Mathematical and Computational Sciences Division

Summary of Activities for Fiscal Year 2008

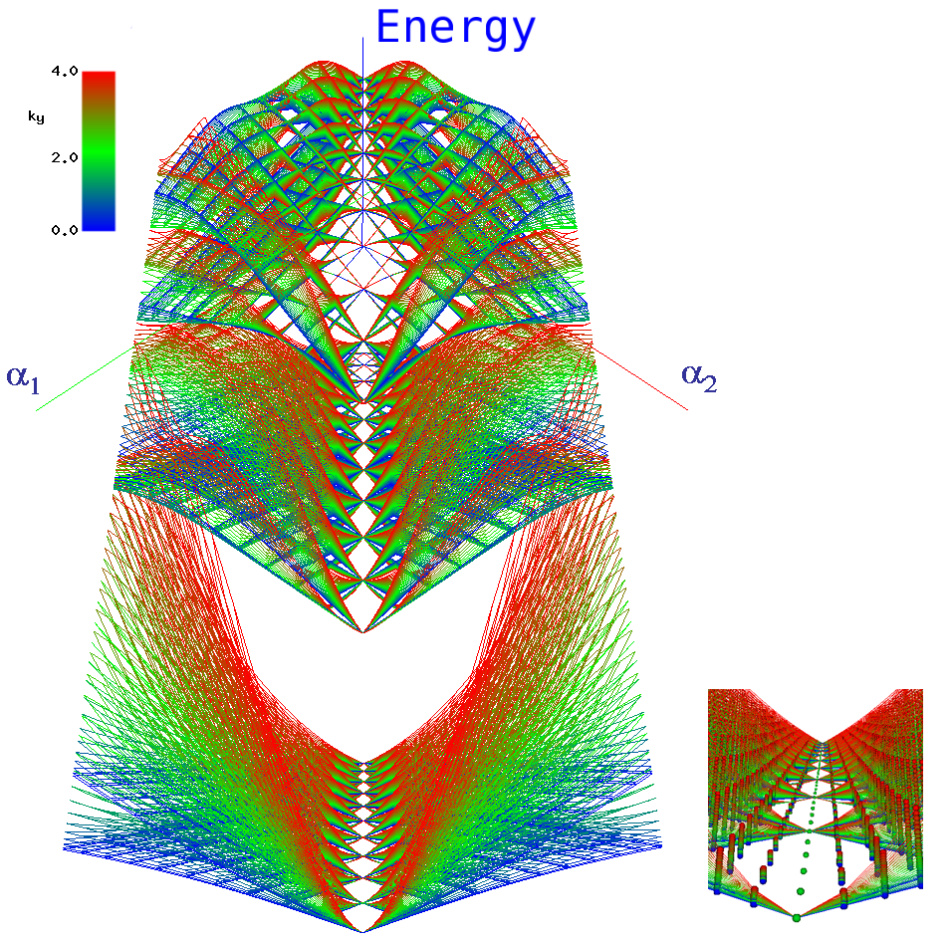

Information Technology Laboratory
National Institute of Standards and Technology
U.S. Department of Commerce

January 2009

INFORMATION TECHNOLOGY LABORATORY

NIST
National Institute of Standards and Technology
Technology Administration, U.S. Department of Commerce

Abstract

This report summarizes the technical work of the Mathematical and Computational Sciences Division (MCSD) of NIST's Information Technology Laboratory. Part I (Overview) provides a high-level overview of the Division's activities, including highlights of technical accomplishments during the previous year. Part II (Features) provides further details on ten projects of particular note this year. This is followed in Part III (Project Summaries) by brief synopses of all technical projects active during the past year. Part IV (Activity Data) provides listings of publications, technical talks, and other professional activities in which Division staff members have participated. The reporting period covered by this document is October 2007 through December 2008.

For further information, contact Ronald F. Boisvert, Mail Stop 8910, NIST, Gaithersburg, MD 20899-8910, phone 301-975-3812, email boisvert@nist.gov, or see the Division's web site at **http://math.nist.gov/mcsd/**.

Editor: Ronald F. Boisvert

Cover Visualization. One of the driving forces behind studying ultracold atoms is the exquisite control and flexibility available in experiments. This flexibility enables us to engineer new states of matter, e.g. cold atom emulators of other physical systems which are not presently accessible to experiment. The figure on the left shows the spectrum of energy levels of a cold atom in a laser configuration designed to make the atom mimic an electron confined in two dimensions, in a generalization of a magnetic field called a *non-Abelian gauge field*. The atom's energy levels are functions of three parameters: α_1 and α_2 are related to the strengths of the artificial magnetic field, and k_y represents the atom's momentum. Along the line $\alpha_1 = \alpha_2$, the field reduces to the usual (Abelian) magnetic field. This type of figure is sometimes called an *Osterloh moth*, after the *Hofstader butterfly* structure that was identified by Douglas Hofstadter in the Abelian system. The figure on the right shows a close up of the line $\alpha_1 = \alpha_2$. The points in the figure are created by directly programming the visualization system's graphics processing unit (GPU) to render the points as fat pseudo-spheres, giving them size and depth. The fat pseudo-spheres size is set interactively by the user.

Cover Visualization by: Terence Griffin, Marc Olano, Judith Terrill

Acknowledgement. We are grateful to Robin Bickel for collecting the information and organizing the first draft of this report.

Disclaimer. All references to commercial products in this document are provided only to document how results have been obtained. Their identification does not imply recommendation or endorsement by NIST.

Table of Contents

Part I

Overview

Introduction

The success of research and development in the sciences and engineering is increasingly recognized as critical to the nation. Such advances are the principal fuel for industrial innovation, and hence for US competitiveness. Over the past 25 years computation has played an increasingly prominent role as a key enabler of science and engineering R&D. Indeed, a 2005 report[1] of the President's Information Technology Advisory Committee (PITAC) concludes that

> *"Computational science is now indispensable to the solution of complex problems in every sector, from traditional science and engineering domains to such key areas as national security, public health, and economic innovation."*

In a similar vein, a 2006 National Science Foundation report[2] states

> *"Simulation-based engineering science … is a discipline indispensable to the nation's continued leadership in science and engineering. It is central to advances in biomedicine, nanomanufacturing, homeland security, microelectronics, energy and environmental sciences, advanced materials, and product development. There is ample evidence that developments in these new disciplines could significantly impact virtually every aspect of human experience."*

Some examples of the context for our work …

"The use of digital technologies, including computation for increasingly complex models and simulations, vast sensor arrays, powerful imaging equipment and detectors, and networked access, interaction, and dissemination tools, has transformed the scientific landscape. … The rate at which these digital data are produced is increasing each year, yielding massive and exponentially-growing data flows in what has been described as a data deluge."– Interagency Group on Digital Data, 2008

"Quantum mechanics may hold a key to sustained growth in computing power in the coming decades, not by continuing along the path of Moore's law of miniaturization but by changing the rules of the game." – NRC Committee on Atomic, Molecular and Optical Physics, 2007

"Modern scientific research requires interdisciplinary collaboration …" – NRC Report on US Mechanical Engineering Basic Research, 2007

" … the (High End Computing) community should develop, maintain, and continually refine a collection of well-engineered software tools and software development environments that facilitate the use of HEC systems. … By doing so, America's scientific and engineering communities and its industries will increase their competitive advantages." – PCAST Report on Networking and Information Technology, 2007

"In the swiftly changing and increasingly competitive global marketplace, innovative design solutions and short product development cycles which rely on integrated product development teams armed with computationally based design, engineering analysis, and manufacturing tools are what give the nation its competitive advantage. A critical missing link in the integrated product development process is a set of predictive computational materials engineering tools. The development of computational tools for materials engineering has lagged behind the development of such tools in other engineering fields." – NRC Report on Integrated Computational Materials Engineering, 2008

[1] *Computational Science: Ensuring America's Competitiveness*, President's Information Technology Advisory Committee, June 2005.
[2] *Simulation-based Engineering Science: Revolutionizing Engineering Science Through Simulation*, Report of the National Science Foundation Blue-Ribbon Panel on Simulation-based Engineering Science, February 2006.

Applied mathematics, statistics, and computer science are the foundation for computational science and engineering. Research in mathematical and statistical analysis, numerical algorithms, software tools, high performance computing, and visualization provide the basis for mathematical modeling, computational simulation, and data analysis in all fields. Thus, success in this arena requires close cooperation between mathematicians, computer scientists, and application scientists. As the PITAC report observes, "the 21st century's most important problems … are predominantly multidisciplinary, multiagency, multisector, and collaborative."

In this report we describe the recent technical work of the Mathematical and Computational Sciences Division, the principle organization within NIST promoting advances in applied mathematics and computational science in support of both NIST internal measurement science research as well as for the computational science community at large.

Organizational Context. The National Institute of Standards and Technology (NIST) plays a central role in the infrastructure for science and technology. Its mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life. In particular, the NIST Measurement and Standards Laboratories conduct research that advances the technology infrastructure needed by U.S. industry to continually improve products and services.

As one of the major operating units at NIST, the Information Technology Laboratory (ITL) promotes US innovation and industrial competitiveness by advancing measurement science, standards, and technology through research and development in information technology (IT), mathematics, and statistics. The importance of

ITL Programs in Brief

Complex Systems
Complex Systems are composed of large interrelated, interacting entities which taken together, exhibit a macroscopic behavior which is not predictable by examining the components. This Program seeks to understand the fundamental science of these systems and develop rigorous descriptions that enable prediction and control of their behavior.

Cyber and Network Security
Cybersecurity is focused on ensuring three IT system objectives: confidentiality, integrity, and availability. This Program addresses long-term scientific issues in some of the building blocks of IT security — cryptography, security testing and evaluation, security metrics, and security properties —providing a more scientific foundation for computer and network security, while maintaining a focus on near-term issues.

Enabling Scientific Discovery
Modern scientific research has become more and more dependent on mathematical, statistical, and computational tools. This Program promotes the use of these tools to dramatically advance our ability to predict the behavior of a broad range of complex scientific and engineering systems and processes. In particular, this Program focuses on inter-disciplinary projects that involve novel computational statistics and simulation methods and software.

Identity Management Systems
Identity management systems are responsible for the creation, use, and termination of electronic identities used to access logical and physical resources. This Program is pursuing the development of common models and metrics for identity management, critical standards, and interoperability of electronic identities. These efforts will improve the quality, usability, and consistency of identity management systems while protecting privacy.

Information Discovery, Use, and Sharing
This Program fosters innovation throughout the information life cycle by developing the measurement infrastructure to enhance knowledge discovery, information exchange, and information usability. The Program enables novel computational approaches to data collection and analysis to be combined with improved interoperability techniques to effectively extract needed information from the wealth of available data.

Pervasive Information Technologies
Pervasive information technology is the trend towards increasingly ubiquitous connected computing sensors, devices, and networks that monitor and respond transparently to human needs. This Program facilitates the creation of standards for sensor communication, networking interoperability, and sensor information security.

Trustworthy Information Systems
Trustworthy information systems perform as intended for a specific purpose, with operational resiliency and without unwanted side effects or exploitable vulnerabilities. This Program will improve the ability to model, produce, measure, and assess trustworthiness through innovative technologies, models, measurements, and software tools.

Virtual Measurement Systems
A virtual measurement is a quantitative result and its uncertainty, obtained by a computer simulation or computer-assisted measurements. This Program introduces metrology constructs, standard references, uncertainty characterization, and traceability into scientific computation and computer-assisted measurements to enable predictive computing with quantified reliability.

MCSD Programs in Brief

Digital Library of Mathematical Functions
The special functions of applied mathematics are fundamental tools enabling modeling and analysis in all areas of science and engineering. To make effective use of such functions, practitioners must have ready access to a reliable source of information on their properties. The goal of this work is the development and dissemination of definitive reference data on the special functions of applied mathematics. The DLMF will be a freely available interactive and richly linked online resource.

Fundamental Math Software Development, Testing
Modern science and engineering in general, and modern measurement science in particular require a wide variety of software tools for scientific discovery, exploration, and analysis. As scientific inquiry becomes deeper and more specialized, so must the supporting software tools. The goal of this work is to develop critical software tools that support measurement science at NIST, as well as computational science and engineering at large.

High Performance Computing and Visualization
Computational capability is advancing rapidly. This means that modeling and simulation can be done with greatly increased fidelity (e.g. higher resolution, more complex physics). However, developing large-scale parallel applications remains highly challenging, requiring expertise that application scientists rarely have. In addition, the hardware landscape is changing rapidly so new algorithmic techniques must constantly be developed. We are developing and applying facilities and expertise of this type for application to NIST problems. Large scale computations and laboratory experiments invariably produce large volumes of scientific data, which cannot be readily comprehended without some form of visual analysis. We are developing the infrastructure necessary for advanced visualization of scientific data, including the use of 3D immersive environments and applying this to NIST problems. One of our goals is to develop the 3D immersive environment into a true interactive measurement laboratory.

Mathematics of Metrology
Mathematics plays an important role in the science of metrology. Mathematical models are needed to understand how to design effective measurement systems, and to analyze the results they produce. Mathematical techniques are used to develop and analyze idealized models of physical phenomena to be measured, and mathematical algorithms are necessary to find optimal system parameters. Finally, mathematical and statistical techniques are needed to transform the resulting data into useful information. The goal of this work is to develop fundamental mathematical methods and analytical tools necessary for NIST to continue as a world-class metrology institute, and to apply them critical measurement science applications.

Quantum Information
An emerging new discipline at the intersection of quantum mechanics and computer science, quantum information science (QIS) is likely to revolutionize science and technology in much the same way that lasers, transistors, electronics, and computing did in the 20th century. At the very least, QIS has the potential to provide phenomenal increases in information storage and processing speed and communication channels with extremely high levels of security guaranteed by the laws of physics. However, the needed manipulation and control of quantum states remains a grand challenge. We are developing theoretical underpinning to support the multi-laboratory experimental program in quantum information at NIST. This work is aimed at the development of a measurement science for quantum information processing systems.

applied mathematics and computational science to ITL's work is clear from the four core competencies that it has identified for itself:

1. IT measurement and testing,
2. Mathematical and statistical analysis for measurement science,
3. Modeling and simulation for measurement science, and
4. IT standards development and deployment.

To respond to the needs of its customers in industry, academia, and government, ITL has developed a set of cross-cutting and interdisciplinary R&D programs. These are summarized in the sidebar *ITL Programs in Brief*.

NIST's other Laboratories and research centers are also important customers of ITL. Indeed, NIST's measurement science research program has been transformed by the advent of computational science and engineering. Nearly every NIST project, both theoretical and experimental, typically now has critical computational components. In addition, an increasing number of NIST "products" are techniques, tools, and reference data to enable modeling, simulation, and data analysis in particular application domains. As a result, the expertise of applied mathematicians and computer scientists are in high demand within NIST.

Mathematics and Computational Science at NIST. The Mathematical and Computational Sciences Division (MCSD) is one of six technical Divisions within ITL. MCSD develops and maintains the scientific competencies and facilities necessary to provide NIST programs with leadership and collaborative research in applied mathematics and computational science to enable NIST to maintain its status as a world-class measurement science institute. In addition, we also engage in highly leveraged research and development efforts to improve the envi-

ronment for computational science and engineering at large.

To accomplish these goals, MCSD staff members engage in peer-to-peer collaboration with NIST scientists and engineers in a wide variety of critical applications, targeted outreach efforts with selected external communities to advance the state-of-the-art in their subfield, development and dissemination of unique mathematical and computational tools, and research in targeted areas of applied mathematics and computer science of high relevance to future NIST measurement programs.

The technical work of the Division is organized into a collection of Lab-centric and Division-centric Programs. The ITL Programs are substantial cross-Divisional efforts aimed at focusing ITL's core competencies on critical National priorities. The ITL Programs in which we are currently participating are (in decreasing order of involvement):

Enabling Scientific Discovery
Information Discovery, Use and Sharing
Complex Systems
Virtual Measurement Systems
Pervasive Information Technology
Trustworthy Information Systems
Cyber and Network Security

Very brief descriptions of these program areas are provided in the sidebar. Our Division-centric work is concentrated in the following five areas (in decreasing order of size):

Digital Library of Mathematical Functions
High Performance Computing and Visualization
Quantum Information
Mathematics of Metrology
Fundamental Math Software Development and Testing

Some of these efforts are characterized as mature Division-centric efforts in support of NIST goals that do not fit in any of ITL's Program areas, but are nevertheless critical for NIST metrology. Others fulfill a need to continuously develop and maintain technical competence and facilities in core areas of mathematics and computational science which are broadly applicable.

In Part III of this document we provide descriptions of most ongoing technical projects of the Division. These are organized under the Program areas listed above.

Highlights

In this section we identify some of the major accomplishments of the Division over the past year. We also provide news related to MCSD staff. Details can be found in subsequent sections.

Technical Accomplishments

MCSD has made significant technical progress in a wide variety of areas during the past year. Here we highlight a few examples. Further details are provided in Part II (Features) and Part III (Project Summaries) of this document.

Seeded with $1.3M of funding from a FY 2007 NIST initiative, MCSD has initiated a new program of research on the *foundations of measurement science for information systems*. The goal of this long-term basic research effort is to develop mathematical models, techniques, and

tools which facilitate the fundamental understanding needed to develop a measurement science for information systems comparable to that of the physical sciences and engineering. The initial foci of the program have been the emerging field of network science, including the analysis and computation of properties and dynamics of networks structures, as well as in selected applications in computer security and software testing. The projects underway in this approximately 5-person effort include the following (described elsewhere in this report): *Methods for Characterizing Massive Networks, Analysis of Distributed Protocols for Network Control, TCP Metastability and Cascading Failures on the Internet, Modeling Network Vulnerabilities Using Attack Graphs, and Combinatorial Methods for Software Testing.* The nascent work of this program has already yielded 18 papers and reports, as well as several software packages which have recently been made available to the public. Most of the work supported by this initiative has been aligned with the ITL Complex Systems Program; additional work is associated with the Trustworthy Systems Program and the Cyber and Network Security Program. The staff members involved in this effort are Isabel Beichl, Brian Cloteaux, Fern Hunt, Raghu Kacker, James Lawrence, Vladimir Marbukh, Roldan Pozo, and Anoop Signhal (ITL Computer Security Division).

MCSD's long-running *Digital Library of Mathematical Functions* (DLMF) Project has entered its final stage. The DLMF will be a freely available, interactive, web-accessible knowledge base of technical information on the special functions of applied mathematics. The DLMF is a modern successor to the classic NBS *Handbook of Mathematical Functions* (M. Abramowitz and I. Stegun, Eds., 1964), the most widely distributed NIST publication of all time, and among the most highly cited works in the mathematical literature. The DLMF will be comprised of 36 chapters of data (three methodology chapters and 33 chapters, each focusing on a class of functions). The project is quite complex, involving the cooperative work of some 50 researchers and staff, mostly external to NIST. Thousands of web pages mathematical of information and hundreds of visualizations are being produced. We expect to unveil the web site in mid 2009. Of the 36 chapters in the DLMF, 24 have been finalized as of this writing; the remaining chapters are having their final versions checked by their authors. A publisher for a print version of the DLMF has just been selected. A preview edition of the DLMF[3] was released to the public in June 2008. The preview edition contains five chapters: "Gamma Function", "Airy & Related Functions", "Functions of Number Theory", "3j, 6j, 9j Symbols", and "Asymptotic Approximations."

To determine the quality of quantum control in the context of *quantum information processing*, Manny Knill of MCSD developed a randomized benchmarking strategy for determining an effective error probability per quantum gate and implemented it with one-qubit gates in an ion trap in collaboration with Dave Wineland's Ion Storage Group in PL. Using his strategy, an effective error probability of only 0.5% was observed. While this is good, much small error tolerances are desired; this error probability probably can be greatly reduced by better laser stabilization and other improvements. Research groups at Waterloo and Yale have already applied Knill's benchmarking strategy to liquid-state NMR and solid state qubit systems. Once sufficiently high-quality two-qubit gates are available, the team hopes to repeat the experiment with operations on two or more ion qubits.

Alfred Carasso of MCSD has developed a novel mathematical framework for *quantifying the smoothness of images*. Potential applications of this metrology tool include monitoring of performance of imaging systems, evaluation of image reconstruction software quality, detection of abnormal fine structure in biomedical images, and monitoring of surface finish in industrial applications. The method is based upon the L_1 Lipschitz exponent, a useful tool for characterizing fine-scale content, provided the image is relatively noise free. The use of this index has been

[3] http://dlmf.nist.gov/

enabled by an effective computational method for estimating image Lipschitz exponents developed by Carasso. The technique has been evaluated in collaboration with András Vladár of the Precision Engineering Division of MEL. Using synthetically degraded images, it was shown that Lipschitz exponents can measure the extent to which competing denoising algorithms remove texture along with the noise. Likewise, Lipschitz exponents can measure the ability of competing deblurring algorithms to recover texture. This was illustrated in the case of the NIST APEX blind deconvolution method applied to Hubble space telescope imagery, as well as imagery from SEM, MRI, and PET devices.

First kind integral equations are routinely used to model instrument distortions when trying to identify an unknown function. The solution to such ill-posed problems is inherently challenging; small variations in input data (such as noise), or small errors made in the course of numerical solution, are magnified without bound. To stabilize the solution process, a technique known as "regularization" is necessary. Such procedures require the *selection of regularization parameters*, which is typically done in an ad hoc way. Few metrics for evaluating the success of regularization parameter selection are known. Bert Rust of MCSD, working with MCSD faculty appointee Dianne O'Leary of UMCP, developed such a metric based upon an evaluation of whether the resulting model is statistically plausible. In particular, candidate parameter choices are evaluated by inspecting the model residual along with its periodogram and cumulative periodogram. Based on this criterion, Rust and O'Leary also propose a new method for choosing regularization parameters that makes the residuals as close as possible to white noise using a diagnostic test based on the periodogram. In a paper published last year in the journal *Inverse Problems*, they compared this new method with standard techniques such as the discrepancy principle, the L-curve and generalized cross validation, showing that it performs better on a variety of benchmark problems.

Recently, William Mitchell of MCSD developed a new method for *partitioning adaptive grids for distribution over parallel processors* in the context of adaptive multilevel methods for solving partial differential equations. For triangular (2D) and tetrahedral (3D) grids, his method is guaranteed to produce *connected* partitions; no other partitioning method is known to do this. (Connected partitions lead to reduced inter-process communication.) Mitchell has implemented this method in his PHAML package for the parallel multi-level adaptive solution of partial differential equations and compared it with other popular partitioning methods. His method computes partitions in an amount of time similar to fast load balancing methods (which create inferior grids), and with mesh quality similar to slower methods (which produce well behaved grids). This work received the 2008 ITL Best Journal Paper of the Year Award.

The innovative efforts of a team of researchers from MCSD and the NIST Building and Fire Research Laboratory (BFRL) have been rewarded with two significant *grants of external supercomputing time for the modeling of complex materials*. This year the team was granted 400,000 hours of CPU time on the Columbia supercomputer at the NASA Ames Research Center and 750,000 CPU hours on the IBM Blue Gene/P system at Argonne National Laboratory. The NASA grant is a follow-up to a 1M CPU hour award the group received last year for the modeling the rheological properties of suspensions. The DOE allocation is one of 55 awards of supercomputer time given in a peer-reviewed competition known as the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. The NIST team of William George, Judith Terrill, and John Hagedorn of MCSD and Nicos Martys, and Edward Garboczi of BFRL are using the granted time to study the flow, dispersion and merging of densely suspended, diversely sized and shaped materials under a variety of conditions. The primary application is the study of cement and concrete. In addition, the MCSD researchers have begun work with Jeffrey Bullard of BFRL to study the scalability of a parallel code (called Hy-

dratiCA) which simulates the three-dimensional changes in the structure and chemical composition of aqueous mineral systems, such as those encountered in environmental geochemistry, ceramic processing, and cement-based materials.

A popular NIST benchmark for scientific computing has been included in the recently released SPECjvm2008, an industry standard *benchmark suite for evaluating the performance of Java runtime environments*. SciMark[4], which was developed by Roldan Pozo and Bruce Miller of MCSD, can be used to compare the performance of Java environments for scientific computing applications. The SciMark benchmark consists of five computational kernels characteristic of the most compute-intensive part of large scientific computing applications, each with different resource usage profiles: FFT, sparse matrix multiply, dense LU factorization, Gauss-Seidel relaxation (grid sweep), and Monte Carlo integration. Originally released in 1999, SciMark has tracked performance improvements in Java for scientific computing from early processing rates measured in 10s of kiloflops (thousands of floating-point operations per second) to a recent mark of over 1 gigaflop. SciMark is one of 10 benchmarks included in SPECjvm2008 to represent the wide range of existing Java applications. SPECjvm2008 was released on May 7, 2008 by the System Performance Evaluation Corporation[5] (SPEC), a non-profit open industry cooperative which develops benchmarks that ensure that the marketplace has a fair and useful set of metrics to differentiate system performance. Members of SPEC include AMD, Dell, HP, IBM, Intel, Microsoft, and Sun.

Technology Transfer and Professional Activities

The volume of technical output of MCSD remains high. During the last 18 months, Division staff members were (co-)authors of 34 articles appearing in peer-reviewed journals, one book chapter, and 18 papers in conference proceedings. Twelve additional papers have been accepted for publication, while 13 others are undergoing review. Division staff members gave 11 invited technical talks and presented 30 others in conferences and workshops.

MCSD continues to maintain an active Web site with a variety of information and services, including the Guide to Available Mathematical Software, the Matrix Market, and the SciMark Java benchmark. During calendar year 2008, the division web server satisfied more than six million requests for pages from more than 536,000 distinct hosts. On a daily basis, more than 13,000 pages and 1.6 Gbytes of data were shipped. In total, there have been more than 147 million "hits" on MCSD Web servers since they went online as NIST's first web servers in 1994.

Among our most popular software downloads for calendar year 2008 were: JAMA (linear algebra in Java): 17,780 downloads, Template Numerical Toolkit (linear algebra using C++ templates): 13,025 downloads, LAPACK++ (dense linear algebra in C++) 11,573 downloads, OOMMF (modeling of micro- and nano-magnetic phenomena): 4,700 downloads, and f90gl (Fortran 90 interface to OpenGL graphics): 3,819 downloads. Another indication of the successful transfer of our technology is references to our software in refereed journal articles. For example, our OOMMF software was cited in 109 such papers which were published in calendar 2008 alone (686 such papers have been identified since 2001).

Members of the Division are also active in professional circles. Isabel Beichl has just begun a term as Editor-in-Chief of *Computing in Science & Engineering*. Staff members hold a total of seven associate editorships of peer-reviewed journals. They are also active in conference organi-

[4] http://math.nist.gov/scimark/
[5] http://www.spec.org/

zation, serving on four organizing/steering/program committees. Staff members organized two minisymposia for the 2008 SIAM National Meeting and two symposia at ASME conferences.

Service within professional societies is also prevalent among our staff. For example, Ronald Boisvert serves as Co-Chair of the Publications Board of the Association for Computing Machinery (ACM) and is a member of the ACM Council, the association's board of directors. Fern Hunt serves on the Executive Committee of the Association for Women in Mathematics. Daniel Lozier serves as Secretary of the Society for Industrial and Applied Mathematics (SIAM) Activity Group on Orthogonal Polynomials and Special Functions. Staff members are also active in a variety of working groups. Ronald Boisvert serves as Chair of the International Federation for Information Processing (IFIP) Working Group 2.5 on Numerical Software, Donald Porter is a member of the Tcl Core Team, and Bruce Miller is a member of W3C's Math Working Group. Judith Terrill represents NIST on the High End Computing Interagency Working Group of the Federal Networking and Information Technology Research and Development (NITRD) Program.

For further details, see Part IV (Activity Data) of this document.

Staff News

Vladimir Marbukh, formerly of the ITL Advanced Networking Technologies Division, joined MCSD in March 2008. He will be enhancing the growing MCSD program on Mathematical Foundations of Measurement Science for Information Systems funded by the 2007 Cyber Security Initiative. Vladimir holds a PhD degree in Systems Engineering from the Polytechnic University in Leningrad. Before coming to NIST in 1998, he served as a consultant at Bell Laboratories (Murray Hill). His research specialization is in the performance analysis of large scale communication networks. His mathematical expertise covers a broad range of topics including queueing theory, optimization, control theory, stochastic processes, and game theory.

Three new NRC Postdoctoral Associates began two-year appointments in MCSD this year. Bryan Eastin, a PhD in Physics from the University of New Mexico, arrived in October 2007 to work with Manny Knill in Boulder. He is working on fault tolerance and error correction for quantum computers. Valerie Coffman, a PhD in Physics from Cornell University began her appointment in MCSD in January 2008. She is working with Stephen Langer on the development of a 3D version of the OOF system for modeling materials with complex microstructures. Our newest postdoc, Aaron Lott, arrived in September 2008. Aaron is a recent PhD in Applied Mathematics and Scientific Computation from the University of Maryland College Park. There he worked with Howard Elman and Anil Deane to develop efficient solvers for spectral element discretizations of the steady advection-diffusion and Navier-Stokes equations. In MCSD he is working with Geoffrey McFadden on the modeling and simulation of complex multiphase flows.

Two new guest researchers have begun visiting appointments in MCSD. In January 2008 Rüdiger Kessell arrived from Germany to work on mathematical applications in satellite calibration, radiation measurements, mass metrology, and software testing. Partial support for his stay here is provided by the Optical Technology Division (NIST PL) and the New Brunswick Laboratories (DOE). Gunay Dogan began a guest researcher appointment in September 2008 to do research associated with a NIST grant to Cornell University. He received a PhD in Applied Mathematics and Scientific Computing from the University of Maryland in 2006, where he worked on image segmentation algorithms. Here he is working on imaging problems associated with the OOF system for the modeling of complex material microstructures.

Aurelie Canale, a guest researcher from ISIMA in Clermont-Ferrand, France, ended her term at NIST in the fall of 2008. She worked with Adele Peskin on visualization applications.

This year MCSD supported two graduate students from the University of Colorado, three summer interns in the NIST Student Undergraduate Research Associateship (SURF) program, as well as two high school volunteers. Details are provided in Table 1.

Table 1. ***MCSD student interns for 2008.***

Name	Institution	Level	Program	Mentor	Project Title
William Hess	Montgomery Blair	High School	Student Volunteer	J. Terrill	Visualization of the Deformation of Structures Due to Fire
Adam Leedy	University of Kentucky	Undergraduate	SURF	A. Peskin	Weighting Factors for Image Segmentation
Adam Meier	University of Colorado	Graduate	PREP	E. Knill	Quantum Computer Simulation
Kevin Rawlings	Montgomery Blair	High School	Student Volunteer	J. Terrill	Visualization of Nanostructures
Axel Rivera-Rodrigues	University of Puerto Rico at Humacao	Undergraduate	SURF	J. Terrill	Measurement of Angles in an Immersive Environment
Amelia Tebbe	St. Mary's College	Undergraduate	SURF	D. Gilsinn	Modeling and Volume Estimation of Lung Cancer Nodules
Yanbao Zhang	University of Colorado	Graduate	Guest Researcher	E. Knill	Tests of Local Realism

SURF: NIST Student Undergraduate Student Fellowship Program.
PREP: NIST Professional Research Experience Program

Recognition

Division staff garnered a significant number of professional recognitions during the past year. These are described below.

Internal Awards

Anthony Kearsley was selected as part of a four-person team to receive the Department of Commerce Bronze Medal for work on matrix assisted laser desorption/ionization-time of flight-mass spectroscopy (MALDI-TOF-MS). The team was cited for work which enabled the widespread use of MALDI as a quantitative measurement tool. In particular, they developed a comprehensive online resource that includes an automated, operator-independent data analysis tool (MassSpectator), a consolidated collection of sample preparation methods, and polymer mass spectroscopy workshop reports. The citation states "Because of their enormous value to industrial, academic, and government researchers, the database and the accompanying tools exemplify the highest level of service to those engaged in mass spectroscopy." The co-recipients of the award were William Wallace, Charles Guttman, and Kathleen Flynn of MSEL. The award was presented at the annual NIST awards ceremony in December 2007.

William Mitchell received the ITL Outstanding Authorship Award for best journal paper published during 2007 by an ITL staff member. The paper, entitled "A Refinement-Tree Based Partitioning Method for Dynamic Load Balancing with Adaptively Refined Grids," appeared in the *Journal of Parallel and Distributed Computing* in April 2007.

Abbie O'Gallagher received the ITL Outstanding Contribution Award for "outstanding initiative and dedication in leading the ITL Diversity Book Club." Chris Schanzle received the ITL Outstanding Support Award for "outstanding initiative and dedication in the enhancement of Division computing resources resulting in improved performance and substantially lower cost."

External Awards

Dianne O'Leary, a faculty appointee in MCSD from the University of Maryland, was named the Sonia Kovalesky Lecturer for 2008 by the Association for Women in Mathematics (AWM) and the Society for Industrial and Applied Mathematics (SIAM). The lecture is intended to highlight significant contributions of women to applied or computational mathematics. The lecture, entitled "A Noisy Adiabatic Theorem: Wilkinson Meets Schrödinger's Cat", was delivered as a plenary talk at the SIAM National Meeting held in San Diego in July 2008.

MCSD's Scientific Applications and Visualization Group won a Department of Energy Office of Advanced Scientific Computing Research (OASCR) Award for their visualization demonstrated at the 2008 SciDAC Conference, held in Seattle, Wash., from July 13-17, 2008. The winning visualization presented the results of a high performance computer simulation of the flow of a particular complex suspension: concrete. The visualization showed a flow of a realistic collection of rocks of various sizes whose shapes were obtained by scanning actual rocks. The direct visualization of rocks is augmented with embedded data to provide additional infor-

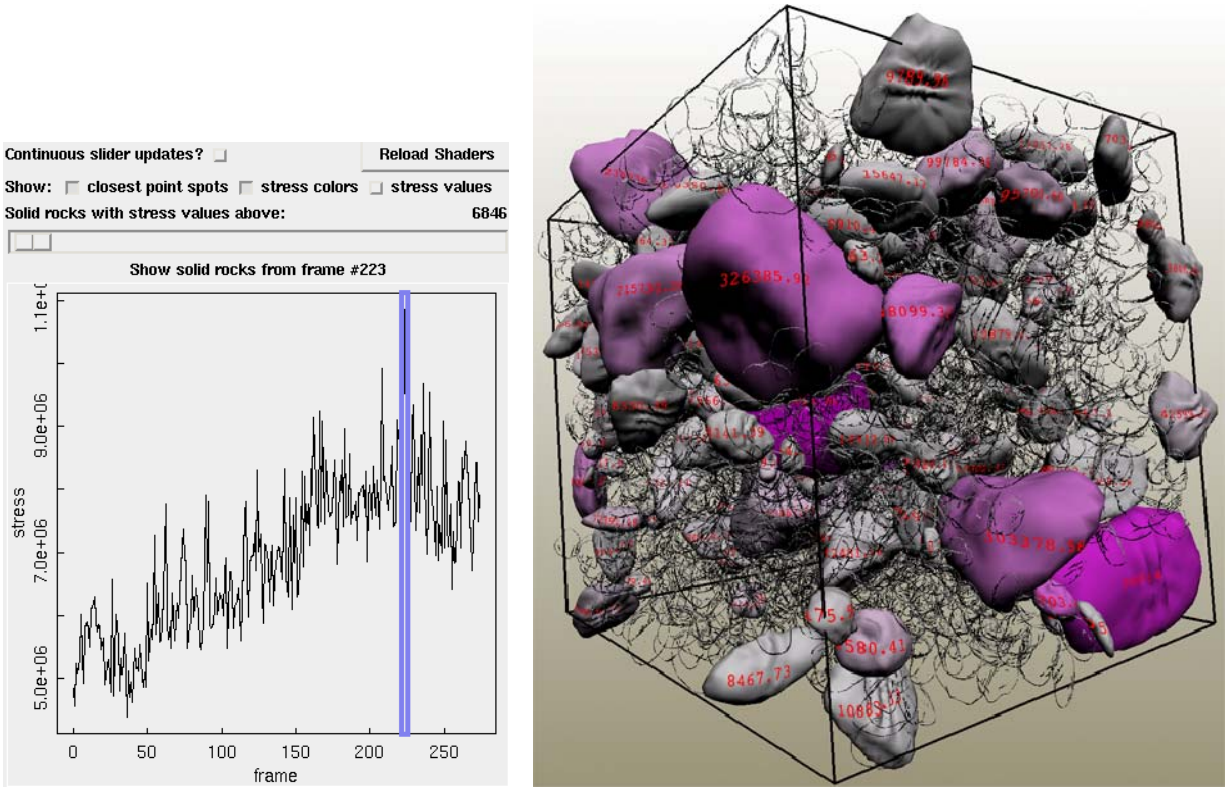

Figure 1. *MCSD visualization that won a DOE Office of Advanced Scientific Computing Research Award at the 2008 SciDAC Conference, held in Seattle, Washington, from July 13-17, 2008. The visualization presents the results of a high-performance computer simulation of the flow of a particular complex suspension: concrete. Illustrated is a flow of a realistic collection of rocks of various sizes whose shapes were obtained by scanning actual rocks. The direct visualization of rocks is augmented with embedded data to provide additional information on rock-rock interactions and per-rock stress.*

mation on rock-rock interactions and per-rock stress. Supported by the Department of Energy Office of Science, the SciDAC (Scientific Discovery through Advanced Computing) program engages computational scientists, applied mathematicians, and computer scientists from across application domains and from universities and national laboratories across the United States. The visualization was created by Marc Olano, Steve Satterfield, John Hagedorn, William George, and Judith Terrill in collaboration with Nicos Martys and Edward Garboczi of BFRL.

Figure 2. *Daniel Lozier received an Honorary Fellowship from ESCMSE.*

Daniel Lozier received the European Society of Computational Methods in Science and Engineering (ESCMSE) Honorary Fellowship award. ESCMSE is a non-profit organization which seeks to enable cooperation among scientists and to disseminate research results related to the construction, development, and analysis of computational, numerical, and mathematical methods and their application in the sciences and engineering. The Honorary Fellowship is the highest distinction offered by ESCMSE. Lozier was cited for "outstanding contribution in the field of computational mathematics and numerical analysis." The award was presented at the Sixth International Conference of Numerical Analysis and Applied Mathematics, which was held in Kos, Greece September 16-29, 2008.

Ronald Boisvert received the International Federation for Information Processing (IFIP) Silver Core Award in recognition of outstanding service to IFIP. IFIP is a non-governmental, non-profit umbrella organization for national professional societies working in the field of information processing. It was established in 1960 under the auspices of UNESCO. Sponsoring more than 100 conferences each year, IFIP links more than 3,500 scientists from academia, industry, and government labs, organized into more than 101 Working Groups reporting to 13 Technical Committees. Boisvert has served as a member of IFIP Working Group 2.5 on Numerical Software since 1996, and has been the Working Group's Chair since 2000.

Figure 3. *Axel Rivera (second from right) receives a Best Paper Award at the Undergraduate Research Symposium held in San Juan, Puerto Rico in September 2008. His paper was based on the project "Principal Components Analysis and Singular Value Decomposition for Automated Angle Measurements," completed in MCSD as part of the NIST SURF Program.*

Axel Rivera, a 2008 participant in the NIST Student Undergraduate Research Fellowship program SURF, received a Best Paper Award, one of three presented at the XIX Undergraduate Research Symposium held in San Juan, Puerto Rico on September 12-13, 2008. Some 150 students attended the conference, which is designed to give students from institutions with minority participation in the science, technology, engineering and mathematics the opportunity to pre-

sent results of their research experiences. Rivera's presentation, entitled "Principal Components Analysis and Singular Value Decomposition for Automated Angle Measurements," reported the results of his 2008 SURF project, completed under the direction of Judith Terrill of MCSD. Riveria is beginning his senior year of undergraduate studies at the University of Puerto Rico, Humacao, where he majors in Computational Mathematics.

Fern Hunt and Vladimir Marbukh received a Best Paper Award for their paper "Dynamic Routing and Congestion Control through Random Assignments of Routes," presented at the 5th International Conference on Cybernetics and Information Technologies, Systems and Applications (CITSA 2008), Orlando, FL, June 2008. Conference participants vote for the best paper in each session. Theirs was tops in the session on Control Technologies and Applications.

Other Significant Recognition

Isabel Beichl has been named the next Editor-in-Chief of *Computing in Science and Engineering* (CiSE), a joint publication of the IEEE Computer Society and the American Institute of Physics. CiSE is a bi-monthly publication that combines peer-reviewed articles, columns, tutorials, and book reviews in a magazine format. According to its statement of purpose, CiSE "aims to support and promote the emerging discipline of computational science and engineering and to foster the use of computers and computational techniques in scientific research and education." Beichl has been the Department Editor for CiSE's Computing Prescriptions department for several years. This section of the magazine publishes tutorial articles which explain mathematical and computational techniques useful in scientific computing. Beichl's appointment as Editor-in-Chief, which was made by the IEEE Computer Society Board of Governors, is for a two-year term commencing in January 2009. She succeeds Norman Chonacky, a physicist from Yale University, in this position.

