a (apparent) physical address, ensuring that the memory semantics are identical to those of the original machine. Our mechanism features a fast algorithm for identifying addresses that are not aliased. The algorithm executes concurrently with the operations already required to maintain coherence and access main memory. Thus, the cost of accessing unaliased locations remains essentially unchanged.

We have designed virtual pinned memory, a novel mechanism for providing true zero-copy message passing. Ordinarily, exchanging data between nodes in cluster-based systems requires memory copies to gather it from and scatter it into the user memory locations. Also, even contiguous user data must be copied into and out of system buffers in main memory or the user memory must be pinned to physical memory so that it can be copied to the network interface. Virtual pinned memory, derived directly from SMP-aware DAO mechanisms, eliminates these requirements. With virtual pinned memory, DAO-based scatter/gather and alias translation mechanisms copy user data directly to the network interface.

Finally, we continued our work on understanding application memory access regularity. Our novel tool that measures an application's regularity gathers statistics that strongly indicate what memory optimizations will improve the application's performance. This research, the subject of a Master's Thesis at the University of Utah, demonstrates that applications with irregular access patterns require our DAO mechanisms.

Visualization Streams for Ultimate Scalability (ViSUS): Monitoring Terascale Simulations in Real Time

Valerio Pascucci

Summary

Modern scientific simulations and experimental settings produce ever-increasingly large amounts of data that traditional tools are not able to visualize in real time, especially on regular desktop computers. Scientists are unable to interactively explore the data sets that they produce, which creates a frustrating slowdown in the overall process of scientific discovery. Use of innovative, high-performance visualization techniques that allow interactive display of very large data sets on simple desktop workstations and the monitoring (or steering) of large parallel simulations will have specific applications to several of the DOE's and LLNL's missions, including stockpile stewardship, energy and environment, nonproliferation, biology, and basic science, that use large-scale modeling and simulations.

The ViSUS system will implement unified scalable solutions allowing large data visualization on a single desktop computer, on a cluster of personal computers, and on heterogeneous computing resources distributed over a wide-area network. When processing terabytes of scientific data, our goal is to demonstrate an effective increase in visualization performance of several orders of magnitude in two major settings: (1) interactive visualization on desktop workstations of large data sets that cannot be stored locally; and (2) real-time monitoring of a large scientific simulation with negligible impact on the computing resources available. The main focus of the ViSUS research is in the development of a novel data-streaming infrastructure based on a suite of progressive and out-of-core visualization algorithms enabling the interactive exploration of scientific datasets of unprecedented size. The methodology aims to globally optimize the data flow in a pipeline of processing modules, with each module reading a multi-resolution representation of the input while at the same time producing a multi-resolution representation of the output. Multi-resolution methods provide the necessary flexibility to trade speed for accuracy in the visualization process. Maximum coherency and minimum delay in the data flow are achieved by extensive use of progressive algorithms (see figure below) that continuously map local geometric updates of the input stream into immediate updates of the output stream. We are developing a prototype implementation of this streaming infrastructure to demonstrate the practical flexibility and scalability of our approach.

During FY2002, we (1) implemented a new progressive slicing technique that computes planar cross sections of rectilinear grids with arbitrary orientation; (2) designed and implemented a new data layout allowing real-time geometric queries for rectilinear grids with up to 0.5 trillion nodes; (3) developed a novel multi-resolution representation that simplifies the real-time extraction of minimal adaptive models; (4) introduced the first subdivision method for unstructured meshes of any dimension and cells of virtually any type; and (5) tested the basic network components of our streaming infrastructure.

Valerio Pascucci

Summary (continued)

In FY2003 we will (1) build a prototype of a remote viewer that interactively slices data sets retrieved from a remote storage system; (2) implement a direct-streaming infrastructure prototype that connects the nodes of a scientific simulation to a set of data servers used for permanent storage; (3) parallelize the preprocessing of the input grids to achieve one order of magnitude improvement in the data reordering into our storage layout; (4) compute in external memory multi-resolution representations of triangulated surface meshes; (5) define and implement multi-resolution wavelet models to be combined with a new volumetric subdivision scheme; and (6) explore new image caching mechanisms for decoupling the display process from the image rasterization process.

Institute for Scientific Computing Research

Student Internship Research Summaries

siMPI: Serial Implementation of MPI

Maria Luisa M. Alano
University of San Francisco

Summary

The need has arisen for creating a sequential implementation of the Message Passing Interface (MPI) that conforms to the high-performance parallel version. Named siMPI, this project enables laboratory users to run MPI programs outside the POE environment of the IBM ASCI machines for debugging and testing purposes. This is convenient for developers and avoids allocation and waste of high-end resources. In addition, from a general perspective, siMPI is a tool that lets developers create parallel programs without the need for actual parallel machines. A single processor suffices, without the overhead resource allocation, trouble, and complexities.

To implement siMPI, the standard C programming language was chosen. The essential calls of MPI (the most common two-thirds of the entire 209 function-interface), such as MPI_Send, MPI_Receive, have been implemented so far. The Fortran equivalents have also been generated.

The behavior and output of the library calls comprising siMPI have been patterned, but not exactly mirrored, after the standard MPICH implementation.

The future for siMPI is promising. Once it is deemed stable and fit, siMPI will be packaged with the standard pyMPI so users may freely explore and use MPI with convenience.

The remaining one-third of the standard MPI calls should not be difficult to complete. Also, as with any software project, siMPI should undergo rigorous testing before release.

Multigrid Performance on Systems of PDEs Using Computational Local Fourier Node Analysis

David Alber

University of Illinois at Urbana-Champaign

Summary

Geometric multigrid is a fast and cheap alternative to Krylov iterative methods for some structured problems, but it is not always clear how well it will perform on a given problem. Surprisingly slow convergence rates sometimes plague a given problem. Multigrid performance on linear systems arising from systems of partial differential equations (PDEs) is understood theoretically for many special problems. This research aims to generalize the systems of PDEs whose multigrid convergence can be understood theoretically.

Local Fourier Mode Analysis (LMA) is a powerful technique that is applied to estimate smoothing and convergence rates of multigrid smoothers and methods. In addition, it is possible to use this technique to determine the spectrum of a preconditioned matrix. This is useful to estimate the effectiveness of using multigrid as a preconditioner on a particular problem. Therefore, LMA was selected as the method to use for gathering data. Unfortunately, very few software packages exist that perform LMA, and doing the analyses by hand is not an option. Therefore, a new software package was written to carry out LMA calculations. Using a library provided by Hiptmair and Metzger we authored an application that calculates smoothing and convergence rates. It also calculates the spectrum of a matrix preconditioned by multigrid. Other tools were added that make using the software more intuitive, such as a function that builds the stencil input files, and a graphical user interface.

This software is now being used within CASC to study different systems of PDEs amid how multigrid performs when trying to solve or precondition them. I hope to apply some of this work to a Masters thesis. I intend to expand the current software to offer more functionality, and, in any case, I plan to continue my collaboration with CASC on the research that has been started. Currently, we are working on a journal article that includes some of this work.

Wrapper Integration within the DataFoundry Bioinformatics Application

John C. Anderson
University of the Pacific

Summary

The DataFoundry bioinformatics application was designed to allow scientists to interact directly with large datasets, gathered from multiple remote data sources through a interactive graphical interface. This type of application is becoming essential, as there are more than five hundred sources of publicly available bioinformatics information available today. The need to access even a small fraction of the publicly available, heterogeneous bioinformatics sources on the World Wide Web motivates the development of a data access solution that facilitates use of external data.

Gathering information from multiple data sources, integrating that data, and providing an interface to the accumulated data is non-trivial. Advanced computer science techniques are required to develop a solution. One possible solution, which was applied to this problem, involves the use of wrappers. Wrappers are specialized information access programs that are able to transmute that information to a form usable by a single application. Wrappers were deemed the most appropriate way to extend the DataFoundry bioinformatics application to support data integration from multiple sources. By adding wrapper support into the DataFoundry application, we expect that this system will provide a single access point to bioinformatics data for scientists.

To date, we have been successful in defining and implementing a communication path between the client and server DataFoundry bioinformatics application programs that includes support for wrapper-based queries. Adding more external bioinformatics sources to the DataFoundry application is now as simple as developing a wrapper for a particular data source.

There are still limitations within the current framework. The most important is that there are currently relatively few wrappers available for bioinformatics data sources. Work continues to develop an automated wrapper generation system able to generate wrappers for any bioinformatics source. The second is performance; with relatively few wrappers, the client and server are able to process the volume of applicable data with ease. Once all of the relevant sources are wrapped, however, performance issues will become pressing.

Implementing the Column Inventory Algorithm

Cheryl Barkauskas
University of Wisconsin-Madison

Summary

One way to evaluate the accuracy of a climate model is to test it against known benchmarks, such as predictions of the amount and flux of dissolved inorganic carbon in the oceans. Calculating the column inventory, the amount of a substance present in the ocean under a certain surface area, is such a test.

Climate model data sets are sufficiently large (often over a hundred thousand points) that speed can be a factor in calculations. This makes the choice of programming language important, to minimize execution time. However, the functionals we are creating are tools for climate scientists, not computer scientists, so the simplicity of the interface to the methods is also a consideration. Therefore, several programming languages - Python, Fortran, and C - are tested to determine the best one in terms of execution speed and ease of implementation.

We first wrote a Python program that read a data file and extracted the necessary variables. This separates the interface from the calculation and provides each of the three implementations with a common standard to follow. We next created a separate set of Python calls to calculate the column inventory and mean. Finally, we translated the code from Python into C, and then from C into Fortran. (To communicate with the Python code, the Fortran module had to be compiled by a tool named Pyfort, a Python-Fortran connector developed by Paul Dubois at LLNL.)

Three working implementations of the column inventory algorithm - in Python, Fortran, and C - were completed. The Python version is most likely to be used by the climate research group at LLNL because most of its members understand Python, so it will be easy to deal with the source code. However, preparing the Fortran was also useful because it taught us the basics in making Python and Fortran communicate. Many of the group's programming routines are already written in Fortran, so the ability to call them from a Python interface is valuable. Speed of execution is not a major factor with current dataset sizes, even though Python generally runs about four times slower than either C or Fortran.

Additional features added to the basic algorithm include correcting volumes of partial bottom cells and applying regional (basin) masks to the data.

The climate group may now compare the output of the column analysis code to ship-track data to determine how well their models match the real world. Also, the versions will be archived for future reference for how to code Python-C and Python-Fortran programs.

We brought our summer research through the lab “review and release” process with a poster presentation at the 2002 student research symposium.

Improving the Network Intrusion Detector

Bridget Benson
Cal Poly San Luis Obispo

Jessica Fisher
Harvey Mudd College

Ian Webb
Colorado State University

Summary

Computer Security has become an increasing concern in today's networked society. The College Cyber Defenders (CCD) program creates a pool of young computer scientists knowledgeable about computer security. Our project was to enhance the Network Intrusion Detector (NID) by adding signatures to its database and updating the web-based Intrusion Detection Exchange Server (IDES). We also worked on finding a reasonable network setup for future CCD programs to use.

First, we created an internal network on which to run attacks, one that was not connected to the lab network. We also modified the Flexnet server developed at Sandia Lab to be a backup server through the Network File System. This system allowed us to recover damaged systems to a known state.

Next, we downloaded exploit code from the internet, wrote it to CDs to bring into our internal network, and attempted to run the exploits against our victim machines. We listened to the traffic with Ethereal, a graphical utility similar to tcpdump. From this traffic, we found packets which contained unique strings to identify the attempted exploits; these strings became new NID signatures. Most of the exploits we downloaded actually did not work, and we were able to obtain only 23 new signatures by this method. We then turned to an open-source intrusion detector, Snort, and wrote a PERL program to convert Snort's "rules" to NID signatures. Through this program, we obtained over 800 more signatures.