Alfred Carasso was awarded a US Patent (number 7,437,012) entitled "Singular Integral image Deblurring Method" on October 14, 2008. According to the abstract,

> *"The present invention provides a method for image deblurring based on correctly specifying the lack of smoothness in the unknown desired sharp image, in terms of Lipschitz (Besov) space classification. The method makes essential use of Singular Integrals and Fast Fourier Transform (FFT) algorithms to estimate the image's Lipschitz exponent α. Such Singular Integrals and Fast Fourier Transforms are then used in a method for regularizing the ill-posed deblurring problem, resulting in a fast, direct, (i.e., non-iterative) computationally effective deblurring procedure."*

Part II

Features

Ion Trap Quantum Computing Benchmarks

Quantum computers have the potential for significantly speeding up many useful algorithms. However, building quantum computers is challenging. Currently available quantum devices are barely sufficient for demonstrating computationally useful registers with two quantum bits (qubits). A leading technology for realizing the first useful quantum computers involves atomic qubits in ion traps. To determine the potential of current ion trap technology for quantum computing, we implement diagnostic quantum subroutines to benchmark the performance of experimental atomic qubits. In previous years, we have prepared quantum states on up to six qubits and measured their fidelity to demonstrate the ability to prepare states with nontrivial quantum features. Recently we developed an efficient, randomized benchmarking strategy to determine the computationally relevant error per operation. We implemented the strategy on one atomic qubit, demonstrating that the error per operation is close to 0.5%.

Emanuel Knill

The goal for experimental efforts in quantum computing is to obtain quantum devices that can be used to build large quantum computers. There are many different physical systems that are being investigated for their utility in obtaining such quantum devices. Examples include systems based on ion traps, superconducting circuits and quantum dots. Although we are still a long way from realizing useful quantum computations, the physical system currently closest to the desired goal is based on ion traps.

In order to determine the suitability of a device for quantum computing and to compare it to other devices, one performs standard "benchmark" experiments that can characterize the error in the physical implementation of small quantum algorithms or that demonstrate a quantum process or protocol useful for scaling up quantum computers. Well-known examples of such benchmarks include quantum teleportation, quantum error correction, versions of the quantum Fourier transform, entangled state preparations and entanglement purification. All of these have been implemented in previous years in the ion trap systems at NIST in the Ion Storage Group of the Time and Frequency Division (PL). Although these benchmarks give an indication of how well the quantum systems can be controlled for specific tasks, they do not give a clear indication of how well the quantum systems would perform in the setting of general quantum computations.

In order to determine how well we can expect a quantum system to perform for quantum computation, we developed a randomized benchmarking method. This method determines an average probability of error for elementary operations called "gates" in a computational context by generating long random sequences of gates, applying them to the initialized quantum system, and making measurements to determine the probability of error. The probability of error can be mapped out as a function of the number of gates used. It is expected to converge exponentially to a limit. The exponent of the exponential is linearly related to an average error per gate.

Useful quantum computations are expected to involve many gates, most of which are dedicated to eliminating errors by means of quantum error correction to ensure fault tolerance. Because of this feature, gates are used in contexts that are essentially random with respect to the parts of the system affected. Randomized benchmarking thus determines the error probability as it affects such quantum computations. This contrasts with other methods for determining the quality of gates, such as quantum process tomography. Unlike process tomography, randomized benchmarking is insensitive to errors in state preparation and measurement, does not require well-characterized "analysis" pulses, determines the gates' behavior in arbitrary computational contexts and is efficient. Randomized benchmarking can be used to convincingly demonstrate that the implemented gates are sufficiently good to realize large quantum computations according to the

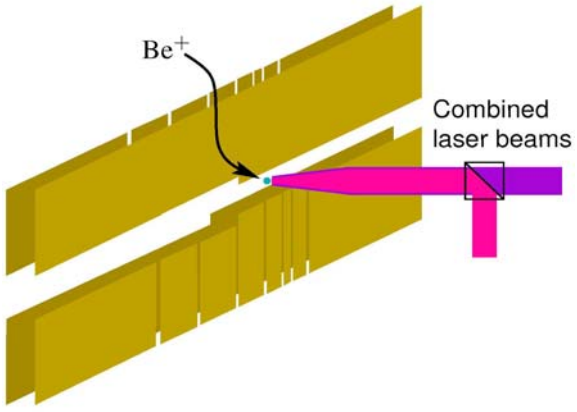

Figure 4. *Schematic of the ion trap setup for an atomic qubit stored in one* Be^+ *ion. The gold electrodes serve to trap the ion in a controlled position. The ion has two internal states that represent the logical basis states of the qubit. Laser pulses are used to "rotate" between the two states. By controlling the phase and duration of the pulses, arbitrary computational gates can be applied. For this experiment, two laser beams are combined on a polarizing beam splitter to provide this capability. Other beams (not shown) are used to trap and cool the motional degrees of freedom of the ion.*

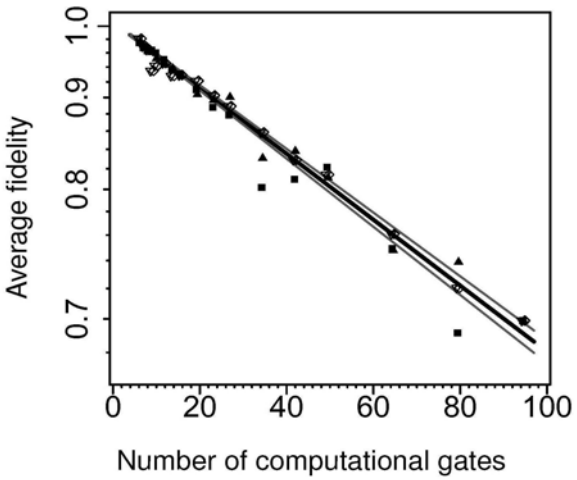

Figure 5. *Average fidelity as a function of the number of steps for gate sequences. The fidelity f is given by 1-p, where p is the probability of error. The different symbols are associated with different sets of randomized sequences of gates. The standard deviations due to measurement noise for the plotted points are smaller than the symbol sizes. The fidelity is plotted on a logarithmic scale. The lines represent the fitted exponential decay with 68% confidence intervals for the fit. The slope implies a probability of error of 0.482(17)% per randomized gate. The deviation of the plotted points from the expected exponential decay may be due to non-probabilistic, coherent errors that increase the variance of the error probabilities over sequence randomizations.*

accepted standard for maximum error probability per gate, which is currently 10^{-4}.

We have implemented randomized benchmarking on one atomic qubit realized by a Be^{+} ion in an ion trap. The atomic qubit's state is manipulated using laser beams as schematically shown in Figure 4. We determined an explicit procedure for combining and randomizing gates and making final measurements. The exponential decay as a function of the number of randomized gates is shown in Figure 5, showing an average error probability of about 0.5% per gate. Although randomized benchmarking cannot diagnose the source of problems contributing the errors, it can reveal the presence of issues of concern. For example, if there are global, slowly varying effects that cause every gate to deviate identically from the intended one, this causes a deviation from exponential decay. Such effects can be caused by fluctuations in a magnetic field or by miscalibration, and they are a source of concern for implementing long computations. Our data does not show a significant deviation from exponential decay. We made independent measurements of known error sources, which include spontaneous decay from excited states of the ion, dephasing due to magnetic field fluctuations, and laser phase and amplitude fluctuations. The estimated net effect of these errors is consistent with the measured average error probability.

Our procedure is given as a recipe that can be implemented with any other technology that may be used in future quantum computers. In fact, it has already been realized on nuclear magnetic resonance (NMR) qubits by Ryan et al. and in superconducting (SC) qubits by Chow et al. For comparison, the single-qubit error probability achieved for NMR and SC is about 0.13% and 1.1%, respectively. Whereas NMR qubits are not likely to be scalable to large systems, SC qubits have recently emerged as the second most advanced quantum computing technology.

Further work is required to develop general-purpose randomized benchmarking strategies for multiple qubits that are both readily implemented and simple to compare for different technologies. The theory of randomized benchmarking also needs further work to formally establish the relationship between the error per gate determined by the benchmark and the computationally relevant probability of error. Currently, the relationship has been proven under strong independence and stationarity assumptions on physical gate errors.

References

[1] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin and D.J. Wineland, Randomized Benchmarking of Quantum Gates, *Physical Review A* **77** (2008) 012307, 7 pages.

[2] C. A. Ryan, M. Laforest, and R. Laflamme, Randomized Benchmarking of Single and Multii-Qubit Control in Liquid-State NMR Quantum Information Processing, arXiv:0808.3973, 2008.

[3] J. M. Chow, J. M. Gambetta, L. Tornberg, Jens Koch, L. S. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Randomized Benchmarking and Process Tomography for Gate Errors in a Solid-state Qubit, arXiv:0811.4387, 2008.

Participants

Emanuel Knill (MCSD); Dietrich Leibfried, David Wineland, and the rest of the Ion Storage Group (PL).

Program area: Quantum Information (Division)

Partitioning for Dynamic Load Balancing with Adaptive Grids

The numerical solution of partial differential equations (PDEs) is the most compute-intensive part of a wide range of scientific and engineering applications. So the development and application of faster and more accurate methods for solving PDEs is a very important field that has received much attention in the past fifty years. Many PDE problems in application areas at the cutting edge of research are extraordinarily challenging to solve. These problems necessitate the use of the most advanced numerical techniques including adaptive grid refinement, multigrid solvers, and the use of parallel computers. To obtain optimal performance of the parallel computer, we have developed a method for dynamically partitioning the data to equidistribute the work load among the processors. The new method provides high quality partitions; for example, it is the only known method for producing partitions that are guaranteed to be connected. In addition, the method's computation time is comparable to the fastest known partitioning methods, which produce inferior partitions.

William F. Mitchell

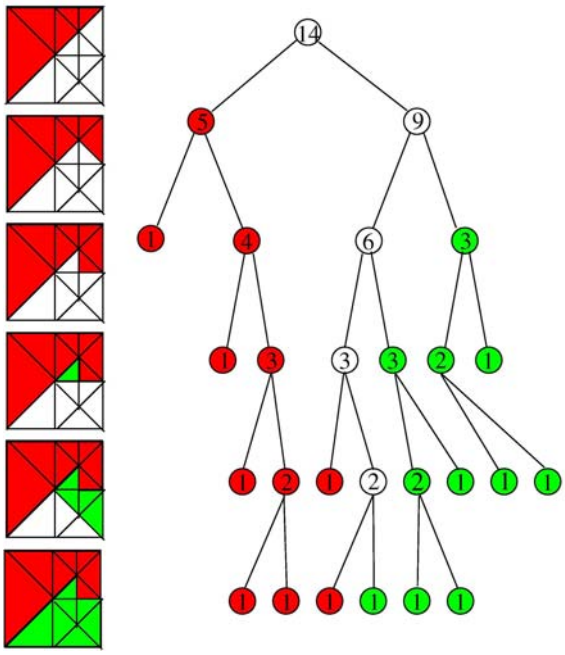

Figure 6. *Partitioning the grid and refinement tree into two sets.*

To solve partial differential equations (PDEs) numerically one converts the continuous problem to a system of algebraic equations defined on small subregions called elements (collectively known as the grid). To resolve steep fronts or highly oscillatory regions in the solution the elements must be small, but if one does this uniformly over the whole domain the number of algebraic equations explodes and the problem becomes intractable. It is necessary, then, to place small elements only where they are needed. Unfortunately, to place elements optimally one must know the solution to the PDE. An alternative is adaptive grid refinement: start with a coarse grid, compute a solution, estimate the error in each element, subdivide elements with unacceptably large error, and repeat. The PDE solver PHAML developed by William Mitchell of MCSD is based on this strategy.

Even with adaptive grid refinement the cost of solving a challenging PDE can be too large and so parallel processing is needed. Here the domain is divided among available processors. Each works on its own subdomain, obtaining data from other processors as needed. To minimize solution time, the partitioning of elements to processors must be done in a way that ensures that the work is divided equally, i.e., that the computational load is balanced. For adaptive grid refinement a sequence of grids are generated, with a new partition computed at each stage; this process is called dynamic load balancing. Recently, we developed a new method for partitioning adaptive grids for distribution over parallel processors in the context of adaptive multilevel methods for solving PDEs. For triangular (2D) and tetrahedral (3D) grids, our method is guaranteed to produce connected partitions; no other partitioning method is known to do this. (Connected partitions reduce inter-process communication.)

The new method is based on the refinement tree, which arises naturally from the adaptive refinement process. The tree begins with a root node with one child for each of the elements in the initial coarse grid. In the case of the elements being triangles that are refined by bisection (see Figure 6), the refinement of an element replaces it with two new elements, which are represented in the refinement tree by two children of the corresponding node. The leaf nodes of the refinement tree correspond to the elements of the final fine grid to be partitioned. We apply weights to the nodes by assigning 1.0 to the leaf nodes, and the sum of the weights of the children to all other nodes. Thus the weight on each node is the number of elements in the final fine grid that are contained within the element corresponding to the node of the tree. Dividing the weight on the root (which is the number of elements in the fine grid) by the number of processors gives the number of elements that should be assigned to each processor to balance the load. Partitions are then constructed while performing a depth-first partial traversal of the refinement tree. We begin by setting the first

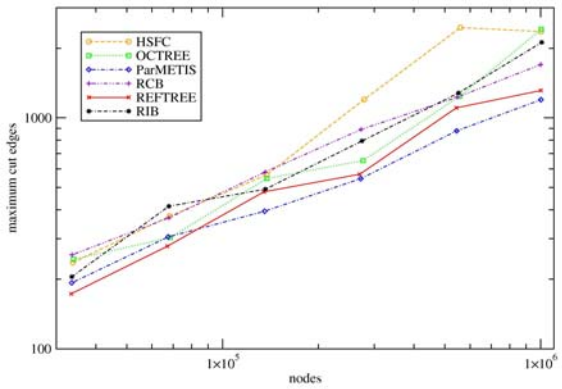

Figure 7. *Quality measure of a partition into 32 partitions for a sample adaptive grid.*

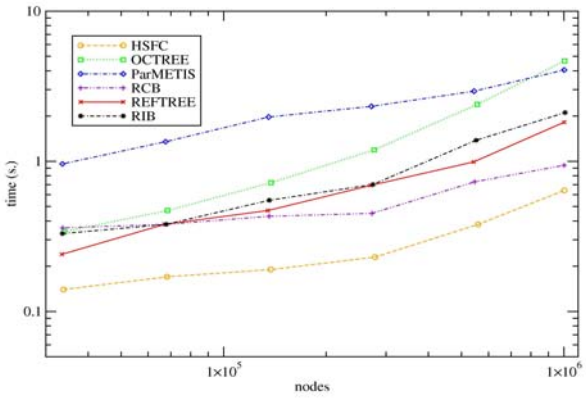

Figure 8. *Computation time to compute 32 partitions for a sample adaptive grid.*

partition to be empty and have a total weight of 0.0, and begin traversing the tree. If the weight on a node is small enough that adding it to the partially constructed partition keeps the total weight of the partition less than the desired number, the entire subtree (and corresponding elements) is added to that partition. Otherwise, the subtree is traversed. When the total weight is equal to the desired number, the partition is declared to be full and the next partition is begun. This method results in equal-sized partitions. Figure 6 provides an example of applying the method to partition a grid into two parts, which are represented by red and green. The sequences of colorations of the grid show the process of constructing the two partitions.

To guarantee that partitions are connected, the order of the children in the tree is crucial. By considering all possible cases, we were able to prove that there exists an order such that the partitions will be connected, and to provide the method for determining that order for all of the most popular element shape and element refinement techniques in 2D and 3D. This includes bisected and quadrisected triangles and quadrisected rectangles in 2D, and bisected and octasected tetrahedra, and octasected hexahedra in 3D. The order of the children is determined by propagating designated vertices called the in-vertex and out-vertex. The out-vertex of one element is the same as the in-vertex of an adjacent element. We proved in [1], with a constructive proof, that such a designation always exists for the stated element types, and in [2] provided the algorithm that propagates the designation to the children when elements are refined. When the grid is traversed by moving from one element to the next by passing through these vertices, a curve known as a space filling curve is created. The refinement-tree partitioning algorithm previously described is equivalent to cutting the space filling curve into segments, which guarantees that the partitions are connected. The designation of the in-vertex and out-vertex determines the order of the children in the tree.

We have implemented this method and compared it with other popular partitioning methods. Our method computes partitions in an amount of time similar to fast load balancing methods (which create inferior grids), and with mesh quality similar to slower methods (which produce well behaved grids). Figure 7 illustrates a measure of the partition quality, the maximum number of edges in the grid that span two partitions (which is related to the volume of interprocessor communication), as a function of the number of elements in the grid. The new method (colored red and labeled REFTREE) is outperformed only by the ParMETIS method. Figure 8 shows the amount of time required to compute these partitions. Here we see that the ParMETIS method takes much longer. Two of the methods are faster than REFTREE, but, relating back to Figure 7, produce inferior partitions. Our results demonstrate that methods that produce high quality partitions typically require excessive computation time. In contrast, our method produces high quality partitions at reasonable cost.

The publication reporting this work [2] received the 2008 ITL Best Journal Paper of the Year Award.

References

[1] W.F. Mitchell, Hamiltonian Paths Through Two- and Three-Dimension Grids, *NIST Journal of Research* **110** (2005), pp. 127-136.

[2] W. F. Mitchell, A Refinement-tree Based Partitioning Method for Dynamic Load Balancing with Adaptively Refined Grids, *Journal of Parallel and Distributed Computing* **67** (2007), pp. 417-429.

Participants

William F. Mitchell

http://math.nist.gov/phaml

Program area: High Performance Computing and Visualization (Division)

Methods for Characterizing Complex Networks

A type of data of increasing importance in information technology is one not based on continuum models but rather on connections between objects which can be modeled as graphs or networks. Complex systems of this type arise in many applications, such as the power grid, the Internet, communication systems, and transportation systems, situations where resources are provided through restricted channels.

In order to quantify the security and reliability of complex IT systems it is important to develop an underlying measurement science based on sound mathematical methods. Unfortunately, the size of real networks often makes exact measurement impossible. Thus, we are building efficient tools to approximate properties of real networks. An important aspect is the creation of realistic models to support research in this field. In particular, it is necessary to understand how real networks differ from random models.

Towards these goals, we have been developing new methods to generate random network models constrained to have a given degree sequence and s-metric. In addition, we are developing Monte Carlo methods and software based on sequential importance sampling for measuring properties of these networks. Applications considered have been counting the number of spanning trees of a graph and approximating the coefficients of the reliability polynomial.

Isabel Beichl and Brian Cloteaux

A *graph* is a collection of nodes (or vertices) and a set of connections between them, called edges. The *degree* of a node is the number of edges that are connected to it. A graph is termed *connected* if there is a path between any two nodes. Graphs are very useful in modeling of interconnected systems in a wide variety of applications from communications to biology.

Network Model Generation. The theory of random graphs is well developed. However, many real world systems cannot be approximated using simple random graphs (the Erdős-Renyi model). One reason for this is that the distribution of nodal degrees of random graphs follows a Poisson distribution, whereas it has been shown that the degree distribution for many real-world systems follows a power law distribution. In other words, the probability p of a node having degree k is $p \sim k^{\alpha}$ for some α. Nevertheless, generating random graphs with a power law degree distribution is still not sufficient to model real networks. For any given degree sequence, there can be a large number of non-isomorphic graphs with that sequence.

We have shown that real networks can have very different properties from random networks with the same degree distribution. We demonstrated this by examining the structure of the autonomous systems (AS) topology of the Internet as measured at UCLA [6]. We compared this with a series of random graphs with the same degree sequence generated using an algorithm of Blitzstein and Diaconis [3]. The conclusion was that simple degree-based models do not capture important details contained in real world data.

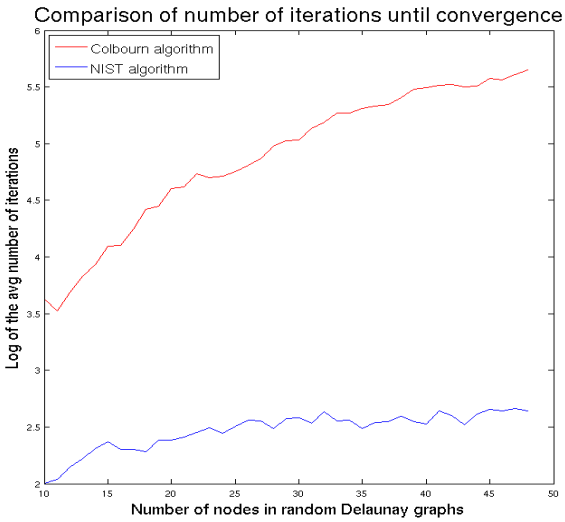

Figure 9. *Comparison of Colbourn and NIST algorithms*

Thus, network models need more structure to be realistic. To be able to further distinguish between networks with identical degree sequences, Li *et al.* [5] introduced the *s-metric*. This measures the connectedness of high degree nodes in networks having a power law distribution. Formally, the s-metric value for a graph G is defined as

$$s(G) = \sum_{(i,j)\in E(G)} \omega_i \cdot \omega_j$$

where $E(G)$ is the set of edges and ω_i is the degree of vertex i.

We have developed a method to generate graphs with a given degree sequence and an s-metric close to a given target [1]. It uses a Monte Carlo technique called threshold acceptance. If we have a graph with a given degree sequence, we use random edge switches that do not change the degree sequence but can increase or reduce the s-metric. While the models generated this way are a definite improvement, they still do not capture certain critical structural details of the AS topology [2]. We are continuing to develop new methods for more accurate network modeling.

Computing the Reliability Polynomial. A simplified measure of the reliability of a network is its reliability

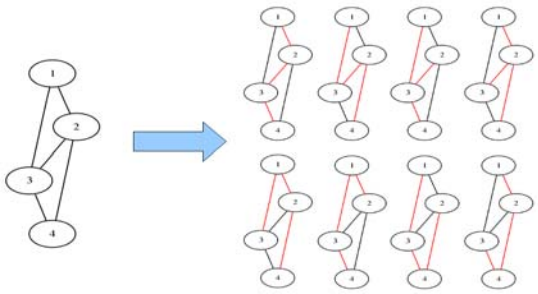

Figure 10. *Illustration of the number of spanning trees in a graph.*

polynomial. This function gives the probability that a network will remain connected given that each edge in it can fail independently with a probability p. One definition for the reliability polynomial R of a connected graph G with n vertices and m edges and for the edge failure probability p is

$$R(G;p) = \sum_{j=1}^{m} a_j (1-p)^j p^{m-j}$$

where a_j is the number of distinct connected subgraphs of G with exactly j edges.

The problem in using the reliability polynomial for network analysis has been the fact that determining the coefficients of the polynomial itself is computationally intractable. In fact, it has been shown that, in general, this problem is #*P*-complete. Thus the best we can expect to do in general is to approximate the polynomial. Our current research has led to a new method to do this. Our approach has two advantages. First, while there are several algorithms that provide pointwise approximations of the polynomial, our approach approximates the polynomial coefficients themselves. Secondly, when compared to the known approximation method of Colbourn *et al.* [4], our method has shown much faster rate of convergence (see Figure 9).

Counting Spanning Trees. A *spanning tree* in a graph is a set of edges such that each node is connected and there is no cycle in the edges (see Figure 10). The number of spanning trees in a graph can be a measure of reliability. As the number of spanning trees increases, the number of edges that need to be cut to disconnect the graph also increases (i.e., more spanning trees means many alternate paths between any two nodes).

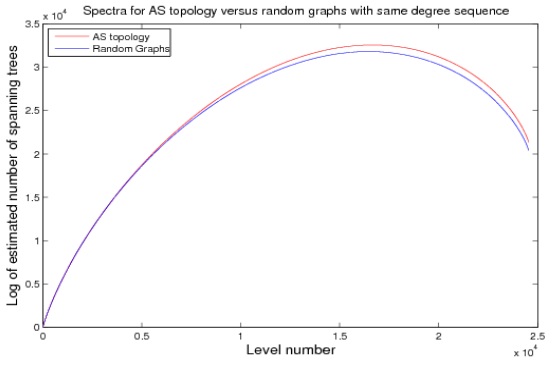

Figure 11. *Comparison of the number of sub-forests of size k (level number) in the AS network and a random model. The number of spanning trees is the value of the right-most point.*

In order to efficiently count the number of spanning trees in a network, we have developed a Monte Carlo method based on sequential importance sampling. The reason for creating an approximation algorithm to count spanning trees is slightly different than for estimating the reliability polynomial. There is an exact method to count all spanning trees of a graph in polynomial time. Unfortunately, this method becomes impractical for the size of the many of networks we measure. An example of a larger network that our new algorithm has let us analyze is the AS topology of the Internet. One instance of the AS topology, which we obtained from the UCLA Internet topology database, contained 24,567 nodes and 102,946 edges. Using our algorithm, we were able to compare random graphs to this AS topology in terms of number of spanning trees. This is shown in Figure 11.

References

[1] I. Beichl and B. Cloteaux, Generating Network Models using the *S*-metric," in *Proceedings of the International Conference on Modeling, Simulation and Visualization Methods*, 2008, pp. 159–164.

[2] I. Beichl and B. Cloteaux, Measuring the Effectiveness of the *s*-metric to Produce Better Network Models," in *Proceedings of the Winter Simulation Conference*, 2008, pp. 1020–1028.

[3] J. Blitzstein and P. Diaconis, A Sequential Importance Sampling Algorithm for Generating Random Graphs with Prescribed Degrees," submitted. 2006.

[4] C. J. Colbourn, B. M. Debroni, and W.J. Myrvold, Estimating the Coefficients of the Reliability Polynomial, *Congressus Numerantium* **62** (1998), pp. 217–223.

[5] L. Li, D. Alderson, J. C. Doyle, and W. Willinger, Towards a Theory of Scale-free Graphs: Definition, Properties, and Implications, Internet Mathematics **2** (2005), pp. 431–523.

[6] UCLA. UCLA Internet topology collection. http://irl.cs.ucla.edu/topology, 2008, accessed 1/14/2008.

Participants

Isabel Beichl and Brian Cloteaux (MCSD); Francis Sullivan (IDA Center for Computing Sciences)

Program area: Complex Systems (ITL)

TCP Metastability and Cascading Failures in the Internet

Computer networks that provide the foundation for modern commerce are astounding in their complexity. Hundreds of millions of nodes interact in real time controlled by lean communication protocols. Unfortunately, unexpected macroscopic behavior, such as waves of network congestion and failure, have been observed in this large complex system. Such events have already led to the loss of opportunity and productivity valued in the billions of dollars. Our fundamental understanding of such behavior remains weak. It remains impossible to reliably predict the emergent collective behavior of the large complex interconnected systems that we can now readily build and deploy. The goal of this project is to develop and investigate realistic models of the Internet subject to flow control protocols. A preliminary analysis of a mean-field model with retransmissions induced by control feedback indicates the possibility of instability in cases thought to be stable from the analysis of simpler models.

Vladimir Marbukh

The Transmission Control Protocol (TCP) controls data rates of the majority of Internet users. While a large number of TCP versions have been implemented, and even more versions have been proposed, all rely on Explicit Congestion Notifications (ECN) as the indication of congestion on the end-to-end path. Under this protocol, each controlled TCP source increases its data rate until it encounters congestion as indicated by an ECN. Once an ECN is received, the source reduces its rate to reduce congestion. In this way the Internet operates as a distributed control system with feedback provided by Explicit Congestion Notifications.

High dimension, non-linearity and delays in the feedback loop make stability and performance evaluation of the Internet under TCP control a very challenging problem. Observations of the Internet, simulations of Internet fragments, and analytical modeling under simplifying assumptions, suggest that the TCP controlled Internet is stable under "normal" operating conditions.

Given the number of flows $N = N_r$ carried on all feasible routes $r \in R$, the equilibrium source transmission rates x_r can be approximated by the weighted (α,w)-fair rate assignment [1],

$$x_r = \left(w_r \Big/ \sum_{l \in r} p_l \right)^{1/\alpha}, \qquad (1)$$

where the parameters α, $w_r > 0$ characterize TCP version-specific fairness, and the parameters p_l represent the average rate of Explicit Congestion Notifications generated by links l on route r. The average transmission rate of each TCP controlled Internet user is determined by the aggregate ECN rate from all links on the user's route r. For most widely used TCP versions, ECNs indicate packet losses and thus the parameters p_l in (1) represent the packet loss probabilities on links l. When $w_r = 1$, the cases $\alpha \to 0$, $\alpha \to 1$, and $\alpha \to \infty$ in (1) correspond, respectively, to a rate assignment which achieves maximum throughput, is proportionally fair, or is max-min fair. The case $\alpha = 2$ corresponds to the so-called TCP-Reno rate assignment in the case of small packet losses $p_l << 1$.

The number of active Internet users fluctuates over time, with each user alternating between active and silent periods, causing arrival and departures of flows carrying files, streaming voice and video. Markov models of fair bandwidth sharing (1) under arriving/departing flows have been proposed and analyzed in [2,3,4]. This analysis suggests that the Internet is stable under arriving/departing flows if the average exogenous file transfer load is less than the link capacity for all network links, that is,

$$\rho_l \overset{def}{=} \frac{1}{C_l} \sum_{r:l \in r} \Lambda_r b_r < 1, \ \ \forall l \in L \,. \qquad (2)$$

In (2) the average arrival rate of files to be transferred on route r is Λ_r, the average size of these files is b_r, the capacity of link l is C_l, and hence the utilization on link l is ρ_l.

Nevertheless, simulations [5] have demonstrated that in the case when ECNs indicate lost packets, and thus parameters p_l in (1) indicate probability of packet loss, the network may be unstable even when conditions (2) are satisfied for all Internet links l. The reason for the inconsistency between analytical modeling [2,3,4] and simulations [5] is that models [2,3,4] do not account for retransmissions of lost packets, which TCP uses to ensure reliable transport.

The goal of our project is to develop and investigate realistic models of TCP controlled Internet under arriving/departing flows. A mean-field model which incorporates packet retransmissions has been proposed in [6,7]. The model accounts for additional link load due to packet retransmissions on links carrying packets dropped downstream and then retransmitted. This additional load increases with increase in the number of flows in progress, creating a positive feedback: sufficient increase in the number of flows in progress due to temporary load increase of fluctuation further increases the number of flows in progress, etc. Consistent with previous simulation results [5], our preliminary stability analysis indicates that this positive feedback may

cause model instability even when conditions (2) are satisfied. Model instability implies that the number of flows in progress grows unbounded in time resulting in congestion collapse. Of course, infinite growth of the number of flows in progress is an abstraction since in the Internet demand drops as user dissatisfaction grows with the increase in the number of flows in progress.

Preliminary analysis [6,7] has also produced new findings, including the following. A large-scale TCP controlled network with exogenous load satisfying (2) has a desirable metastable, i.e., persistent but slowly changing, network state with a finite number of flows in progress. Despite the fact that eventually the network may become unstable resulting in infinitely growing number of flows in progress, on time scales of practical importance the network converges to the desirable metastable state if the initial number of flows in progress is sufficiently small. By analogy with phase transitions of the first kind, one may expect that eventual transition from desirable metastable state to congestion collapse occurs through process of cascading failures: supercritical local overload may spread over a major portion of the network causing global congestion collapse. Our preliminary results indicate that appropriately designed flow admission control can transform the desirable metastable state into a stable state, eliminating instability at the cost of a small flow rejection probability.

Current and future research involves the validation of model results, estimation of the stability margin of the desirable metastable state, and developing practical techniques for preventing cascading failures and the possibility of global congestion collapse. More long-term plans include investigation of the effect of specific TCP parameters, including fairness, on TCP stability and performance. Evaluation and balancing of the numerous trade-offs between TCP fairness, stability, performance and robustness under fixed and fluctuating sets of flows in progress are matters of significant and urgent practical importance for optimization and design of current and future versions of the TCP protocol.

References

[1] J. Mo and J. Walrand, Fair End-to-end Window-based Admission Control, *IEEE/ACM Transactions on Networking* **8** (2000), pp. 556-567.

[2] S. Ben Fredj, T. Bonald, A. Proutiere, G. Regnie, and J. Roberts, Statistical Bandwidth Sharing: a Study of Congestion at Flow Level, SIGCOMM 2001.

[3] G. de Veciana, T. J. Lee, and T. Konstantopoulos, "Stability and Performance Analysis of Networks Supporting Elastic Services," *IEEE/ACM Transactions on Networking* **9** (2001), pp. 2-14.

[4] P. Key, L Massoulie, A. Bain, and F. P. Kelly, Fair Internet Traffic Integration: Network Flow Models and Analysis, *Annales des Telecommunications* **59** (2004), pp. 1338-1352.

[5] L. Massoulié and J. W. Roberts, Arguments in Favor of Admission Control for TCP Flows, in *Proceedings of the 16th International Teletraffic Conference*, Edinburgh, UK, June 1999.

[6] V. Marbukh, Metastability of Fair Bandwidth Sharing under Fluctuating Demand and Necessity of Flow Admission Control, *Electronics Letters* **43** (19) (13 September 2007), pp. 1051-1053.

[7] V. Marbukh, Can TCP Metastability Explain Cascading Failures and Justify Necessity of Flow Admission Control in the Internet? in *Proceedings of the 15th International Conference on Telecommunications*, St. Petersburg, Russia, June 2008.

Participants

Vladimir Marbukh

Program area: Complex Systems (ITL)

Regularization by Residual Periodigrams

First kind integral equations are routinely used to model instrument distortions when trying to identify an unknown function. Examples abound in measurement science, from models of optical systems (i.e., imaging) to temperature probes. The solution to such problems is extremely challenging. Known as "ill-posed" problems in mathematics, solutions to such problems are highly sensitive to small variations in input data (such as noise), or to small errors made in the course of numerical solution. To stabilize the solution process, a technique known as "regularization" is necessary. Such procedures involve the selection of regularization parameters, which in practice is often done in an ad hoc way. Few metrics for evaluating the success of regularization parameter selection are known. We have developed such a metric that assures that the resulting model is statistically plausible. In particular, candidate parameter choices are evaluated by inspecting the model residual along with its periodogram and cumulative periodogram. Based on this criterion, we propose a new method for choosing regularization parameters that makes the residuals as close as possible to white noise, using a diagnostic test based on the periodogram. We have compared this new method with standard techniques such as the discrepancy principle, the L-curve and generalized cross validation, showing that it performs better on a variety of standard problems, as well as two new ones.

Bert W. Rust and Dianne P. O'Leary

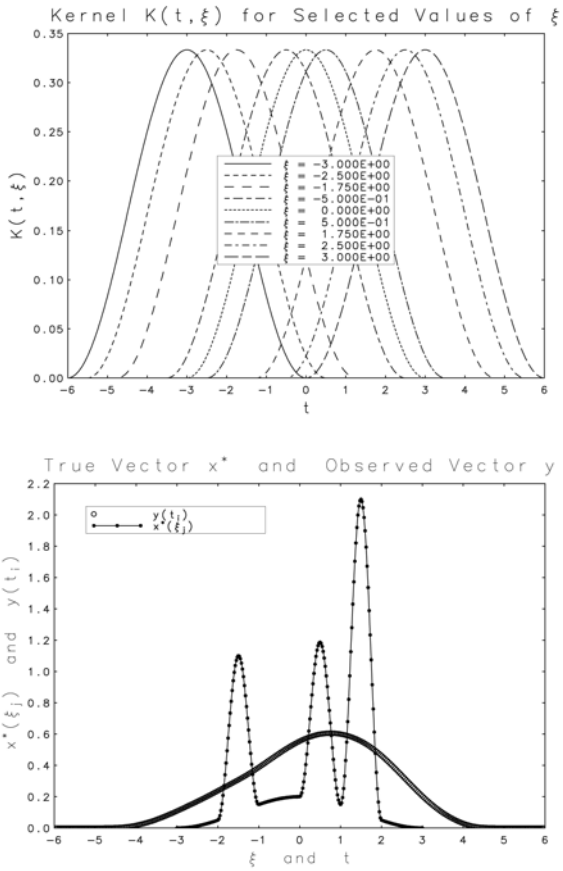

Figure 12. *A variant of the Phillips problem with m = 300, n = 241.*

Many physical measurements $y(t_i)$, such as those in spectroscopy, remote sensing, and image reconstruction, can be modeled by a system of linear, first kind integral equations

$$y(t_i) = \int_a^b K(t_i,\xi)x(\xi)\,d\xi \;+\; \varepsilon_i \quad i = 1,2,\ldots,m \qquad (1)$$

where $x(\xi)$ is the function being measured, the $K(t_i,\xi)$ are known instrument response functions, and the ε_i are random measuring errors which we assume are normally distributed with a known positive definite covariance matrix $\mathbf{S}^2$, i.e.,

$$\boldsymbol{\varepsilon} = (\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_m)^T \sim N(0, \mathbf{S}^2). \qquad (2)$$

Discretizing the integrals with an n-point quadrature rule produces an ill-conditioned linear regression model,

$$\mathbf{y} = \mathbf{K}\mathbf{x}^* + \boldsymbol{\varepsilon}, \quad \text{with } \boldsymbol{\varepsilon} \sim N(0, \mathbf{S}^2), \qquad (3)$$

Where $m \geq n$ and rank($\mathbf{K}$) = n. An example problem is shown in Figure 1. Premultiplying (3) by $\mathbf{S}^{-1}$ and defining

$$\mathbf{b} \equiv \mathbf{S}^{-1}\mathbf{y}, \quad \mathbf{A} \equiv \mathbf{S}^{-1}\mathbf{K}, \quad \boldsymbol{\eta} \equiv \mathbf{S}^{-1}\boldsymbol{\varepsilon} \qquad (4)$$

Gives a model with independent $n(0,1)$ errors, i.e.,

$$\mathbf{b} = \mathbf{A}\mathbf{x}^* + \boldsymbol{\eta}, \quad \text{with } \boldsymbol{\eta} \sim N(0, \mathbf{I}_m). \qquad (5)$$

Because of ill-conditioning, the usual least squares estimate

$$\hat{\mathbf{x}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{b} \qquad (6)$$

gives a wildly oscillating, spurious approximation to $\mathbf{x}^*$.

If $\tilde{\mathbf{x}}$ is any estimate of $\mathbf{x}^*$, then comparing the residual vector $\tilde{\mathbf{r}} \equiv \mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}$ with the error vector

$$\boldsymbol{\eta} = \mathbf{b} - \mathbf{A}\mathbf{x}^*, \quad \boldsymbol{\eta} \sim N(0, \mathbf{I}_m) \qquad (7)$$

suggests that $\tilde{\mathbf{x}}$ is acceptable only if $\tilde{\mathbf{r}}$ is a plausible sample from the $\boldsymbol{\eta}$-distribution. Under our assumptions,

$$\|\boldsymbol{\eta}\|^2 = \|\mathbf{b} - \mathbf{A}\mathbf{x}^*\|^2 \sim \chi^2(m), \qquad (8)$$

whence

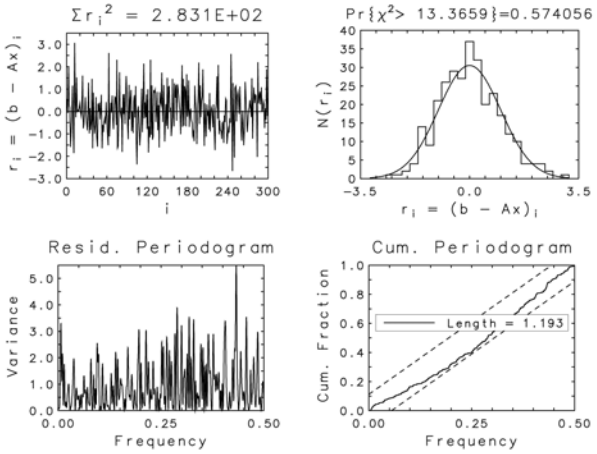

Figure 13. *Residual diagnostics for the true solution x*.*

$$\mathrm{E}\left\{\|\boldsymbol{\eta}\|^2\right\} = m, \quad \mathrm{Var}\left\{\|\boldsymbol{\eta}\|^2\right\} = 2m \tag{9}$$

It follows that the $\tilde{r}_i$ should be independently distributed like $n(0,1)$ with its norm expected to satisfy

$$\left(m - 2\sqrt{2m}\right) \le \|\tilde{\mathbf{r}}\|^2 \le (m + 2\sqrt{2m}), \tag{10}$$

and the sequence $\{\tilde{r}_1, \tilde{r}_2, \ldots, \tilde{r}_m\}$ should comprise a white noise time series. To illustrate these conditions, we added random errors, with

$$\mathbf{S} = (10^{-5})\mathrm{diag}\left(\sqrt{(\mathbf{K}\mathbf{x}^*)_1}, \cdots, \sqrt{(\mathbf{K}\mathbf{x}^*)_m}\right), \tag{11}$$

to the problem in Figure 12, computed the residuals corresponding to $\mathbf{x}^*$, and plotted them (with $\Delta t = 1.0$) in the upper left panel of Figure 13 and as a histogram in the upper right panel. The sum of squared residuals, SSR = 283.1, is well inside the interval (10), and the histogram is well modeled by the renormalized $n(0,1)$ probability density function which is also plotted. The χ^2 test indicates statistically significant consistency between the two curves.

The bottom two plots in Figure 13 give the *periodogram P* and the *cumulative periodogram C* of the residual time series. The former was obtained by zero-padding the series to length 8192, computing its discrete Fourier transform

$$R_k = \sum_{j=1}^{m} \tilde{r}_j \exp\left(-\frac{2\pi \mathbf{i}\, jk}{8192}\right), \quad k = -4096, \ldots, 0, \ldots, 4096 \tag{12}$$

where $\mathbf{i} = \sqrt{-1}$, and then computing

$$P_k = \frac{1}{8192}|R_k|^2, \quad k = 0,1,\ldots,4096. \tag{13}$$

The periodogram gives the distribution of variance with frequency $f_k = k/(8912\Delta t)$. The cumulative periodogram is then computed by

$$C_k = \frac{\sum_{j=0}^{k} P_j}{\sum_{j=0}^{4096} P_j}, \quad k = 0,1,\ldots,4096. \tag{14}$$

The two dashed lines, centered on the diagonal

$$C(f) = 2f, \quad 0 \le f \le 0.5, \tag{15}$$

define a 95% confidence band for white noise. None of the 4096 C_k lie outside this band, so the residuals are an acceptable realization of white noise.

The least squares residuals $\hat{\mathbf{r}} = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}$ do not resemble the true residuals in Figure 13, with $\|\hat{\mathbf{r}}\|^2 = 43.01$ far outside the interval (10) and 2803 of its 4096 C_k lying outside the white noise band. The estimate $\hat{\mathbf{x}}$ is a wildly unstable, oscillating vector. Stabilized estimates can be obtained by appending a set of regularization constraints to the regression model. If the estimate is constrained to lie near the origin, the procedure is called *Tikhonov regularization*, and the augmented model becomes

$$\begin{pmatrix} \mathbf{b} \\ \mathbf{0} \end{pmatrix} = \begin{pmatrix} \mathbf{A} \\ \lambda \mathbf{I}_n \end{pmatrix} \mathbf{x}(\lambda) + \begin{pmatrix} \boldsymbol{\eta} \\ \lambda \boldsymbol{\gamma} \end{pmatrix}, \tag{16}$$

where $\boldsymbol{\gamma}$ is a $N(0, \mathbf{I}_n)$-distributed error vector and λ is a user-chosen parameter that determines how much weight is given to the regularizing constraints. Minimizing the sum of squared residuals gives

$$\tilde{\mathbf{x}}(\lambda) = (\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T\mathbf{b}, \tag{17}$$

for the regularized estimate. The success of the method depends crucially on the choice of the parameter λ. Previously developed methods for making this choice include: (1) the *L-curve* method, (2) minimizing the *Generalized Cross Validation* (GCV) function, and (3) the *Discrepancy Principle* method which, by Eq. (9), can be implemented by choosing λ to satisfy

$$\|\tilde{\mathbf{r}}(\lambda)\|^2 = \|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}(\lambda)\|^2 = m. \tag{18}$$

In 2000, Rust [1] suggested choosing λ to minimize the sum of squared deviations (SSD) of the cumulative periodogram $C(\lambda)$ from the white noise line (15), i.e.,

$$\mathrm{SSD}_{\min} = \min_{\lambda}\left\{\sum_{k=0}^{4096}\left[C_k(\lambda) - 2f_k\right]^2\right\}. \tag{19}$$

In 2008, we proposed several alternative choices, also based on the periodogram [2]. Figure 14 gives a plot of (19) for the problem obtained by adding the errors (11) to the problem shown in Figure 12. The minimum value occurs at $\lambda_{\min} = 24.624$.

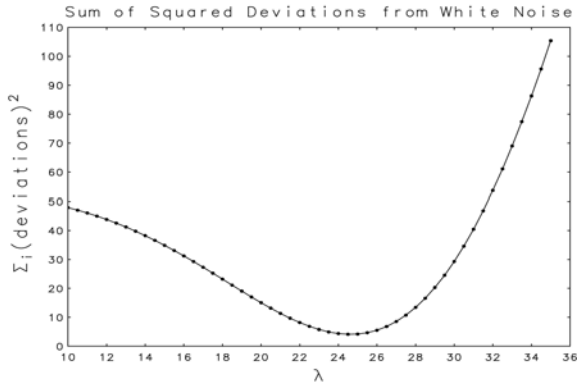

Figure 14. *Sum of squared deviation from white noise.*

Table 2. *Comparison of four different methods for choosing λ.*

Method	λ	$\|\tilde{\mathbf{r}}(\lambda)\|^2$	$\tilde{r}_i$ n.d.	$\tilde{r}_i$ w.n.	*rms*
Cum. Perd.	24.62	265.4	no	yes	0.0085
Discr. Princp.	28.31	300.0	yes	no	0.0084
Min. GCV	13.27	214.1	no	no	0.0127
L-curve Corner	1.61	177.2	no	no	0.1215

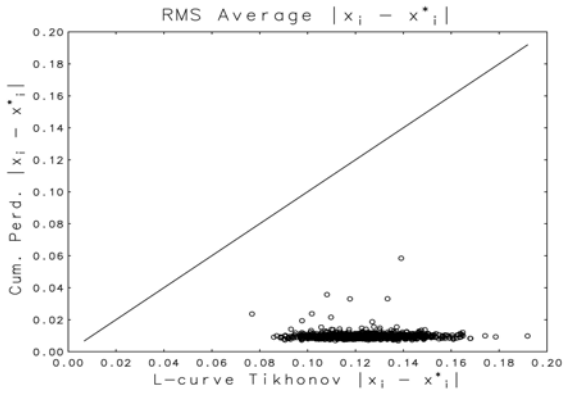

Figure 15. *Cumulative periodogram vs. L-curve.*

Table 2 gives results from the four different methods for choosing λ. In the column headings, "n.d." means "normally distributed", "w.n." means "white noise", and

$$rms\left|\tilde{x}_i - x_i^*\right| = \sqrt{\frac{1}{n}\sum_{j=1}^{n}\left|\tilde{x}_i(\lambda) - x_i^*\right|^2}\,. \qquad (20)$$

There is not much difference between the cumulative periodogram and discrepancy principle methods, but both are clearly superior to the L-curve and minimum GCV methods. Figure 15, Figure 16, and Figure 17 compare the rms errors of the methods for 1000 different realizations of the noise. The L-curve errors all greatly exceed those for the cumulative periodogram, and the minimum GCV errors mostly exceed them, but not by so much. The discrepancy principle errors are mostly only slightly better than those for the cumulative periodogram, so the latter is a reliable alternative to the former, and the cumulative periodogram is a useful diagnostic for any estimate.

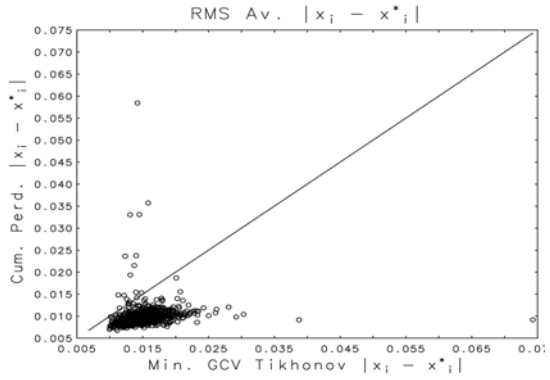

Figure 16. *Cumulative periodogram vs. minimum GCV.*

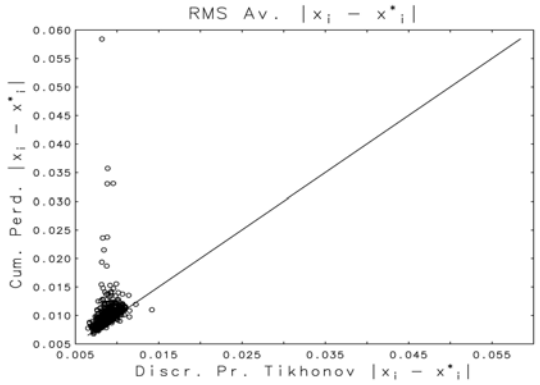

Figure 17. *Cumulative periodogram vs. discrepancy principle.*

References

[1] B. W. Rust, Parameter Selection for Constrained Solutions to Ill-Posed Problems, *Computing Science and Statistics* **32** (2000), pp. 333-347.

[2] B. W. Rust and D. P. O'Leary, Residual Periodograms for Choosing Regularization Parameters for Ill-Posed Problems, *Inverse Problems* **24** (2008), 034005.

Participants

Bert W. Rust and Dianne P. O'Leary

Program area: Mathematics of Metrology (Division)

Calibrating Image Roughness Using Lipschitz Exponents

We have developed a mathematical framework for quantifying the unsmoothness of images. The method is based upon the L^1 Lipschitz exponent α, where $0<\alpha\leq 1$, a useful indicator of fine-scale content, provided the image is relatively noise free. The use of this index has been enabled by an effective computational method for estimating image Lipschitz exponents. Applications of this metrology tool include monitoring of image sharpness and imaging performance in imaging systems, evaluation of image reconstruction software quality, detection of abnormal fine structure in biomedical images, and monitoring of surface finish in industrial applications.

We have studied the use of Lipschitz exponents in the quantitative evaluation of image reconstruction procedures. Using synthetically degraded images, it was shown that Lipschitz exponents can measure the extent to which competing denoising algorithms remove texture along with the noise. Likewise, Lipschitz exponents can measure the ability of competing deblurring algorithms to recover texture. Quantitative confirmation of fine structure recovery is particularly important in blind image deconvolution, where the cause of the blur is unknown. This was illustrated in the case of APEX blind deconvolution applied to Hubble space telescope imagery, SEM, and MRI and PET brain scans.

Alfred S. Carasso

A theorem in pure mathematics characterizing the behavior of Gaussian singular integrals is the basis for a useful image metrology tool that quantifies the degree of unsmoothness in an image. In October 2008, United States Patent #7437012 was awarded for this methodology.

Most natural images $f(x, y)$ display edges, localized sharp features, and other significant fine scale details or texture, and cannot be modeled as smoothly differentiable functions of x and y. In many digital image processing tasks, it is necessary to provide prior information that specifies the degree of unsmoothness in the unknown desired true image. If an image is incorrectly postulated to be too smooth, the processing algorithm may produce a smoothed out version of the true image, in which critical diagnostic information has been lost.

The L^1 Lipschitz exponent α, where $0 < \alpha \leq 1$, is a mathematical index that can capture the fine-scale content and degree of unsmoothness in an image, provided that image is relatively noise free. Images that are of bounded variation (including smoothly differentiable images) have $\alpha = 1$. The value of α decreases systematically with increasing roughness. In most images of interest, $0.1 < \alpha < 0.6$.

There are numerous contexts where a simple practical method of quantifying image roughness can be useful. Many imaging systems suffer performance degradation over time and require periodic maintenance. In scanning electron microscopes, the shape of the electron beam changes with time, often without the user's knowledge. Periodic performance testing can be accomplished by monitoring sharpness degradation in the micrograph of a specially designed test object. Such degradation can be quantified by measuring the increase in Lipschitz exponent. In some imaging systems, a decrease in α might indicate an increase in system noise. Measuring the value of α of a test image can also be used to evaluate and compare the performance of competing image reconstruction packages. Automatic measurement of α might be useful in monitoring the smoothness of surface finishes in certain industrial applications. In some medical applications, measurement of α, along with other appropriate processing, might be helpful in prescreening for certain types of abnormalities.

Calculating the Lipschitz exponent of an image $f(x, y)$ involves the following. For fixed $t > 0$, define the linear operator G^t by means of the Fourier series

$$G^t = \sum_{\xi,\eta=-\infty}^{\infty} \left\{ \begin{matrix} \exp[-t(\xi^2+\eta^2)]\hat{f}(\xi,\eta) \\ \times \exp[2\pi i(x\xi + y\eta)] \end{matrix} \right\}. \quad (1)$$

Let $\mu(t)$ be the L^1 norm relative error in approximating f with $G^t f$

$$\mu(t) = \frac{\left\| G^t f - f \right\|}{\left\| f \right\|}. \quad (2)$$

An image $f(x,y)$ has Lipschitz exponent α if and only if $\mu(t) = O(t^{\alpha/2})$ as t tends to zero. Using FFT algorithms, the Fourier series in Eq. (1) can be evaluated for each fixed $t_n > 0$ in a sequence $\{t_n\}$ tending to zero, together with $\mu(t_n)$ in Eq (2). By plotting $\mu(t_n)$ versus t_n on a log-log scale, one can locate positive constants C and α such that $\mu(t) \leq Ct^{\alpha/2}$ as t tends to zero. This is illustrated in Figure 18 for the Giza Pyramids image. The solid curve A is a plot of $\mu(t)$ vs t, and exhibits a characteristic elbow shape. That curve can be majorized by the two straight lines Γ and Σ. The line Γ reflects behavior at very small values of t, while Σ models behavior for $\log t > -10$. The Lipschitz exponent α is equal to twice the slope of the Σ line. Behavior associated with the Γ line is unrelated to image roughness and must be discarded.

Figure 19 illustrates the behavior of the Lipschitz exponent as an image is degraded and restored. Traces

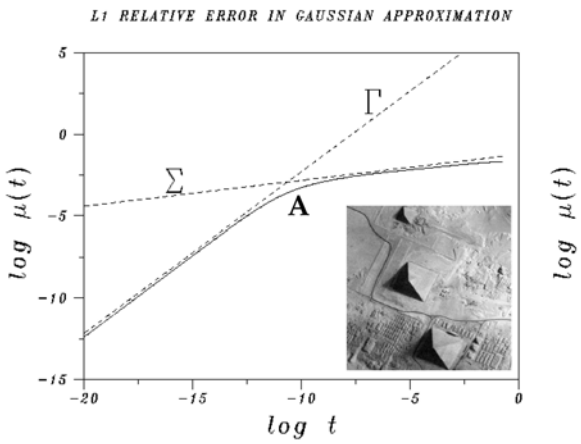

Figure 18. *The pyramids image has a Lipschitz exponent α = 0.318, obtained by doubling the slope of the § line.*

A, B, C, and D correspond to images (a), (b), (c), and (d) respectively. The Σ lines were not drawn to avoid clutter. Sharp image (a) with $\alpha = 0.591$, is degraded by adding salt and pepper noise to form image (b), resulting in a substantially lower $\alpha = 0.302$. Denoising image (b) using nonlinear partial differential equation (pde) methods produces image (c) with substantially larger $\alpha = 0.714$. This reflects the fact that pde denoising eliminates texture along with the noise, producing a lifeless image that resembles a department store mannequin. However, using 2D median filtering produces image (d) with $\alpha = 0.645$, which is close to the original. In this experiment, the Lipschitz exponent can rank order the effectiveness of competing restoration algorithms.

APEX blind deconvolution is a NIST developed technology that has been found useful in a variety of applications. Measurable reconstruction of fine scale information is one of several elements that bear on the plausibility of APEX images. Such sharpening should be accompanied by a quantifiable decrease in Lipschitz exponents. Figure 20 deals with APEX processing of state-of-the-art nanoscale Scanning Electron Microscopy. The original *Au/Pd* decorated magnetic tape sample micrograph in image (a) has trace *A* and Lipschitz exponent $\alpha = 0.303$. The APEX-processed image (b) has $\alpha = 0.234$, a 23% decrease. Figure 21 deals with positron emission tomography in brain research. The original transverse functional PET brain slice in image (a) has $\alpha = 0.502$. APEX processing in image (b) reveals several additional active bright spots, and results in $\alpha = 0.394$, a 21% decrease. Figure 22 deals with a Kitt Peak Observatory image of the Andromeda galaxy M31. The original image (a) has $\alpha = 0.287$. The APEX processed image (b) is noticeably sharper and has $\alpha = 0.153$, a 47% decrease.

References

[1] A. Carasso and A. Vladar, Calibrating Image Roughness by Estimating Lipschitz Exponents, with Applications to Image Restoration, *Optical Engineering* **47** (3) (March 2008), 037012, 13 pages

Participants

Alfred S. Carasso (MCSD); András E. Vladár (MEL)

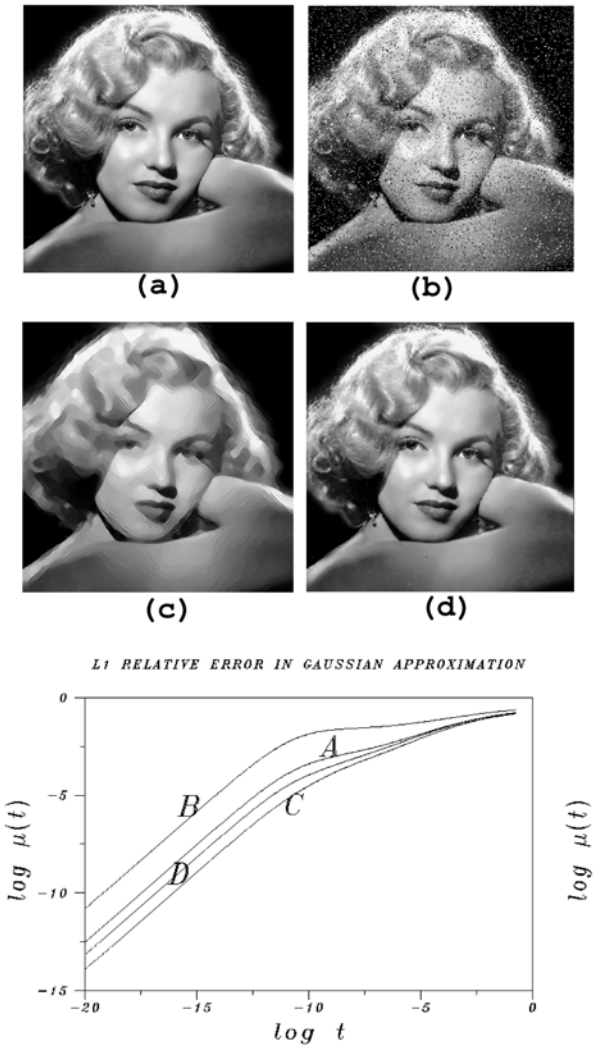

Figure 19. *Comparing image denoising algorithms. Noise addition can artificially lower Lipschitz exponent α while some noise removal algorithms can eliminate texture and increase α. Gaussian traces A, B, C, and D correspond to images (a), (b), (c), and (d) respectively. Original (a) has α = 0.591, (b) has α = 0.302, (c) has α = 0.714, and (d) has α = 0.645.*

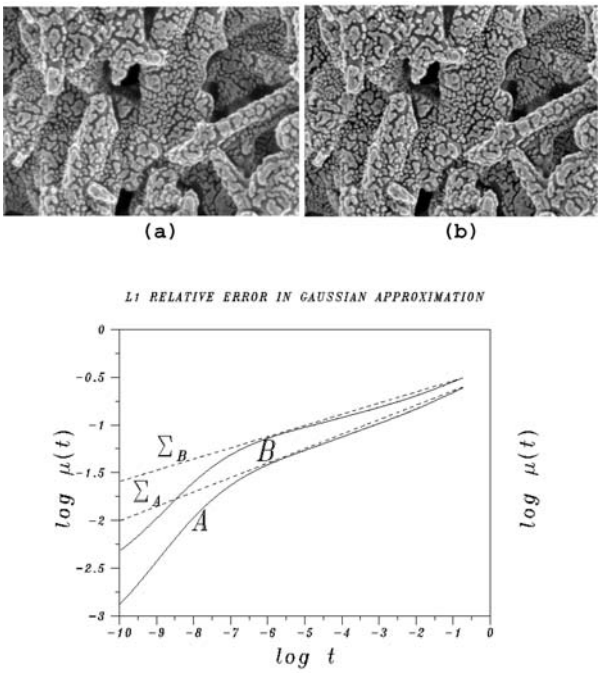

Figure 20. *APEX blind deconvolution in nanoscale Scanning Electron Microscopy. Original micrograph (a) of Au/Pd decorated magnetic tape sample has α = 0.303. APEX processed image (b) has α = 0.234, a 23% decrease.*

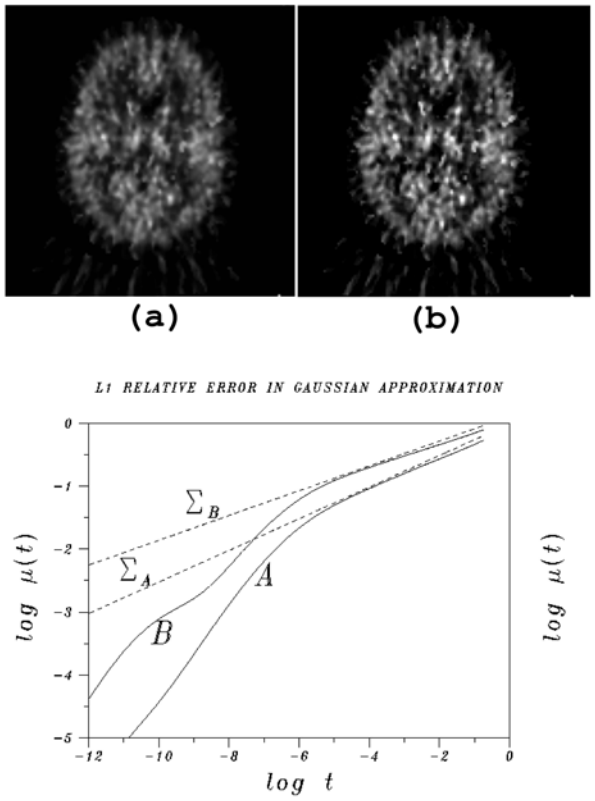

Figure 21. *APEX blind deconvolution of transverse functional PET brain slice. Original (a) has α = 0.502. APEX processed (b) has α = 0.394, a 21% decrease.*

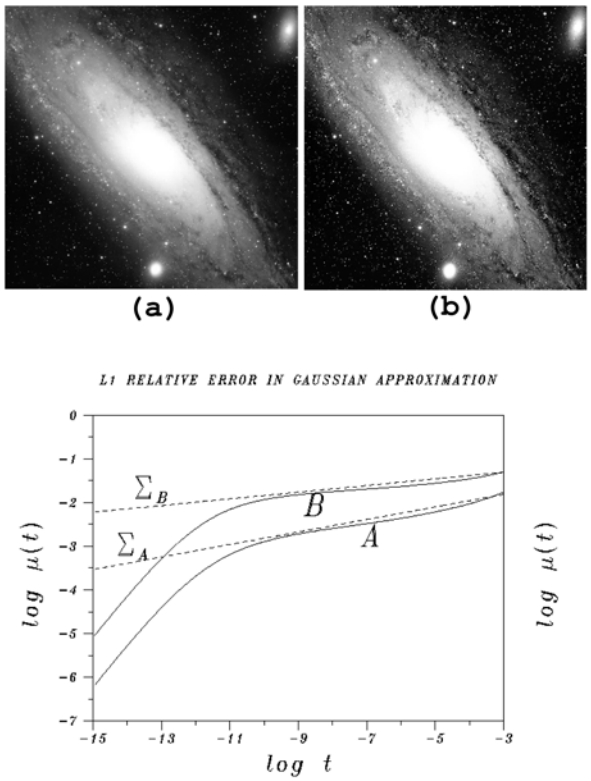

Figure 22. *APEX blind deconvolution of Kitt Peak Observatory Andromeda galaxy image. Original (a) has α = 0.287. APEX processed (b) has α = 0.153, a 47% decrease.*

Program area: Enabling Scientific Discovery (ITL)

Modeling Grain Boundary Premelting in Binary Alloys

The performance of solid-state devices is strongly related to the structure and properties of the underlying materials from which the device is fabricated. Particularly for small-scale devices, it is important to understand how the processing of the material affects the resulting microstructure of the components. Material scientists often cannot process perfect, single-crystal structures, but instead must deal with polycrystalline material whose microstructure includes networks of interfaces ("grain boundaries") that separate the individual domains of single crystal ("grains") in the material.

This project is a theoretical study that is intended to elucidate experimental measurements of the capillary properties of grain boundaries in metallic alloys near the melting point of the material. At these temperatures the incipient liquid phase can intrude into the grain boundary, with a resulting change in the properties of the grain boundaries and, hence, the overall material and its macroscopic properties. This "wetting" of the grain boundary by the liquid phase can occur even when the temperature is below the melting point of the alloy ("premelting"). We study premelting in the context of a continuum theory that models the grain boundary as a diffuse interface of finite width, but with interfacial properties that can be computed in a thermodynamically consistent fashion. This approach complements alternative approaches that employ atomistic descriptions of the underlying structure. The result is a prediction of the extent of premelting as a function of processing temperature and alloy composition, including the associated changes in the surface properties of the grain boundary.

Geoffrey B. McFadden

Grain boundaries are interfaces that occur in polycrystalline solids where the orientation of a crystal lattice changes in passing from one single crystal "grain'' of material to a neighboring grain. Grain boundaries have a variety of material properties that distinguish them from regions of bulk crystal, including surface energy, surface diffusion, surface adsorption, and surface stress. These surface phenomena can all affect the mechanical and/or electrical properties of the material, particularly at small length scales where surface effects compete with volume effects. Over time grain boundaries can also migrate in response to capillary forces or diffusion processes, which can lead to significant dynamic changes in material properties. For these reasons grain boundaries are actively studied by material scientists, with considerable efforts being made using each of experimental, theoretical, and computational approaches.

In this project, researchers in MCSD, the Metallurgy Division in MSEL, and the Department of Physics and Astronomy at George Mason University, are collaborating in a computational study of the properties of grain boundaries in a binary alloy at temperatures near the melting point of the alloy. At the melting point the solid and liquid phases of the material can co-exist in thermodynamic equilibrium, and the phases are separated by a solid-liquid interface where the phase transition occurs. At temperatures below the melting point the solid phase is favored, and regions of bulk liquid phase are unstable to the formation of solid. However, at a grain boundary between two solid grains the liquid phase can preferentially "wet" the grain boundary, even at temperatures below the melting point where the liquid is not the stable bulk phase. Roughly speaking, the liquid wets the grain boundary if the surface energy of the "dry" grain boundary is greater than the combined surface energy of two solid-liquid interfaces, so that wetting lowers the overall energy of the system. This "pre-melting" of the grain boundary by the liquid phase at temperatures below the melting point is another type of surface phenomenon. Whether or not it has been observed experimentally in metallic binary alloys is actually a subject of some dispute among material scientists, and is one of the main reasons for the present study.

Two approaches were used to study grain boundaries theoretically, a continuum [1] and an atomistic [2] description of the grain boundary and neighboring material. The two approaches were found to give consistent results that allow the prewetting behavior of grain boundaries to be modeled. The continuum model that we summarize here is based on a description of the material by using three "order parameters" or "phase field" variables φ_j that indicate the local thermodynamic state of the material at each point in space [3]. The phase-field variables can be interpreted as local phase fractions, with, for example, $\varphi_3(x,t) = 1$ representing a bulk liquid phase at point x and time t. The other two phase field variables each represent grains of given orientation, with all three variables normalized to sum to unity. The composition of the material, $c(x,t)$, is an additional variable in the description, and the temperature, T, is taken to be uniform throughout the sample. A consistent thermodynamic model of the binary allow is then employed [4] that is based on the composition, temperature, and phase-field variables. The specific model was chosen to reproduce the phase diagram of the copper-silver system.

States of thermodynamic equilibrium are computed by minimizing a free energy functional, which

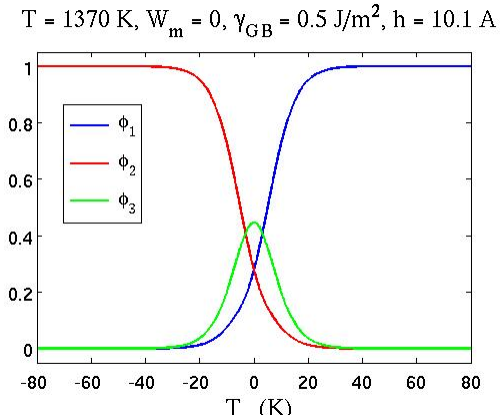

Figure 23. *Phase-field profiles for a grain boundary located near $x = 0$ illustrating wetting of the interface ($\varphi_3 = 0$) in the interfacial region.*

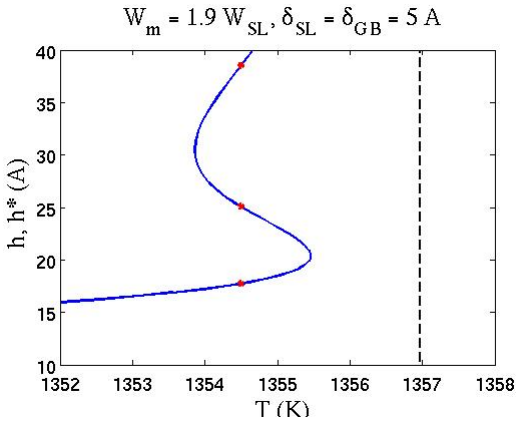

Figure 24. *Computed width of the grain boundary as a function of the system temperature. The melting point ($T_M = 1357$ K) is indicated by the vertical dashed line. The width exhibits hysteresis, suggesting a possible jump in the interface thickness as the temperature varies.*

results in a system of Euler-Lagrange differential equations whose solution provides the desired description of grain boundaries in the material. An example of the resulting phase-field profiles is shown in Figure 23. Here the grain boundary is located near $x = 0$, and the regions where $\varphi_1 = 1$ on the right and $\varphi_2 = 1$ on the left represent the solid grains with different orientations. The liquid order parameter, φ_3, is non-zero in the interfacial region which indicates wetting of the grain boundary.

The desired surface properties can all be determined from the computed phase-field profiles. In these types of models the interface is modeled as a diffuse region, with a characteristic width h that is defined by a convention such as $\varphi_1(h/2) = ½$. In Figure 24 we show the computed grain boundary width h as a function of the system temperature. For these conditions there are multiple solutions that can be interpreted as a first-order phase transition between thin and thick interfaces. For large widths the grain boundary tends to resemble two solid-liquid interfaces with an interior region of liquid near $x = 0$, all coexisting at the melting point.

This study [1], and the complementary study [2], suggests the validity of grain boundary premelting in the copper-silver system. The prewetting occurs over a fairly narrow range of temperatures below the melting point, which may account for the difficulties in the interpretation of experimental studies of premelting.

References

[1] Y. Mishin, W. J. Boettinger, J. A. Warren, and G. B. McFadden, Thermodynamics of Grain Boundary Premelting in Alloys. I. Phase-field Modeling, submitted.

[2] P. L. Williams and Y. Mishin, Thermodynamics of Grain Boundary Premelting in Alloys. II. Atomistic Simulation, submitted.

[3] G. B. McFadden, Phase-field Models of Solidification, in *Contemporary Mathematics*, Vol. 306, *Recent Advances in Numerical Methods for Partial Differential Equations and Applications,* (X. Feng and T.P. Schulze, eds.), American Mathematical Society, Providence, RI, 2002, pp. 107-145.

[4] A. A. Wheeler, W. J. Boettinger, and G. B. McFadden, A Phase-field Model for Isothermal Phase Transitions in Binary Alloys, *Physical Review A* **45** (1992), pp. 7424-7439.

Participants

Geoffrey B. McFadden (MCSD); William J. Boettinger and James A. Warren (MSEL); Yuri Mishin (George Mason University)

Program area: Enabling Scientific Discovery (ITL)

Computational Modeling and Visualization of Cement Paste Hydration and Microstructure Development

The properties, durability, and service life of concrete are directly linked to the structure of the cement paste binder at length scales between 100 nanometers and 100 micrometers. We have created and parallelized a numerical model, HydratiCA, which computes this structure. We have demonstrated that the parallel algorithm scales linearly up to 1,000 processors (the largest number to which we had access). We have run a detailed set of 18 validation runs with excellent results. We have created a set of real-time visualization tools that facilitate the direct analysis of the predicted microstructure. Combined, we have generated new knowledge of the cement hydration curing process.

William George and Judith Terrill

Background. Concrete is the most widely used man-made product in the world. Great improvements in the use of concrete can be attained through the use of sophisticated computer models to enable experimentation with new constituent materials, chemical additives, and manufacturing processes. Models of this type which have been developed at NIST are the result of decades of research in both theory and experiment. These models now employ parallel algorithms which we have developed. Model outputs can now be studied with real-time visualization and analyses that run on the desktop as well as in our immersive visualization environment.

The goal of this research is the moving of concrete testing from the laboratory to the computer, enabling efficient design of cement and concrete with specific properties in mind. The characteristics of cement and concrete that must be predicted include: setting time and flow properties ("workability"), the volume and connectivity of porosity, and the solid phase assemblage as a function of time; shrinkage at early ages, including chemical, autogenous, and drying components; heat evolution; permeability and transport factor of the material with respect to aggressive ions (e.g., chloride); compressive strength development; and susceptibility to common degradation mechanisms, e.g., freeze/thaw. The expected benefits of this research include: concrete with improved (optimized) performance, prediction of in-service performance, performance-based standards, reduced life-cycle cost, increased sustainability, better use of waste stream materials, stimulation of innovation, and reduced cost of testing—all aids to business survival in an era of constrained resources.

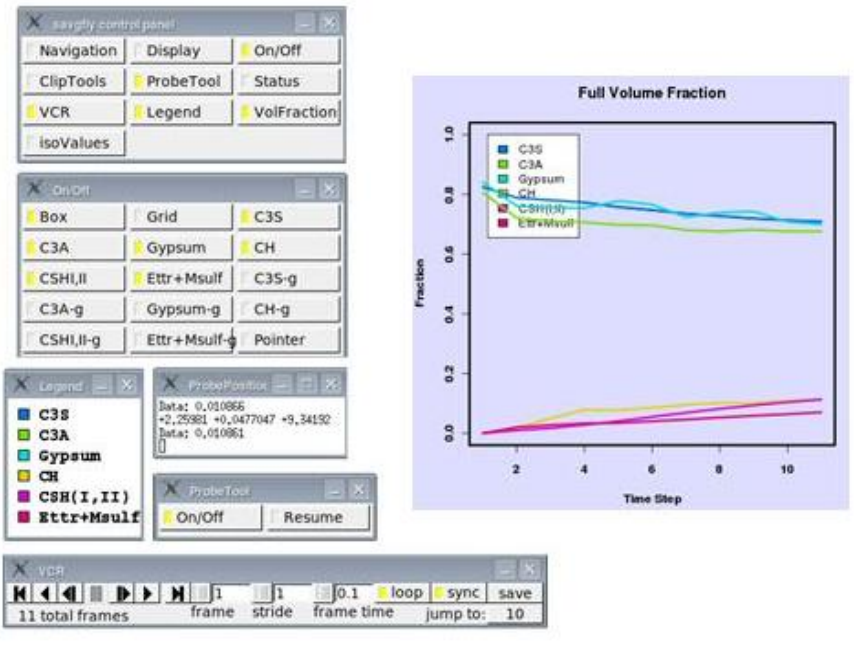

Figure 25. *Visualization of cement paste hydration with interactive controls.*

HydratiCA. The computational modeling and visualization of cement paste hydration and microstructure development is a crucial tool in the move to computational testing and prediction. Understanding the complex chemical changes that occur when cement powder is mixed with water is a longstanding but extremely challenging technological goal. Fundamental computational modeling of the hydration of cement is difficult because it involves a large number of coupled nonlinear rate equations that must be solved on a complex three-dimensional spatial domain.

To address these challenges, we are applying a new computational model called HydratiCA that has several advantages over previous attempts to model cement paste hydration. HydratiCA is based on concepts of transition state theory and uses stochastic cellular automata algorithms to simultaneously model 3-D reaction and transport phenomena. This allows us to track the detailed kinetics and equilibria that occur in a diverse range of cementitious systems. At the length scales required to finely resolve the reaction mechanisms and microstructure changes in cement paste, HydratiCA must take small time steps (approximately 10^{-5} seconds) to remain numerically stable. In other words, tens of millions of time steps are required to simulate the behavior of cement paste for just one hour. Therefore, parallelization of HydratiCA is important to model systems that are large enough to be realistic, to avoid finite-size effects, and to still be able to complete the simulations in a reasonable amount of time.

Through a grant from NASA in 2008, our parallelized version of HydratiCA was run on the Columbia supercomputer at the NASA Ames Research Center with simple single-phase minerals in water, to test how the code speed scaled with number of processors. This preliminary work has verified that HydratiCA scales almost linearly when using between 100 and 1000 processors. This year we also completed a series of 18 validation runs with HydratiCA.

- Portlandite in pure water (examine temperature effect at 293 K, 298 K, 303 K)
- Ettringite in pure water (examine temperature effect at 293, 298 K, 303 K)
- Portlandite in water at 298 K (examine effect of NaOH concentration - 0.01 M, 0.05 M, 0.1 M)
- Water supersaturated with respect to Portlandite at 298 K (testing rate of homogeneous and heterogeneous nucleation as a function of saturation index and nucleation energy barrier)
- CaO in water at 298 K (examine effect of particle size using monodisperse spheres with diameters of 5, 10, or 20 micrometers)

Visualization. Visualization of the output of this model is important both for validation and to enable understanding of the spatial distribution of phases in 3D. The output is a 3D volume of data with percentage values for each of multiple material phases at each lattice site. Over the course of the simulation time, a series of data volumes is produced at the time intervals of interest. During 2008, we built a series of software tools (see Figure 25) that include a Virtual Cement Analysis Probe, represented as a cross-shaped 3D glyph, which allows interactive visualization and interactive measurement of phases in the hydration process. As the user sweeps through the virtual data, a continuous digital read out of the data being measured is displayed. At the press of a button, the current value is saved and displayed at its location within the virtual environment. The researcher is free to fully explore the data and measure concentration at areas of scientific interest.

The phases are represented as isosurfaces. Any combination of phases can be interactively turned on or off in the visualization. A cube representation of the data voxels provides a visual aid to identify the solid portions of the volume enclosed by the isosurface. There are also "VCR" controls that allow the simulation time steps to be played forward, reverse, fast, slow, single step, stop, etc. Additionally, two individual time steps may specified and with a single click the scene jumps from one to the other for an "A vs. B" comparison view.

This year, our simulations clarified the influence of the mineral portlandite and other calcium salts on the kinetics of early-age hydration; revealed the rapidity with which solutions in cement paste have their compositions homogenized; revealed a scaling problem with the implemented nucleation mechanism; visualized chemical "subsystems" within cement paste, i.e. correlation of ettringite crystals with aluminum sources, as opposed to sulfate sources; and confirmed the extent in which temperature and common ions influence the kinetics and equilibrium of mineral salts.

In the coming year, we will be studying the influence of particle size distribution on the setting of tricalcium silicate, and the influence of calcium sulfate on the hydration kinetics and microstructure evolution of tricalcium aluminate in water.

Participants

William George, Edith Enjolras, Steve Satterfield, John Kelso, Judith Terrill (MCSD); Jeffrey Bullard (BFRL)

http://math.nist.gov/mcsd/savg/parallel/hydration/
http://math.nist.gov/mcsd/savg/vis/hydration/

Program area: Enabling Scientific Discovery (ITL)

Quantifying and Characterizing Errors in 3D Rendering

The visual display of three-dimensional (3D) scenes has become commonplace in modern computing. We have developed methods for measuring errors in the rendering of such scenes on pixel-based computer graphics systems. While many rendering applications are intended to provide qualitative experiences to the viewer, others are intended to convey accurate spatial relationship and have important quantitative components. As we use such systems for quantitative tasks it is important to understand how the computer graphics environment contributes to uncertainty in these tasks. Our methods for assessing error in computer graphics rendering is a step toward that understanding.

John Hagedorn

In modern computer graphics systems, a virtual 3D scene is usually described as a set of geometric objects expressed in a 3D coordinate system. These geometric descriptions are then transformed into a set of pixels that are displayed to the user on a monitor. The process of transforming the 3D descriptions to a set of pixels is referred to as *rendering*.

The rendering of virtual 3D scenes has become ubiquitous. There is a vast number of applications such as special effects in movies and TV, realistic computer games, architectural rendering, and scientific visualization. While many of these rendering applications are purely qualitative in nature, others are part of quantitative tasks in which the accuracy of the rendered images is of vital importance; for example, at NIST we have implemented interactive measurement tools within such virtual scenes [1]. Measurement and analysis tools that rely on rendered images are likely to be affect by rendering accuracy. We have recently developed methods for assessing geometric errors in the rendering process. While this is hardly the only possible source of error that may be introduced by computer graphics technology, it is the focus of the work described here [2]

Our approach is to examine geometric errors in the rendering of three 3D geometric primitive forms: points, line segments, and triangles. These represent the simplest and by far the most commonly used geometric forms rendered in current computer graphics systems. Our work starts from the assumption that for each of these 3D geometric forms, there is a correct projection of the form to a two-dimensional (2D) rectangle that represents the viewing area. The viewing area can be thought of as an idealization of (a portion of) the monitor screen, though it is not discretized into a finite set of elements (pixels) as a real screen is. We use this projection from 3D coordinates to the 2D coordinates of the viewing rectangle as an idealized form of the rendering process.

For each primitive type, we create several error metrics that quantify the amount by which the actual rendering of a primitive deviates from its ideal projection. For point rendering, we measure the distance from the center of the rendered pixel to the projection of the point onto the plane of the screen; Figure 26 illustrates this. For line segment rendering, we aggregate the distances from the rendered pixels to the projection of the line segment. In addition, for some of the line segment metrics, we measure the extent to which the set of rendered pixels correctly covers the length of the projected line segment. For triangle rendering, we calculate how the area covered by the rendered pixels relates to the area covered by the projected triangle. This involves measuring the area of rendered pixels that fall outside the projected triangle and the area of pixels inside the triangle that were not rendered. Detailed descriptions of these metrics can be found elsewhere [2].