Our final task was to modify IDES to run with the Apache webserver. IDES provides a way for NID users to interact with developers and suggest improvements to the system, and offers utilities to create configuration files for NID and add signatures to the database. IDES was designed to run with Microsoft's IIS webserver, using Microsoft Access as its database and ASP (Active Server Pages) as its script language. To make IDES less platform-dependent, we modified it to run with the Apache webserver, using MySQL as its database and PHP as its script language. We also modified the PERL program to place the new signatures directly into IDES's database.

Future CCD programs' network setup would be enhanced by the use of a one-way file transfer system, which would help avoid the necessity of using CDs to transfer data. Enhancements to Flexnet to turn it into an easy-to-use backup server would also be beneficial. IDES would be improved by adding functionality to assist more of NID's multiple features.

Improvements to ALE3D

Timothy Campbell
University of Arizona

Summary

We implemented five improvements in the Livermore workhorse code ALE3D.

ALE3D employs a point-sampling algorithm to determine volume fractions in elements composed of mixed materials. We devised and implemented an algorithm that efficiently calculates these volume fractions explicitly, and in comparable run-time. This will result in better simulation accuracy, especially for input with multiple materials that mix.

ALE3D was not capable of material burn across slide surfaces (partially disconnected regions of the mesh). I altered ALE3D's burn routines to detect slides, and burn across, assuming a burn material exists on the other side, and to do so in such a way so as to correct for disparate mesh zoning across the slide.

We introduced Python hooks into ALE3D. The code is now capable of reading in Python source code as input to modify internal data structures at run-time. This allows users, for example, to create boundary conditions using arbitrary functions at arbitrary time steps.

We expanded ALE3D's internal mesh generation capabilities. ALE3D previously relied on externally generated meshes, but can now create simple meshes based on geometries like spheres, cones, blocks, etc. This makes generation of test problems considerably less time-consuming, both in terms of problem generation time and overhead associated with learning how to generate a mesh.

We modified ALE3D's input code to allow it to read in large data files in parallel. Mesh input is read and sent only to nodes that need it, based on an a priori domain decomposition of the problem.

These changes to ALE3D range in type from improved accuracy, greater functionality, greater ease of use, and improved performance.

Helios: Illuminating C++ Memory Management

Karl Chen
UC Berkeley

Summary

Helios is a memory-allocation software tool that provides a flexible and efficient way of debugging large C++ programs and Python extensions in a developer-friendly, user-steerable way. It is intended to be used in debug mode, as well as production code. Helios allows dynamic (run-time) control of optimization and debug options. If necessary, these options may be fixed at compile time for maximum optimization (each option fixed is one less memory access and/or conditional branch at runtime).

Helios contains a “Controller” (C++) that delegates allocation and deallocation to the appropriate “Allocator”. The various “new” and “delete” operators (C++) and malloc() and free() (C) call the Controller's allocate() and deallocate() functions. The Controller can be manipulated directly in C++ or through Python (with Pyffle wrappers). The first time we ran Kull with Helios, Helios immediately caught a certain kind of memory bug (malloc/delete mismatch).

Kull, the project Helios was created to support, and other lab projects are planning on using Helios, and we will personally continue to maintain it. The code for Helios will be put up for software Review and Release.

Draco Help System

John Clark
Northern Arizona University

Summary

Draco is a PMesh physics model program used to create meshes. Draco integrates an OpenGL graphical interface to display the mesh, as well as a command line interpreter used to manipulate the Python objects that control the mesh. The Draco help system provides the user with a localized place for finding documentation on all of the Python objects available. The documentation will inform the user of all the methods, functions, classes, and data that a particular object has, and it is organized in an easy to read website format.

PyDoc is a documentation tool for the Python scripting language, it will get the information about any Python objects and output it into a text format. Unfortunately, PyDoc does not have the ability to create the beautiful html format that we wanted. However, the C++ and C documentation tool Doxygen does have the ability to create html as well as many other desirable formats of output. So we created a bridge between the two. The bridge uses PyDoc to get the information from the Python objects in Draco, then it formats that information to create files that Doxygen can understand. After this, Doxygen can be used to create the html pages.

As another part of the project, we made viewing help pages easy to do while using Draco. Since the pages are in html, a browser is the best way to display them. We allowed for two options. The first option is an integrated browser written with the QT windowing library; the second option is to use a browser integrated into the operating system environment. The integrated browser is very simple and does not allow for things like style sheets, or backgrounds with differentiated highlight colors. The system browser can be anything like Netscape, Mozilla, or Konqueror, and is determined by setting the environment variable DRACOBROWSER to the browser of choice. Either of these browsers may be launched, while running Draco, from the command line interpreter or from the Help pull-down menu in the Draco GUI.

The help system will be maintained and augmented by my supervisor who did the initial work on the bridge between PyDoc and Doxygen.

Computations in Scalar Field Topology

Kree Cole-McLaughlin
University of Utah

Summary

Scalar fields are used to represent data in different application areas like geographic information systems, medical imaging or scientific visualization. One fundamental visualization technique for scalar fields is the display of isocontours - that is, sets of points of equal scalar value, also called the isovalue. For example in terrain models isolines are used to highlight regions of equal elevation. The Contour Tree is a graph that represents the relations between the connected components of the isocontours in a scalar field. Two connected components that merge together (as one continuously varies the isovalue) are represented as two arcs that join in a node of the graph. Each node occurs at a critical point in the scalar field. For a simplicial domain with a piece-wise linear interpolant, Carr et al. present a simple algorithm for computing the Contour Tree with complexity 0(n log n + N), where n is the number of points in the domain, and N is the number of simplices.

One fundamental limitation of the Contour Tree is the lack of additional information regarding the topology of the contours. The topology of an isocontour is fully characterized by its Betti numbers. In 3D fields the isocontours are surfaces and the Betti numbers correspond to the number of connected compo-nents, the number of tunnels, and the number of voids enclosed by a surface. Pascucci has presented a simple and fast algorithm for the construction of the Contour Tree of a 3D scalar field augmented with the Betti numbers of each contour. The complexity of this approach is 0(N log N) in time and 0(N) in storage. The nodes of the Augmented Contour Tree correspond precisely to every critical point of the field. Thus we trade a slightly slower algorithm for more information.

We have implemented both of the above algorithms and produced images of complex contour trees for a few data-sets. The images are produced with the graphiz utility from AT&T. Our data-sets are quickly approaching the limit of this utility's ability. Therefore in the future we will need to develop our own tool for drawing the Contour Tree. We have designed a third algorithm that computes the Augmented Contour Tree in 0(n log n ± N) time. This new algorithm combines the efficiency and simplicity of the scheme of Carr with the augmented information of Pascucci. Thus we have successfully introduced new and useful information to the Contour Tree data structure without slowing down the computation.

Our main result is a completely new approach for computing the Contour Tree for rectilinear grids. This scheme is based on a divide-and-conquer approach, and assumes only that we can construct the tree for a single cube. In this method we are not limited to using piecewise linear interpolation, which can introduce false topology on rectilinear grids. In the time I have spent here I have implemented all of these algorithms, and have collected several practical results. The time complexity of our scheme is further

Kree Cole-McLaughlin
University of Utah

Summary (continued)

improved to 0(n + t log n) where t is the number of critical points of the mesh. We have applied this technique to the methane data set of F. Gigi.

We have several goals for future development of this project. The method of producing the Contour Tree of a cube with a trilinear interpolant involved a careful study of the topology of the trilinear function. For example, we discovered a case missed by earlier papers classifying the possible configurations of the trilinear interpolant. We plan to publish our new result with a formal proof of the completeness of our analysis.

In the possible extension of this project we plan to study the practical scal-ability of the algorithm for very large data-sets. We also plan to implement a parallel version of our new algorithm and develop a new GUI component for the effective display of the Contour Tree.

Late Time Outgoing Distribution from Spheres

Hillary E. Davis
Georgia Institute of Technology

Summary

The dominant radiation transport model used for modeling Inertial Confinement Fusion (ICF) is Implicit Monte Carlo (IMC). Although there exists an acceleration scheme that provides a factor of five increase in performance, its use is restricted to well defined problems. In this project, we are reducing the bias present in the acceleration scheme so that there are no limitations on its use.

The overall project comprises two methodologies: Compressed Walk (a non-recursive solution to the random walk of a constant speed particle) and adjusting diffusion solutions. Our effort was to find the thickness dependent late time outgoing distribution to be used in concert with the diffusion methodology.

We chose a thickness for the media and placed an isotropic point source in the center. Subsequently, using Monte Carlo, we followed photons as they incurred isotropic scatters. As the photons left the surface of the media, tallies were made. These numerical experiments recovered not only the distribution and its error, but also several velocity moments. The distribution was then fitted to a polynomial of the given form and tested against measured moments.

We were able to capture the distribution and moments in less than thirty coefficients. Such a short polynomial makes our scheme practical for use in reducing bias, without incurring a significant runtime penalty.

The angular distribution can be broken down into four regions: early time internal and outgoing, and late time internal and outgoing. We have presented the late time outgoing. A preliminary investigation into the data suggests that the early time outgoing and the late time internal are tractable. Although we can measure the early time internal behavior, the complexity of the distribution makes it difficult to gain runtime advantage over the Compressed Walk solution. Hence, work will continue on the time dependent outgoing and late time internal distributions as well as the Compressed Walk solution.

Tools for Parallel Finite Element Analysis

Paul Dostert
UC Davis

Summary

Computing numerical solutions to finite element problems in parallel can often be difficult task. Our goal was to create tools with which users can simplify the process, from creating a mesh all the way to solving the finite element matrices in parallel.

The natural first step in finite element problems is to define a mesh. For our code, Netgen was used to create triangular or tetrahedral meshes. After a mesh is established, we use the Metis mesh partitioner to split the mesh into multiple subdomains, one per processor. The mesh information, as well as the connectivity between subdomains, is written to a different file for each subdomain.

Each processor reads in the mesh information corresponding to its subdomain, and the finite element matrix is created according to its degrees of freedom on the local subdomain. Next, a matrix corresponding to the local-to-global degree of freedom table is created on each processor. This matrix, which we will call P, as well as the finite element matrix, A, are stored in ParCSR format. The ParCSR format was chosen due to the efficiency which it stores sparse matrices in parallel, and its compatibility with the well established hypre routines. To obtain the full finite element matrix, across subdomains, we take P^T A P. We then apply one of many possible linear solver routines from the Hypre toolkit to generate the solution on each processor.

Results were satisfactory, with some routines scaling at near perfect levels. If input time is included, the total time to input and solve the system was actually greater than two times faster on double the number or processors in some cases. This is an I/O artifact: the input files are smaller if we split a domain into more subdomains, lowering the input time on each processor.

Future development of this code may involve improving compatibility with different mesh generators, including generators that use non-triangular or non-tetrahedral elements. Other generalizations include implementing other parallel matrix storage formats and investigating their benefits.

Efficient Translators for Document Object Model Trees

Roger Elion
Purdue University

Summary

Cyclops is a program written in Java that allows users to utilize a graphical user interface to create or make changes to legacy physics codes. A second program, written in Python, converts a language-specific parameter description file into a data structure that will automatically and dynamically build another data structure that represents an input deck. Upon parsing the input deck, each value and comment is added to a Document Object Model (DOM) tree, which is then output in XML and passed to Cyclops for the user to interface.

This input translator is a far more condensed and efficient than the existing language parser, written in Perl, that was being used. The concept of recycling the parameter description file in order to create a look-up hash automated this portion of the process. (This was the idea of Thomas Richmond, a returning summer student.) In addition to this, the recently released DOM tree structure enables programmers to convert DOM trees to XML by merely calling a function. Overall, the decision to incorporate these fairly young languages with easy to manipulate data structures has condensed thousands of lines of Perl code into hundreds of lines of Python.