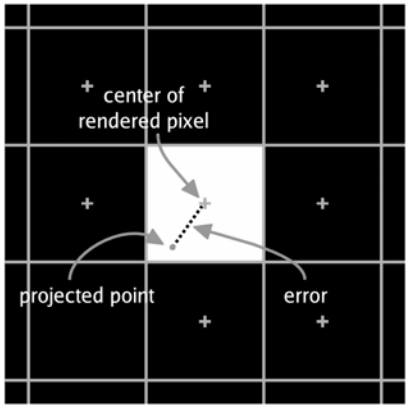

Figure 26. *The point rendering error is the distance from the projection of the 3D point to the center of the rendered pixel. This is the simplest of the rendering error metrics.*

For each of these metrics, we calculate errors in pixel-based coordinates which is a natural coordinate system for the viewing area or screen. In addition, we also define a form for each metric that normalizes the metric based on the pixel dimensions of the viewing rectangle (window). For the line segment and triangle error metrics, we define signed and unsigned versions of the metrics. The unsigned metrics assess the overall deviation of the rendered segment or triangle from the projection of that object, while the signed metric indicates whether there is any bias in the errors.

To show the utility of our methods, we performed a series of rendering experiments on several hardware platforms under a variety of rendering conditions. Each experiment is a run of a program that renders a

set of randomly generated geometric primitives. For each primitive (point, segment, or triangle), the program calculates the projection of the geometry to the idealized viewing plane, looks at the graphics frame buffer to find which pixels were actually rendered, and then uses these to calculate the error metrics.

In our experiments, the programs render the primitives using OpenGL; however our methods are not dependent on any particular rendering software or rendering technique. Indeed, these methods could just as easily be used with other 3D rendering libraries such as DirectX. We chose OpenGL because it is readily available on many hardware platforms and is used in many consumer, industrial, and research applications. Our test software was designed so that it could easily be ported to a variety of hardware platforms and operating systems. We have placed our software on the web at the URL given at the end of this article.

A complete description of the results of these experiments can be found in [2]. As an example of the type of results obtained, we look at the normalized signed line segment errors for one of the hardware platforms that we tested. Table 3 shows the median and inter-quartile range (IQR) for four combinations of viewing conditions. These statistics were obtained by rendering 200,000 randomly generated line segments for each of the four viewing conditions.

Table 3. *Normalized line segment rendering errors ($\times 10^6$).*

Test Case		Normalized Signed Line Segment Errors	
View Angle	Window Size	Median	IQR
30	128	62.71	123.63
150	128	84.55	146.92
30	907	1.67	3.47
150	907	2.24	4.0 0

The most obvious thing that we see is that the smaller windows have much greater errors than larger ones. (This corresponds to our naive understanding that a rendering at a low resolution is inferior to a rendering at a high resolution.) We also see that the 30 degree views have smaller medians and spreads than the corresponding 150 degree views. Note that our methods do not try to explain *why* such differences exist; they are intended only as tools for characterizing and quantifying differences in rendering accuracy.

These error metrics enable us to make quantitative estimates of errors let us compare errors for different rendering cases, including across platforms. For example, we found that point rendering has a median error of about 0.39 in pixel units, that one platform has much less bias error in rendering segments than another platform, and so on. Our normalizations provide a way of quantifying our intuitive understanding that higher resolution renderings produce smaller errors.

These results can provide useful information in real-world viewing situations. These measurements are easily converted to physical dimensions. For example, Table 4 shows how the point rendering error of 0.396 pixel units translates to physical dimensions on three representative screens that we use at NIST.

Table 4. *Point rendering error converted to physical units for three representative displays.*

Display	Error in pixel units	Error in physical units
LCD	0.396	0.101mm
CRT	0.396	0.118 mm
Projection	0.396	0.915 mm

We plan to quantify errors in other aspects of the visual representation of information. For example, we will analyze how these rendering errors result in quantifiable errors in depth representation in stereo displays. We will also be investigating errors that result from the the process of generating renderable 3D geometric primitives from underlying data representations (for example, CT scans). For our work in immersive visualization environments, we are looking at the viewing errors caused by the position tracking systems used in virtual reality systems. Finally, it is our goal to develop methods for aggregating all of these types of errors with uncertainties from other source to form an overall assessment of uncertainty. We believe that these tools enable us to understand better the possible errors in our displays of quantitative information. This understanding is critical for making informed analyses based on these visual representations.

References

[1] J. G. Hagedorn, J. P. Dunkers, S. G. Satterfield, A. P. Peskin, J. T. Kelso, and J. E. Terrill, Measurement Tools for the Immersive Visualization Environment. *Journal of Research of NIST* **112** (5) (2007), pp. 257-270.

[2] J. G. Hagedorn, J. E. Terrill, A. P. Peskin, and J. J. Filliben, Methods for Quantifying and Characterizing Errors in Pixel-based 3D Rendering. *Journal of Research of NIST* **113** (4) (2008), pp. 221-238.

Participants

John Hagedorn, Judith Terrill, Adele Peskin (MCSD); James Filliben (ITL)

http://math.nist.gov/mcsd/savg/vis/metrology

Program area: Virtual Measurement Systems (ITL)

Standardization of Lung Tumor Growth Measurements

Lung tumor malignancy is assessed by measuring the increase in tumor size over time in computed tomography (CT) scans. Current methods for estimating tumor size from CT scans are inaccurate, with results varying greatly from lab to lab. Standardized and accurate methods for estimating tumor volumes are needed. We are developing and assessing methods for such volume calculations. To enable such assessments we are developing realistic, synthetic lung tumor data sets to be used as standards with which to compare volumetric methods.

Adele Peskin, David Gilsinn, and Javier Bernal

Many advances in medicine today depend upon the accurate reading of body images obtained from computed tomography (CT). 3-D CT data is collected as a set of 2D slices approximately 5 mm apart. Within a slice, data is obtained on a set of grid points positioned at distances of about 0.5 mm apart. Radiologists assess growth in CT scans of tumors taken over a period of time by examining the scans and making estimates of volume growth. There are a variety of alternate techniques available for calculating tumor volumes and/or the change in tumor size between two images of the same tumor. We are working to assess the quality of such methods. A critical component of most such methods is the determination of tumor boundaries. This is quite challenging to do in practice. These boundaries are blurred pixel-width regions that have a somewhat wide distribution of pixel intensities. Variations in boundary estimates can have a large effect on computed volumes, making such computed volumes highly unreliable.

Generating Artificial Tumors. The first step was to create sets of artificial tumors of known volumes embedded in real lung data. We create reference spheres as pixel data to insert into an actual set of CT scans. In order to create realistic tumor data it is necessary to vary the signal intensity near the edges of the tumor. A critical component of such a method is a standardized method of calculating "partial volumes," or the volumes of 3-D voxels in the CT data grid that tumor edges only partially fill. To do this, each pixel location in the region of the sphere is defined as the center point of a unit sized voxel. We compute the fraction of the voxel contained inside the sphere by a numerical integration. The voxel is divided into a large number of cubes. If a cube is outside of the sphere, it contributes nothing to the intensity value. In the case of a voxel entirely within the sphere, this sum is equal to 1.0. For a voxel at the edge, this fraction is between 0.0 and 1.0. Once inserted into the lung data, the resulting pixel intensity distributions of lung regions containing the synthetic spheres resemble those from regions containing real tumors. That is, we found that real lung tumor pixel intensity data has certain consistent boundary metrics across data sets that come from the gradient of the pixel intensity fields, and the synthetic tumors, cre-

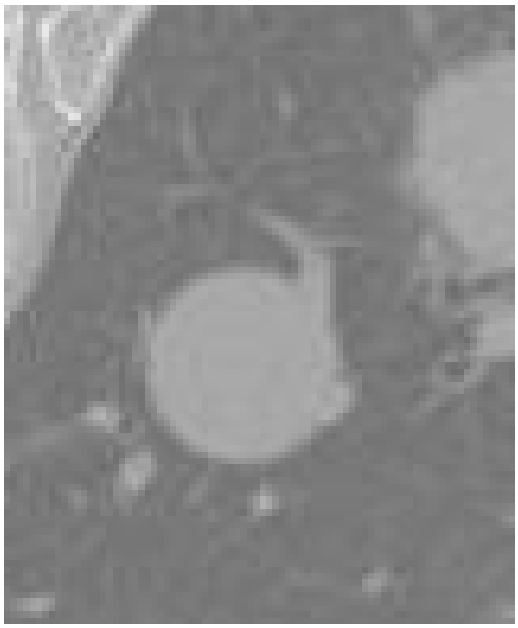

Figure 27. *Sphere embedded in lung data near blood vessel.*

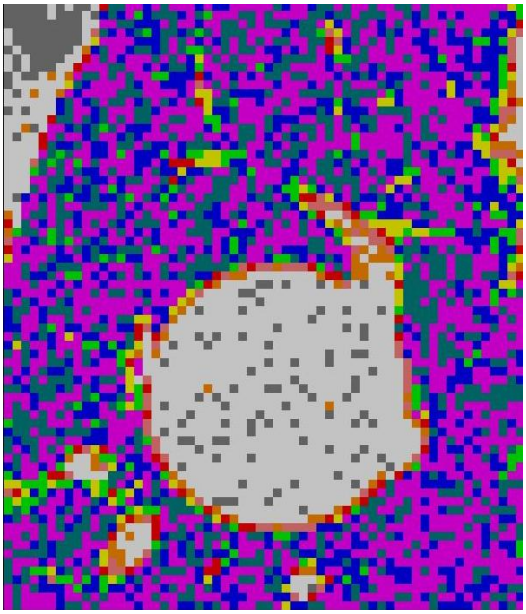

Figure 28. *Same slice of embedded sphere, color coded according to pixel intensity.*

ated by the method above, also share these boundary metrics.

The synthetic data generated in this way can be adjusted to represent different challenges in tumor volume estimation. For example, blood vessels attached to tumors typically result in ambiguous tumor boundaries. Creating synthetic tumors with blood vessels attached provides realistic and challenging standard data to test separation methods. Figure 27 shows a spherical synthetic tumor embedded in a section of the lung surrounded by blood vessels. Figure 28 shows a color-coded pixel intensity distribution of this tumor and the blending of the tumor and blood vessel pixels.

Volume Estimation Methods. Three approaches to estimate volumes of tumors from CT scan data have been investigated. One method is based on determining the probability of each pixel being inside or outside of the tumor, and then adding together contributions from each corresponding voxel based on these probabilities. A second approach involves converting the pixel based data to points in 3-space, called a point cloud, and then generating a tetrahedralization to fill the interior of the point cloud. A third method that accounts for the fact that the scale of the distance between scans is significant compared to the pixel distances within scans was also considered. Here a least squares spline surface model of the point cloud data is constructed and used to estimate the volumes by means of the classical Divergence Theorem. We briefly describe these methods here.

In the first method, we use a statistical approach to weight the intensity of each pixel according to its "likelihood" of being inside or outside of the tumor, and then and sum the contributions of the weighted partial volumes. Each pixel lies in one of three different pixel intensity distributions for the lung pixels: the tumor pixels, the edge pixels (which vary over the entire range), and the background pixels. We fit the data in the region of a tumor to a three-component Gaussian mixture via the Expectation Maximization (EM) algorithm [2]. The challenge is to determine which pixels from the edge pixel distribution are part of the tumor, and which are part of the background; the former receive full weight (1.0) while the latter receive no weight (0.0). Based on the assumption that each pixel has some probability of being tumor or background, we calculate a weight that each individual intensity comes from one or the other of these distributions. Edge pixels receive partial weights in a smooth way, depending upon the distance of their intensities from the tumor pixel distribution. A smooth weight function that satisfied this property is the biweight function:

$$w(u) = (1-u^2)^2 I_{[-1,1]}(u), \quad u = (x-T)/(kS)$$

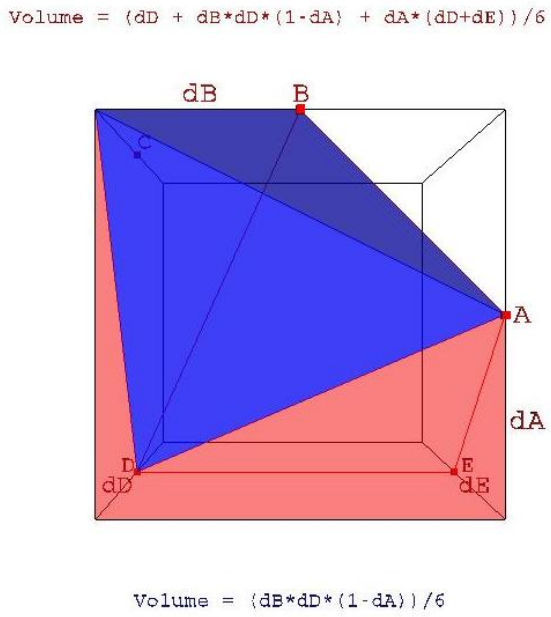

Figure 29. *Marching cube example volume.*

where x is a value for a pixel intensity, T is an estimate of the average of the distribution of tumor pixel intensities, S is an estimate of the scale (e.g., standard deviation) of this distribution, and k is a "tuning constant" [1].

We found that the required constants can be calculated directly from the CT data, with no user input needed. They are derived from the EM mixtures results and the distribution of the gradients of the pixel intensity field. Since part of our goal is to develop standard methods for volume calculations, repeatable from laboratory to laboratory, it is important to use parameters in these volume calculations that are computed internally and not input by the user.

The tumor volume itself is computed by the following method. For each pixel, there is an isosurface that defines a volume. We can determine that volume using a type of marching cubes algorithm in which each partially filled voxel region is computed based on a linear interpolation within that voxel (see Figure 29 for an example). We compute the volume for the lowest pixel intensity for which the computed weight above is equal to 1.0. We also calculate a volume for each pixel intensity in which the weight is between 0.0 and 1.0. Each volume becomes slightly larger as the pixel intensity decreases, forming shells around the original volume. Each shell makes a contribution to the total volume, based on its volume and weight.

The total volume of tumor is then calculated as a function of the individual volume calculations at each of the isovalues (V_i = volume at isovalue i) and their corresponding weights (w_i = weight at isovalue i) as

$$V^* = V_0 + \sum_{i=1}^{n} w_i (V_i - V_{i-1})$$

A second approach to estimating a tumor volume has been implemented based on Canny's algorithm [3] for detecting edges in a 2-D image. This method results in a set of points on the surface of the tumor that can be used for estimating the volume. The program identifies pixels on the edges of the images in the scans using a gradient based algorithm applied to the pixel intensities. The distances between pixels and scans are known, so that the Cartesian coordinates of each identified pixel can be computed. This results in the desired set of points on the surface of the tumor. These points form what is known as a point cloud. Using a method based on Voronoi diagrams, the interior of the point cloud is filled with tetrahedrons whose volumes can be summed. The boundary surface of the sphere formed by the tetrahedrons is a triangulated surface. The method is known as a power crust method.

The third approach chosen for this study is based on a fundamental theorem in calculus, the Divergence Theorem, a generalization of Green's Theorem for two dimensional space. In the Divergence Theorem a particular volume integral is shown to be equal to a surface integral. The surface points from the point cloud data used for the second method above is also used here. The data is modeled using a set of basis functions called B-splines. These are functions formed by joining low order polynomials to create a smooth function. The piecewise polynomials are joined at points called knots. B-splines of order four are cubic splines that are non-zero only over four adjacent knot intervals. The notation for a B-spline basis function is $N_{i,k}(x)$, where $N_{i,k}$ is zero everywhere except in the range $x_{i-k} < x < x_i$, where $k = 4$ for cubic splines. The surface spline is given by the tensor product [4, 5]

$$s(x,y) = \sum_{i=1}^{m+4} \sum_{j=1}^{n+4} c_{ij} N_i(x) N_j(y) .$$

where here x and y are angular coordinates.

The unknown coefficients c_{ij} are determined by a least squares fit to the point cloud data. Here the standard normal equation approach cannot be used due to the fact that the resulting system matrix is often rank deficient. (There are often zero rows or columns due to the sparse distribution of data.) Thus, in order to compute the least squares solution in this case we rely on the singular value decomposition (SVD).

Once the surface is fit by a tensor product of B-splines, the surface integral portion of the Divergence Theorem is computed by dividing the tumor surface into triangular surface patches and computing the surface integral for each patch and summing these to give the total surface integral. The Divergence Theorem then gives the volume integral directly [6].

Estimates of the embedded sphere volumes for the case of a radius of 20 pixels show clearly that the volume is highly dependent of the method of determining the sphere boundary. In our test, the side of each voxel was 0.57 mm, with 1.25 mm between slices. The simulated spherical tumor had a radius of 11.4 mm and a volume of 6205.87 cubic mm. The volume estimate by the first method described was 6196.74 cubic mm, approximately a 0.15% error. The radius estimate in this case would be 11.39 mm. Thus a 0.09% radius error produced a 0.15% volume error. When the Canny algorithm, based on pixel gradient fields, was applied to the pixel data to generate a boundary point cloud, the volume estimate by the power crust method was 6596.88 cubic mm with a radius of 11.63 mm. In this case a 2% error in radius estimate produced a 6.3% volume error. Finally, using a grid of 30 latitudes and 60 longitudes a volume estimate by the B-spline approach of 6626.57 cubic mm was based on a radius estimate of 11.65 mm. In this case a 2.2% error estimate for the radius produced a 6.8% volumetric error.

Further work is needed in developing a more robust boundary algorithm for producing boundary point cloud values in Euclidean coordinates. Future work will include greatly expanding our synthetic lung tumor standard data set, testing a wide variety of volumetric algorithms, and increasing the accuracy of the new volumetric methods we are developing.

References

[1] K. Kafadar, The Efficiency of the Biweight as a Robust Estimator of Location, *Journal of Research of the NBS* **88** (2) (1983), pp. 105-116.

[2] T. Hastie, R. Tibshivari, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, And Prediction*, Springer, New York, 2001.

[3] R.C. Gonzales and R.E. Woods, *Digital Image Processing*, Pearson, Upper Saddle River, NJ, 2008.

[4] C. de Boor, *A Practical Guide to Splines*, Springer-Verlag, New York, 1978.

[5] D.F. Rogers and J.A. Adams, *Mathematical Elements for Computer Graphics*, McGraw-Hill, Boston, 1990

[6] P.J. Schneider and D.H. Eberly, *Geormetric Tools for Computer Graphics*, Morgan Kaufmann Publishers, Amsterdam, 2003.

Participants

Adele P. Peskin, David Gilsinn, Javier Bernal (MCSD); Alden Dima, Walter Liggett (ITL); Karen Kafadar (Indiana University)

Program area: Information Discovery, Use and Sharing (ITL)

Part III

Project Summaries

Information Technology Laboratory Programs

Enabling Scientific Discovery

Modern scientific research has become more and more dependent on mathematical, statistical, and computational tools for enabling discovery. The Enabling Scientific Discovery Program promotes the use of these tools to dramatically advance our ability to predict the behavior of a broad range of complex scientific and engineering systems and enhance our ability to explore fundamental scientific processes. This Program focuses on interdisciplinary scientific projects that involve novel computational statistics and the development of simulation methods and software. These efforts will have a foundational impact on scientific discovery throughout U.S. industry, government, and academia. – Anthony Kearsley, Program Manager

Calibrating Image Roughness Using Lipschitz Exponents

Alfred S. Carasso
András E. Vladár (NIST MEL)

See feature article, page 34.

Modeling Grain Boundary Premelting in Binary Alloys

Geoffrey B. McFadden
William J. Boettinger (NIST MSEL)
James A. Warren (NIST MSEL)
Yuri Mishin (George Mason University)

See feature article, page 37.

Computational Modeling and Visualization of Cement Paste Hydration and Microstructure Development

William George
Edith Enjolras
Steve Satterfield
John Kelso
Judith Terrill
Jeffrey Bullard (NIST BFRL)

See feature article, page 39.

Funding: *Virtual Cement and Concrete Testing Laboratory Consortium*

Adaptive Finite Element Modeling of Two Confined Interacting Atoms

William F. Mitchell
Marjorie A. McClain
Eite Tiesinga (NIST PL)

We have implemented a parallel adaptive refinement and multigrid finite element code for the solution of partial differential equations, PHAML, and applied it to a two-dimensional Schrödinger equation in order to study the feasibility of a quantum computer based on extremely cold neutral alkali-metal atoms. In this physical system, quantum bits (qubits) are implemented as motional states of an atom trapped in a single well of an optical lattice. Quantum gates are constructed by bringing two atoms together in a single well, with the interaction between them causing entanglement. Quantifying the entanglement reduces to solving for selected eigenfunctions of a Schrödinger's equation that contains a Laplacian, a trapping potential and a short-ranged interaction potential. The behaviors of the eigenfunction due to the trapping and interaction potentials are on very different scales, which require adaptive refinement and high order finite elements.

Use of high degree piecewise polynomials has proved to be crucial to get the high accuracy solutions required for correctly modeling the atom interactions. We are studying the use of a new technique, *hp*-adaptive refinement, to improve on this even further. With *hp*-adaptive refinement one adapts both the element size h and polynomial degree p locally. Properly done, this technique has a faster convergence rate than fixed-degree adaptive methods. A number of *hp*-adaptive strategies, i.e., methods for determining whether a particular element should be refined by h or by p, have been proposed in the literature. We have designed and performed a benchmark experiment to compare the convergence rates of 10 *hp*-adaptive strategies with very high order, fixed degree, *h*-

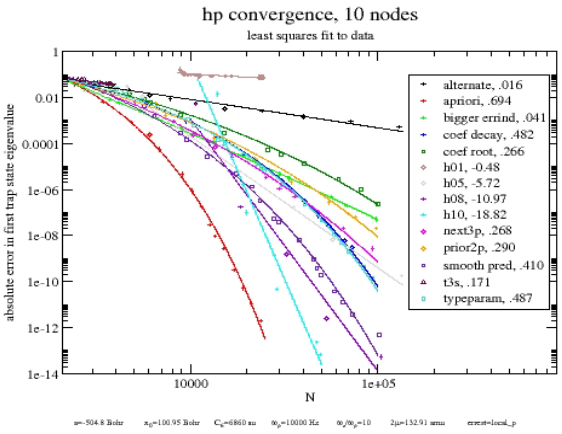

Figure 30. *Computational experiment comparing the convergence rates of 10 hp-adaptive strategies with very high order, fixed degree, h-adaptive refinement, using a 10-node model problem.*

adaptive refinement, using a 10-node model problem (see Figure 30). The results show the *a priori* method to superior, and for the accuracy we are interested in (5 or 6 digits); the smooth prediction strategy and *h*-adaptive refinement with elements of degree 8 and 10 are virtually tied for second.

We have begun studying methods for solving the time-dependent Schroedinger equation, which will be necessary for understanding the time-dependent behavior of the interaction between atoms. We found that standard methods for time-dependent parabolic partial differential equations, such as the backward Euler and Crank-Nicholson methods, are not stable enough for computing interior eigenvalues. With these methods, the slightest round-off error destroys the unitary propery of the operator, and all modes except the ground state (smallest eigenvalue) are quickly damped away. As a different approach, we have implemented a method that uses eigenvectors of the time-independent problem as a basis set for the time-dependent problem. In our initial tests, this approach appears to be giving solutions with the expected behavior, but further tests are needed to verify this. We will also consider other possibly more efficient approaches. Future work will include the solution of multi-channel problems.

[1] W. F. Mitchell, "A Parallel Multigrid Preconditioner for High-order and hp-Adaptive Finite Elements," 13th SIAM Conference on Parallel Processing for Scientific Computing, Atlanta, GA, March 12-14, 2008.

[2] W. F. Mitchell, "Application of a Parallel Adaptive Finite Element Code to Confined Interacting Atoms," SIAM Annual Meeting, San Diego, CA, July 7, 2008.

[3] W. F. Mitchell, "Strategies for *hp*-Adaptive Refinement," 6th Int. Conference of Numerical Analysis and Applied Mathematics, Kos, Greece, Sept. 16-20, 2008.

Nano-structures, Nano-optics, and How to Squeeze the Light out of Quantum Dots

James S. Sims
Kevin Rawlings
John Kelso
Garnett W. Bryant (NIST PL)

Research and development of nanotechnology, with applications ranging from smart materials to quantum computation to biolabs on a chip, has the highest national priority. Semiconductor nanoparticles, also known as nanocrystals and quantum dots, are one of the most intensely studied nanotechnology paradigms. Nanoparticles are typically 1 nm to 10 nm in size, with a thousand to a million atoms. Precise control of particle size, shape and composition allows one to tailor charge distributions and control quantum effects to tailor properties completely different from the bulk and from small clusters. As a result of enhanced quantum confinement effects, nanoparticles act as artificial, man-made atoms with discrete electronic spectra that can be exploited as light sources for novel enhanced lasers, discrete components in nanoelectronics, qubits for quantum information processing, and enhanced ultrastable fluorescent labels for biosensors to detect, for example, cancers, malaria or other pathogens, and to do cell biology. We are working with the NIST Physics Lab to develop computationally efficient large scale simulations of such nanostructures. We are also working to develop immersive visualization techniques and tools to enable analysis of highly complex computational results of this type.

We study the electrical and optical properties of semiconductor nanocrystals and quantum dots. In the most complex structures this entails modeling structures with on the order of a million atoms. Highly parallel computational and visualization platforms are critical for obtaining the computational speeds necessary for a systematic, comprehensive study.

Often it is easy to define the simple subsystems that make up a complex, heterogeneous nanosystem [1]. However, it may be difficult to explicitly define the entire structure. A novel feature of our code is the ability to link together heterogeneous nanostructures (nanosubsystems). Since we can do multiprocessor runs routinely, we have the basic building blocks for making larger structures by "stitching" together disparate subsystems into composite structures, each separate component subsystem being a smaller multiprocessor run. Now that we can routinely model nanosystems [2], our focus has shifted to modulating and controlling the optical properties of these self-assembled quantum dots using external strain [3].

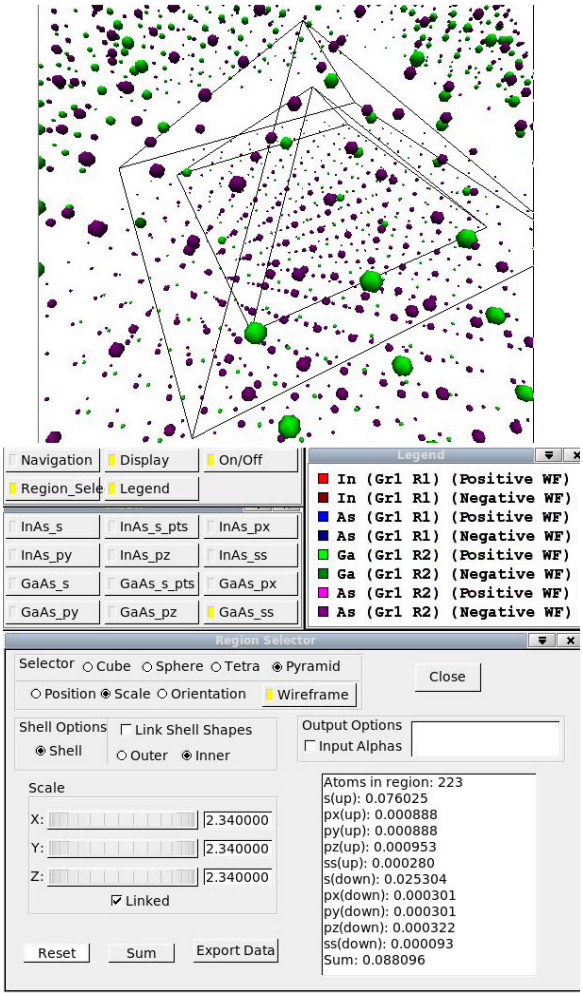

Figure 31. *Multiple geometric structures can be selected and the charge density summed up within (or between, as appropriate) them. In this case, the charge is being summed up between two pyramids.*

To understand the coupling between externally imposed strains and the electronic states of the self-assembled quantum dots, we have developed an atomistic tight-binding theory of the confined states in quantum dots that incorporates the local strain due to atomistic lattice mismatch and the externally imposed strain due to applied stressors or to the bend in a vibrating nanomechanical oscillator. Both local strain and the externally imposed strain are important, so we include them on an equal footing via an atomistic valence force field approach, with the externally imposed strain modeled as a distortion of the lattice on the boundary of structure. A full tight-binding model including an sp3s*d5 orbital model and spin-orbit effects is used [2, 3, 4].

To understand how applied stress can be used to actively control dot optical properties, dots buried at different points in a nanobridge oscillator or cantilever are considered. The dependence of the quantum dot electronic states and optical transitions on the coupling to bending modes of the oscillator is studied. Both isolated and coupled dots are considered. Level crossings for whole states are found. State symmetries can be significantly distorted by the applied stress. Transitions can be strongly polarized parallel or perpendicular to the external strain or suppressed, depending on how the dots are stressed. The dependence on applied stress can be understood as a competition between the internal and external stress that can either enhance or suppress the local strain at the dot, depending on how the external stress is applied.

As the computational model has been extended to handle more complex and larger systems by including not only the nanostructures but also the substrate and environment around them, parallel processing has become even more of a necessity. In the most recent parallelizations both shared memory OpenMP parallelization and distributed memory MPI are employed. For a typical calculation with 10-11 nodes, each using 4 CPUs (for a total of 40 - 44 processors), a structure of 10,000 atoms can be solved in 5.3 days of wall clock time. A 100,000 atom structure of the same calculation would take 530 hours (about 22 days), compared to six months without the parallelization. A paper on the latest developments has been recently delivered at the 13th Advanced Heterostructures and Nanostructures Workshop [5].

To study the increasingly larger systems we are working with, we have expanded the visualization capability with multiple selection tools that can sum up charge density in different types of regions. This is illustrated in Figure 31.

[1] G. W. Bryant and W. Jaskolski, "Electronic Properties of Quantum-dot Molecules," *Physica E* **17** (2003), pp. 40-41.

[2] J .S. Sims, W. L. George, T. J. Griffin, J. G. Hagedorn, H. K. Hung, J. T. Kelso, M. Olano, A. P. Peskin, S. G. Satterfield, J. Devaney Terrill, G. W. Bryant and J. G. Diaza, "Accelerating Scientific Discovery through Computation and Visualization III. Tight-Binding Wave Functions for Quantum Dots," *Journal of Research of NIST* **113** (3) (May-June 2008), pp. 131-142.

[3] G. W. Bryant, W. Zielinski, W. Jaskolski, J. Aizpurua, and J. S. Sims, "Controlling the Optical Properties of Self-assembled Quantum Dots Using External Strain," Fall 2007 Meeting of the Materials Research Society, Boston.

[4] G. W. Bryant, M. Zielinski, W. Jaskólski, J. Sims, and J. Aizpuru, "Squeezing the Light Out of Quantum Dots Using External Strain," 5th International Conference on Semiconductor Quantum Dots (QD2008), Gyeongju, Korea May 11-16, 2008.

[5] G. W. Bryant, M. Zielinski, J. Sims, W. Jaskolski, and J. Aizpurua, "Squeezing Light Out of Quantum Dots with Externally Imposed Strain," 13th Advanced Heterostructures and Nanostructures Workshop (AHNW2008), Kohala Coast, Hawaii, December 2008.

Modeling Fluid Flow in Materials Processing

Geoffrey B. McFadden
P. Aaron Lott
Sam R. Coriell (NIST MSEL)
Daniel Anderson (George Mason University)
Katherine Gurski (Howard University)
Bruce Murray (SUNY Binghampton)

The study of the stability of a fluid-fluid interface is important in a number of scientific and technological applications. In this project we consider two fluid layers separated by a horizontal planar interface subject to a vertical temperature gradient. The effects of various driving forces on the stability of the system can be taken into account, including buoyancy (known as Rayleigh-Benard convection), bulk density differences (known as Rayleigh-Taylor instabilities), and surface tension gradients along the interface (known as Marangoni instabilities). If the two layers represent different phases of the same material the stability results for a two-phase bi-layer system are quantitatively and even qualitatively different than for those of an immiscible system.

To estimate the relative importance of these types of instabilities for a two-phase system, we have considered a bi-layer geometry in which a horizontal fluid-fluid interface separates two semi-infinite layers of a single component fluid [1, 2]. We have performed linear stability calculations for horizontal fluid bi-layers that can undergo a phase transformation [1], taking into account both buoyancy effects and thermocapillary effects. We compare the familiar case of the stability of two immiscible fluids in a bi-layer geometry (the "inert" case [2]) with the less-studied case in which the two fluids represent different phases of a single-component material. The two cases differ in their respective interfacial boundary conditions.

Recently we have also studied a novel version of the two-phase system in which the onset of the instability is oscillatory in time [3]. This instability was first noted by researchers at the University of Minnesota, although the physical origins of the instability are poorly understood. We have been able obtain both numerical and analytical approximations which help explain the physical mechanism of the instability. We are currently working on an extension of the two-phase, bi-layer problem to the case of a binary alloy system. In a related study, we are considering the effects of buoyant convection on the stability of a mushy zone in a ternary alloy. The mushy zone is an idealized model of a two-phase region where liquid and solid co-mingle.

[1] G. B. McFadden, S. R. Coriell, K. F. Gurski and D. L. Cotrell, "Onset of Convection in Two Liquid Layers with Phase Change", *Physics of Fluids* **19** (2007), 104109.

[2] G. B. McFadden, S. R. Coriell, K. F. Gurski and D. L. Cotrell, "Convective Instabilities in Two Liquid Layers," *Journal of Research of NIST* **112** (2007), pp. 271-281.

[3] G.B. McFadden and S.R. Coriell, "Onset of Oscillatory Convection in Two Liquid Layers with Phase Change," submitted.

Funding: *NIST/NRC Postdoctoral Associateship Program*

Analysis of Diffusion and Stress in Binary Alloys

G. B McFadden
W. J. Boettinger (MSEL)

The *Kirkendall effect* refers to the motion of lattice planes in a crystalline binary alloy that is caused by a difference in the rates of diffusion of the two species. The difference in diffusion rates leads to a stress-free stain in the sample. This process is an example of the coupling between diffusion and elasticity during phase transformations. Current interest in stress generation in thin films due to diffusion processes has stimulated new interest in this classical problem. The effect can lead to measurable shape changes perpendicular to the direction of diffusion [1]. Such changes can often be detected by a shift of inert markers parallel to the diffusion direction; that is, deformation with pure axial displacement [2].

In current work we are analyzing the time-dependent bending of a bimetallic strip due to interdiffusion in a binary alloy. In this case the deformation is due to diffusion across the short dimension of the strip, so that the deformation and diffusion are co-linear. The model, which employs a Fourier method, couples simple beam theory and diffusion. The results are in good agreement with previously reported experiments.

[1] W. J. Boettinger, G. B. McFadden, S. R. Coriell, R. F. Sekerka, and J.A.Warren, "A Model for the Lateral Deformation of Diffusion Couples," *Acta Materialia* **53** (2005) 1995-2008.

[2] W. J. Boettinger, J. E. Guyer, C. E. Campbell, and G. B. McFadden, "Computation of the Kirkendall Velocity and Displacement Fields in a One-dimensional Binary Diffusion Couple with a Moving Interface," *Proceedings of the Royal Society* **463** (2007), 3347–3373.

[3] W. J. Boettinger and G. B McFadden, "Bending of a Bimetallic Beam Due to the Kirkendall Effect," in preparation.

Modeling of Self-Similar Grain Size Distributions

Geoffrey B. McFadden
C. S. Pande (Naval Research Laboratory)
K. P. Cooper (Naval Research Laboratory)

Polycrystalline solids typically consist of a large number of grains that differ in the orientation of their crystallographic axes. These grains have a distribution of sizes, and are separated by curved grain boundaries that, in turn, meet at grain boundary junctions. Driven by combinations of elasticity, diffusion, or capillary forces, the grain boundaries migrate in time, and the distribution of grain sizes consequently evolves dynamically.

A statistical description of the process is given by a grain size distribution function, which is assumed to depend on a mean grain radius and time. Grain growth under usual circumstances is known to approach a quasi-stationary distribution of grain sizes after a transient period. The quasi-stationary state in a wide variety of materials exhibits a scaling property, such that the grain size distribution has an invariant form when expressed in terms of the grain size scaled by its mean value. An accurate description of the spatial and temporal evolution of a polycrystal from an initial stage, through the transient period, and finally to the quasi-stationary state is still only poorly understood.

In this work, a size-based continuum stochastic formulation is presented based on topological considerations. This analysis leads to a Fokker-Planck equation for the size distribution which admits a similarity solution that describes the long time behavior of the grain growth. The solution for the case of two-dimensional grain growth recently appeared [1]. We are also addressing the three-dimensional version of this problem [2]. The resulting grain size distributions are shown to be in agreement with those obtained from computer simulations, indicating the validity of the stochastic approach.

[1] C.S. Pande, K.P. Cooper, G.B. McFadden, "Grain Size Distribution in Two dimensions in the Long Time Limit," *Acta Materialia* **56** (2008), pp. 5304–5311.

[2] C.S. Pande and G.B. McFadden, "Self-similar grain size distribution in three dimensions," in preparation.

Computational Modeling and Visualization of the Flow of Concrete

William George
Julien Lancien
Steve Satterfield
Marc Olano
John Kelso
John Hagedorn
Nicos Martys (NIST BFRL)

`http://math.nist.gov/mcsd/savg/parallel/dpd/`
`http://math.nist.gov/mcsd/savg/vis/concrete/`

Understanding the mechanisms of dispersion or agglomeration of particulate matter in complex fluids, such as suspensions, is of technological importance in many industries such as pharmaceuticals, coatings, and concrete, and is essential for improving our nation's infrastructure. These fluids are disordered systems consisting of a variety of components with disparate properties that can interact in many different ways. Modeling and predicting the flow of such systems represents a great scientific and computational challenge requiring large scale simulations at a fundamental level.

In collaboration with scientists in the Building and Fire Research Laboratory (BFRL), we are developing an application, called QDPD, a dissipative particle dynamics code which is capable of performing large scale simulations of suspensions. QDPD is highly parallel and has been shown to efficiently scale up to at least 1000 processors when running on the NASA supercomputer Columbia. Our goal in this project is to advance our understanding of the flow properties of a specific suspension, fresh concrete, a suspension composed of cement, water, sand, and stones.

This year we were awarded an additional 400,000 hours of computer time on the NASA Columbia supercomputer. In addition, we won 750,000 hours on the

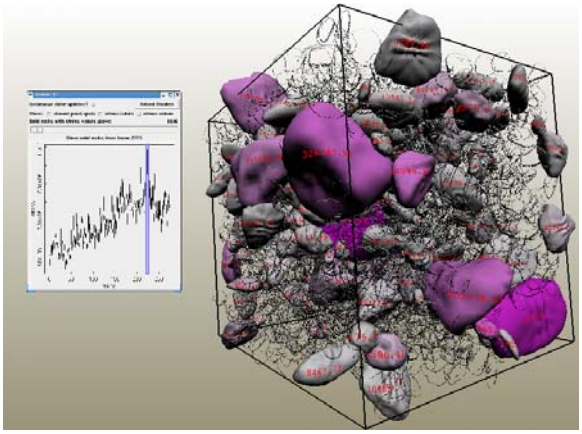

Figure 32. *Flow of Suspensions, with interactive controls.*

Argonne Blue Gene/P supercomputer at the Argonne Leadership Computing Facility, a system comprising 163,840 processors and 80 terabytes of main memory. The allocation was one of 55 awards of supercomputer time given in a peer-reviewed competition known as the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. Our work was recognized by a Department of Energy Office of Advanced Scientific Computing Research (OASCR) Award for a visualization demonstrated at the 2008 SciDAC Conference, held in Seattle, Washington, from July 13-17, 2008. The winning visualization presented the results of a high-performance computer simulation of the flow of a particular complex suspension: concrete. The visualization showed a flow of a realistic collection of rocks of various sizes whose shapes were obtained by scanning actual rocks. The direct visualization of rocks is augmented with embedded data to provide additional information on rock-rock interactions and per-rock stress. See Figure 32.

This year we modified the QDPD simulation code to include non-Newtonian fluids, that is, fluids in which the viscosity changes as a function of the applied strain rate. We improved the computation of curvatures and principal direction vectors that are calculated for the rocks that are input to the simulation. The code was ported to the Blue Gene/P supercomputer. The simulation is regularly run on 2,048 or more processors. Substantial changes are in progress to optimize the I/O for QDPD on systems of this size. Other modifications to the simulator are also in progress to allow QDPD to run efficiently on larger numbers of processes.

We have applied the QDPD code to study the onset of flow in colloidal suspension. These studies have significantly improved our understanding of the physical mechanisms controlling the yield stress of dense suspensions. Yield stress is a well-known property, crucial for applications, but is not well-understood. It is often thought of as the minimum force needed to break bonds in colloidal materials. However, we have shown that an important cause of the measured yield stress is a supplementary stress that develops due to the organization of the aggregate contacts in the compression quadrant as stress is distributed throughout the suspension when it is sheared; see Figure 33. This important discovery will affect how chemical admixture companies design their products to increase the workability of fresh concrete.

Our simulations have also shown that at low volume fractions, a slowly applied strain had a tendency to facilitate aging of the suspension. Stress growth curves could collapse onto a universal curve. The maximum stress was not dependent on the aging and could be reached sooner with the applied strain. At volume fractions close to 50 %, the stress growth curves did not exhibit a universal behavior. As the suspension aged, a barrier had to be overcome to begin flow.

This year we plan to continue large-scale simulations of colloidal suspensions, focusing more on the effect of size distribution and particle shape. We also plan to examine the effect of having a non-Newtonian fluid matrix in a suspension. Demonstration of concept is very important so we will be collaborating with experimentalists from industry to help validate our approach of modeling suspension with a non-Newtonian fluid matrix.

[1] N. S. Martys, D. Lootens, W. L. George, S. G. Satterfield, and P. Hébraud, "Spatial-Temporal Correlations in Concentrated Suspensions," 15th International Congress on Rheology, Monterey CA, Aug 3-8, 2008.