Increasing the Transfer Speed of SSH to Utilize Full Network Bandwidth

Jason R. Estrada
Baylor University

Summary

This report presents a performance analysis and describes several techniques to obtain more performance from the secure networking tool SSH, including implementing a high performance protocol, using multiple processors to perform encryption, and making use of multiple sessions. A solution, using multiple SSH sessions to multiplex data transfer over the network has been developed in an effort to gain faster transmission speed over the network, while still using strong encryption ciphers.

SSH is a protocol for secure remote login and other secure network services over an insecure network. SSH has become the de facto standard for obtaining a remote shell or transmitting data across the network securely in a day in which security is an increasing concern. However, we would also like to accomplish this very quickly.

Most of the processing time spent within SSH is in the actual encryption of the packets, and not in the transfer of the packets across the network. It, of course, takes time to send the packets over the network, but that time pales in comparison to the time spent in the encryption algorithms. Since SSH (at least in the case of OpenSSH) is not multithreaded, CPU cycles are often dormant on multiprocessor machines that could be used to perform encryption/decryption on the outgoing/incoming packets.

All testing and benchmarks were performed in a non-routable subnet containing two host machines with a round trip time (RTT) of 116 usec, each running Redhat Linux 7.3. Neither host was mounting NFS or AFS, in order to minimize the network traffic over the line. A set of benchmarks was performed on the systems and the network to discover performance and scaling metrics. Netperf was used to measure the maximum achievable throughput of the network

The Secure Shell used in the testing was from OpenSSH. At the time of testing, the most recent version of OpenSSH was 3.2.3pl. OpenSSH can be obtained at www.openssh.org. Only the SSH2 protocol was used and the SSL library used for encryption was provided by OpenSSL's distribution of SSL. At the time of testing, the most recent version of OpenSSL was 0.9.6d. OpenSSL can be obtained at www.openssl.org. Public/Private key authentication was used since SSH was compiled and run within user space, instead of running as the system. Although not recommended, a null key phrase was used on the keys so it would not be required to enter a password. This made the testing of SSH more feasible.

Netperf is a benchmark utility that can measure end-to-end network latency and throughput. The latest version at the time of benchmarking was 2.lpl3. Netperf can be found at www.netperf.org. The results from these two test show that the maximum available throughput of the 100Mb Ethernet line is approximately 94Mb/sec with 99%

Jason R. Estrada
Baylor University

Summary (continued)

confidence. No TCP optimizations were applied to the network, such as TCP tuning, since no root access to the machines was available.

Testing the transfer speed of SSH was achieved by using two shell scripts, ssh2_test. sh and full_ssh2_test. sh. The shell script full ssh2_test. sh was used to test all the supported SSH2 ciphers. The shell script logged the results to individual log files for later examination. Minor additions to the SSH source code were made to time and calculate the effective transfer rate of data across the network without using an external program. The shell script assumes that it is executing the modified binary of SSH; otherwise no output will be displayed in the log files.

Efficient Approximation of Solutions to Mass Matrix Equations

Aaron Fisher
UC Davis

Summary

The finite element method approximates solutions to partial differential equations on arbitrary domains tesselated with a mesh. The bulk of the CPU time involved in computing this approximation comes from repeatedly solving a linear system, Ax=b, composed of a mass matrix A, a vector of the degrees of freedom x, and a right-hand side b. The mass matrix, which is dependent on the geometry of the mesh and the order of the approximation, tends to be sparse, but can be quite complicated for unstructured meshes. Furthermore, in problems where the mesh is allowed to deform, the mass matrix changes at each time step to accommodate the deformation.

Since mass matrix solutions dominate the execution time, their CPU requirements make a natural target for amelioration in finite element problems. Even relatively small improvements in these solves prove valuable since they appear in each time step.

Our project focused on approximations to the mass matrix equations for an optical fiber waveguide problem. By comparing the results to an analytical solution to the optical waveguide problem, error values were obtained for the various approximations. We studied both this discretization convergence behavior and algebraic convergence behavior for the iterative solution of the resulting system.

Specifically we studied the preconditioned conjugate gradient with block Jacobi, banded, and ILU preconditioners, with variations in block size, band radius, and fill levels. We also studied a well-known approximation to the mass matrix known as mass lumping.

Through this work we gained a valuable understanding of the error levels involved in this class of approximations. We also found that the convergence behavior of the methods is sufficiently well behaved to consider loosening up the convergence tolerances in order to save CPU time. Finally, we found that by adjusting the block size/fill levels in the block Jacobi/ILU preconditioners we could save a moderate amount of CPU time.

Despite some modest successes, this remains a work in progress. Future plans include propagation of linear system error into the overall simulation in order to determine how much is tolerable, and experimenting with more esoteric solution methods.

Compression of Scalar Functions

Ilja Friedel
Caltech

Summary

Scientific simulations generate huge data sets, which need to be transmitted and stored for future analysis. To reduce size and transmission time of the generated files, compression schemes are employed. Current methods are often based on wavelets defined on grids. These methods perform well on smooth data but loose efficiency near discontinuities. It is known that meshes that adapt size and shape to the data have better approximation properties than regular grids. Unfortunately the storage overhead for arbitrary meshes is high compared to regular grids. It is our goal to create a progressive, triangle based compression method for 2d slices of data. Progressivity permits browsing of data sets at coarse resolution without need of transmitting the whole file.

It was our goal to avoid the overhead associated with the transmission of connectivity information. For this reason in our model the receiver makes data dependent decisions (based on the already received data points) on the construction of the mesh. The sender monitors these decisions by running the same construction.

At each iteration of the algorithm
- the receiver selects the triangle with the expected highest error
- a point inside this triangle is chosen, such that the insertion is expected to provide the largest error reduction
- the difference to the prediction of the z-coordinate of this point is transmitted from the receiver
- the selected triangle is split using the new point
- based on this new knowledge the receiver re-triangulates the mesh to obtain a good approximation for the next step

We examined different strategies for each of the steps in this algorithm and compared them using rate distortion curves. Most of these strategies require the construction of a local quadratic model of the surface based on discrete curvatures. Point insertion based on a modified farthest point strategy combined with edge flipping based on bending energy where shown to be efficient for many smooth input data sets. With the new algorithm approximation errors of 35% to 50% of the magnitude of grid based errors where achieved when using the same number of data points.

To get a practical method the residuals of the transmitted data points need to be compressed. We are considering different bit-encoding schemes.

The current algorithm can be classified as a lossy method, because it assumes continuous data. In many applications the data is defined on a grid. By choosing the inserted points from this grid the algorithm will operate in a lossless mode.

In our implementation only residuals are transmitted. We plan to examine heuristics for the sender to guide the decisions made by the receiver.

Exploiting Workload Distributions for Application-Level Dynamic Load Balancing

Karen Glocer
UC Santa Cruz

Summary

A large volume of research in distributed computing suggests that load balancing is a primary concern, and it will inevitably become more difficult as applications get larger and more complex. The goal of our work is to allow automated load balancing at the application level. Target applications are typically implemented in MPI. MPI has been described as the assembly language of parallel programming. Like assembly language, it is both efficient and difficult to manage. For years people wrote in assembly for its efficiency, but eventually it was superseded by high-level languages through the advent of optimizing compilers. MPI is currently the standard for high performance scientific applications, but as problems increase in complexity and size, some difficulties will become more evident and it will benefit developers to look at handling these difficulties with higher-level parallel programming models.

Most previous load balancing research falls into three categories: scheduling between independent jobs, workload characterization, and balancing the load within an application, something that is usually handled by the application programmer. Only relatively recently has automated load balancing been proposed for balancing load within the application, and not just between independent applications. The two most notable examples of this are Charm++ and Cilk, both of which use a programming model in which work is divided into discrete chunks where the number of chunks is much greater than the number of processors. These chunks can then be moved around based on their estimated computational and communication costs.

In these systems, computation and communication costs are estimated heuristically and no attempt is made to characterize the workload distribution of real scientific applications. Our work fills this gap. Using a new tool, the Sequoia Tracer, we gathered traces of four applications: sPPM, SWEEP3D, High Performance LINPACK, and FLASH. These are each representative of programs run on the ASCI clusters and are common in scientific computing. In particular, FLASH uses adaptive mesh refinement, a technique that is growing in popularity and which requires a great deal of internal load balancing.

The traces contain each entry and exit from a computational state as well as communication. A trace is generated for each processor used by the application in question. Because tracing can itself affect timing information, it makes sense to use alternative metrics to measure the amount of work each processor is doing. In this case, a count of floating point operations is used. From this information, we were able to get the distribution of work done by each processor between communications. Each instance is analogous to a chunk of work in a programming model such as Charm++, so the distribution can serve as a way to estimate computational costs that were only estimated with heuristics in the past.

Karen Glocer
UC Santa Cruz

Summary (continued)

The results are surprising. The workload distributions for these chunks appear to be strongly bimodal. Most of the chunks perform very little computation, a few are very long, but chunks in between are infrequent. The small computational loads seem to fit a Pareto distribution in three of the four applications we tested. The distribution of chunks with large computational load is harder to classify because the numbers are smaller. Most importantly, these results show that it is very easy to predict which chunks will perform the most work, and these can be preemptively moved to either a dedicated processor or to one with a much smaller load, in a method similar to the one proposed by Harchol-Balter and Downey.

Interactive Exploration of Large Isosurfaces in Volume Datasets

Benjamin Gregorski
UC Davis

Summary

Numerical simulations performed on LLNL supercomputers are generating unprecedentedly large amounts of data. A standard method for visualizing, exploring, and understanding this data is to examine various isocontours. The problem with directly viewing the isocontours is that each consists of hundreds of millions of elemental triangles. Surfaces of this size cannot be viewed at interactive frame rates on conventional desktop machines. Intelligent preprocessing techniques and runtime algorithms are needed for interactive exploration and visualization. The goal of this project is to develop a system for dataset exploration on desktop workstations.

Large datasets are divided into subsets called bricks and stored on disk or on a remote system. These bricks are loaded and unloaded by the application as they are needed so that system resources are effectively utilized.

Interactive frame rates can be achieved using algorithms that employ multiresolution data structures with view-dependent rendering. Multiresolution data structures make better use of available storage space and computational power by representing important areas with more detail and less important areas with less information. View-dependent rendering selects what portion of the data to render based on where the user is looking and on the user's perception of the data. Objects outside the field of view do not need to be rendered, and objects that are far away can be rendered at lower resolutions. View-dependent rendering takes advantage of the multiresolution data structure to efficiently select what to draw. Our multiresolution data structure is a recursive tetrahedral mesh based on longest edge bisection. The mesh structure supports fast, local refinement necessary for view-dependent rendering. In addition, it supports an efficient method for ensuring mesh continuity required for iso-surface extraction. View-dependent rendering calculates the distortion on the view screen; it adjusts the refinement of the multiresolution mesh, selecting finer levels where more detail is needed and coarser levels where less detail is needed. Our continuing work is focused on working with large datasets, accurate representation of the data, and computational optimizations.

A paper documenting our summer work has been published at the 2002 IEEE Visualization Conference. It is also available as UCRL-JC-146819. Our continuing work is focused on working with large time-varying datasets, accurate representation of the data, and computational optimizations including occlusion culling, asynchronous rendering and compression.

A Parallel, Adaptive Implementation of the Immersed Boundary Method using SAMRAI

Boyce Griffith
New York University

Summary

The immersed boundary (IB) method provides a mathematical and computational framework for addressing problems involving fluid-structure interaction and has proved to be especially useful in simulating biological fluid dynamics. Realistic simulations that use the IB method require both high spatial resolution and very small timesteps. Consequently, high performance computing is an important component of simulation research employing the IB method.

Currently, all high performance IB software is designed to run on shared-memory vector machines. There is a clear need to develop IB software which will run efficiently on distributed-memory computers. We are harnessing the expertise available at CASC to develop such software.