[2] N. S. Martys, D. Lootens, W. L. George, S. G. Satterfield, and P. Hébraud, "Stress Chains Formation Under Shear of Concentrated Suspension," 15th International Congress on Rheology, Monterey CA, August 3-8, 2008.

[3] N. S. Martys, W. L. George, S. G. Satterfield, and M. Olano, "Spatial-Temporal Correlations in Start-up Flows of Colloidal Suspensions," *Physical Review Letters,* submitted.

[4] N. S. Martys, W. L. George, and B.-W. Chun, "Modeling the Flow of a Suspension with a non-Newtonian Fluid Matrix," in process.

[5] N. S. Martys, C. F. Ferraris, V. Gupta, J. H. Cheung, J. G. Hagedorn, A. P. Peskin, E. J. Garboczi, "Computa-

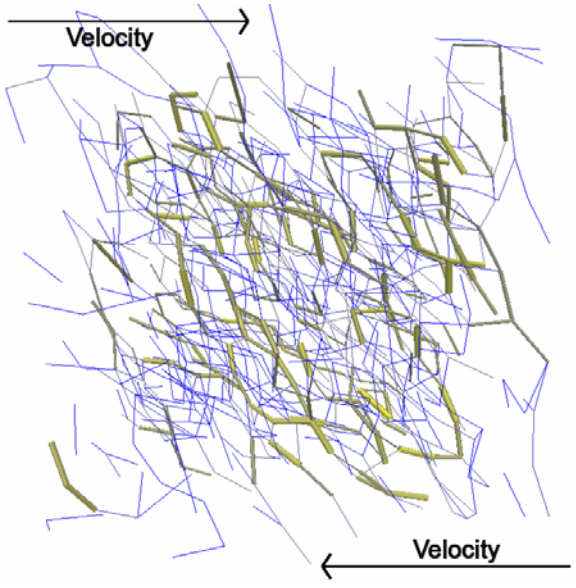

Figure 33. *Stress chains visualization. The representation consists of lines connecting aggregate center-to-center. The stress values (low to high) are color encoded (blue to yellow) as well as size (thin to thick line segments) encoded to emphasize the high stress chains. Here, the suspension is straining to the right at the top and to the left on the bottom. Note the lines are directed along the compression quadrant (oriented towards the bottom right). Stress is transmitted diagonally through a network of interparticle contacts. This mechanism is responsible for a supplementary stress.*

tional Model Predictions of Suspension Rheology: Comparison to Experiment," 12th International Conference on the Chemistry of Cement, Montreal, Canada, July 11, 2007.

[6] N. S. Martys, C. F. Ferraris, W. George, J. Lancien, J. Terrill, E. Garboczi, A. P. Peskin, J. Hagedorn, M. Olano, J. Kelso, S. Satterfield, H. Zhu, J. H. Cheung, B.-W. Chun, V. Gupta, D. Lootens, R. Flatt, and B. Descheneaux, "Computational Modeling of Suspensions: Recent Advances and Future Research Directions," VCCTL Bi-Annual Meeting, Gaithersburg MD, November 13-14, 2007.

[7] W. George, J. Lancien, N. Martys, J. Terrill, and E. Garboczi, "Large Scale Simulations of Suspensions," NASA Booth, Supercomputing 2007, Reno, NV, November 10-16, 2007.

[8] N. S. Martys, W. L. George, M. Olano, J. Kelso, John Hagedorn, S. G. Satterfield, and J. Terrill, "VCCTL Rheology Update", VCCTL Bi-Annual Meeting, Dec 3-4, 2008, Gaithersburg, MD.

[9] W. George, S. Satterfield, M. Olano, and J. Terrill, "Flow of Suspensions," movie at the Argonne National Laboratory booth, Supercomputing 2008, Austin, TX, November 15-21, 2008.

Funding: *Virtual Cement and Concrete Testing Laboratory Consortium*

Materials Data and Metrology for Applications to Machining Processes, Frangible Ammunition, and Body Armor

Timothy Burns
Steven Mates (NIST MSEL)
Richard Rhorer (NIST MEL)
Eric Whitenton (NIST MEL)
Debasis Basak (Orbital Sciences Corporation)

The NIST Pulse-Heated Kolsky Bar has been developed primarily for the measurement of the dynamic properties of metals. While it has been used for a number of other interesting applications over the past few years, including work on frangible bullets and bullet-proof vests, it was originally developed and its main use has been to obtain improved material constitutive response data for the modeling and simulation of high-speed machining operations. Two high strength steel bars, each of 1.5 m length and 15 mm diameter, are mounted on bearings to enable easy sliding of the bars in the axial direction and to resist bending in other directions. A cylindrical sample of the material to be tested is inserted between the two long bars, with care taken to align the system for axial symmetry. Ignoring radial effects, the data taken with this system can be analyzed using one-dimensional wave theory. One of

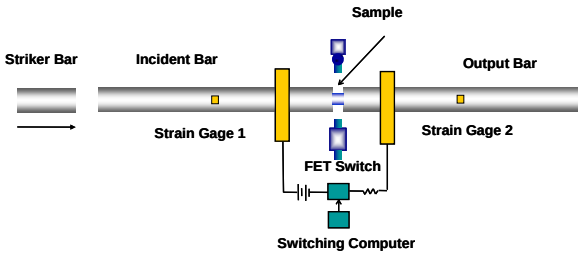

Figure 34. *Schematic drawing of the NIST Kolsky Bar, with electrical pulse-heating of the sample.*

the long bars, called the incident bar, is impacted by a striker, launched by an air gun. The striker is a much shorter bar made from the same steel, and with the same diameter, as the two long bars. In this way, the sample is rapidly loaded by a compressive wave. The corresponding stress vs. strain response of the sample material is obtained by means of two strain gauges, one placed at the midpoint of each of the long bars.

What makes the NIST Kolsky bar unique is that the system has been combined with an existing resistive-heating facility, the NIST Subsecond Thermophysics Laboratory. This Laboratory was originally developed to measure physical properties of metals at high temperature, such as the critical point at melting of a pure metal. It has the capability to pre-heat an experimental sample extremely rapidly, in situ, using a precisely controlled DC electrical current.

With this combined system, a metal sample can be pulse-heated from room temperature to a significant fraction of its melting temperature in less than a second. Following this rapid pre-heating, the temperature

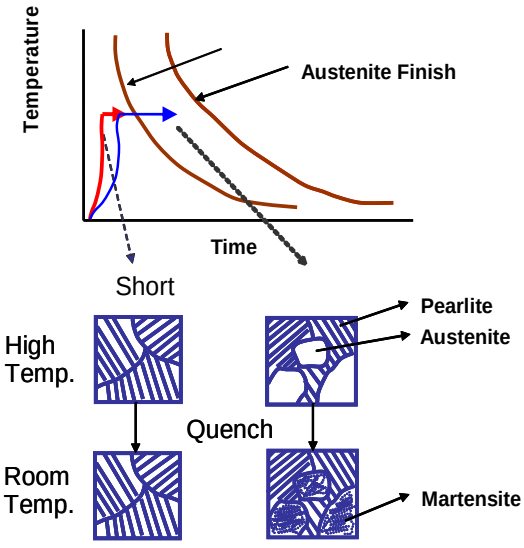

Figure 35. *Schematic illustration of hypothetical microstructural state of AISI 1045 during "short" and "long" rapid heating tests, and after rapid quenching in air immediately following Kolsky Bar impact testing.*

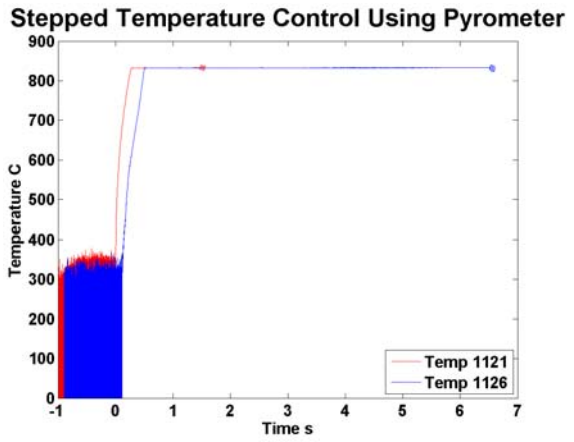

Figure 36. *Examples of pulse-heating of samples prior to impact. To minimize arcing and local hot spot formation, the current is applied in two steps. It appears to be noisy below approximately 350 °C, because the pyrometer used for temperature control only becomes sensitive at higher temperatures.*

can be precisely controlled to remain constant for up to several seconds, if desired, and then the current shut off extremely rapidly by the control system just as the air gun that launches the striker bar is fired to initiate a Kolsky bar test. We now have the capability to explore the interesting question of whether or not there are differences in the mechanical response of alloys at high temperature that depend upon the rate of heating and the time at temperature of a sample of the material prior to loading it in rapid compression.

Because we received no funding from the NIST Office of Law Enforcement Standards this year, work in the Kolsky Bar Laboratory has focused on studies of carbon steels of interest in manufacturing. Results we have obtained support the hypothesis that some thermally activated changes in the material's microstructure have insufficient time to run to completion under the rapid heating conditions that are present in high-speed machining operations, with the result that higher cutting forces are required for these processes. This has significant implications for the modeling of machining operations using finite-element software.

[1] S. P. Mates, R. Rhorer, E. Whitenton, T. Burns, and D. Basak, "A Pulse-heated Kolsky Bar Technique for Measuring the Flow Stress of Metals at High Loading and Heating Rates," *Experimental Mechanics*, to appear.

OOF: Finite Element Analysis of Material Microstructures

Stephen A. Langer
Andrew C.E. Reid
Valerie R. Coffman
Gunay Dogan (Cornell University)

http://www.ctcms.nist.gov/oof/

The OOF Project, a collaboration between MCSD and MSEL, is developing software tools for analyzing real material microstructure. The microstructure of a material is the (usually) complex ensemble of polycrystalline grains, second phases, cracks, pores, and other features occurring on length scales large compared to atomic sizes. The goal of OOF is to use data from a micrograph of a real material to compute the macroscopic behavior of the material via finite element analysis. OOF is intended to be a general tool, applicable to a wide variety of microstructures in a wide variety of physical situations.

During FY 2008 the OOF team released seven new versions of the OOF2 program, namely version 2.0.4 and six nominally "alpha" versions of 2.0.5. The 2.0.5 versions were called alpha because some of the new features were incomplete, but important improvements to old features made a release worthwhile. OOF2 was downloaded over 1,100 times from the NIST website during the year.

The OOF customer base continues to be broad and difficult to define, since users aren't required to register before using the software. The developers received e-mail from 62 different users who needed help installing or running the program. Besides being available for download from NIST, OOF2 can be run on-line through Purdue University's nanoHUB site for nanoscience simulation[6]. OOF2 had 160 nanoHUB users in 1,360 interactive sessions in the past 12 months.

Most of the work on OOF2 this year was behind the scenes. The new releases this year use memory and CPU more efficiently. The majority of the development effort, however, was spent on features that will appear in future versions, particularly time dependence and three dimensions.

OOF2 now has the infrastructure for solving time-dependent problems, including a variety of time stepping methods (both implicit and explicit, linear and nonlinear, and first and second order in time). It can cache data from previous time steps either in memory or on disk. The user can specify a schedule of output quantities to be computed during the time evolution and displayed on the screen (either as text or graphics) or saved in a file. With the addition of adaptive time

[6] http://www.nanohub.org/

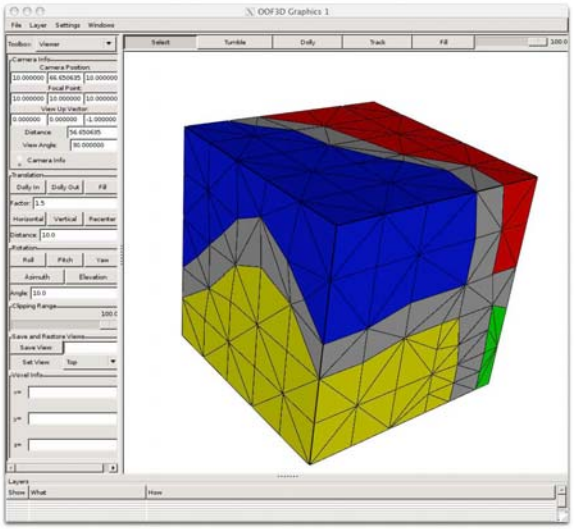

Figure 37. *A screenshot of OOF3D, showing a finite element mesh created from an artificial test microstructure.*

stepping and an improved user interface for time evolution, a time dependent version of OOF2 should be available in FY 2009.

Development of the three dimensional version, OOF3D, is proceeding in parallel with OOF2. The program can now load a 3D microstructure, and create and solve a finite element mesh, although it only contains a subset of the OOF2 tools, as well as new tools required by the presence of a third dimension. OOF3D will be released once the full OOF2 toolset is in place and various memory and user interface issues have been addressed.

[1] C. E. Reid, S. A. Langer, R. C. Lua, V. R. Coffman, S.-I. Haan, and R. E. Garcia, "Image-based Infinite Element Mesh Construction for Material Microstructures," *Computational Materials Science* **43** (2008), pp. 989-999.

Funding: *NIST/NRC Postdoctoral Associateship Program*

Micromagnetic Modeling: Software and Applications

Michael Donahue
Donald Porter
Robert McMichael (NIST CNST)
June Lau (NIST MSEL)

http://math.nist.gov/oommf/

Advances in magnetic devices such as recording heads, field sensors, magnetic nonvolatile memory (MRAM), and magnetic logic devices are dependent on an understanding of magnetization processes in magnetic materials at the nanometer level. Micromagnetics, a mathematical model used to simulate magnetic behavior, is needed to interpret measurements at this scale. MCSD is working with industrial and academic partners, as well as with colleagues in the NIST CNST, MSEL, and EEEL, to improve the state-of-the-art in micromagnetic modeling.

Michael Donahue and Donald Porter in MCSD have developed a widely used public domain computer code for doing computational micromagnetics, the Object-Oriented Micromagnetic Modeling Framework (OOMMF). OOMMF serves as an open, well-documented environment in which algorithms can be evaluated on benchmark problems. OOMMF has a modular structure that allows independent developers to contribute extensions that add to its basic functionality. OOMMF also provides a fully functional micromagnetic modeling system, handling both two and three-dimensional problems, with sophisticated extensible input and output mechanisms. During fiscal year 2008 alone, the software was downloaded more than 4,700 times, and use of OOMMF was acknowledged in more than 125 peer-reviewed journal articles. OOMMF has become an invaluable tool in the magnetics research community.

Key new developments in OOMMF over the last year include:

1) Increasing support for multi-threaded computation to exploit the concurrency opportunities made possible by the increasing availability of multi-core hardware platforms. Common micromagnetic simulations can be computed about three times faster with these enhancements.
2) Collaborative work with the Polish Academy of Sciences to compute magnetic energies according to assumptions of periodic boundary conditions. This is a popular approach to greatly extend the size of devices that can be simulated.
3) Creation of additional modules tailored for the simulation of polycrystalline materials, common in practical device designs.
4) Development of a norm-preserving semi-analytic numerical integration scheme that improves numerical stability and reduces problem stiffness in some critical cases.

In addition to the continuing development of OOMMF, Donahue and Porter also engage in collaborative research using OOMMF. They have recently provided modeling support for a four-laboratory (EEEL, MSEL, PL and ITL) NIST Innovation in Measurement Science (IMS) project on high sensitivity magnetic sensors, and also play an instrumental role in the 2007 IMS project (with CNST, MSEL and EEEL) entitled "Magnetic nanostructures for post-CMOS electronics." They also provide modeling and simulation consultation for work in EEEL on spin momentum

transfer. The MCSD micromagnetic project produced three journal papers, two conference presentations and two invited talks this past year.

Work in progress includes a continuing effort to develop techniques to minimize or control errors introduced when the spatial resolution of the computation grid does not exactly represent the boundaries of the simulated material. Further progress on multi-threaded computation is focused on partitioning energy calculations to more effectively use limited memory bandwidth on common multi-processor architectures. Another focus area is the modeling of spin momentum transfer to magnetic domain walls arising from electric current flow. This spintronic effect plays an increasing role in the near-term designs of magnetic memory and logic devices.

[1] K. M. Lebecki, M. J. Donahue, and M. W. Gutowski, "Periodic Boundary Conditions for Demagnetization Interactions in Micromagnetic Simulations," *Journal of Physics D-Applied Physics* **41** (2008), Article No. 175005 (10 pages).

[2] D. G. Porter and M. J. Donahue, "Precession Axis Modification to a Semi-analytical Landau-Lifshitz Solution Technique," *Journal of Applied Physics* **103** (7) (February 25, 2008), pp. D920-1 to D920-3.

[3] J. M. Shaw, S. E. Russek, M. J. Donahue, T. Thomson, B. D. Terris, M. Olsen, and M. Schneider, "Reversal Mechanisms in Perpendicular Magnetic Nanostructures," *Physical Review B* **78** (2008), Article No. 024414 (5 pages).

Funding: *NIST Innovations in Measurement Science Program*

Sparse Representations in High Dimensional Geometry

Bradley Alpert
Yu Chen (New York University)

Laboratory and clinical diagnostic systems often produce images that lack quantitative information or comparability across imaging systems, which can lead to difficulties in diagnosis and long-term management of disease that requires patient monitoring. Elaborate measurements are taken, digitally processed, then reported in essentially qualitative form. For MRI, functional MRI, and some other processes, formation of the images relies on filtering and Fourier transformation.

Although the computational recovery of functions from their Fourier representations is considered routine, aside from possible efficiency concerns arising from non-grid data, standard approaches assume that the functions are smooth and can be recovered with the discrete Fourier transform (computed via the FFT). When this assumption fails, as is evident from Fourier data that are not small at the highest frequencies measured, some sort of attenuation, or belling, scheme is typically used. This procedure tends to blur sharp features and, being more art than science, can lead to different functions (images) in different systems. Although this problem, and the Gibbs phenomenon, have received considerable attention from mathematicians in recent years, recently developed methods tend to work well only in an asymptotic sense; they do not exploit available data efficiently.

Despite inherent limits on resolution that can be obtained from truncated Fourier data, a change of assumptions from smooth to piecewise smooth can lead to significantly improved recoveries. Procedures implementing this idea are not completely established even for functions (signals) in one dimension; they are yet more challenging in two and three dimensions, where discontinuities may be expected to occur along mostly smooth curves or surfaces. This constraint, which magnifies the advantage of the piecewise smoothness assumption, must be appropriately reflected in the methods used. Bradley Alpert of MCSD and Yu Chen of New York University are conducting extensive numerical experiments to understand this environment and to develop reliable procedures for these problems. They have recently developed procedures for locating discontinuities, or even a linear combination of delta-functions, from Fourier data taken at well below Nyquist sampling rates.

A related, yet more general, challenge is parsimonious, or sparse, representation and recovery of functions under assumptions appropriate to an application. Although there is considerable interest in the mathematical community in these problems (L_1-norm minimization, compressed sensing, sparse representation in high-dimensional geometry), and much recent progress (initiated in pivotal work by Candes, Romberg, and Tao) most of the methods being explored are limited to linear spaces. While linearity is a natural starting point, strong evidence suggests that image recovery cannot be done this way yet nevertheless is within reach. It is the goal of this project to develop procedures for robust recovery of piecewise smooth, or otherwise constrained, functions from Fourier data.

Information Discovery, Use and Sharing

Society is awash in data – our ability to amass data has outpaced our ability to use it. Extracting knowledge, information, and relationships from this data is one of the greatest challenges faced by the scientists in the twenty-first century. The data can be as diverse as biological research data, medical images, automated newswire, speech, or video. The Information Discovery, Use, and Sharing Program fosters innovation throughout the information life cycle by developing the measurement infrastructure to enhance knowledge discovery, information exchange, and information usability. The Program enables novel computational approaches to data collection and analysis to be combined with improved interoperability techniques to effectively extract needed information from the wealth of available data. – Mary Brady, Program Manager.

Standardization of Lung Cancer Growth Measurements

Adele Peskin
David Gilsinn
Javier Bernal

See feature article, page 43.

Funding: *2007 NIST ACI Initiative on Bioimaging*

Figure 38. *Simulated Lung Cancer Nodules.*

Estimating Volumes of Simulated Lung Cancer Nodules

David E. Gilsinn
Bruce R. Borchardt (NIST MEL)
Amelia Tebbe (St. Mary's College of Maryland)

Lung cancer is a disease of uncontrolled cell growth in tissues of the lung. It is the most common cause of cancer related deaths in men and second most in women and is responsible for an estimated 560,000 deaths in the US in 2007. Lung cancer lesions may be seen in Computerized Tomography (CT) images. Modern CT is a medical imaging method that employs digital image processing techniques to generate three dimensional images from a large sequence of two-dimensional X-rays made around a single axis.

The Food and Drug Administration (FDA) is conducting research on the development of reference cancer lesions, called phantoms, to test CTs and their underlying proprietary software. Two sample phantoms were loaned to NIST to estimate volumes. The material of which the phantoms are composed simulates lung cancer material. They can be inserted into a simulated body torso for the testing of CT scanning equipment. The two phantoms are shown in Figure 38. Although they seem spherical they are slightly non-spherical.

The approach chosen to determine precise volumes for this study is based on a fundamental theorem in calculus, the Divergence Theorem, which is an analogue of Green's Theorem in two dimensional spaces. In the Divergence Theorem a volume integral is shown to be equal to a surface integral. Thus, if a model of the surfaces of the phantoms could be developed, we could use the Divergence Theorem to estimate their volumes. In order to develop a surface model we needed to obtain data about the surface. This was done using a coordinate measuring machine (CMM) in the Manufacturing Engineering Lab (MEL) at NIST; see Figure 39.

Figure 39. *Coordinate Measuring Machine.*

This machine produced a set of precisely measured (x, y, z) points on the surface of each phantom.

The data were then transformed to spherical coordinates and modeled using a set of basis functions, called B-splines. After computing a best fit of the B-spline model to the data in the least squares sense, the model was used to generate a grid of values on the surface. These values were used to form surface triangles that were then used to compute the necessary surface integrals and finally the volume.

The quality of the volume estimates depended on the surface grid sizes. The estimates of the volumes of the two simulated lung cancers were 4337.1 cubic mm with an expanded uncertainty of 3.5 cubic mm for the nodule on the right of Figure 38 and 3922.1 cubic mm with an expanded uncertainty of 6.6 cubic mm for the smaller nodule on the left in Figure 38. The larger uncertainty for the smaller nodule was due to its more non-spherical shape.

Characterizing and Improving Segmentation Methods for Images of Biological Cells

Adele Peskin
Javier Bernal

Precise identification of individual cells in an image, a process known as segmentation, is necessary in order for biologists to characterize cell lines. Because of the nature of this image data, segmentation is a complex problem. Image quality varies depending on the underlying imaging techniques employed, which differ widely across laboratories. The goal of this project is to improve segmentation methods for use in cell-based image analysis. In particular, we seek to characterize the quality of cell imagery in order to be able to predict which segmentation techniques will work well. In addition, we wish to predict, based on this quality, what kind of segmentation errors are to be expected with different segmentation techniques.

As part of this project an effort is under way to develop and implement mathematically sound reference implementations of known segmentation methods. This effort includes an analysis of the underlying mathematical principles on which these methods are based in order to better characterize their behavior. We have found that a single segmentation method gives varying results depending upon the details of its implementation. For example, we have several different implementations of a KMeans thresholding segmentation method. Repeated measurement of a single image gave consistent segmentation results using our group's implementation, but a range of different results using an open source implementation.

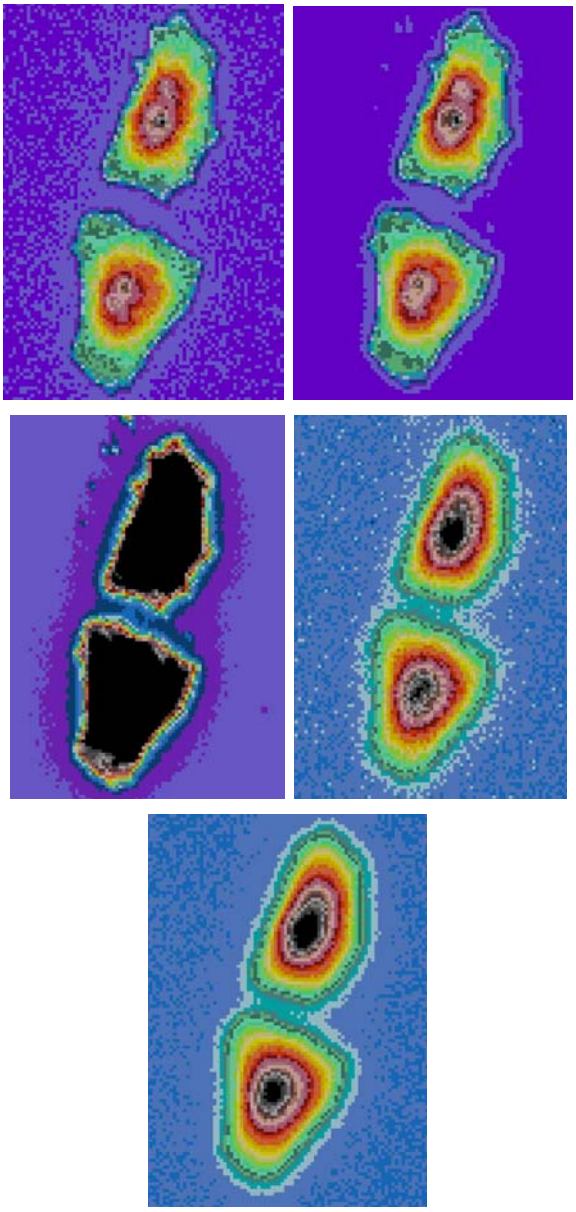

Figure 40. *Five pictures of the same biological cells taken with five different set of imaging factors, color-coded with 40 colors over the range of each image, to show the differences in pixel intensity gradients in the five images.*

One of our tasks involves devising a method to assess the quality of an image in order to select an appropriate segmentation method and predict its accuracy. We have sets of images each of which were taken on identical groups of cells. For example, Figure 40 shows pictures of identical cells, taken with five different image settings, and then color-coded pixel by pixel to see the difference in quality. We are studying the pixel intensities and the corresponding gradient fields of well characterized data in order to be able to assess image quality.

Cultivating (Legacy) Mathematical Data

Bruce R. Miller
Michael Kohlhase (Jacobs University, Bremen)
Heinrich Stamerjohanns (Jacobs University, Bremen)

http://dlmf.nist.gov/LaTeXML/

LaTeX may not be the first document format reached for by a generation raised on word-processors. Nevertheless, due to its support of extensive, complex, mathematics and high-quality typesetting, it has become the *de facto* standard for most mathematically oriented scientific fields. It is the preferred format for the preprint archive at arXiv[7], the American Mathematical Society, the American Physical Society, the Association for Computing Machinery, and others. It was the clear choice for authoring in NIST Digital Library of Mathematical Functions[8] (DLMF) project.

Ideally, a "digital library" will go far beyond print — LaTeX's primary focus — and beyond a mere archive, to support search, connections between related documents, and layers of information appropriate to different audiences. Tools are needed to convert material from the source LaTeX into XML and MathML that can be analyzed and transformed for use on the Web and beyond. For all its benefits, LaTeX and TeX are notoriously quirky and programmable, leading to a number of difficulties in conversion. There being no available tool sufficiently powerful and robust for the job, we created LaTeXML for converting LaTeX to XML. A period of parallel development of the DLMF and LaTeXML ensued, with each project leading to improvements in the other.

Meanwhile, a project at Jacobs University in Bremen (formerly International University Bremen), was carrying out research in Mathematical Knowledge Management (MKM), particularly in mathematically aware indexing and search. Desiring to get beyond toy problems, they needed a large corpus of material to work with, preferably covering a variety of styles and subjects. What better resource than the arXiv, with currently over half a million articles in fields of astronomy, mathematics, physics?

Thus began a collaboration for the mass conversion of documents in the arXiv from LaTeX to XHTML+MathML. The usage scenario here is quite different from the DLMF. In the DLMF we own the sources and can adjust the markup, add declarations and fine tune the system to completely control the transformation to get exactly the output we desire. With the arXiv, authors use a vast array of styles of markup and many obscure add-on packages. Indeed, many articles in the arXiv are just drafts that are not even valid LaTeX! Clearly, we cannot devote hours of attention to each article to get a proper conversion even if it were allowable to modify the sources.

To deal with this task, a system for distributed processing and reporting was constructed. This allows for relatively quick reprocessing of the entire body after changes are made. Several students at Jacobs are also involved in developing implementations of the more commonly used style files. At present, we are working with a set of over 400,000 documents. Almost half have been converted into relatively clean form; most of the remaining articles are likely to be useable for some purpose. Statistics on the articles processed are available online[9].

This exercising of LaTeXML has led to many improvements in coverage and robustness, leading to a tool of much broader application outside of the DLMF. Moreover the rich variety of mathematical notations encountered is leading to improvements in the mathematical grammar that can be recognized. With this material now available to them, researchers at Jacobs are starting to look into ways to infer semantic content from primarily presentational markup.

Related to this work, we have given presentations at the Mathematical Knowledge Management conference at Birmingham in June. Finally, Bruce Miller and Michael Kohlhase have been active participants in the W3C Math Working Group developing the third version of the MathML standard.

Funding: *NIST Systems Integration for Manufacturing Applications (SIMA) Program.*

[7] http://arxiv.org/
[8] http://dlmf.nist.gov/
[9] http://www.kwarc.info/projects/arXMLiv/

Complex Systems

Complex Systems are composed of large interrelated, interacting entities which taken together, exhibit a macroscopic behavior which is not predictable by examination of the individual entities. The Complex Systems Program seeks to understand the fundamental science of these systems and develop rigorous descriptions (analytic, statistical, or semantic) that enable prediction and control of their behavior. Initially focused on the Internet and Grid Computing, this Program will facilitate predictability and reliability in these areas and other complex systems such as biotechnology, nanotechnology, semiconductors, and complex engineering. – Sandy Ressler, Program Manager

Methods for Characterizing Complex Networks

Isabel Beichl
Brian Cloteaux
Francis Sullivan (IDA Center for Computing Sciences)

See feature article, page 27.

Funding: *2007 NIST ACI Initiative on Cyber Security*

TCP Metastability and Cascading Failures in the Internet

Vladimir Marbukh

See feature article, page 29.

Funding: *2007 NIST ACI Initiative on Cyber Security*

Analysis of a Distributed Protocol for Network Control

Fern Y. Hunt
Vladimir Marbukh

Congestion control is critical to the (smooth) functioning of modern communication networks in particular and the Internet in general. The goal of congestion control protocols is to allow many users to share network resources without causing congestion collapse. In addition, it is desirable to use as much of the available network capacity as possible and ensure some fairness among users. In was shown by Frank Kelly of the University of Cambridge, and later many others, that achievement of these objectives can be represented as a problem of global optimization of system utility where some notion of fairness is expressed by a choice of utility function. Congestion control can then be seen as a distributed iterative solution to this problem (Lin and Shroff).

Recently Vladimir Marbukh considered the problem of appropriately routing traffic networks within the TCP/AQM protocol and proposed the introduction of controls that randomly assign routes to traffic in a way that minimize the mean link load. We formulated a utility optimization problem that incorporates a random route allocation scheme, defined by a probability distribution whose degree of randomness is controlled. We seek to use the distribution entropy to modulate between the two extremes of path specification; minimum cost path selection, as used in the Open Shortest Path First (OSPF) routing protocol on one end and equiprobable path allocation on the other. Earlier, in simulations of a stochastic model, Marbukh and Klink showed that networks operating under a random allocation scheme displayed oscillatory behavior when the entropy approached zero — corresponding to minimum path cost routing, and non-oscillatory (called stable) behavior as the entropy increased passed a critical value.

We derived a general model of network operation under this congestion/routing control protocol as a solution of our optimization problem. We then examined the behavior of this model for two sample topologies. Here we demonstrated the existence of equilibria that are determined by the link capacities and the topology of the network. Given a topology and set of link capacities there is a distinguished critical route distribution where one sees stable equilibrium behavior. This is maintained for distributions with larger entropy but unbounded behavior appears for distributions with smaller entropy. This is analogous to the simulation observations of Marbukh and Klink. We show that under smoothness assumptions on our objective function, the equilibria are solutions of the original optimization problem.

One of our examples, a five-link, three-route network, as simple as it is, shows a new phenomenon that is not present in a two link network — the first and simplest possible network. Using the route distribution entropy as a parameter it can be shown that the problem changes character once the entropy of the three

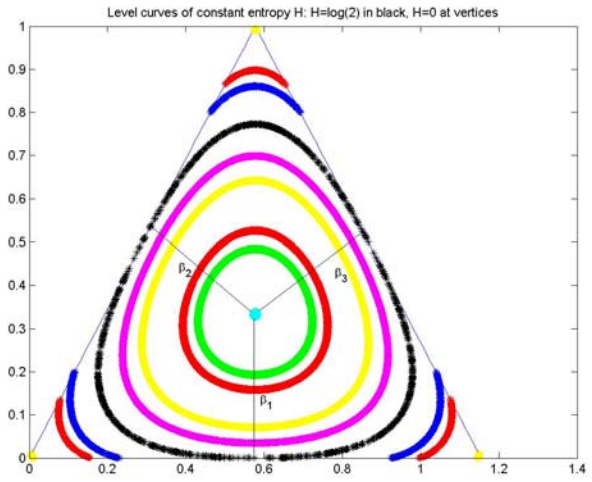

Figure 41. *In order to control congestion and maximize network utility, a variety of routing protocols can be used to allocate traffic. When multiple paths linking a source and destination are available but only a single path, e.g. a shortest path or path of minimum cost, is selected we have the OSPF protocol. On the other hand, traffic could be allocated uniformly when the paths have the same cost. A spectrum of path specification strategies lies in between these extremes. The accompanying figure represents these possibilities in the case of a single source destination pair with 3 possible connecting paths.*

route network reaches log 2, the largest entropy achievable by a two route network. This is illustrated in Figure 41. For entropies less than log 2, the region defined by the optimization problem is no longer convex. Thus for the problem we consider the range of route distributions is restricted to those that are not achievable by a simpler topology with two links.

Each allocation can be associated with a point in or on the equilateral triangle with an altitude height of 1. Here β_1, the perpendicular distance of the point to the horizontal base of the triangle is the fraction (or probability) of traffic sent to route 1. β_2, the perpendicular distance from the point to the left side of the triangle is the fraction allocated to route 2 and β_3, the fraction allocated to route 3 is the analogous distance for the right side. The entropy associated with an allocation,

$$H(\beta_1,\beta_2,\beta_3) = -(\beta_1 \log(\beta_1) + \beta_2 \log(\beta_2) + \beta_3 \log(\beta_3))$$

provides an important quantitative description of its degree of randomness. In OSPF ($H = 0$), traffic is sent along a path defined by one of the vertices of the triangle. Equal cost multipath allocation is assigned to the center of the triangle and has the maximum value of $H = \log(3)$. Bands of points with the same value of H (up to accuracy .001) are given the same color. The black band is associated with the critical value $H = \log(2)$, the maximum entropy for a network with 2 paths.

[1] V. Marbukh, "Decentralized Control of Large-Scale Networks as a Game with Local Interactions: Cross-layer TCP/IP," in *Proceedings of the 2nd International Conference on Performance Evaluation Methodologies and Tools*, Nancy, France, October 2007.

[2] F. Hunt and V. Marbukh, "A Model of Joint Congestion Control and Routing through Random Assignment of Paths," in preparation.

[3] F. Hunt and V. Marbukh, "Dynamic Routing and Congestion Control Through Random Assignment of Routes," in *Proceedings of the 5th International Conference on Cybernetics and Information Technologies, Systems and Applications (CITSA 2008)*, Orlando FL, July 2008, pp. 161-164.

Funding: *2007 NIST ACI Initiative on Cyber Security; NIST Innovations in Measurement Science Program*

Grid Simulation using Piecewise Homogeneous Markov Chains

Fern Y. Hunt
Christopher Dabrowski (NIST ITL)

This past year, F. Hunt and C. Dabrowski have collaborated on the development of a Markov chain model of the operation of a large scale computational grid system. Typically, discrete event simulation is used to model the interaction of service providers and consumer clients as they enter in service agreements for a variety of computer services. Such simulations, particularly those that seek to accurately reproduce system structure and component behavior, are computationally expensive; even simulating a single day takes many hours of CPU time. Nevertheless, there is a critical need for models that can provide rapid, and potentially scalable, simulation of large scale grid systems. Such models can be perturbed to produce alternative system execution paths and identify scenarios that lead to system breakdown or other anomalous behavior.

By compressing the more than 38 states associated with the grid model into seven states, Dabrowski was able to convert the simulation model into a piecewise homogeneous Markov chain with a small number of pieces. The nature of these states enabled us to develop perturbation schemes that are associated with three scenarios involving failure to meet service guaranties. The simulations with the piecewise Markov chain compared very favorably with the results of large scale grid simulations of the same scenarios. We were able to produce an envelope of Markov chain behaviors into which the grid simulations fell, and moreover we could identify thresholds in variables that lead to system failures (as measured by a decrease in job completion rate). Our formalism enabled us to propose possible reasons for this failure in terms of the state variables. The Markov chain simulations (which take a most a couple of hours) are good predictors of the re-

sults of grid simulations of substantially longer duration.

[1] C. Dabrowski and F. Hunt, "Using Markov Chain Analysis to Study Dynamic Behavior in Large-Scale Grid Systems," in *7th Australasian Symposium on Grid Computing and e-Research*, Wellington, New Zealand, 20 - 23 January, 2009, to appear.

[2] C. Dabrowski and F. Hunt, "Markov Chain Analysis for Large-Scale Grid Systems," in preparation.

Funding: *NIST Innovations in Measurement Science Program*

Standard Reference Data for Complex Network Research

Roldan Pozo

The study of complex or "self-organizing" networks is a relatively young, yet active research area. Much of the focus is to determine how graphs from "real" phenomena (such as protein interactions, social networks, structure of the World Wide Web, etc.) differ from synthetic or "engineered" graphs that are constructed by design. Clearly they are different, but how? Three key metrics (clustering coefficient, graph diameter, and degree distribution) are commonly used in the literature, and these techniques have served as encouraging first steps, but further research is needed to yield more meaningful distinctions and truly understand the nature of these complex networks.

One fundamental question to ask, however, is what are the precise "real" graphs that are being used to drive these theories? Clearly, the conclusions being reached can be only as good as the quality and availability of the original data. Most papers provide little information about this, making any kind of comparison or verification of approaches difficult.

For example, a commonly studied collaboration network in the literature is the Movie Actor graph, where two movie actors are connected if they have worked together on the same movie. This information is usually culled from the Internet Movie Database[10]. However, there is no single Movie Actor graph because this database is a changing snapshot of the entertainment industry at any given time. Furthermore, what one chooses to count in their particular search (foreign films, documentaries, TV movies, specific epochs or genres) will generate different graphs. In fact, there are several versions of these Movie Actor graphs cited in the literature, so a citation is not really meaningful unless it contains a timestamp and documentation of the specific queries used in its generation. A similar argument can be made for other association networks, such as co-author or citation graphs.

Another example is networks based on the structure of the World Wide Web. In this case, each web page is a node and two pages are connected if there a hypertext link from one to the other. These are typically generated by web crawlers and can vary greatly depending on where they crawl and what decisions they make about what constitutes a "link." For example: Do different anchors to the same page count as separate links? How are non-static pages handled? What file types are processed? How robust it is with broken links and sloppy HTML pages? This list can be quite long, and each decision path will yield a different graph. Aside from this, the Web is, of course, not a static entity. Many of the Web graphs cited provide little of these details and usually present only the topological structure (i.e., URLs removed) so it is difficult to confirm or validate them independently. To address these issues, we have developed and deployed our own set of web crawlers that report not only topological structure, but also record the node URLs for verification and validation, as well as generate a unique checksum to ensure that pages which contain identical material (e.g., aliased, or mirrored pages) show up as unique. Furthermore, these web crawls can be fined-tuned to record and characterize data along a precise decision path, thus allowing us to see what differences are generated when we alter the definition of "what is a link" and "what is a node?"

Similar concerns exist for graphs taken from other application domains. Often the data is incomplete, poorly documented, or missing pieces crucial to its description. There are ongoing debates about whether particular networks are really scale-free or not, or whether the conclusions of some papers are actually valid. Usually these concerns lie at the heart of how data was collected and analyzed.

Our motivation for developing a standard reference data set for complex networks is to provide a collection of public and well-documented "real-world" graphs that serve as a test bed of reasonably validated data which researchers can utilize to (1) verify and compare the algorithms and analysis of various approaches in the literature, (2) to ensure that everyone is using the same reference data in their studies, (3) to have one convenient location where one can browse and search for network graphs of various characteristics, (4) to provide a focal point for the research community to contribute and exchange network graphs from various application domains, and (5) to provide a testbed for development of software analysis tools (e.g., graphical viewers, graph partitioners, clustering algorithms) that will aid in the further development and research of complex networks.

In further aiding the understanding of such network systems, we have developed a set of software

[10] http://www.imdb.com/

tools for working with mathematical graphs (Ngraph++) and processing and transforming data collections into various formats (e.g., Matrix Market, GraphViz, Matlab) for output and analysis. Ngraph++ computes clustering coefficients and degree distributions for directed/undirected graphs. Although a small and simple library, it provides intuitive graph operations (subgraphs, unions, differences) and complex-network features not commonly found in more complex packages. Together with various Python and scripting tools, and a reference dataset of network graphs, this forms a kernel package for analyzing complex network data in a portable and universal fashion.

Funding: *2007 NIST ACI Initiative on Cyber Security*

Visualization of Network Dynamics

John Hagedorn
Cedric Houard
Dong Yeon Cho
Judith Terrill
Fern Hunt
Kevin Mills (NIST ITL)
E. Schwartz (NIST ITL)

http://math.nist.gov/mcsd/savg/vis/network/
http://math.nist.gov/mcsd/savg/software/DiVisa

As part of the ITL Complex Systems program, we are creating visualizations and developing a software system that enables researchers to view visualizations and analyze data from their model output. The purpose of these visualizations is to provide better understanding of the grid computing and network models. Figure 42 shows a screen shot of a software tool DiVisa that we developed to do this.

We have been using machine learning techniques to study the model outputs. Figure 42 shows a study of congestion control algorithms.

In the coming year we will be visualizing routing protocols, as well as continuing with our data analyses.

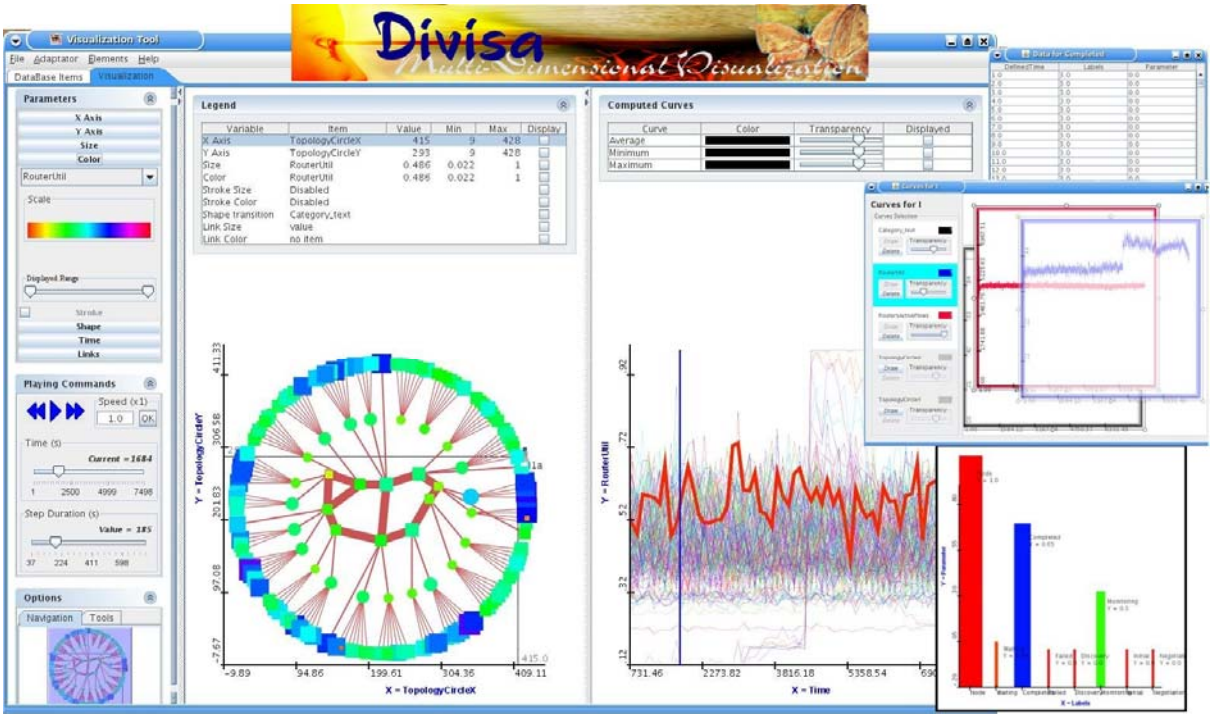

Figure 42. *One frame from an animated visualization of a simulation of the "Abilene" network. On the left panel, the innermost nodes are the backbone routers, the next ring represents the subnet routers, and the outermost ring shows the leaf routers. On the right panel, the time series curves of the routers (one is currently selected) are displayed. On the right side, different options are available: numeric data display (top), superposed time series curves (middle); and diagram view (bottom).*

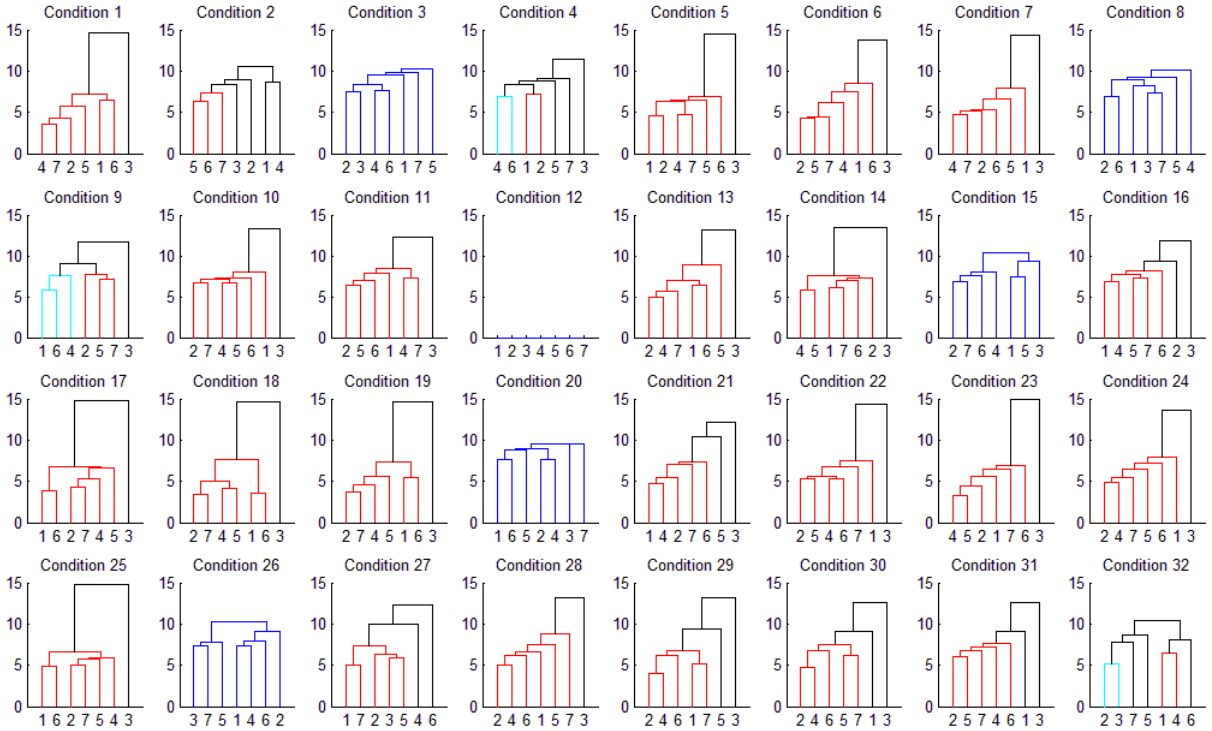

Figure 43. *A cluster analysis of the differences in one time period from among 45 responses. The X axis represents y congestion control algorithms. Each "small" figure represents one of 32 conditions applied to the 7 algorithms. Note that in condition 12 we see an identical response from each algorithm (due to no congestion on the network). Also note that in many of the conditions the largest distance graphed is from algorithm number 3 and the rest of the algorithms. Number 3 is the "FAST" TCP congestion control algorithm which performs markedly different.*

Virtual Measurement Systems

A virtual measurement is a quantitative result and its uncertainty, obtained primarily by a nontrivial computer simulation or computer-assisted measurements. Examples of virtual measurements include computational models of physical systems and visualizations of the results. The Virtual Measurement Systems Program introduces metrology constructs — standard references, uncertainty characterization and propagation, and traceability — into scientific computation and computer-assisted measurement technologies. As with physical measurement systems, development of a virtual metrology infrastructure will result in predictive computing with quantified reliability. In turn, this will enable improved decision making contingent on virtual measurement systems. – Andrew Dienstfrey, Program Manager

Quantifying and Characterizing Errors in 3D Rendering

John Hagedorn
Judith Terrill
Adele Peskin
James Filliben (NIST ITL)

See feature article, page 41.

Uncertainties in Virtual Measurements from Quantum Chemistry Models

Raghu Kacker
Rüdiger Kessel
Karl Irikura (NIST CSTL)
Russell Johnson (NIST CSTL)

By a virtual measurement we mean a prediction along with its associated uncertainty determined from a computational model as an alternative to a physical measurement for the value of a quantity. An important application is prediction of molecular properties from computational quantum chemistry models for spectroscopy, thermo-chemistry, and reaction kinetics. Generally, predictions from computational quantum chemistry models are reported without uncertainties. As the technology improves, the need and importance of reliable uncertainties in virtual measurements is being recognized.

This project is focused on developing and applying methods to quantify the uncertainties associated with virtual measurements from computational quantum chemistry models. The benefits accrue to research and development of chemical processes, materials development, and drug discovery. Predictions from computational quantum chemistry models seldom agree with corresponding high-quality physical measurements. The differences are not random but systematic. Therefore, a common practice is to apply an empirical scaling factor to computational predictions to bring them closer to the true quantity values. The empirical scaling factor carries uncertainty. We have developed a methodology to quantify the uncertainty associated with a scaling factor. This approach is based on the Guide to the Expression of Uncertainty in Measurement, which is an international standard. The uncertainties for scaling factors lead to corresponding uncertainties for the virtual predictions.

In 2005 we reported uncertainties in the scaling factors for *ab initio* vibrational frequencies from 40 models (combinations of theory and basis sets) intended for predicting fundamental frequencies from computed harmonic frequencies. The uncertainties were larger than generally acknowledged. Subsequently, in 2006-2007 we determined a new reference data set of 60 diatomic and 15 polyatomic experimental vibrational zero-point energies that includes estimated uncertainties. Most recently, in 2008 we refined the methodology for calculating uncertainties and reported uncertainties for 32 models intended for predicting vibrational zero-point energies from computed harmonic frequencies. In 2009 we plan to determine uncertainties for virtual measurements from computed an-harmonic vibrational frequencies.

Standard Problems for Micromagnetic Modeling

Michael Donahue
Donald Porter

http://math.nist.gov/oommf/

Advances in magnetic devices such as recording heads, field sensors, magnetic nonvolatile memory (MRAM), and magnetic logic devices are dependent on an understanding of magnetization processes in magnetic materials at the nanometer level. Micromagnetics, a mathematical model used to simulate

magnetic behavior, is needed to interpret measurements at this scale. MCSD is working with industrial and academic partners, as well as with colleagues in the NIST CNST, MSEL, and EEEL, to improve the state-of-the-art in micromagnetic modeling.

Michael Donahue and Donald Porter in MCSD have developed a widely used public domain computer code for doing computational micromagnetics, the Object-Oriented Micromagnetic Modeling Framework (OOMMF). OOMMF serves as an open, well-documented environment in which algorithms can be evaluated on benchmark problems.

OOMMF is part of a larger activity, the Micromagnetic Modeling Activity Group (muMAG), formed to address fundamental issues in micromagnetic modeling through two activities: the development of public domain reference software, and the definition and dissemination of standard problems for testing modeling software. MCSD staff members are involved in development of the standard problem suite as well. There are currently four standard problems in the suite, testing both static and dynamic magnetization properties.

Efforts have recently begun to draft a new standard problem to compare models and techniques for simulating the effects of spin momentum transfer with the support of ITL's Virtual Measurement Systems program. Also in the realm of testing, an extensive regression test suite has been added to OOMMF so that any changes to computed results that arise from ongoing code development are noticed and verified as they happen.

High Precision Hy-CI Variational Calculations on Small Atomic Systems

James S. Sims
Stanley A. Hagstrom (Indiana University)
Maria Belen Ruiz (University of Erlangen, Germany)

Few-electron atomic systems continue to play an important role as a testbed for understanding fundamental physical and chemical theories. Non-relativistic ground state energies and wavefunctions for few-electron atoms are primary inputs for computations of atomic properties such as electron affinities and ionization potentials, as well as for computing perturbative corrections due to relativity and quantum electrodynamics [1]. Furthermore, fundamental quantum theory dictates that, in general, uncertainties in wavefunctions scale as the square root of uncertainties in corresponding energies. Therefore, when using convergence of energies as a metric for convergence of associated wavefunctions, extremely high accuracy in the former must be achieved to guarantee accuracy in the later. The combination of the foundational nature of these non-relativistic ground states and the accuracies required for downstream scientific research position these computations at the level of generalized fundamental physical constants.

The ultimate goal of this work is to produce wave functions and energy levels of few-electron atoms and ions with a precision that goes well beyond what can be achieved experimentally. In the past two decades, there have been breathtaking improvements in computer hardware and innovations in mathematical formulations and algorithms, opening up the real possibility that theoretical successes for the two-electron atomic problem (helium) and the two-electron molecular problem (H_2) can be extended to atoms (and molecules) with more than two electrons. There are formidable bottlenecks to accomplishing this goal, as evidenced by the fact that in going from helium (two electrons) to lithium (three electrons) to beryllium (four electrons), the situation degrades to the point that there are no calculations of the ground or excited states with an error of less than present day experimental uncertainty. This project attempts to remove the bottlenecks by combining mathematical formulations and algorithms revolving around a mathematical technique pioneered by Sims and Hagstrom many years ago (the Hylleraas-Configuration Interaction (Hy-CI) technique [2]), and the use of extended precision, parallel computing, and efficient techniques for integral calculation and storage.

Previous accomplishments of this project have been at the two-electron level. NIST computations of the non-relativistic ground state energies and wave functions for members of the helium (He) isoelectronic sequence using the Hy-CI technique are among the most accurate in the world, achieving an accuracy greater than one part in 10^{20} [3]. Computations on the ground state of dihydrogen (H_2) represented the highest level of accuracy ever achieved (10^{-12} hartree) in molecular quantum computations (except for trivial one-electron cases) [4].

Work this year has moved to three-electron atomic states, as we have endeavored to develop a three-electron atomic code capable of achieving high precision results for the ground state of atomic lithium (Li) as well as several of its excited states. By focusing on Li we are testing our three electron integral routines [5, 6] and we are also testing Hy-CI against the large number of very accurate calculations that exist for this system.

The major theoretical issue addressed in these calculations is whether linked terms like $r_{12}^m r_{13}^n$ with m and n odd are really necessary in either Hy-CI or in the formally equivalent competitive technique, Hylleraas (Hy), or is it sufficient to use only a single

odd r_{ij} power in linked r_{ij} products. Preliminary results of our Li work show that these $r_{12}^{m} r_{13}^{n}$ n13 odd-odd terms are unimportant at the nanohartree level of accuracy, leading us to believe that we can achieve nanohartree accuracy for all the Li states we compute, resulting in predictions for Li that are more accurate than have been done before, and more accurate than modern experiments.

The combination of this year's accomplishments and previous efforts culminating in an extended-precision parallel dense matrix generalized eigenvalue solver, as well as our work on Hy-CI four-electron integrals [7]), should lead to a benchmark calculation on the four electron Beryllium atom (the prototype for high precision atomic calculations for atoms with four or more electrons), resulting in theoretical predictions as accurate (or better) than modern experiments.

[1] F.W. King, "High-Precision Calculations for the Ground and Excited Excited States of the Lithium Atom," *Advances in Atomic, Molecular, and Optical Physics* **40** (1999), pp. 57-105.

[2] J.S Sims and S.A. Hagstrom, "Combined Configuration-Interaction-Hylleraas-Type Wave-Function Study of the Ground State of the Beryllium Atom," *Physical Review A* **4** (1971), pp. 908-916.

[3] J.S. Sims and S.A. Hagstrom, "High Precision Hy-CI Variational Calculations for the Ground State of Neutral Helium and Heliumlike Ions," *International Journal of Quantum Chemistry* **90** (6) (2002), pp. 1600-1609.

[4] J.S. Sims and S.A. Hagstrom, "High Precision Variational Born-Oppenheimer Energies of the Ground State of the Hydrogen Molecule," *Journal of Chemical Physics* **124** (2006), 094101.

[5] J.S. Sims and S.A. Hagstrom, "Math and Computational Science Issues in High-precision Hy-CI Calculations I. Three-electron Integrals," *Journal of Physics B: Atomic, Molecular, and Optical Physics* **37** (7) (2004), pp. 1519-1540.

[6] J.S. Sims and S.A. Hagstrom, "Math and Computational Science Issues in High-precision Hy-CI Calculations II. Kinetic Energy and Electron-nucleus Interaction Integrals," *Journal of Physics B: Atomic,Molecular, and Optical Physics* **40** (2007), pp. 1575-1587.

[7] J.S. Sims and S.A. Hagstrom, "Math and Computational Science Issues in High-precision Hy-CI Calculations III. Four-electron Integrals," *Journal of Physics B: Atomic, Molecular, and Optical Physics*, to appear.

A Metrological Approach to Uncertainty Estimation of Finite Element Simulations

Jeffrey T. Fong
James Filliben (NIST ITL)
Alan Heckert (NIST ITL)
Roland DeWit (NIST MSEL)

The finite element method (FEM) has been a widely used tool for simulating complex physical systems for more than 50 years. What has been lacking in a certain class of FEM-based simulations is an easy-to-use but rigorous method of verification. Of interest are estimations of uncertainties in the simulation results due to modeling (physics), discretization (mesh geometry) and solution approximation algorithm (computational software). The problem of estimating such the FEM simulation uncertainties is fraught with problem-dependent issues, which is one reason why endeavors in this field have not produced decisive results.

Motivated by a recent NIST-authored paper [1], we developed a systematic uncertainty estimation methodology for FEM-based simulations by adopting a 4-part metrological approach as described below.

Part 1. We consider a virtual measurement laboratory where each FEM simulation is considered an experiment. We then introduce a classical method in statistics known as the orthogonal factorial design of experiments (DOE), where a properly "balanced" set of experiments, real or virtual, can be conducted to estimate uncertainties due to known or guessed variability in the physical model, and discrete changes in FEM mesh design and the solver algorithm. Here, each variability or change is known as a *process factor*, or simply, a factor, and Part 1 of our methodology is a listing of k such factors.

Part 2. With the listing of *factors* in Part 1, we begin DOE for FEM-based simulations with choosing a specific objective of our virtual experiment, e.g., the peak stress or strain at a specific location of a specimen, or the first resonance frequency of a vibrating cantilever beam. It is worth noting that DOE is an efficient procedure for planning experiments so that the data obtained can be analyzed to yield valid and objective conclusions. The statistical theory underlying DOE begins with the concept of the so-called *process model.* A process model of the "black box" type is formulated with several discrete or continuous input factors that can be controlled, and one or more measured output *responses*. The output responses are assumed continuous. Real or virtual experimental data are used to derive an approximate empirical model linking the outputs and inputs such that an

estimate of the uncertainty of the output response (the experimental objective) due to prescribed variability or change in each of the k input factors can thus be estimated. Part 2 of our methodology is to introduce a specific virtual experimental design known as the 2-level full or fractional orthogonal factorial design, and the associated choice of n runs within the constraint of an orthogonal design. Part 2 is completed when all FEM simulations associated with n runs ($n > k$) are executed, and the specific objective in each simulation as chosen by the user is identified. *Note*: A full factorial orthogonal design with k factors requires n runs, where $n = 2^k$. A fractional factorial orthogonal design could get away with less than 2^k runs, i.e., $n = 2^{k-i}$, where $i = 1, 2, 3$, etc., as long as the number of runs is not less than k. For example, if $k = 5$, a full factorial design requires 32 ($= 2^5$) runs, but a fractional factorial design may need only 16 ($= 2^{5-1}$), or, 8 ($= 2^{5-2}$) runs, but not 4 ($= 2^{5-3}$), because 4 is less than k ($= 5$). Depending on the interaction properties of the input factors and the confounding nature of the analysis results, an 8-run fractional factorial design ($n = 8$) can, in principle, take care of problems with up to 7 factors (here $k = 7 < n$), for which a full design would have required 2^7, or, 128 runs. For the number of factors higher than 7, a 16-run fractional factorial design is needed and is good for problems with up to 15 factors, for which a full orthogonal factorial design would require 2^{15}, or, 32,768 runs.

Part 3. Using a software package named DATAPLOT [2] and applying a ten-step uncertainty analysis algorithm also due to Filliben and Heckert (2002), which was specifically designed for 2-level, k-factor, n-run, fractional orthogonal factorial designs, we complete Part 3 of this methodology by analyzing the set of n FEM-based simulation responses (objective) and obtaining the results in the form of ten plots:

Plot 1. Ordered data plot.
Plot 2. DOE scatter plot.
Plot 3. DOE mean plot.
Plot 4. Interaction effects matrix plot.
Plot 5. Block plot.
Plot 6. DOE Youden plot.
Plot 7. Effects plot.
Plot 8. Half-normal probability plot.
Plot 9. Cumulative residual std deviation plot.
Plot 10. DOE contour plot of 2 dominant factors.

Part 4. The final plot allows us to reduce a very complex multilinear model of k factors into a simpler empirical (approximate) model of two factors that are known to be "dominant." Part 4 of our methodology is to estimate the FEM-based simulation uncertainty by calculating the prediction confidence intervals of the chosen objective of the virtual experiment using standard statistical data analysis techniques.

An expository paper of this methodology was presented at the 2008 Annual Meeting of the American Society for Engineering Education [3]. Two examples were included: (1) The free vibration of an isotropic elastic cantilever beam with a known theoretical solution, and (2) the calculation of the first resonance frequency of the elastic bending of a single-crystal silicon cantilever beam without known solutions. In each example, the FEM-simulated result is accompanied by a 95% confidence interval.

This method of DOE has a built-in mechanism for the user to quantify the problem-dependent factors and issues, thus advantageously employing the knowledge and skill of the scientist or engineer and making the methodology "global" in nature.

In summary, the methodology, along with a future project to develop a "plug-in" for automatic FEM uncertainty analysis, will allow an FEM user to translate his or her knowledge of the variability of problem-dependent factors into error bounds for the final FEM simulation results via a semi-turnkey "black box" which is systematic, computationally-efficient, and statistically-rigorous

[1] D. C. Hurley, et al., "Atomic force acoustic microscopy methods to determine thin-film elastic properties," *Journal of Applied Physics* **94** (2003), pp. 2347-2354.

[2] J. J. Filliben and N. A. Heckert, 2002, DATAPLOT: A Statistical Data Analysis Software System, public domain software downloadable from http://www.itl.nist.gov/div898/software/dataplot.html.

[3] J. T. Fong, J. J. Filliben, N. A. Heckert and R. deWit, "Design of Experiments Approach to Verification and Uncertainty Estimation of Simulations based on Finite Element Method," in *Procedings of the Annual ASEE Conference,* June 22-25, 2008, Pittsburgh, PA, Paper AC2008-2725.

Pervasive Information Technology

Pervasive information technology is the trend towards increasingly ubiquitous connected computing sensors, devices, and networks that monitor and respond transparently to human needs. The Pervasive Information Technologies Program facilitates the creation of standards for sensor communication, networking interoperability, and sensor information security. The Program enables the use of pervasive information technologies to enhance personal and professional productivity and quality of life. – Kamran Sayrafian-Pour, Program Manager

Visualization of Radio Frequency Propagation for Implant Communication Systems

John Hagedorn
Wenbin Yang
Kamran Sayrafian (NIST ITL)
Judith Terrill
Ansoft Corporation
National Inst. of Information and Comm. Tech. (Japan)

Future medical implants will have the ability to communicate data to external monitoring systems. To enable such data transfer, short range, wireless communication to and from antennas implanted inside the human body is required. Currently no standards for such devices exist. Such a standard should use frequency bands approved by national medical and/or regulatory authorities.

A necessary input for such standards is the knowledge of RF propagation from an implant device. There are currently no models or data available to characterize this propagation. As physical experimentation is nearly impossible, we are using simulation models to study this problem. The impact would be an improvement in the quality of care, wellness, protection, safety, and vendor interoperability among wearable and implantable devices. Customers include standards organizations (e.g., IEEE 802), public safety users, first responders, healthcare providers, military personnel, and the public.

This year we built real time 3D visualizations that run on the desktop and in the immersive visualization laboratory to observe the results of simulations of RF propagation from implant devices. These visualizations have enabled researchers to investigate the results of simulations for both scalar and vector fields with a variety of visual representations. Interactive tools have been implemented to enhance researchers' understanding of the data; see the Figure 44.

Additionally, we have streamlined data processing between HFSS (the software that models RF propagation), HBM (the human body model) and our 3D immersive platform. We obtained estimates of path loss versus distance between an implant and an on-body node (for multiple directions and/or antenna orientations), and path loss versus distance between two implants. We made presentations to the IEEE 802.15.6 Task Group regarding the application of the 3D immersive system for the channel modeling activity.

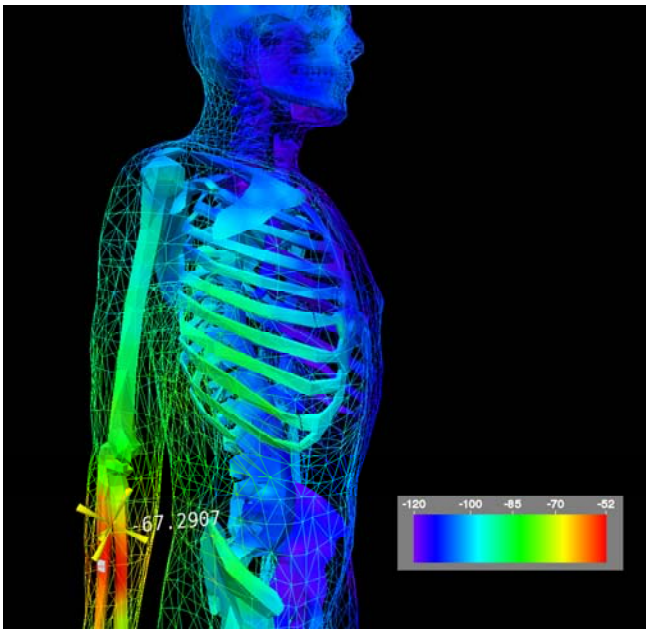

Figure 44. *Still image from an interactive visualization of a simulation of a transmitting antenna implanted in the right forearm. The antenna is visible as a white rectangle in the forearm and the color variations show RSS (received signal strength) at the surface of the skeleton and on the surface of skin (which is depicted as wireframe). Other organs and structures within the body can be turned on and off interactively. The cross-shaped object that appears near the antenna is a 3D cursor that functions as an interactive probe of the data. The legend in the lower right shows the correspondence between color and RSS.*

In the future we plan to investigate appropriate interactive 3D data visualization for various quantities (e.g., RSS, E-field, H-field, heat, etc.) and to collaborate to obtain the appropriate antenna design for various on-body or implant applications. We want to characterize path loss from an implant to another implant or on-

body node and to investigate the possibility of driving a statistical channel model for MICS (Medical Implant Communication System) operation. We will contribute the results to the IEEE 802.15.6 Task Group on body area networking and the corresponding channel modeling document. We will also study coexistence and interoperability issues with other wireless technologies.

Mobile Sensor Networks

Vladimir Marbukh
Kamran Sayrafian-Pour (NIST ITL)

Mobile sensor networks are envisioned for detecting and tracking potential targets and events for civilian as well as military purposes. The location of sensors in a mobile sensor network affects not only network detecting and tracking ability, but also its ability to communicate the relevant information to the intended recipients. Communication reliability can be improved if sensors are capable of self-organizing into a multihop mobile network where some sensors relay the other sensors' information in addition to transmitting their own information. Since the detecting and tracking requirements could potentially compete with the core communication responsibilities of a node, optimal realization of mobile sensor networks requires the ability to balance these competing requirements using local, and typically incomplete, information. Energy conservation needs impose constraints on the sensor network's ability to complete its mission due to its possible impact on the sensor network lifespan.

The goal of this project is to develop strategies for self-organized mobile sensor networks capable of adjusting to target movement and other environmental changes under energy constraints. Achieving this goal will require solving complex optimization problems, including balancing numerous inherent trade-offs in a distributed fashion. As a first step in achieving this goal, we have proposed a tractable extension to the Network Utility Maximization (NUM) framework aimed to address some of these algorithmic challenges [1]. The conventional NUM framework assumes elastic users, sources, or applications capable of adjusting their bandwidth requirements in response to network congestion. Cross-layer NUM optimization includes flow control, routing, scheduling and power control, and is achieved by a decentralized, adaptive, closed-loop algorithm with feedback signals representing resource congestion prices.

In [1] we propose Location-aware Network Utility Maximization (L-NUM), which assumes that the aggregate sensor utility is a function not only of the sensor information rates but also sensor physical locations. L-NUM operates in two time scales. The "fast" time scale L-NUM is identical to conventional NUM, and yields the optimal cross-layer network design for given sensor locations. In the "slow" time scale mobile sensors adjust their locations in attempt to further increase the aggregate utility. In practice, reaching the optimal locations by mobile sensors may be infeasible due to inaccessible terrain and/or limitations on the mobile sensor energy supply. L-NUM proposes to account for these factors through utility production, which is a function of both sensor locations and speeds. The sensor speeds are selected to maximize utility production, given the current sensor locations.

Current and future efforts will concentrate on two major issues: quantification of location-aware sensor utility and distributed implementation of the aggregate utility maximization. From the algorithmic perspective, the main difficulty in resolving these issues is possible redundancy of the information obtained by different sensors. Indeed, while it is natural to assume that the utility of sensor information increases as the sensor approaches the target, proximity to other sensors may deem this information redundant. Marginal utility of sensor information depends on the details of sensor information exchange and processing, and is difficult to quantify. The ability to account for sensor information redundancy while keeping the information exchange sufficiently low is critical for mobile sensor network self-organization.

[1] Vladimir Marbukh and Kamran Sayrafian-Pour, "A Framework for Joint Cross-Layer and Node Location Optimization in Mobile Sensor Networks," in *Proceedings of 7th International Conference on Ad-hoc Networks and Wireless* (ADHOC-NOW'08), Sophia-Anipolis, September 2008.

Trustworthy Information Systems

A trustworthy system is one that performs as intended for a specific purpose, when needed, with operational resiliency and without unwanted side effects, behaviors, or exploitable vulnerabilities. The Trustworthy Information Systems (TIS) Program conducts research, development, and testing to improve the ability to model, build, test, measure, and assess information system trustworthiness through the development and application of new and innovative technologies, models, measurement methods, and tools. The aim of the TIS program is to reduce the risk and uncertainty associated with information systems that must be trusted, and to improve the ability to build in and evaluate trustworthiness in information applications and systems. – Tom Rhodes, Program Manager

Automated Combinatorial Testing for Software Systems

Raghu Kacker
Rüdiger Kessel
Stephen Langer
James Lawrence
Yu (Jeff) Lei
Andrew Reid
Rick Kuhn (NIST ITL)
Vincent Hu (NIST ITL)
Tom Rhodes (NIST ITL)
Renee Bryce (Utah State University)
Sagar Chaki (Carnegie Mellon University)
Michael Forbes (MIT)
Arie Gurfinkel (Carnegie Mellon University)
Sreedevi Sampath (UMBC)
Tao Xie (North Carolina State University)

Combinatorial testing is an approach to detect combinatorial interaction faults in complex software systems. A combinatorial interaction fault is a defect that may cause the system to fail when particular values of certain parameters are combined with particular values of certain other parameters in the system. Interaction faults remain dormant until the unfortunate combination of parameter values is encountered. Many critical system failures in recent times have been traced to interaction faults in the underlying software. Thus, combinatorial testing, which is designed to detect such interaction faults, is a fundamental technology for ensuring trustworthy software in business, industrial, medical, scientific, and transport systems.

It is often difficult or impossible to characterize all possible combinatorial interactions, let alone test them. Combinatorial testing is based on the observation that not all parameters are involved in every interaction fault, and that most interaction faults involve only a few parameters. A NIST study of actual faults in software showed that most faults involved a single parameter, a smaller proportion of faults resulted from interactions between values of two parameters, and progressively fewer interaction faults involved 3, 4, 5, and 6 parameters. So far, a fault involving more than six parameters has not been seen.

Most combinatorial testing has so far been limited to pair-wise (2-way) testing, in which all interactions involving two parameters are incorporated. Development of efficient test methods to detect higher order interaction faults has so far not been available. The goal of this project is to advance the technology from pair-wise to higher order (strength) testing and to demonstrate successful applications.

In 2008 we issued version 1 of a publicly available research tool, FireEye, to generate test suites for high strength testing. We have distributed over 115 copies of FireEye to US government agencies and corporations, including Microsoft, Lockheed Martin, Cisco, SAP, AT&T, Intel, and SAIC. We published four papers in refereed publications.

In 2009 we plan to continue to enhance FireEye and to investigate applications in computer security, GUI based applications, web-applications, and for optimization and the testing of simulations of complex systems.

[1] R. Kuhn, Y. Lei, and R. Kacker, "Automated Combinatorial Test Methods – Beyond Pairwise Testing" *CrossTalk*, June 2008, pp. 22-26.

[2] R. Kuhn, Y. Lei, and R. Kacker, "Practical Combinatorial Testing: Beyond Pairwise," *IT Professional*, May/June 2008, pp. 19-23.

[3] Y. Lei, R. Carver, R. Kacker, and D. Kung, "A Combinatorial Testing Strategy for Concurrent Programs," *Software Testing, Verification and Reliability* **17** (4) (December 2007), pp. 207-225.

[4] Y. Lei, R. Kacker, D.R. Kuhn, V. Okun, and J. Lawrence, "IPOG/IPOG-D: Efficient Test Generation for Multi-way Combinatorial Testing," *Software Testing, Verification and Reliability* **18** (2008), pp. 125-148.

[5] Y. Lei, R. Kacker, D.R. Kuhn, V. Okun, and J. Lawrence, "IPOG: A General Strategy for T-way Software Testing," in *Proceedings of the 14th Annual IEEE International Conference and Workshop on the Engineering*

of Computer-Based Systems (ECBS'07), March 26-29, 2007, Tucson AZ.

[6] W. Wang, S. Sampath, Y. Lei, and R. Kacker, "An Interaction-based Test Sequence Generation Approach for Testing Web-applications," in *Proceedings of 11-th IEEE High Assurance Systems Engineering Symposium*, Nanjing, China, December 3-5, 2008.

[7] R. Kuhn, Y. Lei, and R. Kacker, "Advances in Combinatorial Testing," *IEEE Spectrum,* submitted.

[8] J. Lawrence, R. Kacker, Y. Lei, D.R. Kuhn, and M. Forbes, "Two-valued Covering Arrays," *Designs, Codes and Cryptography*, submitted.

[9] J. Lawrence, R. Kacker, Y. Lei, D.R. Kuhn, and M. Forbes, "Binary Covering Arrays," *Designs, Codes and Cryptography*, submitted.

Funding: *2007 NIST ACI Initiative on Cyber Security*

Division Programs

Digital Library of Mathematical Functions

The special functions of applied mathematics are fundamental tools enabling modeling and analysis in all areas of science and engineering. To make effective use of such functions, practitioners must have ready access to a reliable source of information on their properties. The goal of this work is the development and dissemination of definitive reference data on the special functions of applied mathematics. The DLMF will be a freely available interactive and richly linked online resource. – Daniel Lozier, Program Coordinator

Digital Library of Mathematical Functions

Frank W. J. Olver
Daniel Lozier
Ronald Boisvert
Peter Ketcham
Marjorie McClain
Bruce Miller
Bonita Saunders
Abdou Youssef
Charles Clark (NIST PL)
Qiming Wang (NIST ITL)
Brian Antonishek (NIST ITL)
Richard Askey (Univ. of Wisconsin, Madison)
Michael Berry (University of Bristol, UK)
Leonard Maximon (George Washington Univ.)
Morris Newman (U. California, Santa Barbara)
Ingram Olkin (Stanford University)
Peter Paule (J. Kepler University, Austria)
William Reinhardt (University of Washington)
Nico Temme (CWI, Amsterdam)

38 authors and validators at universities and research institutions

http://dlmf.nist.gov/

Mathematics and science go hand in hand. Historically, purely mathematical developments find application in the description of scientific phenomena, seemingly by serendipity, and, in the other direction, cutting-edge science is a major motivator and driver of mathematical research. Often the mathematical objects at the intersection of math and science are mathematical functions. Thus there has always been an interest in, and a need for, accurate and comprehensive compilations of the many intricate properties of mathematical functions widely used in applications.

Traditionally, massive published handbooks, such as the 1000-page *NBS Handbook of Mathematical Functions*, created in 1964 under the editorship of Milton Abramowitz and Irene Stegun, have served this purpose. The NBS handbook was an immediate success, especially with physical scientists and engineers, and remains a standard reference today. In the current project NIST aims to supersede the NBS handbook by creating a completely new compilation of the scientifically useful properties of functions, and also to create a new means of disseminating the compiled information. The new compilation is the Digital Library of Mathematical Functions (DLMF). The content of the DLMF is nearly complete. This content includes:

- Methodology chapters covering parts of algebra, real and complex analysis, asymptotics, and numerical analysis related to the analysis and computation of mathematical functions.
- Definitions, notations, and important properties of many classes of elementary and special functions.
- Graphics, including many line graphs and surface plots.
- Methods of computation and approximation, including links to available software.
- Large bibliography and index.

The 36 chapters planned for the initial release have been written, validated, and approved for publication by their authors and validators. A demanding editorial process, resulting from repeated written communications with the Mathematics Editor, was followed to reach this stage. The result is a well-integrated presentation of very technical material that is uniform across the whole content, and not just a collection of individually written chapters gathered together hurriedly. What remains is completion of front and back matter, such as introductions and indexes, and resolution of small issues that have been raised by authors and/or validators that must be reconciled.

Similarly, the new means of dissemination via the DLMF Web Site is in an advanced state. A five-chapter preview was released for public review in June 2008. The preview shows how the DLMF has succeeded in using the latest Web technology to:

- Deliver text and math in MathML to Web browsers that support these technologies, or text in HTML and math in GIF images to less capable Web browsers, and to detect such browser capabilities automatically at the Web server.
- Provide zoomable and rotatable color visualizations of surfaces of real or complex functions of two real variables to browsers that support VRML.
- Support search based on mathematical queries, as well as text queries; for example, the query "$\sin^2 + \cos^2$" locates many identities that contain the sum of squares of sines and cosines.

Figure 45 shows the home page of the preview release. Chapters 2, 5, 9, 27 and 34 are fully functional and contain full mathematical content.

The complete DLMF is scheduled to appear in mid 2009. It will include Chapters 1-36; Chapters 37 and 38 have been deferred to a later release.

The Web-based DLMF will be accompanied by a commercially published print edition, the *NIST Handbook of Mathematical Functions*, or *HMF*. A solicitation for a publisher was announced on the DLMF Web Site on November 7; see Figure 45. Four proposals were received before the cutoff date of December 12. A publisher will be selected in January 2008. The HMF will appear at about the same time as the DLMF is released.

Several future developments related to the DLMF project are planned or under consideration. One is to further the semantic capabilities of the XML tools and techniques developed for the DLMF. Another is to develop tools and techniques for generating validated numerical values and graphs on demand. A third is to extend the power and applicability of the DLMF search tools and techniques. Also, resumed development of Chapters 37 and 38 is planned.

[1] M. E. Altamimi and A. Youssef, "A Math Query Language with an Expanded Set of Wildcards," *Mathematics in Computer Science* **2** (2) (published online December 10, 2008).

[2] D. W. Lozier, "Asymptotics in the NIST Digital Library of Mathematical Functions," AMS-SIAM Special Session on Asymptotic Methods in Analysis with Applications, Joint Mathematics Meetings, San Diego, CA, January 6, 2008.

[3] D. W. Lozier, "Numerical Analysis in the NIST Digital Library of Mathematical Functions," International Conference of Numerical Analysis and Applied Mathematics, Kos, Greece, September 18, 2008.

[4] A. Youssef, "Relevance Ranking and Hit Description in Math Search," *Mathematics in Computer Science* **2** (2) (published online December 2008).

Funding: *National Science Foundation, NIST Systems Integration for Manufacturing Applications (SIMA) Program, NIST Standard Reference Data Program.*

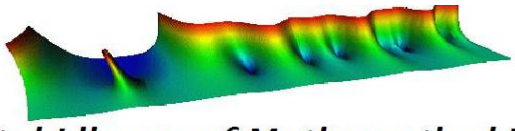

NIST Digital Library of Mathematical Functions

Preview Release

This is a preview of the Digital Library of Mathematical Functions. Not all of the content is present, nor has all of the intended functionality been implemented. Given those limitations, we welcome feedback about the current status, so that we can improve the final result. Note that citations to specific items (such as figures or equations) should be deferred until final release, as it is still possible that these may change. Please consult About the Project for more information.