The IB method specifies the interaction of a fluid, described as an Eulerian variable, and an elastic material, described as a Lagrangian variable. Consequently, the fluid is typically discretized on a Cartesian grid, while the elastic material is described by a network of Lagrangian points. A smoothed approximation to the Dirac delta function is used to connect these two quantities. Through the discrete delta function, quantities such as velocity may be interpolated from the Cartesian grid to the Lagrangian points, and quantities such as force or density may be spread from the Lagrangian mesh to the Cartesian grid.

High spatial resolution is required near the boundaries in order to capture important boundary layer flow. Since the current discretization of the IB method is very well suited for use of structured AMR, we are developing IB software using the SAMRAI library.

After one summer's work, single level IB software using SAMRAI and hypre is nearly complete. This software should allow for IB computations which make effective use of distributed-memory computational facilities.

Upon completion of the single level IB software, we will begin incorporating AMR into the solvers. Currently, there is no parallel implementation of the IB method that includes AMR. We are also interested in extending the IB method to coupled electrical-mechanical biological models.

We are also interested in developing new computational approaches for fully implicit temporal discretiza-tions of the IB equations. The current semi-implicit discretization typically requires very small timesteps due to the stiffness of the IB equations. An efficient implicit timestepping scheme could potentially ease this timestep restriction and allow for more realistic simulations.

Undecimated Wavelet Transforms for Image De-noising

Aglika Gyaourova
University of Nevada, Reno

Summary

Almost every kind of data contains noise. Noise reduction is a required step for any sophisticated algorithm in computer vision and image processing. This problem has existed for a long time and there is no general-purpose solution for it. A tradeoff always exists between the removed noise and the blurring in the image. The use of wavelet transforms for signal de-noising has been started in last decade. The capability of wavelets to render detail about spatial-frequency information is the main reason for this investigation. This property promises a possibility for better discrimination between the noise and the real data. Successful exploitation of wavelet transforms might lessen the blurring effect or even overcome it completely.

There are two main types of wavelet transforms – continuous and discrete. Because of the discrete nature of computers, computer programs exploit the discrete wavelet transform. The discrete transform is very efficient from the computational point of view in operation count and in compression of storage through decimation – the discarding of fine-scale data that is not needed to represent regions of coarse scale variation only. Its only drawback is that it is not translation invariant. Translations of the original signal lead to different wavelet coefficients. In order to overcome this and to get more complete characteristics of the analyzed signal, the undecimated wavelet transform was proposed. It carries out the full transform without decimating the signal. Thus, it produces more precise information for the frequency localization. From the computational point of view the undecimated wavelet transform has larger storage space requirements and involves more computations.

There are two classical algrorithms for computing the undecimated wavelet transform exist – “algorithm a trous” and Beylkin's algorithm. These algorithms approach the problem from different directions. Another class of undecimated algorithms has been discovered in a search for completely different characteristics. We have constructed three new undecimated algorithms and tested their performance for image de-noising. The experimental results have shown that these algorithms have better performance in the terms of noise the removal/image blurring ratio.

Mesh Quality Analysis

Matthew Haddox
University of the Pacific

Summary

Mesh generation plays a vital role in physics simulations. The quality of the mesh strongly correlates to the quality of the simulation: A problematic mesh can lead to inaccurate answers in simulations. It is therefore worth investigating the mesh quality before simulation, to avoid wasting time and resources on a simulation only to find inaccurate results – or worse, not to notice such results.

Three different classifications of algorithms were used. Individual element metrics from the Verdict library were incorporated, for both two- and three-dimensional meshes. Such metrics include aspect ratio, skew, taper, volume uniformity, area uniformity, stretch, diagonal ratio, oddy, condition, jacobian, scaled jacobian, shear, shape, relative size, shape and size, and others. Topological checks were also used. Measurements of node degree, and Eulerian conformal test algorithms were coded and incorporated. Interference checks were also developed to ensure the mesh was not tangled and elements did not overlap. Such tests were incorporated into the Visit visualization project, and the Draco mesh generation project. In such programs, elements of poor quality can be assessed, and then visually identified.

An investigation of poor metric correlations is being considered, and may be performed in the future. The current mesh quality tests will be used in multiple simulation projects to prescreen meshes before simulation.

Limited-view X-ray Tomography with the Stochastic Engine

Keith Henderson
Purdue University

Summary

HADES is a high-fidelity radiographic simulation code that can handle both mesh files and combinations of simple volumes such as spheres, ellipsoids, and cylinders. The Stochastic Engine is a code that works with a forward simulation code and experimental data to perform inverse reconstructions. It uses powerful statistical techniques to "guide" a random walk through the parameter space specified by the user. Together, HADES and the Stochastic Engine grant more powerful tomographic capabilities than traditional techniques, especially when the number of experimental radiographs is limited by experimental conditions.

A full Python installation, including several extension packages, was installed on the local Linux machine. The Stochastic Engine (a set of Python modules) statistical optimizer was also installed. Python modules were developed to vary input parameters to the forward model. The radiographic simulation code HADES was ported to the system and installed as the forward model into the Engine.

Concurrently, several metrics (e.g., Euclidean distance) for image comparison were developed and tested. The metrics are available to be incorporated into the Engine as user-selectable Likelihood Functions.

The combined simulation-optimizer was successfully tested with a single sphere, an ellipse-sphere Boolean object, and a helical post. Each object had at least two model parameters that were examined during the optimization process, and which were accurately reconstructed.

The integrated package will be used to analyze several experiments in which only two or three radiographs were taken. It is hoped that some expertise can be gained using the engine to estimate variances in various radiographic parameters, so that the corresponding image analysis can be made more quantitative.

An Adaptive Grid Method for an Ocean General Circulation Model

Aaron Herrnstein
UCDavis

Summary

The continuously increasing rate of CO2 emissions into the atmosphere is a growing concern of climatologists. It is hypothesized that an excess of such gases will produce a greenhouse effect thus increasing global temperatures and ultimately distorting global climate. As an alternative to releasing greenhouse gases into the atmosphere, it has been proposed to deposit them in the ocean. Major concerns of such "carbon sequestration" are: (1) the amount of CO2 that escapes back into the atmosphere and (2) any biological impacts brought about by changes in oceanic pH levels.

Ocean General Circulation Models (OGDM) are used to model sequestered carbon over a time scale of centuries. The OGCM used by LLNL's Climate and Carbon Cycle Modeling group shows discrepancies in pH changes between fine and coarse grids. An even finer grid is needed to determine convergence, but refining the entire grid will be quite costly in terms of computer time (on the order of months). Adaptive Mesh Refinement (AMR) allows desired sections of the grid to be refined, thus reducing run time considerably.

It is the goal of this research to produce an OGCM that uses AMR. Such a tool will be very useful for areas such as carbon sequestration. For implementing AMR, the LLNL-developed software package SAMRAI (Structured Adaptive Mesh Refinement Application Infrastructure) is used. Refinement of the grid on which the tracer field is hosted is based on gradients and possibly other criteria as needed. In addition, regions of critical topography may be refined if necessary.

Current tools developed are time integrators common to OGCMs but uncommon to AMR. The numerics implemented by these integrators include leapfrog, predictor-corrector, and Runge-Kutta. Tests were formed in 2D on a convection/diffusion model and on a shallow water model. Data for both cases was defined at cell centers. Work for the immediate future involves modifying the integrators to advance data on a staggered grid. Such grids are more typical of OGCMs which define momentum at nodes and all other quantities (pressure, temperature, salinity, etc.) at cell centers.

Surface Patch Trimming for Isosurface Visualization

W. Taylor Holliday
UC Davis

Summary

Very large time-varying datasets, such as the ASCI Turbulence PPM Simulations, will require temporally aware compression algorithms to facilitate interactive isosurface visualization. This may be accomplished, in part, by re-parameterizing the isosurface mesh to a form more compatible with compression and interactive display: a multiresolution mesh structure. The final steps in this process are the fitting of a subdivision surface to the re-parameterized mesh, and the compression of the surface. Surface patch trimming was explored as a way to “cut” into an isosurface, selectively revealing hidden features.

The fitting algorithm is based on the quasi-interpolation method of M. Duchaineau, which is simple and local. The algorithm has been extended to adaptively sacrifice locality for accuracy (a user-specified error tolerance is met) at lower levels of resolution, where improving the detail coefficients is computationally inexpensive.

The patch trimming algorithm is a variant of Warnock's algorithm, extended for view-dependant multiresolution trim curves. A toy implementation of the trimming tool was implemented for two-dimensional domains. A compression algorithm remains to be implemented.

The fitting algorithm should be tested with CIPIC's (http://graphics.cs.ucdavis.edu) re-meshing pipeline, and on more realistic data. A comparison with the wavelet approach of N. Litke should be made and filtering of the surface data should be explored. Creating temporally coherent parameterizations is also a topic of future research. Lastly, an interface for interactive trimming of surfaces should be implemented.

PMESH - CommData Checks

Bryan Hunter
Allegheny College

Summary

The goal of this project is to provide certain checks in the communication data output in the Kull-Lite data structure. There are certain problems with this aspect of Kull-Lite when running serial, on one processor, with multiple domains. The goal was to take the output and provide several checks on the data to locate where the possible error is present. This would give some insight on to where the problem would be located and hopefully aid in debugging.

The first task at hand was to understand the parallel programming aspect of the code and how to communicate between processors using nonblocking communication with the Message Passing Interface (MPI) in C++. Once a reasonable understanding was obtained of how that process worked, a study of the Mesh design and its representation in the code was needed. To accomplish this objective we used MPI in C++ for the parallel check, and were also capable of running these checks in serial. The first step is to verify that all sent and received data on the different domains matched in content, which consists of checking all the zones, faces, and nodes. The next priority is to check each domain to reassure that it only had data that it should own – no more and no less. Next, we confirm that all the data to be sent is on the domain boundary, adjacent to the data to be received. The final check is similar to the previous, except from the receive data side, checking to make sure all its data is on the domain boundary, adjacent to send data.

From some trial tests, this procedure has shown progress in finding regions of code where communication bugs occur. These checks will continue to be used and will provide a guide status report for when the data is created, displaying any possible errors located within it. The PMESH team has already found and fixed an error that was revealed by the checks program implemented as described above.

4D Compression

Lorenzo Ibarria
Georgia Institute of Technology

Summary

Contemporary scientific simulations frequently produce large datasets consisting of voxels, representing a 4D volumetric set of values. Such datasets may consume several terabytes of space, and researchers need to inspect them, but usually in practice in only a couple of senses: some high-level views through time, and small sections through time. To facilitate these common modes of interacting with terabyte datasets, we propose a compression method that reduces the size of the dataset, and allows small sections of the dataset to be viewed through time slices.

To address these requirements three different approaches have been tried. All are extensible to 4D, to take advantage of the coherence in time. We tried wavelets, interpolation, and an extrapolating predictor. Of the different options present in the techniques, the predictors work better in the L-infinity metric, while wavelets outperform in L-2 metric. These techniques are complementary, for instance, wavelets do not allow for a small subset of the data to be decompressed without requiring traversing the whole dataset; however, with the extrapolating predictor this can be avoided. The best method we found to compress the corrections of the predictor was a geometric context arithmetic encoder, which is a set of arithmetic encoders indexed by the previous prediction in the dataset.

We plan to extend the scheme to multiple processors, something not easily accomplished because of the propagation of the error in the predictions, and the requirements of the arithmetic encoder. This would allow for a fast compression, and a guided decompression that chooses only the important parts of the dataset, not the whole. This advance is motivated by the new prevalence of multiprocessor machines, and the perspective that there is a lot to gain.

Algebraic Multigrid Methods for Contact Problems in Linear Elasticity

Ana Iontcheva
University of California, Davis

Summary

The reliable simulation of contact problems is of great importance in many applications. In virtually every structural and mechanical system there exists a situation in which one deformable body comes in contact with another. Signorini's problem, describing the contact of a linearly elastic body with a rigid frictionless foundation, is a classical problem and a basis for many generalizations. The construction of fast and reliable solvers is a challenging task due to the intrinsic nonlinearity of the problem – prior to the application of loads to a body the actual contact surface is unknown so a free boundary problem is obtained. Most of the methods derived for Signorini's problem assume that the body that comes in contact has a regular shape so can be approximated by a structured mesh. The problem that we are solving is a Signorini's problem with a body with complex geometry approximated by an unstructured mesh.

The finite element method is used for the discretization of the problem, so the element matrices are explicitly available. We need a method of optimal order for the solution of the discretized problem, namely multigrid methods. The standard geometric mutigrid can be applied only to problems with structured grid, so we are interested in an algebraic multigrid approach. Possessing the element matrices, we can use AMGe (element based algebraic multigrid), but we need to take into account specific difficulties of Signorini's problem – the existence of a contact surface. A special coarsening away from it is appropriate.

Summarizing, we are creating and applying a nonlinear element based algebraic multigrid method with special coarsening away from the contact boundary. The code for the Projected Block Gauss Seidel Relaxation for Signorini's problem and the code for the coarse grid solution procedure – based on Dostal's algorithm – has been developed.

An extension to the visualization tool GlVis for the visualization of the grids created by the coarsening procedure in AMGe has been accomplished in both 2D and 3D. The next step is the incorporation of these codes in the full code for the nonlinear AMGe for Signorini's problem. Our future plans include solving the two-body contact problem.

Accelerating Volume Rendering by Cache Optimization for Graphics Hardware

Ming Jiang
Ohio State University

Summary

The motivation for this work stems from the need to perform volume rendering of unstructured datasets at interactive frame rates. This level of performance is essential for exploring and visualizing large-scale scientific datasets. In order to improve the performance of the volume renderer, we propose a scheme that allows for more efficient usage of the graphics hardware, In particular, we introduce a preprocessing stage to the volume renderer in which the unstructured dataset is first sorted into Hilbert order and then compressed by eliminating duplicate vertices.

At the current stage of development, an unstructured dataset is first sampled, slice-wise, along the X, Y, and Z directions. For each slice, the sampled data points are triangulated using a Marching Cubes algorithm. Depending on the viewing direction, the actual rendering is accomplished by compositing one of the three sets of slices. It is at this point that we introduce the preprocessing stage to improve the rendering speed.

Our goal for this project is to improve the overall rendering speed of the volume renderer by compressing the unstructured dataset and optimizing cache performance for graphics hardware. Our compression scheme is simple yet effective. For each slice, it eliminates all the duplicate vertices generated from the Marching Cubes algorithm. There are two types of duplicate vertices: numerical and theoretical. For most cases, the coordinates of duplicate vertices share the exact same numerical values. However, due to finite precision, the Marching Cubes algorithm can produce duplicate vertices whose coordinates may not match numerically. In order to efficiently eliminate the duplicated vertices without 0(N••2) comparisons, we utilize a 2D hashing scheme in which only vertices within the same hash bin are compared. To ensure that theoretical duplicates are eliminated, we also compare vertices in the immediate neighboring bins as well. The number of hash bins is determined adaptively for each slice, based on the density of the vertices (i.e. number of vertices divided by their spatial extent). Using this method, we were able to achieve 80-85% data reduction.

Our cache optimization scheme involves sorting the triangles on each slice based on the Hilbert order. A Hilbert curve is a space-filling curve that can be drawn from one point on a plane to another without any intersections. Essentially, this curve fills up an entire region before proceeding to the next. In a similar manner, we sort the triangles on each slice so that all the triangles within a region precede the ones in the following region. Using the Hilbert order, we achieve near optimal spatial coherence for the triangles on each slice, which results in near optimal cache utilization for the graphics hardware. In this case, we conducted our experiments on NVIDIA GeForce3 graphics cards.

Currently, all the preprocessing steps are implemented with respect to each slice, which is necessary for the compositing step. However, since the number of vertices per slices can vary drastically, the preprocessing and rendering times can vary drastically as well. To address this issue, we are considering dividing each slice into blocks of fixed size and render the preprocessed blocks instead.

Parallel Algebraic Multigrid Method for Finite Element Problems Based on Domain Decomposition

Tzanio Kolev
Texas A&M University

Summary

The goal of this project was to develop an efficient parallel solver for finite element problems posed on large unstructured grids. Such problems arise naturally in simulations, where only a very fine discretization of the domain is can resolve the physics of interest. In order to achieve performance comparable with multilevel methods for the geometrically refined case, one can use an algebraic method based on sequence of coarsened meshes. Our objective was to develop a parallelization of one such algorithm, namely, the agglomeration-based algebraic multigrid for finite element problems (AMGe).

Our method starts with a partitioning of the original domain into subdomains with a generally unstructured finite element mesh on each subdomain. The agglomeration-based AMGe is then applied independently in each subdomain. It needs access to the local stiffness matrices which are reconstructed after coarsening by the variational principle. Note that even if one starts with a conforming fine grid, independent coarsening generally leads to non-matching grids on the coarser levels. We use an element-based dual basis mortar finite element method to set up global problems on each level. Since AMGe produces abstract elements and faces defined as lists of nodes, the mortar multiplier spaces are also constructed in a purely algebraic way. This construction requires inversion of the local mass matrices on each interface boundary shared between two subdomains. This is possible because of the way AMGe agglomerates the faces. This completes the (non-nested) spaces.

The algorithm was implemented in a general, object-oriented MPI code that uses parts of the HYPRE preconditioning library. Specifically, the program constructs the stiffness matrix A and the mortar interpolation matrix P on all levels and stores them in the ParCSR parallel format. The global matrix for the mortar method is P^T A P and is computed using a “RAP” procedure from HYPRE. We need explicitly the entries of this matrix, since we use parallel Gauss-Seidel with processors coloring as a smoother on each level of the multigrid. Our experience shows that this is a good choice, because its implementation is independent of the spectrum of the matrix. A software framework was developed, where the geometric information provided by a mesh generator is converted to an algebraic one by a problem generator. This is in turn read by the solver, which is independent of the coordinates, the dimension, and the type of finite element basis functions used. Currently, we have problem generators for model two-dimensional and general three-dimensional problems. They run in parallel, refining independently in each subdomain and allowing for general geometry including non-matching fine grids in three dimensions.

A number of tests were performed in order to investigate the properties of the method. Results were obtained on ASCI Blue Pacific running jobs with more than 500 processors.

Tzanio Kolev
Texas A&M University

Summary (continued)

The general observation is that the method is reasonably scalable when the number of processors is increased, while the size of the problem in each processor is kept constant. The setup cost (as with many algebraic methods) is high, but remains bounded as we use more processors. The solution time also scales well if we ignore the time for the exact solve on the coarsest level of the multigrid. The latter starts to dominate for large number of processors and dealing with it through some type of processor agglomeration is a topic of future research.

Another project I was involved in was the numerical testing of negative-norm based least-squares algorithms for div-curl systems. Two-dimensional computations, demonstrating a uniform convergence rate for a method involving finite element space enhanced with “face-bubble” functions were performed. This study is motivated by our interest in robust solvers for the Maxwell equations and is still in its initial stage.

A Discrete Differential Forms Framework for Wave Equations

Joseph M. Koning
UC Davis

Summary

The discrete differential forms framework for wave equations defines a mimetic finite element method for discretizing three-dimensional scalar and vector wave equations.

Scalar and vector basis functions are used to define the discrete differential forms. The method is valid on unstructured grids and conserves all relevant physical quantities such as energy, charge, momentum, and mass. For the lowest order basis functions, the method is second-order accurate in space and time.

Full-wave, parallel simulations in the fields of electrodynamics, linear elasticity, linear acoustics and linear magnetohydrodynamics have been completed. Currently, the research is focused on simulating optical structures such as optical fibers and photonic band gap devices. One of the major accomplishments of this framework is a mimetic discretization of the electrodynamic Helmholtz equation that shifts all zero eigenvalues to calculate the extremal eigenvalues.

In the future the method will be enhanced with higher order basis functions (in progress), h/p-adaptivity and complex materials.

Generating Library-Specific Optimizing Preprocessors using ROSE

Markus Kowarschik
University of Erlangen-Nuremberg

Summary

Large application codes in scientific computing are usually based on the use of libraries, which provide a variety of data structures and functions. In the case of object-oriented applications, these libraries implement high-level abstractions such as array classes, grid classes, particle classes, etc. Unfortunately, the use of high-level abstractions, which are defined in underlying libraries, cannot be optimized by any standard compiler since the semantics of these abstractions are user-defined and therefore unknown to the compiler. This lack of knowledge is a major reason for the poor efficiency of the majority of object-oriented scientific codes: the Mflop/s rates which can be measured at runtime are just tiny fractions of the theoretically available peak performances that the vendors claim for their machines.

ROSE is a software infrastructure for generating library-specific preprocessors which perform source-to-source transformations; e.g., eliminating the need for creating temporary objects, fusing and blocking loops, and introducing several other kinds of cache-based transformations into numerically intensive application codes. Internally, a preprocessor (which has been built using ROSE) parses the original code (currently C++ code) and assembles the corresponding abstract syntax tree (AST). Automatically generated tree traversal routines based on inherited and synthesized attributes are then employed in order to both recognize library-specific high-level abstractions and introduce transformations by replacing old AST fragments by new ones. These new AST fragments are generated by building context-specific code strings and then passing them to the compiler front-end. After the AST has been modified according to the transformations specified by the library writer(s), it is unparsed. This final step yields the optimized C++ source code.

The work during this summer has focused on two issues. The first task was the handling of preprocessor directives such as comments, “include” directives, etc., and their introduction into the generated C++ source code. The problem with these directives is based on the fact that they are not part of the C++ language itself but processed by the standard C++ preprocessor before the actual C++ compiler front-end is called. The second task was to implement the generation of C++ source code which performs cache-optimized array assignment statement transformations for the A++/P++ preprocessor. This preprocessor is used to enhance the performance of applications based on the A++/P++ array library that was also developed at CASC.

Extending the KULL Automated Testing System

Dedaimia Kozovsky
University of Wisconson-Madison

Summary

The Automated Testing System (ATS) is a tool for testing KULL executables. The ATS runs KULL job scripts, then compares the results to previous runs or analytical data to determine if an executable has passed or failed. However, determining correctness is a challenge for two reasons: (1) job scripts do not have a standard type of output or output format – users can have scripts handle output in almost any way they want and (2) there is no single standard for determining when output is "correct". New output is generally compared to some kind of reference data, but this can be done with many different comparison routines – including having a human manually review results. A main focus of development of the ATS this summer was designing a system to overcome these two difficulties. I also added a few other features including a GUI and a new method for storing results and references in the ATS database.

Rather than try to restrict "test" job scripts to certain types of output and predefined comparison routines, the ATS was designed to be extremely flexible by using several libraries:

1) The AtsOutputTypes library defines all types of output that the ATS will accept from a job script. All built-in Python types are in the library, and users can easily add new types.
2) The AtsReturns library defines functions to return output from a test script to the ATS. Output is serialized using Python's pickle module, dumped to standard out, then unpickled by the ATS. Since both the ATS and the test script can access the AtsOutputTypes library, any object type defined in the library can be returned successfully.
3) A third library, AtsComparisons, contains comparison routines that can be accessed either within the job script or by using the "review" command in the ATS. New comparisons can be added to the library at any time. The ATS also allows users to interactively choose comparison routines and set arguments to the comparison.

In addition to the flexibility of these libraries, users can easily change what reference data job results are compared to by adding new references either from a run result or from a file.

Put together, all of these features make the ATS is extremely flexible and extendible. It can handle literally any type of output, as long as the type has been registered in the AtsOutputTypes library, and perform any comparison, as long as the comparison has been added to the AtsComparisons library. Although this amount of flexibility might be overkill in many testing systems, it is essential for thorough testing of a package as complex as KULL.