Project News

2008-11-07 Call for proposals to publish a printed version of the DLMF (Updated 2008-11-21)

1 Algebraic and Analytic Methods
2 Asymptotic Approximations
3 Numerical Methods
4 Elementary Functions
5 Gamma Function
6 Exponential, Logarithmic, Sine and Cosine Integrals
7 Error Functions, Dawson's Integral, and Fresnel Integrals
8 Incomplete Gamma and Related Functions
9 Airy & Related Functions
10 Bessel Functions
11 Struve and Related Functions
12 Parabolic Cylinder Functions
13 Confluent Hypergeometric Functions
14 Legendre Functions
15 Hypergeometric Function
16 Generalized Hypergeometric Functions
17 q-Hypergeometric Functions
18 Orthogonal Polynomials
19 Elliptic Integrals
20 Theta Functions
21 Multidimensional Theta Functions
22 Jacobian Elliptic Functions
23 Weierstrass Elliptic Functions
24 Bernoulli and Euler Polynomials
25 Zeta and Related Functions
26 Combinatorial Analysis
27 Functions of Number Theory
28 Mathieu Functions and Hill's Equation
29 Lamé Functions
30 Spheroidal Wave Functions
31 Heun Functions
32 Painlevé Transcendents
33 Coulomb Functions
34 3j, 6j, 9j Symbols
35 Functions of Matrix Argument
36 Integrals with Coalescing Saddles
37 Computer Algebra
38 Statistical Methods & Distributions
Bibliography
Index
Notations
Software

Figure 45. *Home page of the preview release of the Digital Library of Mathematical Functions.*

Visualization of Complex Function Data

Bonita Saunders
Qiming Wang (NIST ITL)
Brian Antonishek (NIST ITL)
Sandy Ressler (NIST ITL)
Daniel Lozier
Frank Olver

Clear and informative 3D visualizations can substantially enhance a researcher's understanding of high level mathematical functions that arise as solutions to problems in the mathematical and physical sciences. However, the poor quality of 3D graphs generated by many commercial packages suggests the difficulty of providing accurate visualizations. Consequently, we have devoted a considerable effort to developing state-of-the-art interactive 3D visualizations for the NIST Digital Library of Mathematical Functions (DLMF). Using techniques from numerical grid generation and web-based visualization technologies such as the Virtual Reality Modeling Language (VRML) and X3D, we have completed more than 200 3D visualizations for the digital library. Samples can be viewed in the gamma and Airy function chapters available in the DLMF Preview Edition[11].

With all visualizations essentially completed for the first release of the DLMF, we are now focusing on the following: improving the presentation quality of the visualizations on the website, validating plot data, and making the visualizations accessible on multiple platforms. Presentation quality can be enhanced by improving the underlying grid used in the function computation. We have used numerical grid generation techniques to create computational grids that successfully capture key surface features such as zeros, poles, and branch cuts, but we can improve the color map of surfaces with steep gradients by using grids with cell spacing adapted to function gradient and curvature data. We have made significant progress in this area, creating grids that adapt to simply defined curves; however, the specialized nature of some high level functions may mean that accessing and linking the codes needed to compute gradient and curvature data is a nontrivial task.

Validation of 2D data files is almost complete and validation of 3D files has been started. To validate data accuracy we compute the function values using codes from at least two independent sources, including widely available computer algebra packages such as Mathematica and Maple, routines from repositories such as the NIST Guide to Available Mathematical

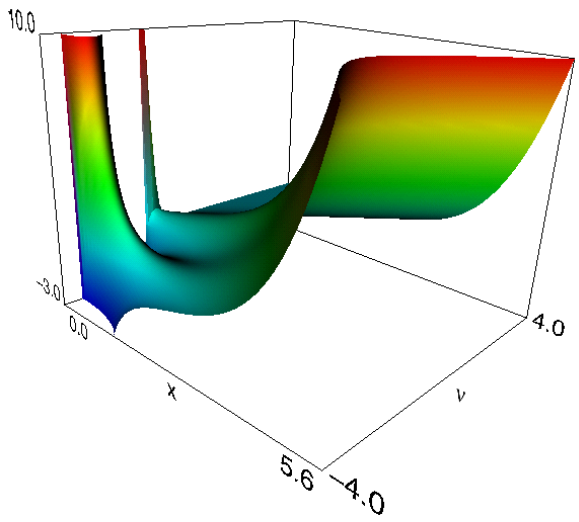

Figure 46. *Plot of Struve function $L_v(x)$.*

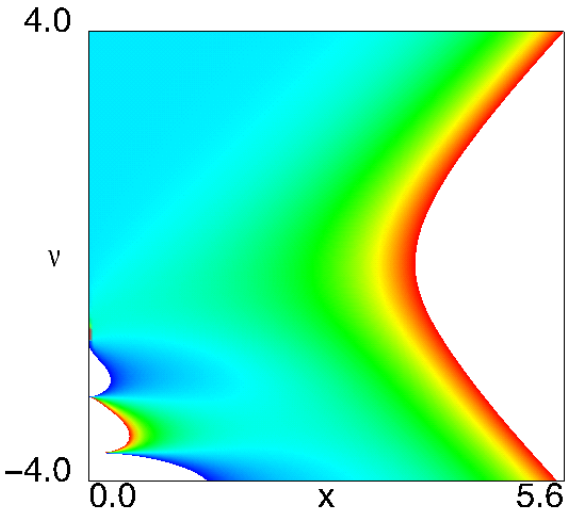

Figure 47. *Density plot of Struve function.*

Software[12] (GAMS), and personal Fortran and C codes provided by the chapter authors themselves. To date, most conflicts have not resulted in a change in the visual appearance of a plot, but some inconsistencies found between Mathematica and Maple data are being examined more carefully.

One of the more problematic issues has been making sure our visualizations are accessible on multiple platforms. The performance of VRML and X3D browsers/plugins can be very inconsistent even if they claim to handle codes that follow the official VRML 97 standard. Our goal is to make sure we can find at least one reliable browser or plugin for users in Windows, Apple Macintosh, and Linux environments. Since it appears that on some platforms most VRML/X3D browsers are stand-alone applications that often have

[11] http://dlmf.nist.gov/

[12] http://gams.nist.gov/

trouble finding external routines and gif files, we have designed a self-contained VRML file which includes imbedded gif images and all routines needed to render the VRML image. Furthermore, we have also developed prototypes using alternative technologies such as embedded 3D and video capture to make the visualizations more accessible to users on platforms where no suitable VRML/X3D plugin is currently available.

As the visualization work for the current phase of the NIST DLMF winds down, we will concentrate more on extending our work to on-demand generation of mathematical graphs and tables. This may be a stand-alone capability or ultimately be incorporated as part of the DLMF. We are searching the web for sites that may offer some of these capabilities and looking at the possibility of extending work on the live generation of tables started by Marjorie McClain, Joyce Conlon and former NIST postdoc and faculty appointee Bruce Fabijonas.

[1] B. Saunders and Q. Wang, "Tensor Product B-Spline Mesh Generation for Accurate Surface Visualizations in the NIST Digital Library of Mathematical Functions," in *Proceedings of the Seventh International Conference on Mathematical Methods for Curves and Surfaces*, Tonsberg, Norway, June 26 - July 1, 2008, to appear.

[2] Q. Wang , B. Saunders, and S. Ressler, "Dissemination of 3D Visualizations of Complex Function Data for the NIST Digital Library of Mathematical Functions," *CODATA Data Science Journal* **6** (2007) (Supplement on the *Proceedings of the 20th International CODATA Conference*, Beijing, China, October 2006*),* pp. S146, S154.

[3] B. Saunders and Q. Wang, "Spline Mesh Generation to Effective Visualizations for the NIST Digital Library of Mathematical Functions," *Curve and Surface Design: Avignon 2006, Proceedings of the Sixth International Conference on Curves and Surfaces,* Avignon, France*,* pp. 235-243.

Quantum Information

An emerging new discipline at the intersection of quantum mechanics and computer science, quantum information science (QIS) is likely to revolutionize science and technology in much the same way that lasers, transistors, electronics, and computing did in the 20th century. At the very least, QIS has the potential to provide phenomenal increases in information storage and processing speed and communication channels with extremely high levels of security guaranteed by the laws of physics. However, the needed manipulation and control of quantum states remains a grand challenge. We are developing theoretical underpinning to support the multi-laboratory experimental program in quantum information at NIST. This work is aimed at the development of a measurement science for quantum information processing systems. — Ronald Boisvert, Program Coordinator

Ion Trap Quantum Computing Benchmarks

Emanuel Knill
Dietfrich Leibfried (NIST PL)
David Wineland (NIST PL)

See feature article, page 23.

Funding: *2007 NIST ACI Initiative in Quantum Information Science*

Quantum Computing Theory

Bryan Eastin
Scott Glancy
Emanuel Knill
Adam Meier
Yanbao Zhang
R. Somma (Los Alamos National Laboratories)
H. Barnum (Los Alamos National Laboratories)
S. Boixo (California Institute of Technology)
G. Ortiz (University of Indiana)

The implementation of quantum computers is based on a substantial body of theoretical work that shows the utility of quantum algorithms and provides techniques for the protection of quantum devices from inevitable noise. Our contributions to quantum computing theory include work on better understanding the limitations as well as the power of quantum computers.

Most quantum algorithms are based on applications of a small number of powerful quantum computing strategies. One such strategy, adiabatic quantum computing, exploits the property that states tend to follow eigenstates of slowly varying Hamiltonians. A related strategy called "quantum walks" can be used in optimization problems or for sampling from thermal probability distributions. We contributed to work that shows how to use adiabatic quantum computing and quantum walks to improve the convergence of classical simulated annealing algorithms. For this purpose we simplified adiabatic quantum computing by using randomization of time when evolving the slowly varying Hamiltonians. We are currently working on improving the complexity of quantum circuit emulations of adiabatic quantum computing.

The applications of quantum information to communication include quantum cryptography and quantum secret sharing. Such applications make possible protocols that cannot be realized using only classical information processing. One reason for the power of quantum information in this setting is the presence of quantum non-locality, which consists of correlations with no classical analogue. Experiments that verify the presence of quantum non-locality have been performed, but are still subject to "loopholes," that is, reasonable but unverifiable assumptions are required to ensure that there is no classical explanation for the data. Removal of the loopholes requires better detectors and/or better choices of experimental configurations. We are using a computational framework developed over the last two years to determine conditions on detector efficiency and on other sources of experimental noise that are required to eliminate the most prominent loopholes. See also "Optical Quantum Metrology and Quantum Computing," page 80. So far we have confirmed the conditions for unbalanced Bell states, and compared the requirements for number-resolving detectors to those for photon detectors when entanglement is obtained by combining two independent photons on a polarizing beam splitter. Although photon number-resolving detectors outperform photon detectors, the difference is not as large as we hoped.

In previous years we developed and simulated fault-tolerant quantum computing architectures that can tolerate surprisingly high error rates per quantum gate. We are in the process of implementing new simulation algorithms for fault-tolerant architectures. The goal of our effort is to greatly improve the efficiency of the simulations, and to enable simulation of architectures based on quantum systems with more than two dimensions. Previously, we implemented C++ libraries for the basic operations of such quantum systems and for

data structures needed to keep track of the necessary quantum gates and their errors. We are refining the general formalism needed for simulating architectures that use quantum codes based on general Galois fields. An important feature of the formalism and its implementation is that it can be used not only for simulating architectures, but also for the decoding algorithms used in feedforward from physical measurements in implementations. This requires keeping track of the relationship between abstract error syndromes and eigenvalues of the observables used in the physical measurements. For architecture simulations, only the error syndromes need to be tracked. Future research is aimed at simulation of universal gate sets and non-stochastic error models.

[1] R. Somma, S. Boixo, H. Barnum, and E. Knill, "Quantum Simulations of Classical Annealing Processes," *Physical Review Letters* **101** (2008), 130504, 4 pages.

[2] S. Boixo, E. Knill, and R. D. Somma, "Quantum Computing through Decoherence," in preparation.

Funding: *2007 NIST ACI Initiative in Quantum Information Science*

Optical Quantum Metrology and Quantum Computing

Scott Glancy
Emanuel Knill
Yanbao Zhang
Tracy Clement (NIST EEEL)
Shellee Dyer (NIST EEEL)
Jingyun Fan (NIST PL)
Thomas Gerrits (NIST EEEL)
Alan Migdall (NIST PL)
Richard Mirin (NIST EEEL)
Sae Woo Nam (NIST EEEL)
Sergey Polyakov (NIST PL)
Marty Stevens (NIST EEEL)

MCSD is contributing to the development of an experimental research program in optical quantum metrology and quantum computing. This project has developed expertise in the preparation, manipulation, and measurement of exotic quantum states of light, such as entangled states of N photons and Schrödinger Cat states. The entanglement properties of these states can be exploited for high precision interferometry, quantum communication, and quantum computation. All of these technologies require the ability to control and measure very delicate and sensitive quantum states. This project will significantly expand NIST capabilities in quantum optical metrology, enabling us to expand our position as the global leader in measurement and enabling technology as applied to quantum optics.

In the last year we created Schrödinger Cat states by subtracting either one or two photons from squeezed light pulses. The two photon subtraction experiment uses superconducting transition edge sensors developed in EEEL to count photons. Members of the MCSD designed software to reconstruct a quantum state from thousands of individual measurements. The quantum state reconstructed by our software clearly shows the quantum interference pattern characteristic of high fidelity Schrödinger Cat states. We have also been improving methods to calibrate the measurement system and to measure and improve the purity of the squeezed light.

Measurements of Bell's inequalities with entangled states have demonstrated the usefulness of quantum theory and the falsity of a large class of theories, known as local-realistic theories. However loopholes have been present in all previous tests of local-realism. Because of their high efficiency, NIST's photon number resolving detectors may provide a method to finally these loopholes. MCSD members have been helping to design this experiment. We are also developing a general theory that can assign a confidence with which an experiment has ruled out local realism. This is accomplished by comparing the probability that a set of measurements can be predicted using quantum theory to the probability given by the optimal local-realistic theory.

[1] S. Glancy and H. Vasconcelos, "Methods for Producing Optical Coherent State Superpositions," *Journal of the Optical Society of America B* **25** (5) (April 16, 2008), pp. 712-733.

[2] S. Glancy, "Diagnosis of Pulsed Squeezing in Multiple Temporal Modes," Quantum Estimation in Theory and Practice Workshop, Perimeter Institute, Waterloo, Canada, September 25-30, 2008.

[3] T. Gerrits, T. Clement, S. Glancy, S. W. Nam, R. Mirin, E. Knill, "Generation of Optical Cat States by squeezed photon subtraction," Southwest Quantum Information Technology Network Workshop, Santa Fe, New Mexico, February 14-17, 2008.

[4] S. Glancy, "Calibration for Slightly Unbalanced Homodyne Detection," Southwest Quantum Information Technology Network Workshop, Santa Fe, New Mexico, February 14-17, 2008.

Funding: *NIST Innovations in Measurement Science Program*

Fault-tolerant Quantum Computation

Bryan Eastin
Emanuel Knill

Advances in the theory of fault-tolerant quantum computation are important for the development of quantum computers as the hardware employed in the first generation of such devices is likely to be highly prone to error. The theory of fault tolerance addresses the problem of implementing encoded gates and error correction in a way that precludes a catastrophic failure from resulting from the malfunction of a few isolated components. Primarily, this is achieved by preventing errors from spreading within individual blocks, that is, independent error-correcting units, of a quantum code.

This year we settled a long-standing question in quantum information theory, namely, whether there exist nontrivial quantum codes for which all logical gates can be implemented transversally. Transversal encoded gates never cause subsystems within a code block to be coupled and are, consequently, fault tolerant. Their simplicity and robustness makes transversal encoded gates valuable tools for quantum computation, but, for all known codes, such gates are inadequate for achieving full computational universality. We showed that no non-trivial transversal encoded gate set could be universal.

In the coming year we plan to investigate ways of simplifying quantum error correction and reducing the propagation of errors. Specifically, we hope to extend work by DiVincenzo and Aliferis on using unverified ancillae in quantum error correction and measurement; explore the possibility of constructing quantum codes with the property that encoded gates are individually, but not as a group, transversal; study the efficacy of proposed passive quantum error correcting architectures; and seek new architectures supporting passive correction. The success of any of these research projects would move us closer to the goal of constructing a quantum computer by reducing the overhead and/or precision required.

[1] B. Eastin and E. Knill, "Restrictions on Transversal Encoded Quantum Gate Sets", arXiv:0811.4262.

Funding: *NIST/NRC Postdoctoral Associateship Program*

Mathematics of Metrology

Mathematics plays an important role in the science of metrology. Mathematical models are needed to understand how to design effective measurement systems, and to analyze the results they produce. Mathematical techniques are used to develop and analyze idealized models of physical phenomena to be measured, and mathematical algorithms are necessary to find optimal system parameters. Finally, mathematical and statistical techniques are needed to transform the resulting data into useful information. The goal of this work is to develop fundamental mathematical methods and analytical tools necessary for NIST to continue as a world-class metrology institute, and to apply them critical measurement science applications.

Regularization by Residual Periodigrams

Bert W. Rust
Dianne P. O'Leary

See feature article, page 31.

Comparing MSU T2 Time Series with Ground-based Measurements

Bert W. Rust

NOAA's Microwave Sounding Unit's (MSU) T2 time series of monthly average global temperatures anomalies in the middle troposphere is measured from polar orbiting satellites. A question currently of great interest is whether or not these measurements are consistent with those made at the earth's surface, and, in particular, whether both records indicate the same rate of global warming. Fitting straight lines to the period 1987.00 – 2006.75 in both records give slopes that are statistically consistent, but a more remarkable demonstration of the consistency of the two records is obtained by comparing the T2 record with a model for the UK Climactic Research Unit's (CRU) measurements of annual average global temperature anomalies developed recently by Rust and Thijsse [1]. The essentials of this model are shown in Figure 48. The demonstrated linear dependence of temperatures on atmospheric carbon dioxide concentration can be regarded as a "smoking gun" for global warming. The section of the solid curve from 1987 to the end of the record is compared with the T2 time series in Figure 49. Although there was no adjustments of free parameters to make the carbon dioxide driven model fit the T2 data, it nevertheless fits the data better than the least squares straight line. The sum of squared residuals for the regression line was 7.290 while it was 7.228 for the replotted carbon dioxide driven model.

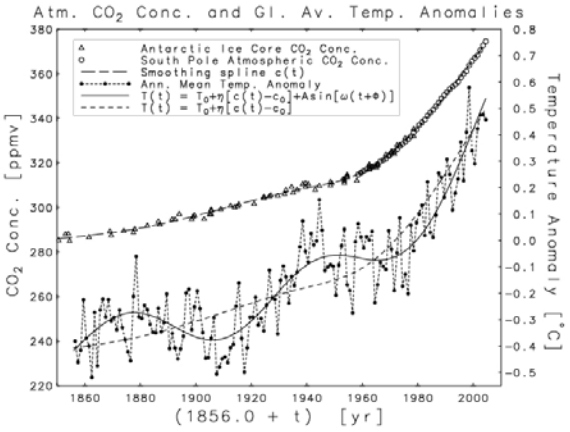

Figure 48. *The upper plot gives an optimal smoothing spline fit to the record of atmospheric carbon dioxide concentrations obtained by combining atmospheric measurements at the South Pole with reconstructions from Antarctic ice cores. The spline c(t) was then used to model the CRU record of annual average global temperature anomalies in the lower plot. The two smooth curves were obtained by fitting a model having a simple linear dependence of the temperatures on the carbon dioxide concentration. The solid curve is the fit of the full model, which also includes an ~70 year sinusoidal cycle, whose existence was first discovered in 1994, and the dashed curve is just the baseline part which is linearly dependent on the carbon dioxide level.*

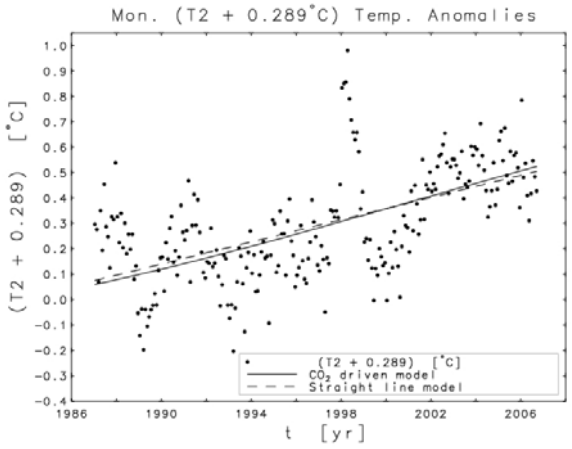

Figure 49. *The discrete data points are the measured T2 anomalies shifted to have the same zero point as the CRU anomalies. The dashed curve is the best fitting straight line, and the solid curve is just the final segment of the full model shown in Figure 48, replotted here with no adjustments to optimize the fit.*

[1] Bert Rust and Barend J. Thijsse, "Data-based Models for Global Temperature Variation", in *Proceedings of the 2007 International Conference on Scientific Computing*, CSC 2007, June 25-28, 2007, Las Vegas Nevada, CSREA Press 2007, pp. 10-16.

Statistics of Electromagnetic Fields

Andrew Dienstfrey
Brian Rider (University of Colorado at Boulder)
Ronald Wittmann (NIST EEEL)

Statistical descriptions of random electromagnetic (EM) fields are increasingly significant in diverse research areas. In electromagnetic compatibility (EMC), EM field statistics provide the foundation for new compliance testing techniques. For example, in concerning unintentional radiators one wishes to understand their directivity statistics. Similar statistics describing cavity quality factors underlie tests for EMC in cluttered environments. In both settings, statistical testing methods are more practical than direction-by-direction predecessors. Further, the consequences of incorrect testing can be considerable, such as adverse biomedical effects and poor device performance. It is therefore essential to put the statistical models on a solid footing as contrasted to the heuristic approach which is often employed at present.

By way of introduction, let $f(\hat{\mathbf{r}})$ be a complex-valued function on the sphere $S^2 \subset \mathbf{R}^3$. For any such function, the associated directivity function and directivity are defined by, respectively,

$$D_f(\hat{\mathbf{r}}) \equiv \frac{|f(\hat{\mathbf{r}})|^2}{\|f\|^2}, \qquad (1)$$

$$D_f \equiv \max_{\hat{\mathbf{r}}\in S} D_f(\hat{\mathbf{r}}).$$

In electromagnetics, given a radiator *f*, the directivity function represents the ratio of the time-averaged power to total power that is emitted in the far-field in the direction $\hat{\mathbf{r}}$. Modern electronic devices often have internal EM fields oscillating at microwave frequencies. Some of these devices, for example cell phones, contain antennas that are designed to transmit and receive these signals. For other devices, such as desktop computers and microwave ovens, these fields are meant only for the internal operations and any radiation outside the device is unintentional.

As microwave radiation can have potential non-trivial interactions with living tissue, as well as interfere with designed emitters, regulatory structures exist to control the quantity of unintended radiation. In some instances these regulations impose restrictions on the maximum directivity, D_f, that *f* may emit. To verify compliance, a test device is mounted on a stage and detailed measurements of the emitted power, i.e., the directivity function, are made over the 4π-steradian surface of a sphere. The directivity is then estimated accordingly. This is a costly measurement process, leading some to consider the possibility of using statistical methods to estimate the directivity of unintentional radiators. Preliminary results were derived using heuristic methods in [5]. In 2008 we continued our investigation of these statistics using more rigorous techniques.

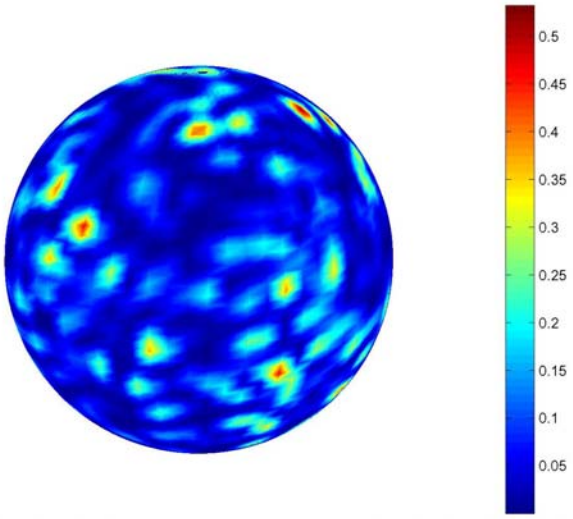

Figure 50. *A sample directivity function,* $D_f(\hat{\mathbf{r}};N)$ *drawn from the random model.*

We require a statistical model for random functions on the sphere $S^2 \subset \mathbf{R}^3$. Let $Y_{n,m}(\hat{\mathbf{r}})$ denote a spherical harmonic, we introduce a class of random functions on S^2 by

$$f(\hat{\mathbf{r}}) \equiv \sum_{n=0}^{N} \sum_{|m|\le n} \hat{f}_{n,m} Y_{n,m}(\hat{\mathbf{r}}), \qquad (2)$$

where $\hat{f}_{n,m}$ are i.i.d. standard complex Gaussian. Although a simple model, rigorous statements on the distributions of derived random variables require new results in probability theory. For example, the isotropy of (2) is a surprisingly subtle issue, see [2]. For any *N* the above sum can be viewed as representing the angular variation of a radiating solution to the scalar Helmholtz equation in the far field. In such case, by standard scaling arguments, the value of *N* corresponds to the electromagnetic size of the radiator. Using a heuristic approach, based on a discrete sampling of the field which was then treated as i.i.d., [5] provided an upper bound on the expectation of $\max_{\hat{\mathbf{r}}\in S} D_f(\hat{\mathbf{r}})$. In fact, such samples are not independent, and the consequences of this for the heuristic arguments are not well

understood. By contrast, we rigorously proved the following tail estimate.

Theorem 1 [3]. *Let* $f(\hat{\mathbf{r}})$ *be as in* (2). *Then for N large and all s>0, it holds*

$$\mathrm{Prob}\left(\|f\|_\infty^2 \geq \mathrm{E}\left(\|f\|_2^2\right)(s + 2\log(N)) \right) \leq C\exp(-s).$$

This bound on tail probability is a substantial improvement over a treatment restricted to the mean, and provides useful information for testing. Still, one is immediately prompted to ask whether the bound is sharp. We are considering analytic approaches to this problem using topological and statistical techniques, see for example [1], [4], and [6].

As an alternative, it is also possible to perform numerical experiments. For different values of *N*, we generate a sample from the space of random functions (2) and compute the associated maximum directivity (1). We show a plot of a single random directivity function with $N = 20$ in Figure 50. Drawing N_{sam} = 10,000 such samples, it is possible derive empirical statistics for the random variable as a function of electromagnetic size, $D_f = D_f(N)$. In Figure 51 we show the empirical probability density functions observed over a range of values for *N*. In Figure 52 we show the scaling of $\mathrm{E}(D_f(N))$ and its variance as a function of *N*. The scaling law $\mathrm{E}(D_f(N) \approx \log(N)$ is expected from general arguments as well as the heuristic arguments put forth in [5]. However the constant terms are inconsistent with [5] and require further analysis to be explained. In Figure 14 we also find that the variance of the directivity is constant independent of *N*. This result is surprising and again awaits further analysis. Ultimately, our goal is to provide rigorous foundations for statistical testing frameworks applied to problems in electromagnetic compatibility and related fields.

[1] R.J. Adler and J.E. Taylor, "Euler characteristics for Gaussian fields on manifolds," *Annals of Probability* **31** (2) (2003), pp. 533-563.

[2] P. Baldi and D. Marinucci, "Some characterizations of the spherical harmonics coefficients for isotropic random fields," arXiv: Pr/0606709, 2006.

[3] A. Dienstrey and R. Wittmann, "Extrema of random isotropic fields," 2006, submitted.

[4] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd edition. Springer, 1998.

[5] P. F. Wilson, D. A. Hill, and C. L. Holloway, "On determining the maximum emissions from electrically large sources," *IEEE Transactions on Electromagnetic-Compatibility* **44** (1) (2002), pp. 79-86.

[6] J. E.Taylor, K. J. Worsley, and F. Gosselin, "Maxima of discretely sampled random fields, with an application to `bubbles'," *Biometrika* **94** (2007), pp. 1-18.

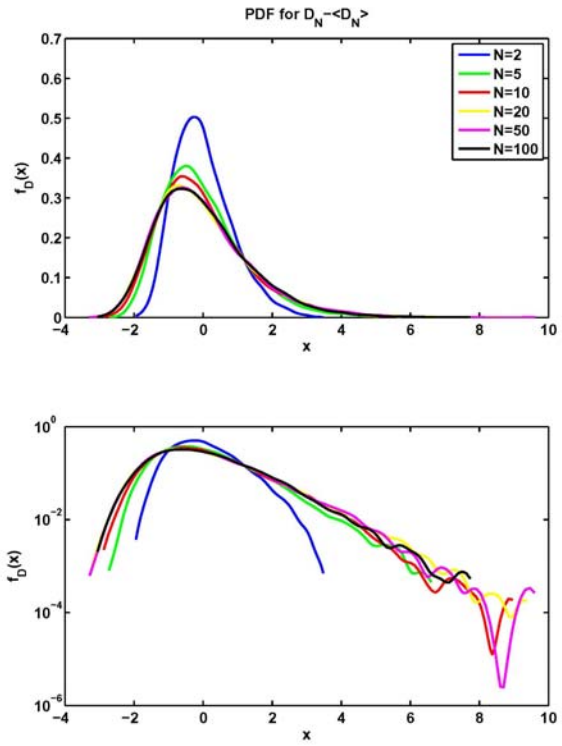

Figure 51. *The empirical probability density functions (PDF) for $D_f(N)$ for different values of N. The PDF's are shown on both linear and log scales to examine the bulk and tail behavior.*

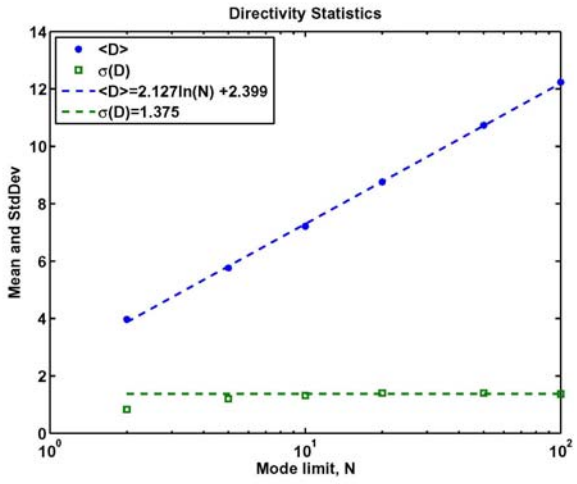

Figure 52. *Scaling of the mean and standard deviation,* $\langle D_f(N)\rangle$ *and* $\sigma(D_f(N))$*, as functions of N. The growth in expectation as* log(*N*) *is expected however the constant factors are unexplained. The variance saturation is surprising.*

Boundary Slip Effects on Instability in Pipe Flow

Geoffrey B. McFadden
P. Aaron Lott
David L. Cotrell (Lawrence Livermore National Lab)

This study explores a classical problem in the stability of the flow of fluid through a cylindrical pipe. This type of flow is relevant to a number of important applications, ranging from material transport in pipelines to blood flow in veins and arteries. The flow is driven by pressure gradients along the direction of the pipe, and is retarded by the effects of friction on the walls of the pipe. In general, the nature of the flow depends on the ratio of inertial forces to viscous forces in the fluid, as measured by the dimensionless Reynolds number. For small Reynolds numbers (high relative viscosity) the flow is laminar. For large enough Reynolds numbers, however, the laminar flow is observed to become unstable, leading to more complicated ("secondary") flows. The resulting turbulence in the flow field gives rise to additional phenomena such as enhanced mixing and an increase in the pressure gradient required to drive the flow.

The prediction of the onset of instability is often possible by employing linear stability theory to determine critical modes that grow temporally or spatially. Such studies have been successful for a number of flow geometries, but not for the case of pipe flow. For an idealized pipe of infinite length and circular cross-section a simple one-dimensional flow can be found from the underlying equations of motion whose stability can be examined through numerical means. There is a long known contradiction between the resulting linear stability results and the experimental observation that such flows become unstable at a Reynolds number of about 2000 for ordinary pipes. A hint that wall roughness may be important can be gathered from experiments which show that for smoothed pipes the onset of the instability can greatly exceed 2000.

We have investigated the possibility that the use of no-slip boundary conditions in modeling the flow is responsible for the discrepancy between theory and experiment, because such models ignore the amplitude variations associated with the roughness of the wall [1]. A rough surface is sometimes modeled though the introduction of a slip boundary condition, in which flow tangential to the boundary is allowed: the magnitude of the tangential flow is proportional to the local shear at the boundary. We currently are investigating the linear stability of pipe flow with a slip boundary condition using a standard normal mode analysis. The no-slip version of this problem has also been studied from the viewpoint of its pseudospectra, since the differential equations are non-normal. We are also computing the associated pseudospecta of the problem with slip boundary conditions.

[1] D. L. Cotrell, G. B. McFadden, and B. J. Alder, "Instability in Pipe Flow," *Proceedings of the National Academy of Science* **105** (2008) 428–430.

Funding: *NIST/NRC Postdoctoral Associateship Program*

Simulation of Bioregulatory Networks Involved in Cell Cycle Control

Geoffrey McFadden
Mirit Aladjem (National Institutes of Health)
S. Kim (National Institutes of Health)
Kurt Kohn (National Institutes of Health)
A. Luna (National Institutes of Health)

We are developing models of bioregulatory networks that are involved in cell cycle control. The models consist of systems of nonlinear ordinary differential equations or delay differential equations that typically exhibit switching behavior, limit cycles, and other types of bifurcations. In this effort, Geoffrey McFadden of MCSD is a co-advisor to S. Kim, a postdoctoral fellow, and A. Luna, a graduate student from Boston University, who are both doing research at the Laboratory of Molecular Pharmacology in the National Cancer Institute at NIH; their co-advisors are NIH researchers M. Aladjem and K. Kohn.

Proper cell growth depends on a network of interacting molecules that monitors cellular metabolism and environmental signals. This network ensures that cells halt their growth in response to unfavorable conditions such as the absence of sufficient nutrients or the presence of potentially damaging agents. When cells escape these controls, the results are developmental abnormalities, genomic instability, and cancer. Much current work focuses on a protein known as p53, a tumor suppressor that causes cell cycle arrest or programmed cell death in response to stress signals such as DNA damage. Regulating the appropriate levels of p53 is essential for cell survival. Two associated proteins, Mdm2 and Mdmx, are known regulators of p53. Mdm2 can facilitate degradation of p53, whereas the mechanism of the regulatory interaction of Mdmx with p53 is not clear. It is also not obvious how those three molecules will operate together under various conditions.

To address those questions, a mathematical model has been developed to investigate the interactions of these three partner molecules by numerical simulations [1]. An interesting feature of this system is the experi-

mental observation of time-periodic behavior of the measured amounts of p53 and Mdm2 in the system under certain conditions. These results show the stabilizing effect that Mdmx has on the system at long times: for large enough amounts of Mdmx in the system the entire branch of steady state solutions is found to be linearly stable. This modeling work is intended to guide experimental investigations of the role of Mdmx in the cell cycle that are being performed by Dr. Kim at NIH. A. Luna has recently joined the team to conduct his Ph.D. research on a combination of stochastic and deterministic modeling of cell cycle regulation.

[1] S. Kim, M. I. Aladjem, G. B. McFadden, and K. W. Kohn, "A Role for MdmX in Fine Tuning the Response of p53 to DNA Damage, submitted.

High Performance Computing and Visualization

Computational capability is advancing rapidly. This means that modeling and simulation can be done with greatly increased fidelity (e.g. higher resolution, more complex physics). However, developing large-scale parallel applications remains highly challenging, requiring expertise that application scientists rarely have. In addition, the hardware landscape is changing rapidly so new algorithmic techniques must constantly be developed. We are developing and applying facilities and expertise of this type for application to NIST problems. Large scale computations and laboratory experiments invariably produce large volumes of scientific data, which cannot be readily comprehended without some form of visual analysis. We are developing the infrastructure necessary for advanced visualization of scientific data, including the use of 3D immersive environments and applying this to NIST problems. One of our goals is to develop the 3D immersive environment into a true interactive measurement laboratory.

Partitioning for Dynamic Load Balancing with Adaptive Grids

William Mitchell

See feature article, page 25.

Visualization of the Energy Spectrum of Ultracold Atoms in a Synthetic Magnetic Field

Terence Griffin
Judith Terrill
Jay Vaishnav (NIST PL)
Charles Clark (NIST PL)

http://math.nist.gov/mcsd/savg/vis/moth/index.html

Among the driving forces behind the study of ultracold atoms is the exquisite control and flexibility available in experiments. This flexibility enables the engineering of new states of matter, e.g. cold atom emulators of other physical systems which are not presently accessible to experiment.

When an electron moves around in a two-dimensional crystal lattice, turning on a magnetic field makes the electron's energies generate a special kind of fractal known as the Hofstadter butterfly – one of the first fractals discovered in physics. With cold atoms, one can not only mimic this system, but can also manipulate it in exotic ways that might be difficult or impossible to do for a real electron. In 2005, scientists in Europe engineered a laser configuration where the cold atoms precisely mimicked the physics of these electrons, and then extended it so that instead of a real magnetic field, the lasers created an exotic artificial magnetic field known as a non-Abelian gauge field. The scientists named the energy spectrum the "Hofstadter moth." The NIST Physics Laboratory and their collaborators went one step further and removed the crystal lattice, generating a slightly different species of Hofstadter's moth.

Scientific visualizations can make physical phenomena easier to understand — especially, in quantum mechanics, where almost nothing can be easily seen. These data are complicated three-dimensional energy surfaces, and are difficult to interpret. Using visualization, certain features of the data jump out. For example,

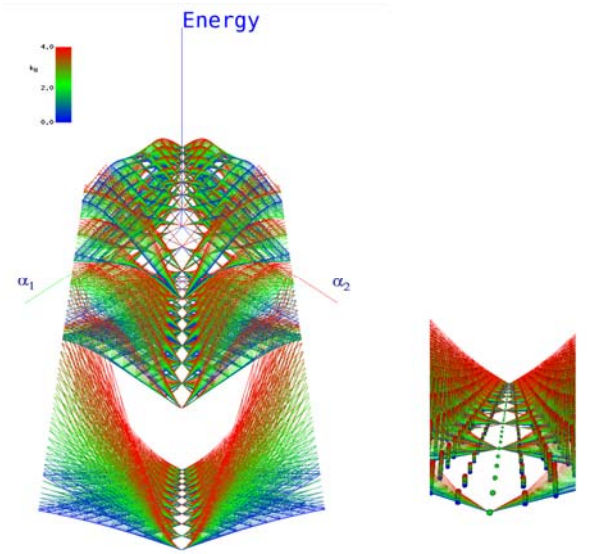

Figure 53. *Visualization of the Energy Spectrum of Ultracold Atoms in a Synthetic Magnetic Field. The figure on the left shows the spectrum of energy levels of a cold atom in a laser configuration designed to make the atom mimic an electron confined in two dimensions, in a generalization of a magnetic field called a "non-Abelian gauge field." The atom's energy levels are functions of three parameters: α_1 and α_2 are related to the strengths of the artificial magnetic field, and k_y represents the atom's momentum. Along the line ($\alpha_1 = \alpha_2$), the field reduces to the usua l (Abelian) magnetic field. This type of figure is sometimes called a "Osterloh moth," after the "Hofstadter butterfly" stru cture that was identified by Douglas Hofstadter in the Abelian system. The figure on the right shows a close up of the line ($\alpha_1 = \alpha_2$).*

that the "moth" has a "backbone" where the energy surfaces all collapse, indicating that for certain laser configurations, the artificial magnetic field is not non-Abelian.

Tere Griffin of MCSD has created a visualization of the moth, which runs both on the desktop and in the immersive visualization laboratory; see Figure 53, as well as the cover of this document. The data contains three independent variables: α_1 and α_2, which are related to the strengths of the artificial magnetic field, and k_y, which represents the atom's momentum. For each triple, there are six dependent variables, representing energy. For the visualization, α_1 and α_2 are represented as the x and y axes. k_y is represented by color using a linear color ramp where the lowest value, zero, is blue, the middle value, 2.0 is green and the highest value, 4.0 is red. The six dependent variables are plotted in the z direction. Structure was added to show the surface relationship for each dependent variable by building a grid on each surface. All grid lines are composed of actual data points, therefore there are no interpolation artifacts. There is a button on the user interface to individually turn on or off the points for each value of k_y.

Finally, a fatpoints shader was used to generate pseudo-spheres at each point that are of constant (but user adjustable) size with respect to the virtual coordinates. Normally points occupy a constant number of pixels on the display so the visual density of a point set gets lower as the view point gets closer. The fatpoints generate billboards that are shaded to convincingly appear as spheres. Since this is done on the GPU, there is no use of CPU resources. If the size is set to zero, the fatpoints shader is turned off, and non-fatpoints are generated. Using these features, scientists can easily visualize and highlight different relationships between the data points.

Future work will examine measurements on the moth.

Virtual Cement and Concrete Testing Laboratory

Edith Enjolras
William George
Terence Griffin
John Hagedorn
John Kelso
Julien Lancien
Adele Peskin
Steve Satterfield
James Sims
Judith Terrill
Jeffrey Bullard (NIST BFRL)
Clarissa Ferraris (NIST BFRL)
Edward Garboczi (NIST BFRL)
Nicos Martys (NIST BFRL)

The NIST Building and Fire Research Laboratory (BFRL) does experimental and computational research in cement and concrete. MCSD has an ongoing collaboration with them to develop highly efficient simulations and in creating visualizations of their data. This work is done in the context of the Virtual Cement and Concrete Testing Laboratory (VCCTL) Consortium which we helped form in 2001. The NIST-led consortium consists of seven industrial members: BASF Admixtures, Ready Mixed Concrete Research and Education Foundation (RMC), National Stone Sand and Gravel Association (NSSGA), W.R. Grace, Sika Technology AG, Mapei, and the Federal Highway Administration Turner-Fairbank Laboratory. The overall goals of the consortium are to develop a virtual testing system to reduce the amount of physical concrete testing and to expedite the overall research and development process. It is expected that this will result in substantial time and cost savings to the concrete construction industry as a whole. MCSD continues as an active participant in the VCCTL.