Future plans for the ATS include: (1) adding more routines to the comparisons library, which is somewhat sparse at the moment, (2) changing the run manager to allow jobs to be run on the batch system as well as interactively, and (3) adding an interface to Perforce to obtain version information on test scripts and executables.

Time-parameterized Contour Trees

Ajith Mascarenhas

University of North Carolina at Chapel Hill

Summary

A common and useful technique for visualizing simulation data is to compute iso-surfaces. Contours of elevations found in topographic maps, where points on the same height on a terrain are shown as curves, are a familiar example of iso-surfaces. We are interested in computing iso-surfaces for volumetric data that is time varying. Such datasets are produced by simulations of physical phenomenon like fluid-flow, mixing of liquids like oil and water, and so forth.

Some questions that come up during visualization of iso-surfaces are how many components are present in the iso-surface? What components merge, split, appear or disappear if one changes the iso-value by a small amount? A contour tree is a useful structure that can be used to answer these questions. We are interested in computing a variant of the contour tree structure that can be applied to time-varying volume data.

Our technique is based on a topological approach to the problem of constructing time-parameterized contour tree. Given the fact that the nodes of a contour tree can be mapped to the critical points - maxima, minima, saddles - of the data, we aim to compute the trace of critical points as a function of time. We call this trace the “gamma” curve. Given the contour tree for time t=0, and the “gamma” curve we aim to compute the contour tree for any subsequent time t=t'.

During the summer we have implemented the construction of “gamma” for a data sampled on a uniform grid. We use a canonical connectivity scheme wherein each grid cell is split into simplices about the main diagonal. We have computed gamma for a few small datasets of size 64x64x64x20 (i.e., a cube of size 64 on a side through 20 timesteps) and plan to work of larger data sets.

In continuation to the work done this summer we hope to develop the theory and algorithms for constructing time-parameterized contour trees. This work will require significant effort during the coming year and next summer.

Detecting and Exploiting Spatial Regularity in Memory References

Tushar Mohan
University of Utah

Summary

Rising processor speed, unaccompanied by corresponding reductions in memory access latency, has caused the performance of codes to be limited by memory accesses. Strided memory accesses, or streams, can be a significant source of memory stalls in loops in practical applications, if undetected. If known to exist in an application, they can be targeted by a host of optimizations, such as stream prefetching, relocation, remapping and vector loads.

Existing stream detection mechanisms either require special hardware, which may not gather enough stream statistics for subsequent analysis, or are confined to limited compile-time detection of array accesses in loops. Formally, little treatment has been accorded to the subject; the concept of locality fails to capture the existence of streams in a program's memory accesses.

In this project we define spatial regularity as a means to depict the presence of strided memory accesses. We develop measures to quantify spatial regularity, and design and implement an on-line, parallel algorithm to detect streams, and hence regularity, in running applications. We identify critical program sections for regularity measurements by using PAPI – a performance measurement API – to access hardware performance counters portably. Dyninst's dynamic binary translation infrastructure is leveraged to perform selective and transitory instrumentation in the application. This allows the user to limit the stream detection overhead at the cost of measurement accuracy. We use examples from real codes and popular benchmarks to illustrate how stream information can be used to effect profile-driven optimizations.

MR-File: A New Multi-Resolution Spatial Index for High-Dimensional Scientific data

Mohamed F. Mokbel
Purdue University

Summary

Existing multi-resolution models for large-scale simulation data can support a wide range of spatial range queries. However, queries about non-spatial variables are not necessarily efficient. Non-spatial variables are stored in the multi-resolution model in a statistical form, where the minimum, maximum, mean, and standard deviation for each variable are stored. The main objective is to build an indexing scheme on top of the multi-resolution model described in to efficiently answer non-spatial queries while keeping the efficiency of the spatial queries.

Our approach is to map the problem into a spatial access method (SAM). The non-spatial v variables in the multi-resolution model are mapped into non-zero sized hyper-rectangles in the v-dimensional space. The state of the art software for spatial indexing, GiST, is used. However, results from GiST are completely unsatisfactory. GiST does not scale with the massive size and the high-dimensionality of scientific data, a phenomenon known as the curse of dimensionality.

Reviewing the literature of spatial indexing methods, it is clear that the curse of dimensionality has not been investigated for spatial access methods. A similar but easier problem, which is indexing zero-sized objects in the high-dimensional space, is well investigated with some solutions for the curse of dimensionality. In this research, we aim to solve the curse of dimensionality for spatial access methods. As a result, a new spatial indexing scheme, termed the Multi-Resolution file, MR-File for short, is proposed. MR-File is scalable in terms of dimensionality and data size.

Parallel Optimizations of Parallel Algorithms

Evan Moran-Bernard
Carnegie Mellon University

Summary

Parallel computing normally calls to mind hundreds of processors working away at some big problem: modeling a galaxy of stars under gravitation, modeling the weather over Kansas, calculating pi to fantastic precision, or analyzing trends in the stock market. Where feasible, parallel computing delivers these results much more rapidly than serial. Why, then, don't all personal computers have several cheap processors instead of one expensive one? Why can't one's palm pilot link to every palm in a room, rather than just one at a time? In general writing parallel programs is more difficult, and takes more time. Our main interest in this project is to optimize parallel code: to make it be "more parallel."

Making programs "more parallel" involves several levels of parallel-algorithm design. Within the file system I developed parallel algorithms to sort and parse enormous files, and to distribute the data to an arbitrary number of processors. To handle per-processor balancing issues I used several randomizing techniques, which appeared to be successful and scalable. Several constraints occurred at the per-system memory level, where it was necessary to keep a minimal amount of data in memory during the sorting process. Specifically, I used an Oct-Tree data structure, well known in graphics research for improving rendering performance, to spatially sort files and to improve locality in the parallel algorithms. At the distributed systems level I used MPI, the Message Passing Interface, to allow inter-process communication. This, in particular, became a conceptual hurdle because of the inherent difficulties with deadlock and test case reproduction.

Additional efficiency in programming parallel systems may reside in more automated time analysis techniques used to analyze time performance. This information would be useful for finding performance bottlenecks in working code. Analysis also informs the programmer of how close to the theoretically best performance his or her program has achieved. This is perhaps not as important to performance as adding more processors, but the potential is definitely interesting.

Parallel AMG Performance Issues

Arne Naegel
University of Heidelberg

Summary

Driven by the need to solve linear systems arising from problems posed on extremely large, unstructured grids, algebraic multigrid (AMG) methods have been recognized as an efficient way to treat sparse systems of the form Ax=b. The advent of high performance computers with large numbers of processors sparked interest in parallel versions of AMG that are capable of employing the available sheer computational power to solve large systems of equations in a robust fashion. Although multigrid methods in general are known for optimal (i.e., O(n)) performance in the solution phase and AMG particularly shows fast convergence rates for a wide variety of problems, in general, results are still sometimes dissatisfying. The reason is that the non-geometric approach — due to the lack of knowledge about an underlying grid structure — requires a setup phase that is often remarkably expensive, and its performance usually degrades further in a parallel environment.

Abstracting the implementation of the BoomerAMG, the parallel AMG solver integrated in the software package hypre (High Performance Preconditioners), we derived a parallel computing model for the setup phase and described the performance of its three major components in terms of the degree of the related matrix graph, the number of unknowns, and of the processors involved. As it involves critical message passing components, one key part was the analysis of one of the coarsening strategies used. Additionally, several numerical tests were run and its results were compared to the developed theory.

The aim of this work was to analyze the behavior of the AMG setup phase in a (massively) parallel environment on a distributed-memory architecture like ASCI Blue Pacific at LLNL, and to gain insight as to those stages of AMG code / algorithm where improvements can be achieved. The close collaboration with Rob Falgout and the Scalable Linear Solvers Group at CASC assures that the results will lead to setup phase improvements in the future.

Regression Testing for Lustre's Distributed Locking Subsystem

James Newsome
University of Michigan

Summary

Lustre is a high performance distributed file system being developed by Cluster File Systems. The project is being assisted by Lawrence Livermore, Sandia, and Los Alamos labs so that it may be used in the next generation of supercomputing clusters.

One of the innovations in the Lustre file system is a distributed lock manager. Rather than having a central server that each node must contact to lock resources, each node in the cluster is capable of managing resources. This has two distinct advantages: First, it eliminates the lock server as a single point of failure. Second, by allowing the node that uses a particular resource the most to manage that resource, we eliminate the need for many lock requests to be sent over the network.

Thorough testing is extremely important in any software project. It is especially important in a system as complex as Lustre's distributed lock manager. My role in the project was to develop a regression test that would verify that the locking subsystem works correctly under a large volume of requests.

One goal of the test was to make it as independent as possible from other parts of Lustre. If the test made heavy use of other parts of Lustre, it would be more difficult to find the source of a failure.

Another goal was to balance testing as much functionality as possible with making it easy to identify problem areas. In a test that tests everything, passing the test assures us that everything is working correctly. However, failing such a test does not provide much information in itself, since it could be any part of the locking subsystem.

In order to isolate the test from other parts of Lustre, it was put inside the Linux kernel. This is currently the only way to directly access the locking subsystem. Another option would have been to have a user-space test work with filesystem objects, putting various resources under contention. However, this would defeat the goal of isolating the test from the rest of Lustre. The test would be dependent on the whole system rather than just the locking subsystem.

The test exercises the following functions of the locking subsystem:
– Enqueueing an extent lock
– Matching an extent lock
– Using completion callbacks to keep track of locks
– Using blocking callbacks to cancel locks

The test first starts some number of threads. Each thread then runs in a loop. In each iteration of the loop, each thread randomly either tries to match/enqueue a lock on a

James Newsome
University of Michigan

Summary (continued)

random extent of a random resource, or decrement one of the lock references it already holds. The test runs indefinitely until a user stops it. Short runs of the tests, on the order of several hundred enqueues and matches, pass. Unfortunately, I have observed some strange behavior when the test is run longer. It is currently unclear whether this is caused by a bug in the lock manager, or a bug in the test itself.

For the future, the first thing that needs to be done is to find what is causing the test to fail on long runs. Once this is done, there are several enhancements that could be added to the test.

First, the test should also test lock conversions. The code for this is partially written, but I left it disabled for now in favor of getting the simpler version of the test to work. Once this is added, the test will be able to verify the functionality of all the basic parts of the distributed lock manager.

The next step after that would be to make the test more distributed. Currently it is designed to run only on a single node. A more realistic and thorough test would be to have several nodes all trying to work with the same set of resources.

After these features are added, I believe the test would be complete. The only other thing to add to the test would be boundary conditions, such as cancelling a lock that hasn't been granted yet, or having a node fail during the test.

Multigrid for Linear Hyperbolic PDEs

Luke Olson
University Of Colorado at Boulder

Summary

Multigrid is firmly established as the method of choice for systems of elliptic and parabolic partial differential equations and the Laplacian-type terms that dominate the ill-conditioning of more general systems of PDEs. Many practical codes have been deployed in ASCI and other applications based on multigrid methods, leading to increased desires to extend multigrid to additional regimes. The First-Order System Least-Squares formulation extends multigrid theory to many such systems. However, the purely hyperbolic case has received relatively little attention (or successful attention) from the multigrid community.

In this project we consider First-Order System Least-Squares formulations for purely convective partial differential equations and analyze the multigrid perfor-mance for solving the linear system arising from a finite element discretization. Prior research exposed poor multigrid convergence results under certain boundary conditions with Algebraic Multigrid (AMG). Local Mode Analysis (LMA) was applied to the resulting finite element stencil and we explored a geometric multilevel framework to further reveal the unexplained inefficiencies in the solver. The LMA predictions of relaxation and two-grid performance were then compared to the actual implementation of the Geometric Multigrid algorithm.

Through our analysis and numerical tests, we found both poor smoothing properties and unsatisfactory coarse-grid approximations to cause the degraded performance. Based on these results, we proposed adjustments on improving the smoothing and coarse-grid components of the multigrid algorithm. Still, these new algorithms need further investigation. This research will be continued at the University of Colorado, where we intend to address issues such as scalability and robustness of the method.

A Hierarchical Shrink-Wrapping Approach to Surface Reparameterization

Serban Porembescu
UC Davis

Summary

Our goal is to develop a robust, efficient, and easy to implement surface reparameterization technique. Parameterized surfaces are amenable, among other things, to Continuous Level of Detail Control, Multiresolution Editing, Texture Mapping, and most importantly, Compression.

Our technique begins with an arbitrarily triangulated surface and constructs a hierarchy of fine resolution to coarse resolution surfaces. As the hierarchy is constructed, mappings between adjacent resolutions are maintained.

We reparameterize the surface by applying, in combination, a set of atomic operations. These operators are Subdivide, Smooth, and Snap. The Subdivide operator refines the base mesh by performing bilinear subdivision. This gives our surface a quadrilateral grid-like structure. The Smooth operator minimizes the mesh distortion by trying to make all quadrilaterals equally sized. Using the Snap operator we can project the base mesh onto any level of the hierarchy.

Given these atomic operators we construct a set of Super operators: Smooth-Snap and Subdivide-Snap. The Subdivide-Snap introduces new vertices in the base mesh and projects these vertices onto the current hierarchy level. The Smooth-Snap reduces distortion by pulling vertices off of the current level of the hierarchy and then projecting the modified vertices back onto the current hierarchy level. The end result is a reparameterization of the original surface.

To make this algorithm truly useful for large-scale data visualization we need to make the algorithm work out-of-core. Extremely large data sets are being generated by improved resolution scanners, sensors, and computer simulations. Often this data is too large to fit in the main memory of a high-end workstation and in some cases even in the main memory of the supercomputer that generates the data. The current algorithm has been designed to generalize to out-of-core implementation.

High Order Discrete Differential Forms for Computational Electromagnetics

Robert N. Rieben
UC Davis

Summary

We are concerned with the finite element solution of Maxwell's equations on unstructured hexahedral grids. We focus our attention on the high-order discretization of this problem, specifically the recently proposed differential forms based approach for constructing curl-conforming and divergence-conforming vector bases. In the computational electromagnetics (CEM) community, it is well known that in order to accurately perform electrically large simulations (i.e., simulations whose characteristic dimension is many wavelengths long), high-order methods must be used. Low-order finite difference and finite element methods are ineffective for such problems due to numerical dispersion. Recently, several advancements have been made in the field of (what are presently called) Vector Finite Element Methods. While it is known that in general higher-order methods have less dispersion and should be superior to low-order methods, this has not yet been demonstrated for Vector Finite Element solutions of Maxwell's equations. In addition, there are few high-order implementations of this method; and of these, none present a general technique for the construction of hierarchical bases, which are necessary for any sort of p-refinement strategy. In conjunction with my thesis project, my research this summer has been the development and implementation of a method for solving Maxwell's equations using Discrete Differential Forms (DDF's) of arbitrary order on unstructured hexahedral grids.

In the language of differential forms, we can unify several existing finite element methods using a clear and concise notation. Traditional finite elements (scalar nodal) as well as vector finite elements (edge and face elements) and discontinuous scalar elements are combined into one set of discrete differential forms. We begin by constructing DDF's of arbitrary order on a reference (or unit) hexahedron. We then define a set of linear functionals (referred to as Degrees of Freedom) that map the DDF basis functions to real numbers. These are necessary for the construction of element mass and stiffness matrices. A major accomplishment of this summer was the development of a general method for constructing hierarchical vector basis functions using an appropriate set of Degrees of Freedom. A hierarchical basis is necessary for any p-refinement method. DDFs on the reference element are then transformed via conformal mappings to global mesh elements of arbitrary location, orientation and distortion.

Over the summer, the mathematical framework of this method was developed and implemented in the FEMSTER framework. FEMSTER is a modular finite element class library for solving three-dimensional problems arising in electromagnetism. The software consists of a set of abstract interfaces and concrete classes, providing a framework in which the user is able to add new schemes by reusing the existing classes or by incorporating new user-defined data types.

Future work will include the integration of the FEMSTER library into a general parallel matrix assembly program. Once this is complete, work will begin on a time stepping algorithm in conjunction with the HYPRE library for time domain simulation of various electromagnetic transmission devices such as optical fibers and Photonic Band Gap devices.

A Web Services Framework for Bioinformatics

Daniel Rocco
Georgia Institute of Technology

Summary

The transition of the World Wide Web from a paradigm of static web pages to one of dynamic web services provides new and exciting opportunities for bioinformatics with respect to data dissemination, transformation and integration. However, the rapid growth of bioinformatics services coupled with non-standardized interfaces diminish the potential that these web services offer.

For instance, the BLAST family of bioinformatics data sources allow biologists to compare DNA and protein sequences with an existing body of knowledge to find similar sequences in other organisms. BLAST provides data management and query services to genomic data sources and allows research groups to set up local repositories or specialized sequence databases. One catalog lists over 500 such data sources, some subset of which will contain data relevant to an arbitrary query.

Unfortunately, there is no common interface or data exchange mechanism for these sites. When performing a search, the scientist chooses a set of known sites, enters a query into each site, and integrates the result by hand.

The problems with this approach are numerous: the scientist may not have the most current or most relevant site for the search at hand, the search must be entered multiple times, and the results of the search must be merged together by hand to obtain an integrated set of results.

Our challenge is to provide a common interface to the vast and dynamic assortment of BLAST data sources. This problem can be divided into three sub-problems. First, how will new sources of relevant information be integrated into the interface? Next, how can the interface determine which discovered sources are relevant to a particular query? Finally, how can the interface mediate between sources and produce an intelligently integrated result set?

We propose the notion of a service class description as a solution to the problem of discovery, classification, and integration of bioinformatics sources. A service class description describes the relevant portions of the service from the point of view of the intended application. The description includes the various data types used by the service, example queries and output, and a graph representation of how service class members are expected to operate. For example, a simplified view of the DNA sequence BLAST service class includes a DNA sequence input type, a BLAST result output type, and descriptions of the intermediate pages. The control flow graph might show the input page being connected to a result page, possibly through a delay page.

The service class provides an abstract view of bioinformatics sources that allows developers to reason about the class rather than being concerned with the intricate details of each member. Using the service class description, we are creating an integrated BLAST service that will operate via a common application to an automatically discovered set of BLAST sources, each wrapped to provide a transparent interface between the application and the data source.

The Generic Simulator

Sadik Gokhan Caglar
University of San Francisco

Summary

The GenSim project is a generic simulator, that is mainly used to simulate a cluster of machines that do distributed computing using message passing libraries. GenSim is developed under C++ using object oriented techniques, for the purpose of reusability, modularity, and maintainability. The reusability principle also goes hand-in-hand with the generic nature of the project.

GenSim is a multi-layered simulator. At the very bottom level, there are seven different classes. These base classes provide the general framework on which the simulator runs. From a different perspective, these base classes can be seen as the engine of the simulator. The different layers on top of this layer inherit properly to model specific applications.

The naming scheme of these base classes might suggest that the framework is designed to run only cluster models. In fact, other environments, such as a distributed database, or a specific architecture can be modeled with the base classes, by implementing appropriate interface layers.

The project currently has two layers implemented. The first layer is used to model a programming language behavior. It provides tools like conditionals, deterministic and probabilistic loops, system messaging, creation and termination of jobs, and computation.

A system heap for the variables is not yet implemented, so the conditionals are only probabilistic for the time being. We are able to have deterministic loops, due to the internal structure of the loops. The statements (events) that are in the loop have no idea about the current state of the loop. For instance, if a given loop body executes three times, a computation event has no idea to which execution the current one belongs; only the loop conditional event is aware of it.

The second layer is a subset of MPI, the message passing library. The layer provides means to model blocking and non-blocking communication, waiting for non-blocking communication, and collective communication. Blocking communication is modeled using basic send and receive events. The events use a tag, and for every send event there is a matching receive event in the peer, and vice versa. Non-blocking communication works the same way; however a wait must be called in order to make sure that the communication has been completed. This forces a rule on our system. The loop bodies have to be sound in terms of the non-blocking events they have. In other words, they have to call the matching waits before restarting the loop. If the loop is restarted without knowing that the non-blocking communication has been terminated successfully, the action may cause the event to be discarded. Collective communication

Sadik Gokhan Caglar
University of San Francisco

Summary (continued)

in its nature is blocking, so it doesn't suffer from this possibility. Since all of the collective communication routines have the same requirements for execution (that is, all the processes in the communicator call the same collective, with the same parameters), they can be implemented in the same way. The differentiating property is the type of the collective communication, which is handled by the tags, and the amount of time it is likely to take to run.

Since in real runs the order of collective operations is not always fixed, a random number generator package, randlib.c from University of Texas is integrated into the system. The computational, collective communication and point-to-point communication events take advantage of different distribution functions of this package.

There are two front ends that the users can use to model algorithms; they can use C++ directly and link with the existing classes, or they can use the GenSim runtime environment. The latter is encouraged. The Jacobi iteration algorithm is implemented as an example how the suite can be used to model algorithms for the runtime environment.

The two layers are examples of the capability of the simulator. New classes can be inherited from the existing classes to achieve more functionality. The runtime environment has to be updated with the possible newly added clusters, events, constraints, messages, jobs or processors.

Transformations for the Cache Optimization of Applications using Object-Oriented Abstractions

Sunjeev Sikand
UC San Diego

Summary

A common problem with object-oriented C++ scientific computing is that the high level semantics of abstractions introduced (e.g., parallel array objects) are ignored by the C++ compiler.

ROSE is a programmable source-to-source transformation tool for the optimization of applications using high-level abstractions commonly found in C++ object-oriented frameworks. In our work we target the A++/P++ array class library and implement three cache-efficient optimizations. These transformations exploit both the spatial and temporal locality of memory references exhibited by stencil statements. Experiments have shown a performance improvement of 1.5 to 2.2 time beyond that of previous transformations of array statements into lower level C code associated with the stencil operations (which provided a 2 to 4 times improvement). While the C code transformations are relatively complex, the cache-based transformations are beyond the complexity that is reasonable to expect programmers to write by hand. Thus we present the case that the use of the semantics of high-level abstractions can be expected to lead to faster code than what can likely be produced by hand.

Our results are encouraging and leave the possibility open for further refinement of these transformations to achieve greater performance gains.

RVR, A Genome Comparison System

Jonathan Strasser
UC Davis

Summary

Within the human genome project, it is important to see the similarities between genomes of different species. The most common approach to do this is to perform a BLAST search of one genome against the other. There are two main problems with this approach, both of which come from the fact that the entire genomes are very large. The first problem is that the searches will take weeks or even months. The second problem is that the computer may run out of memory and crash. Our objective is to develop a system that will not have these problems, but at the same time give us the same results.

Sam Rash and Paramvir Dehal came up with the main idea behind our approach. This is to create a map of regions on one genome to regions on the other genome. In our case we wanted to map regions on Fugu Rubripes to regions on the human genome. To do this we decided to first map the IPI protein set to both Fugu rubries and human. Since this a much smaller data set, the comparisons can be completed in a much smaller time period. We were also able to make use of a program called BLAT, which performs BLAST-like searches in less time. In our case we were able to complete these maps over the course of a few days. Once both maps were complete we used the transitive property to finally map Fugu Rubripes to human. Now we can perform a BLAST search using much smaller regions rather than entire genomes.

I implemented this method through a series of scripts, written in perl, that I call the RVR (region vs. region) system. It contains eight main scripts along with four subscripts that take the user through every step of the process. From the map creation using IPI to the running and parsing of region BLAST searches. The RVR system takes advantage of Sam Rash's job distribution program and is able to split the process into thousands of jobs and send them to the JGI servers.

After using this system several times, I have determined that this process can be completed in around a week using the RVR system. The results of the region BLAST were a mixed blessing. On one hand they had much less noise (bad hits that mean nothing and take up space), but on the other hand there were some hits that the RVR system had clearly missed. This can be expected, however, since regions that did not map to an IPI would have not been included in the final BLAST. Since a large amount of time and computing power is saved, the slightly sub par output is acceptable.

Now that the system is complete and tested, it will be used with future versions of Fugu Rubripes and human. It may also be used to compare other genomes as well. The RVR system will also continue to be debugged and tweaked if needed.

Multi-Processor Isosurface Calculation

Mark Stuppy
University of Missouri-Rolla

Summary

The calculation of an isosurface on a single computer is a time-consuming task. The time required to complete this task can be greatly reduced by splitting the work up amongst multiple processors. The objective of this summer's project was to implement a multi-processor isosurface calculation program and test its efficiency and scalability.

The program itself was written using many different languages that worked together to create the final output. Interpreted Python, with the pyMPI and NumPy modules, was used to set up the environments on each processor and to pass messages between processors, a combination of compiled Python and C++ code was used to do the actual isosurface calculation, and the Yorick visualization tool was used to view the isosurface.

To make the original isosurface calculation process work on multiple processors in parallel, each processor was given a small portion of the full problem space to work on. Message passing, using the pyMPI module, was used to send necessary boundary cell information to adjacent processors. The results from the calculation were then dumped to multiple files (one for each processor), which could be read into and viewed by Yorick at a later time. The resulting program ran in a time which is inversely proportional to the number of processors used in the calculation.

Future improvements include finding more efficient data structures in which to keep intermediate information, finding a less-restrictive scheme for passing multiple messages at the same time, and finding a more structured file format in which to create dump files.

Embedded Boundary Technique in Structured Adaptive Mesh Refinement CFD

Ryuta Suzuki
University of Minnesota

Summary

The embedded boundary technique is a means of extending structured mesh discretizations, including adaptively refined structured meshes, to completely general geometries, such as those encountered in computational fluid dynamics (CFD). Motivated by the current development of micro- and nano-technology in aerospace applications, we implement a Navier-Stokes solver for incompressible flow simulation. Efficient computations of this type of flow are increasingly important aids in the design of micro- and nano-scale vehicles.

We implemented the incompressible Navier-Stokes solver in an adaptive Cartesian mesh setting. In order to handle the incompressibility constraint we used the projection method. The projection method is based on the Helmholtz-Hodge decomposition theorem where the arbitrary velocity field with appropriate support and no-slip boundary condition is decomposed into an incompressible velocity field and the gradient of some scalar function. The scalar function is computed from the Poisson equation with Neumann boundary condition and it well approximates the pressure field we seek. The Poisson equation solver is the most intensive part in incompressible flow simulation. Brian Gunney in the Center for Applied Scientific Computing (CASC) implemented the Fast Adaptive Composite Grid (FAC) Method for the Poisson equation in Structured Adaptive Mesh Refinement Application Infrastructure (SAMRAI) framework. FAC achieves fast convergence of the iteration using the multigrid philosophy on composite grids.

One of the drawbacks of the Cartesian mesh is its difficulty in handling the boundaries that are not aligned on the meshes. Thus, it is important to develop the embedding boundary technique to handle more complex geometry in structured meshes. We employed the so-called “masking” technique to embed impermeable wall boundary conditions. Based on a no-slip and no-flux boundary condition, we interpolate the velocity and flux to satisfy the above two conditions. Although this is a rough approximation it works well thanks to the multi-level meshes provided by SAMRAI.

Future work will be dynamic adaptation, utilizing the SAMRAI adaptative capabilities. Devising an a posteriori error estimator for goal-oriented adaptation is also an important issue.

An Implementation of ARCHES into a Spatially Adaptive Framework for Pool Fire Simulations

Jeremy Thornock
University of Utah

Summary

The solution of a hyperbolic partial differential equation may be achieved by casting the spatially semi-discretized equation as a high-dimensional ordinary differential equation system through a method of lines strategy. Time integration may then be performed with various standard techniques. However, some methods prove unstable when discontinuities or high gradients in the field are present. This can result in overshoot, oscillations, and general instabilities. One cure for such problems lies in a class of time integrators, namely the Runge-Kutta "Strong Stability Preserving" (SSP) integrators. This type of integrator maintains the stability properties provided by a forward first-order Euler explicit integration yet increases the time accuracy without instabilities.

The main objective was to implement a Runge-Kutta SSP time integrator in a spatially adaptive framework, SAMRAI. This included validation and verification of the integrator within the framework using both adaptivity and non-adaptivity to ensure correct implementation. Upon successful completion of this task, a CFD code (ARCHES) was to be implemented into the adaptive framework using the SSP time integrator with an implicit pressure solve.

In order to simplify the solution procedure, time refinement was not used when implementing the Runge-Kutta SSP integrator. The time step was calculated as the minimum of all time steps over all levels and patches, which met with the Courant-Friedrichs-Levy (CFL) condition. This simplifies the algorithm by negating the need for flux-corrections at coarse-fine interfaces. This approach can be justified quantitatively through experience of those familiar with spatially adaptive algorithms. In general, it has been found that the bulk of the calculation in an adaptive environment is spent on the more refined levels, especially when a higher refinement ratio is used. Therefore, little efficiency is lost in stepping the coarser levels at the small time steps. No study was performed here to validate this conventional wisdom.

The algorithm was tested using Burger's equation with Riemann initial conditions and the linear advection equation using several different initial conditions. The MUSCL scheme was used to discretize the spatial elements of Burger's equation while a monotonicity-preserving scheme (MP5) was used for the linear advection problem. They were both compared with an equal order, non-SSP time integrator. The solution was also verified using the theoretical wave speed to mark the solution for refinement. In other words, the cell tagging routine was in no way connected with the state of the solution. It was found that the integrator performed well in both an adaptive and non-adaptive environment. It was also found that the solution followed the theoretical wave speed exactly, both with and without adaptivity. A write up on the subject of the time integrator was produced and can be referenced for more detailed information.

Jeremy Thornock
University of Utah

Summary (continued)

The next step involved implementing ARCHES into the SAMRAI framework using the same time integration technique, again without time refinement. This is work in progress. The Navier Stokes equations present a special challenge when dealing with a multilevel calculation due to the link between the pressure and mass preservation. It is proposed to use a canned FAC solver that is available in SAMRAI to tackle this issue.

Future plans include finishing the ARCHES/SAMRAI code and testing it with a simple, laminar helium plume case. This should help validate a successful implementation of the scalar transport equation, mixing model, pressure solver, and momentum update using the Runge-Kutta time integrator. Turbulence models are to be added later.

Using a unique formulation which is evolving through the CSAFE project at the University of Utah for handling multiphase flow, it is also proposed that research into so called “cut-cell” methods are investigated using the ARCHES/SAMRAI formulation. Current studies include using such a formulation to model reacting, multiphase mixing tanks.

Distributed Linda Tuplespace

Aaron Wegner
Baylor University

Summary

We tested a Distributed Linda Tuplespace implementation called OpenLinda with a small Monte Carlo simulation. OpenLinda proved to be a weak implementation in that scaling to use multiple machines was very inefficient. We therefore wrote a new Distributed Linda-style Tuplespace implementation with better scaling ability. We began with Postgres Database as the server, but later adopted the the more efficient MySQL database.

The Linda tuplespace technique is a simple way to do distributed computing. Shared data is stored in the tuplespace, and other computers connected to this distributed tuplespace can access it. The three main functions by which Linda is defined are:

set() - insertion of a tuple into the tuple space

get() - destructive retrieval of a tuple from the tuplespace

read() - nondestructive retrieval of a tuple from the tuplespace

My Distributed Linda Tuplespace API also includes getall() and readall(), which allocate and return large arrays of all matching tuples rather than just a single one.

Each function in the API can take any number of arguments up to 256. Types accepted are int, long int, double, char*, and BLOB. The BLOB type is primarily used for storing other types of data than the standard types, such as user created structures. The general format for each function is:

func(TSconn *conn, char *types, var, var, ...)

- where *conn is a pointer to a connection of type TSconn to the Linda Tuplespace.

- *types is a string specifying the types of the following variables, and

- var, var, ... is the list of variables to get or set.

The syntax for the specifying data types in the types string is simple. For each variable, put a '%' followed by the first letter of the desired data type. For example, to insert a tuple with types 'int, int, double, blob', the types string would be

Aaron Wegner
Baylor University

Summary (continued)

"%i%i%d%b". It is also acceptable to use the full data type, delimited with %'s. "%int %int %double %blob" would work fine. For return values, precede the type with lower case 'r'. "%i%ri%d%rb" will return and int and a blob for a tuple that matches the given int and double.

A MySQL database must be installed to use this Linda implementation. The Linda Tuplespace must have privileges to select, insert, update, delete, create, drop, and index.

The main limitation of the efficiency of this software is the database backend. Triggers and stored procedures are not supported in the current version of MySQL (v4.0). Once these are implemented, Linda could be modified to use them to run somewhat faster. Additionally, it should be fairly easy to adapt this source code to take advantage of a different, more efficient database than MySQL. One more possibility would be to write a tuplespace server from scratch that does not have the overhead of a relational database. In the future, I would like to work out some solution to get better performance out of the tuplespace server.

I would also like to add XML support to enable using the tuplespace across multiple platforms simultaneously. Data would be encapsulated into a standardized format and viewable from any machine. For functionality, it would be nice to be able to specify exactly how many tuples to grab from the tuplespace instead of just being able to get one tuple or all tuples. It also might be possible to take advantage of the high efficiency of getall() to improve the speed of get().

This new implementation of a Linda Tuplespace is a success, both in and of itself, and as a step along a development path. Since it handles an arbitrary number of connections and uses BLOB data, it is versatile. It scales well because the main limitation to scaling is the database backend used. MySQL does a good job of handling many connections, and a fairly good job with speed in relation to comparable databases. This Linda Tuplespace software would be most useful on a system that does a large amount of computation and has many connections to the tuplespace.

Building a Prototype System for the Automation of the Acquisition of Earthquake Bulletins

I-Hsuan (Sandy) Wu
UC Davis

Summary

To support ongoing research, the GNEM program maintains a large and growing database of seismic waveform and parametric data. Although we have largely automated the process of loading data into the database, acquisition of the data still requires nearly full-time attention of two people. Since much of the data are available electronically, we intend to automate a large portion of the data acquisition. As a first step, we have decided to automate the acquisition of earthquake bulletins. We have decided that the system must be distributed among several computers. However, we have not settled on an implementation language or on a technology for handling distributed objects. This summer's project was to build a prototype FTP agent that could be used by the system to automatically acquire earthquake bulletins and forward them to a processing queue. This served as an experiment to help us evaluate the implementation difficulty and efficiency of available technologies.

Several programs were written to compare the capabilities and functions available in C++ and Java for performing FTP connections. Then, we did some research into the available technologies to determine the best way to acquire a distributed system. In the end, Java was selected as the programming language to be used in order to take advantage of the Java Remote Method Invocation (RMI), which enabled distributed programming. The FTP package used is NetComponents. The resulting prototype is a stand-alone system that runs an http server to receive RMI requests from clients. The system retrieves the requested data from a remote location with FTP and returns the result to the client. By using Java, the system becomes portable. And the RMI allows data requests to be made over networks and does not require the client to reside on the same machine with the FTP agent.

Based on the work accomplished, we have decided to implement the FTP agents and their manager using Java and RMI. This will involve evolving the prototype to forward error handling to the agent manager and to process additional messages that the final system will require.

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551