For further details of work done under the auspices of the VCCTL, see *Computational Modeling and Visualization of Cement Paste Hydration and Microstructure Development*, page 39, and *Computational Modeling and Visualization of the Flow of Concrete,* page 53.

Funding: *Virtual Cement and Concrete Testing Laboratory (VCCTL) Consortium.*

Three-D Desktop

John Kelso
Terence Griffin
John Hagedorn
Marc Olano
Adele Peskin
Steven Satterfield
Judith Terrill

http://math.nist.gov/mcsd/savg/vis/desktop

Immersive visualization systems have many advantages, but not all users have ready access to them. Additionally, sometimes it's easier to look at new data on a desktop in the office rather than go down the hall to use the immersive system, and sometimes the immersive system is not available at all due to distance, time or other constraints. To this end, SAVG is developing the "3D desktop", a set of tools to help users more effectively use our immersive software in a desktop environment. These tools can not completely bridge the gap between desktop and immersive systems, but they can help ameliorate the situation. Some of the tools will be general and can be used in any application, while others will be geared towards a specific application or type of application.

Our immersive software development API, DIVERSE, automatically allows applications to be used in both immersive and desktop systems. But, the API does not address user interface questions that are raised due to the differences between the two types of environments. Therefore our tools do not need to modify the underlying immersive application, but merely replace the immersive interaction techniques with ones more suitable for desktop use, and augment these techniques with additional tools to make up for the lack of intuitive presence necessitated by the desktop system.

This year we developed a measurement capability that works with a geometric constraint, i.e., the measurement of length in a two dimensional plane. This works in the immersive environment and also on the desktop. The plane measurement tool is very flexible. The plane can be put anywhere and is visualized as a grid of any shape, size, and color. The measurement glyph can be any defined 2D glyph of any shape, size, and color. When a measurement is made, a marker is deposited that can also be any 2D shape, size, or color. On the desktop, the measurement plane is positioned with rollers that can be moved in the *x*, *y*, and *z* directions (in data or screen coordinates), and also can be moved in heading, pitch, and roll. In the immersive environment the plane is moved with a position-tracked wand. The text can be made any size or color. See Figure 54 for an example.

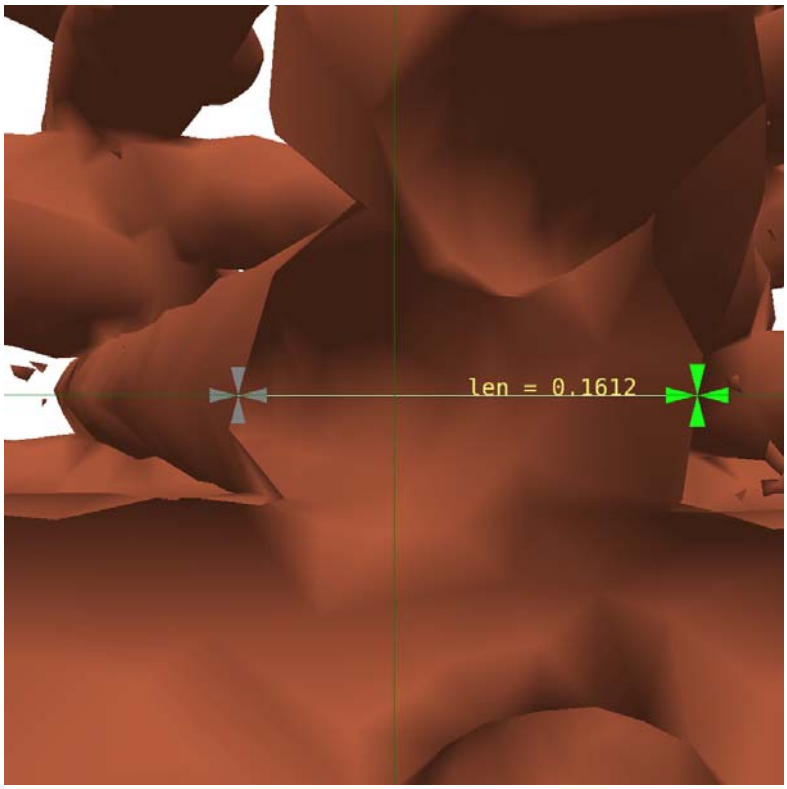

Figure 54. *Example of measurement in a two dimensional plane with a green rectangular grid and a green cross for a selection glyph. The marker glyph is the same shape as the measurement glyph, but shown in gray.*

Parallelization, Visualization, and Analysis of Hurricane Storm Surge

William George
Audrey Lemoussu
Alexandre Marie
Judith Terrill
Long Phan (NIST/BFRL)

http://math.nist.gov/mcsd/savg/vis/StormSurge

We are working with the Building and Fire Research Laboratory (BFRL) to contribute to a collaboration with NOAA's National Weather Service (NWS) and the Oceanic and Atmospheric Research Office (OAR) to develop methodology for use in deriving site specific, risk-based design criteria for coastal structures subjected to the combined effects of hurricane wind speed and storm surge. The methodology utilizes an integrative, interdisciplinary approach that incorporates state-of-the-art knowledge in hurricane science, hydrology, probabilistic methods, and structural engineering needs. It involves the following steps: (1) selection of a stochastic set of hurricane storm tracks affecting the region of interest, (2) hydrodynamic simulation of the region of interest using NOAA's SLOSH (Sea, Lake, and Overland Surges from Hurricanes) model to obtain storm tracks to generate time histories of wind speeds and corresponding time histories of storm surge heights at sites within the affected region, and (3) use of the resulting hurricane wind speed and storm surge height data to develop probabilistic information on joint wind speed/storm surge height events, and ultimately risk-consistent structural design criteria.

This year we parallelized the runs of SLOSH on the NIST Linux cluster, and are working on incorporating the probabilistic data post-processing routines developed by BFRL in the proof of concept phase into a map-based graphical user interface (GUI) that allows users to graphically select multiple points of interest, extract wind speed and storm surge height data output by SLOSH and linked geographically to those points, and perform probabilistic analyses of data for joint probabilities of different combinations of wind speed and storm surge heights. The map-based GUI, preliminarily named *SLOSH Statistical Map Tool*, is based on the NASA World Wind map. Figure 55 shows the current version.

With this map-based tool, one can browse a SLOSH Basin database (a collection of time step files). The program will automatically detect the basin on the earth component, save all information about the basin (e.g., basin name, number of time step files, number of locations), display this information, and save all points of interest. Within the tool, software for computing and plotting joint histograms, probabilities of exceedance, and mean recurrence intervals of different wind speed and storm surge height events for any points of interest can be activated.

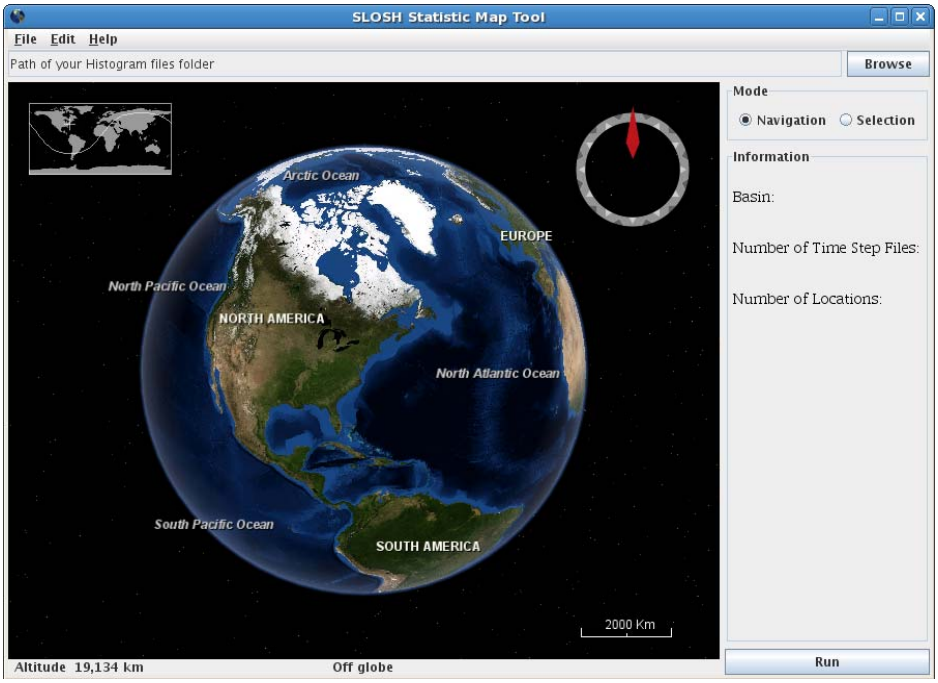

Figure 55. *SLOSH Statistical Map tool.*

There are thirty nine hurricane basins currently available in the SLOSH system. These are being updated by the National Weather Service due to changes in coastal topography over time. We have developed a separate program to read, triangulate, and visualize arbitrary basin surfaces based on the SLOSH basin data format. All of the basins have been visualized with this program. The basins contain additional information such as channels and obstructions. In the coming year, we will be making modifications to our program to incorporate this information.

We have also developed a program to read the output of SLOSH and triangulate the surface of the storm surge. Additional work will be done in the coming year to add realism. The Figure 56 shows an example storm surge visualized using out tools.

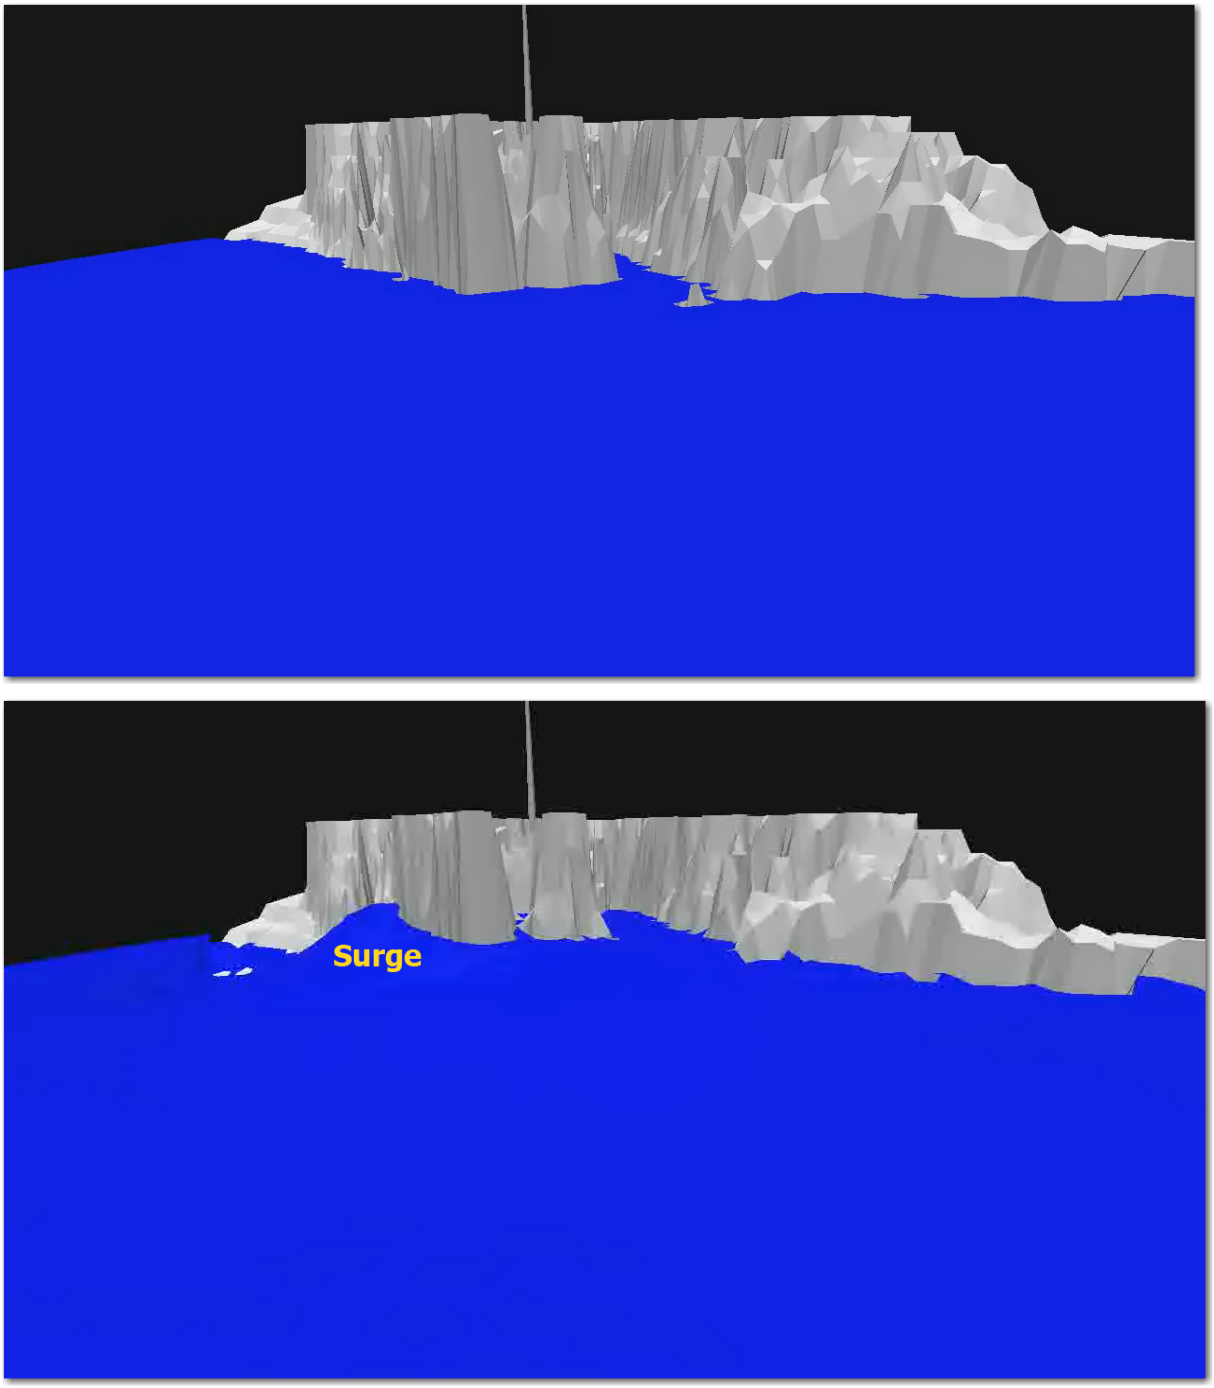

Figure 56. *Tampa Bay at start of storm (top) and during storm surge (bottom).*

Visualizing Fire and the Deformation of Structures Due to Fire

William Hess
Marc Olano
Judith Terrill
Dilip Banerjee (NIST BFRL)

Researchers at the Building and Fire Research Lab (BFRL) at NIST are currently studying the effects of fire on building structures to advance structural engineering and building safety. We are using a sequential process in which the BFRL Fire Dynamics Simulator (FDS) is used to simulate the start and propagation of fire in a room, after which a second code calculates how the temperature computed by FDS diffuses into the structural members. A third program takes the output of the second code and computes how the structure deforms over time due to combined effects of temperature and loads. We used a test case with a single beam in a room. We link the fire, thermal, and structural data with a separate real-time visualization capability that can be run on the desktop and in the NIST 3-D immersive environment. This provides interactive analysis of the resulting data.

We have explored two methods to visualize the fire. The first uses a polygonal representation and the following volumetric algorithm: each cube in the mesh was subdivided according to a level of detail coefficient, and the heat release rate per unit volume values of interior points were interpolated linearly. For each of the resultant cubes, a random vertex was selected, and a tetrahedron was drawn between that point and the three adjacent points. The color and transparency of these tetrahedra were related to the heat release rate per unit volume value of the points. This is shown in Figure 57. Algorithms to visualize the smoke and the deformation of the beam were developed. The application also includes interactive controls to select a point and numerically integrate the fluid flow velocities for a distance and display the generated curve. This allows the user to see what is important to them, rather than look at every single flow vector.

The second method to visualize the fire uses a graphics processing unit (GPU), but rather than use it for traditional polygon rendering, we use it's parallel processing ability to trace rays through the volumetric data, computing the accumulated color and opacity of each ray. This results in a high-quality interactive visualization for interactive exploration, either on the desktop or in the immersive RAVE. See Figure 58.

In the coming year, we will look at additional ways to visualize the fire and the deformation as well as measurement and analysis on the results.

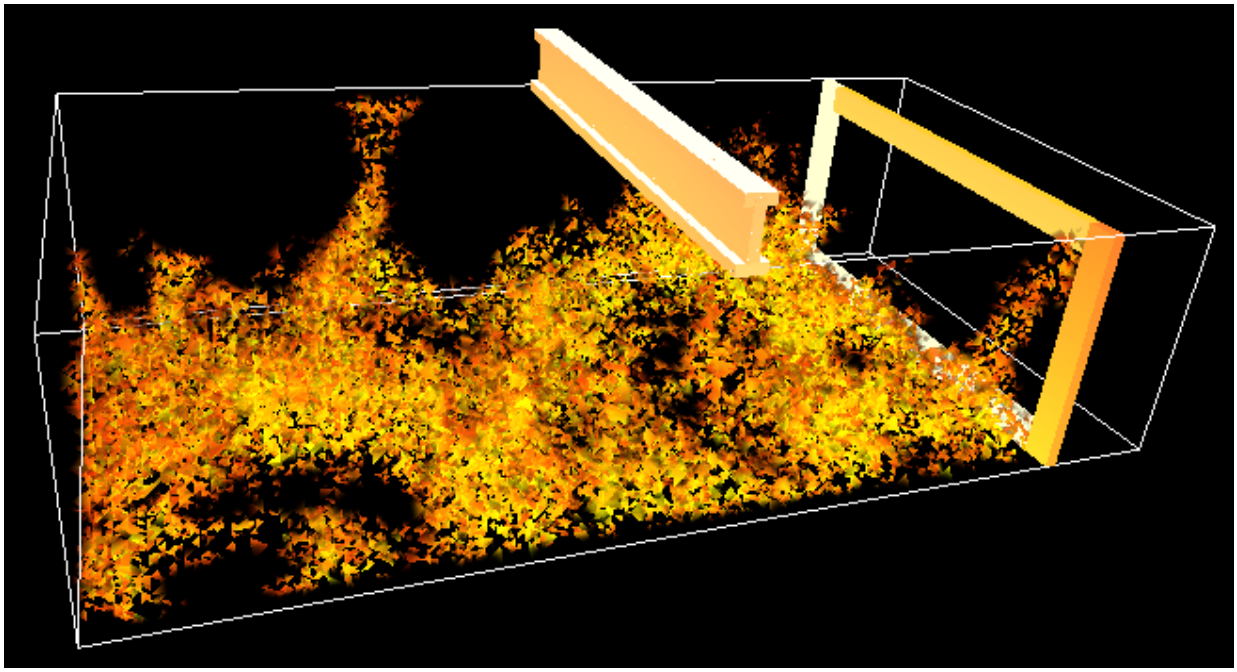

Figure 57. *Polygon based visualization of fire in a room with an exposed beam.*

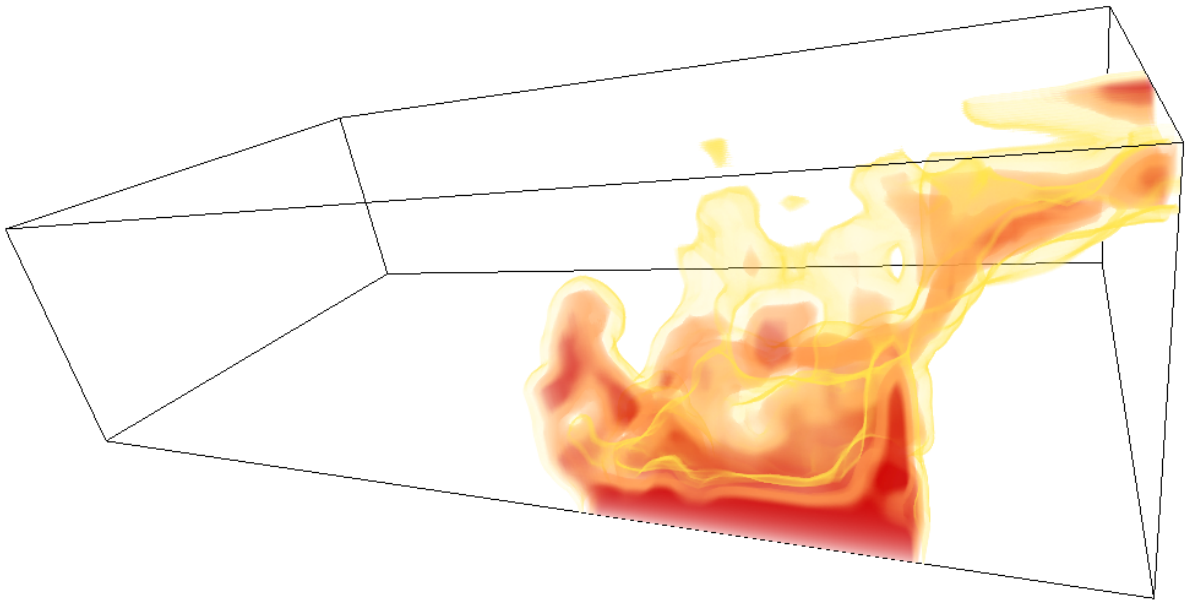

Figure 58. *This image shows one frame from a visualization of a 2,000 frame fire dynamic simulation.*

Fundamental Math Software Development and Testing

Modern science and engineering in general, and modern measurement science in particular require a wide variety of software tools for scientific discovery, exploration, and analysis. As scientific inquiry becomes deeper and more specialized, so must the supporting software tools. The goal of this work is to develop critical software tools that support measurement science at NIST, as well as computational science and engineering at large.

SciMark, a Web-based Benchmark for Numerical Computing in Java

Roldan Pozo
Bruce Miller

http://math.nist.gov/scimark

The NIST SciMark benchmark continues to be one of the most widely used Java scientific benchmarks and is now an official component of the industry standard SPECjvm2008. SciMark consists of computational kernels for FFTs, SOR, Monte Carlo integration, sparse matrix multiply, and dense LU factorization, comprising a representative set of computational styles commonly found in numerical applications. SciMark can be run interactively from Web browsers, or can be downloaded and compiled for stand-alone Java platforms. Full source code is provided, in Java and C programming languages for comparison under different compilers and execution environments. The SciMark result is recorded as a series of megaflop rates for the numerical kernels, as well as an aggregate score for the complete benchmark. The current results database lists results of more than 3,000 submissions from computational platforms ranging from palm devices to high-end servers, representing nearly every operating system and virtual machine environment currently in use, including Solaris, FreeBSD, MacOS, Sun OS, IRIX, OSF1, Linux, OS/2, and Windows 95, 98, 2K, ME, NT, and XP.

As of December 2008, the record for SciMark is over 1.1 Glops, with some of the kernel benchmarks (notably the LU factorization) running over 2.1 Gflops on a single-processor PC. A recent search on Google for "java numerical benchmark" returned SciMark as the #1 result out of 87,000 sources.

SciMark and its kernel components have become a *de facto* standard in industry and academia; for example, they have been adopted by the Java Grande Benchmark Forum, and Sun Microsystems has used SciMark 2.0 to demonstrate floating-point improvements to their JVM. This year SciMark was included in the newly released SPECjvm2008, an industry standard benchmark suite for evaluating the performance of Java runtime environments. SciMark is one of 10 benchmarks included in SPECjvm2008 to represent the wide range of existing Java applications. Other benchmarks included in the suite assess performance in areas such as cryptographic operations, audio decoding, graphics processing, database logic, and XML processing. SPECjvm2008 was released on May 7, 2008 by the System Performance Evaluation Corporation[13] (SPEC), a non-profit open industry cooperative which develops benchmarks that ensure that the marketplace has a fair and useful set of metrics to differentiate system performance. For example, results of the well-know SPEC CPU benchmarks (such as SPECint, SPECfp, and SPECrate) are typically cited on specification sheets used in the marketing of computer processors. Members of SPEC include AMD, Dell, HP, IBM, Intel, Microsoft, and Sun.

Sparse BLAS Standardization

Roldan Pozo
Iain Duff (Rutherford Appleton Labs)
Michael Heroux (Sandia National Laboratory)

http://math.nist.gov/spblas
http://www.netlib.org/blas/blast-forum

MCSD has played a leading role in the standardization effort for the Basic Linear Algebra Subprograms (BLAS) and continues to be a major contributor for the design and development of reference software and documentation. The BLAS are kernels for computational linear algebra subroutines comprising of fundamental matrix/vector operations common to most scientific computing applications. By standardizing such interfaces, computer manufacturers and software vendors can provide high-performance implementations especially suited to a specific hardware platform. By developing their applications in terms of BLAS, computational scientists can achieve high levels of performance and portability. The original BLAS, which were developed for dense vector and matrix operations from the late 1970s through the early 1990s, achieved this goal very well. To this, the BLAS Technical Forum (an international consortium of industry,

[13] http://www.spec.org/

academia, and government institutions, including Intel, IBM, Sun, HP/Compaq/Digital, SGI/Cray, Lucent, Visual Numerics, and NAG) formed an updated BLAS standard which addresses these extensions.

Among the most significant components of the updated BLAS standard is support for sparse matrix computations. R. Pozo of MCSD served as chair of the Sparse BLAS subcommittee during the standardization process, and NIST was first to develop and release a public domain reference implementation in ANSI C for early versions of the standard, which helped shape the final specification. After the standard was formally approved and accepted, the complete technical specification was published, and a special issue of the ACM Transactions of the Mathematical Software (TOMS) was devoted to the new BLAS standard, including a paper co-authored by R. Pozo and other subcommittee members providing an overview of the sparse matrix interface.

There are existing interfaces for the sparse BLAS in C and Fortran. The latest C++ interface, developed at NIST, uses a simplified approach which reduces the number of functions and parameters, while retaining the generality of sparse storage structures and precision. This interface specification is less than 600 lines of C++, yet captures the complexity of Level 1, 2, and 3 vector/matrix operations with basic types (real and complex, single or double precision) and supports general types (e.g. intervals, multiprecision variables) beyond the original BLAST specification.

TNT: Object Oriented Numerical Programming

Roldan Pozo

http://math.nist.gov/tnt/

NIST has a history of developing some of the most visible object-oriented linear algebra libraries, including Lapack++, Iterative Methods Library (IML++), Sparse Matrix Library (SparseLib++), Matrix/Vector Library (MV++), and most recently the Template Numerical Toolkit (TNT). This package has been downloaded by several thousand developers and is currently in use in several industrial and commercial applications.

TNT incorporates many of the ideas we have explored with previous designs, and includes new techniques that were difficult to support before the availability of ANSI C++ compilers. The package includes support for both C and Fortran multidimensional arrays, vector, matrices, and application modules, such as linear algebra.

The design of TNT separates the interface specification from the actual implementation. This allows library developers to create specialized modules that take advantage of particular hardware platforms, utilize vendor-specific libraries, or implement different C++ strategies, such as expression templates, or instrumented versions for debugging sessions.

Developments in the latest design of TNT (version 3.0.11) provide support for both multidimensional arrays and integrate a linear algebra module of fundamental algorithms (LU, Cholesky, SVD, QR, and eigenvalues), as well as sparse matrix support (with computational operations via the Sparse BLAS). Example codes to interface with LAPACK high performance linear algebra library are also provided.

The TNT web site provides a basic implementation for testing and development, as well as links to other library packages that utilize the TNT interface. Full documentation and source code for all TNT components are available online.

Part IV

Activity Data

Publications

Appeared

Refereed Journals

1. M. E. Altamimi and A. S. Youssef, "A Math Query Language with an Expanded Set of Wildcards," *Mathematics in Computer Science* **2** (2) (published online December 10, 2008).
2. W. J. Boettinger, J. E. Guyer, C. E. Campbell, and G.B. McFadden, "Computation of the Kirkendall Velocity and Displacement Field in a 1-D Diffusion Couple with a Moving Interface," *Proceedings of the Royal Society of London* **463** (2007), pp. 3347-3373.
3. A. Carasso and A. Vladar, "Calibrating Image Roughness by Estimating Lipschitz Exponents, with Applications to Image Restoration," *Optical Engineering* **47** (3) (March 2008), 037012, 13 pages.
4. D. L. Cotrell, G. B. McFadden, and B. J. Alder, "Effect of an Axially-periodic Radius on the Linear Stability of Pipe Flow," *Proceeding of the National Academy of Sciences*, **105** (2008), pp. 428-430.
5. T. Dennis, S. D. Dyer, A. Dienstfrey, G. Singh, P. Rice, "Analyzing Quantitative Light Scattering Spectra of Phantoms Measured with Optical Coherence Tomography," *Journal of Biomedical Optics* **13** (2008), 024004, 9 pages.
6. A. Dienstfrey, T. L. Oreskovic, L. T. Hudson, H. S. Bennett, "ISCD-NIST DXA Survey: Preliminary Report," *SCAN* **IX (**Quarter 1) (2008), pg. 14.
7. R. J. Epstein, S. Sedelin, D. Leibfried, J. H. Wesenberg, J. J. Bollinger, J. M. Amini, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland, "Simplified Motional Heating Rate Measurements of Trapped Ions," *Physical Review A* **76** (2007), 033411, 5 pages.
8. D. E. Gilsinn, "Computable Error Bounds for Approximate Periodic Solutions of Autonomous Delay Differential Equations," *Nonlinear Dynamics* **50** (2007), pp. 73-92.
9. S. Glancy and H. Vasconcelos, "Methods for Producing Optical Coherent State Superpositions," *Journal of the Optical Society of America B* **25** (5) (April 16, 2008), pp. 712-733.
10. R. Kacker, "Comments on 'Bayesian Evaluation of Comparison Data," *Metrologia* **44** (December 2007), pp. L57- L61.
11. R. Kacker, K.-D. Sommer, and R. Kessel, "Evolution of Modern Approaches to Express Uncertainty in Measurement," *Metrologia* **44** (December 2007), pp. 513-529.
12. R. Kacker, A. Forbes, R. Kessel, and K.-D. Sommer, "Classical and Bayesian Interpretation of the Birge Test of Consistency and Its Generalized Version for Correlated Results from Interlaboratory Evaluations," *Metrologia* **45** (3) (April 2008), pp. 257-264.
13. E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin and D.J. Wineland, Randomized Benchmarking of Quantum Gates, *Physical Review A* **77** (2008) 012307, 7 pages.
14. R. Kuhn, Y. Lei, and R. Kacker, "Automated Combinatorial Test Methods – Beyond Pairwise Testing" *CrossTalk*, June 2008, pp. 22-26.
15. R. Kuhn, Y. Lei, and R. Kacker, "Practical Combinatorial Testing: Beyond Pairwise," *IT Professional*, May/June 2008, pp. 19-23.
16. K. M. Lebecki, M. J. Donahue, and M. W. Gutowski, "Periodic Boundary Conditions for Demagnetization Interactions in Micromagnetic Simulations," *Journal of Physics D-Applied Physics* **41** (2008), 175005, 10 pages.
17. Y. Lei, R. Carver, R. Kacker, and D. Kung, "A Combinatorial Testing Strategy for Concurrent Programs," *Software Testing, Verification and Reliability* **17** (4) (December 2007), pp. 207-225.
18. Y. Lei, R. Kacker, D. R. Kuhn, V. Okun, and J. Lawrence, "IPOG/IPOG-D: Efficient Test Generation for Multi-way Combinatorial Testing," *Software Testing, Verification and Reliability* **18** (2008), pp. 125-148.
19. D. Leibfried, E. Knill, C. Ospelkaus, and D. J. Wineland, "Transport Quantum Logic Gates for Trapped Ions," *Physical Review A* **76** (2007), 032324, 12 pages.
20. Z. H. Levine, M. Li, A. P. Reeves, D. F. Yankelevitz, J. J. Chen, E. L. Siegel, A. Peskin, and D. N. Zeiger, "A Low-Cost Density Reference Phantom for Computed Tomography," *Medical Physics* **36** (2009) , pp. 286-288.
21. G. B. McFadden, S. R. Coriell, K. F. Gurski, and D. L. Cotrell, "Onset of Convection in Two Liquid Layers with Phase Change," *Physics of Fluids* **19** (2007), 104109, 13 pages.

22. D. G. Porter and M. J. Donahue, "Precession Axis Modification to a Semi-analytical Landau-Lifshitz Solution Technique," *Journal of Applied Physics* **103** (7) (February 25, 2008), pp. D920-1 to D920-3.

23. A. C. E. Reid, S. A. Langer, R. C. Lua, V. R. Coffman, S.-I. Haan, and R. E. Garcia, "Image-based Infinite Element Mesh Construction for Material Microstructures," *Computational Materials Science* **43** (2008), pp. 989-999.

24. B. Rust and D. O'Leary, "Residual Periodograms for Choosing Regularization Parameters for Ill-Posed Problems," *Inverse Problems* **24** (2008), 034005, 30 pages.

25. B. Rust and B. Thijsse, "Freestyle Data Fitting and Global Temperatures," *Computing in Science & Engineering* **10** (1) (January/February 2008), pp. 49-59.

26. C. A. Ryan, C. Negrevergne, M. Laforest, E. Knill, and R. Laflamme, "Liquid State NMR as a Test-bed for Developing Quantum Control Methods," *Physical Review A* **78** (1) (July 2008), pp. 012328-1 – 012328-14.

27. J. M. Shaw, S. E. Russek, M. J. Donahue, T. Thomson, B. D. Terris, M. Olsen, and M. Schneider, "Reversal Mechanisms in Perpendicular Magnetic Nanostructures," *Physical Review B* **78** (2008), 024414, 5 pages.

28. R.D. Somma, S. Boixo, H. Barnum, and E. Knill, "Quantum Simulated Annealing," *Physical Review Letters* **101** (13) (September 2008), pp. 130504-1 – 130504-4.

29. A. S. Youssef, "Relevance Ranking and Hit Description in Math Search," *Mathematics in Computer Science* **2** (2) (published online December 5, 2008).

Journal of Research of NIST

1. I. Beichl, S.S. Bullock and D. Song, "A Quantum Algorithm Detecting Concentrated Maps," *NIST Journal of Research* **112** (6) (November-December 2007), pp. 307-311.

2. J. Hagedorn, J. Dunkers, S. Satterfield, A. Peskin, J. Kelso, and J. Terrill, "Measurement Tools for the Immersive Visualization Environment: Steps Toward the Virtual Laboratory," *NIST Journal of Research* **112** (5) (September-October 2007), pp. 257-270.

3. J.G. Hagedorn, J.E. Terrill, A.P. Peskin, and J. J. Filliben, "Methods for Quantifying and Characterizing Errors in Pixel-Based 3D Rendering," *Journal of Research of NIST* **113** (4) (July-August, 2008), pp. 221-238.

4. G. B. McFadden, S. R. Coriell, K. F. Gurski, and D. L. Cotrell, "Convective Instabilities in Two Liquid Layers," *Journal of Research of NIST* **112 (5)** (September-October 2007), pp. 271-281.

5. J. S. Sims, W. L. George, T. J. Griffin, J. G. Hagedorn, H. K. Hung, J. T. Kelso, M. Olano, A. P. Peskin, S. G. Satterfield, J. Devaney Terrill, G. W. Bryant, and J. G. Diaz, "Accelerating Scientific Discovery through Computation and Visualization III: Tight-Binding Wave Functions for Quantum Dots," *NIST Journal of Research* **113** (3) (May-June 2008), pp. 131-142.

Book Chapters

1. J. Terrill, W. George, T. Griffin, J. Hagedorn, J. Kelso, M. Olano, A. Peskin, S. Satterfield, J. Sims, J. Bullard, J. Dunkers, N. Martys, A. O'Gallagher, and G. Haemer "Extending Measurement Science to Interactive Visualization Environments," in *Trends in Interactive Visualisation,,* (E. Zudilova-Seinstra, T. Adriaansen and R. van Liere, eds.), Advanced Information and Knowledge Processing Series, Springer, London, 2008, pp. 287-302.

Conference Proceedings

1. I. Beichl and B. Cloteaux, "Generating Network Models Using the S-Metric," in *Proceedings of the International Conference on Modeling, Simulation and Visualization Methods* (2008), Las Vegas, NV, July 14-16, 2008, pp. 159-164.

2. M. J. Cohn, J. T. Fong, and P. M. Besuner, "A Quantitative Approach to a Risk-based Inspection Methodology of Main Steam and Hot Reheat Piping Systems," in *Proceedings of the ASME Pressure Vessels & Piping Conference*, July 27-31, 2008, Chicago, IL, Paper No. PVP2008-61242.

3. J. T. Fong, J. J. Filliben, N. A. Heckert, and R. deWit, "A Design-of-Experiment Approach to Verification and Uncertainty Estimation of Simulations based on Finite Element Method," in *Proceedings of the 2008 Annual Conference of the American Society of Engineering Education*, June 22-25, 2008, Pittsburgh, PA, Paper No. 2725.

4. J. T. Fong, J. J. Filliben, N. A. Heckert, and P. V. Marcal, "An Intelligent Flaw Monitoring System: From Flaw Size Uncertainty to Fatigue Life Prediction with Confidence Bounds in 24 Hours," in *Proceedings of the 8th World Congress on Com-*

putational Mechanics, Venice, Italy, June 30 - July 5, 2008.

5. J. T. Fong, J. J. Filliben, N. A. Heckert, R. deWit, and B. Bernstein, "Robust Engineering Design for Failure Prevention," in *Proceedings of the ASME Pressure Vessels & Piping Conference*, July 27-31, 2008, Chicago, IL, Paper No. PVP2008-61602.
6. J. T. Fong, O.F. Hedden, J. J. Filliben, and N. A. Heckert, "A Web-based Data Analysis Methodology for Estimating Reliability of Weld Flaw Detection, Location, and Sizing," in *Proceedings of the ASME Pressure Vessels & Piping Conference*, July 27-31, 2008, Chicago, IL, Paper No. PVP2008-61612.
7. C. G. Interrante, J. T. Fong, J. J. Filliben, and N. A. Heckert, "Uncertainty Estimate of Charpy Data Using a 5-Factor 8-Run Design of Experiments," in *Proceedings of the ASME Pressure Vessels & Piping Conference*, July 27-31, 2008, Chicago, IL, Paper No. PVP2008-61565.
8. R. Kessel and R. Kacker, "Framework for Evaluation of Uncertainty with Test of Linearity Using Covering Arrays," in *Proceedings of the Conference on Advanced Mathematical and Computational Tools in Metrology and Testing* (AMCTM), Paris, France, June 23-25, 2008.
9. Y. Lei, R. Kacker, D. R. Kuhn, V. Okun, and J. Lawrence, "IPOG: A General Strategy for T-way Software Testing," in *Proceedings of the 14th Annual IEEE International Conference and Workshop on the Engineering of Computer-Based Systems (ECBS'07)*, March 26-29, 2007, Tucson AZ.
10. V. Marbukh and K. Mills, "Demand Pricing and Resource Allocation in Market-Based Compute Grids: A Model and Initial Results," in *Proceedings of the 7th International Conference on Networking*, Cancun, Mexico, April 2008.
11. V. Marbukh, "Can TCP Metastability Explain Cascading Failures and Justify Necessity of Flow Admission Control in the Internet?," in *Proceedings of the 15th International Conference on Telecommunications*, St. Petersburg, Russia, June 2008.
12. V. Marbukh and D. Genin, "Towards Understanding Metastability in Cellular CDMA Network: Emergence and Implications for Performance," in *Proceedings of the 12th World Multi-Conference on Systemics, Cybernetics and Informatics* (WMSCI 2008), Orlando, Florida, June 2008.
13. V. Marbukh and F. Hunt, "Dynamic Routing and Congestion Control through Random Assignments of Routes," in *Proceedings of the 5th International Conference on Cybernetics and Information Technologies, Systems and Applications* (CITSA 2008), Orlando, Florida, July 2008, pp. 161-164.
14. V. Marbukh and K. Sayrafian, "A Framework for Joint Cross-Layer and Node Location Optimization in Mobile Sensor Networks," in *Proceedings of ADHOC-NOW*, Nice, France, September 10-13, 2008.
15. V. Marbukh, "Towards Modeling Network Self-Organization/Evolution: Effects of Fairness, Risk Averseness, Competition, and Economic Pressures," in *Proceedings of the 4th European Conference on Complex Systems (ECCS 2008)*, Jerusalem, Israel, September 14-19, 2008.
16. N. S. Martys, D. Lootens, W. L. George, S. G. Satterfield, and P. Hébraud, "Spatial-Temporal Correlations in Concentrated Suspensions," in *Proceedings of the 15th International Congress on Rheology*, August 3-8, 2008, Monterey CA.
17. N. S. Martys, D. Lootens, W. L. George, S. G. Satterfield, and P. Hébraud, "Stress Chains Formation Under Shear of Concentrated Suspension," in *Proceedings of the 15th International Congress on Rheology*, August 3-8, 2008, Monterey CA.
18. X. Tang, L. Ma, A. Mink, T. Chang, H. Xu, O. Slattery, A. Nakassis, B. Hershman, D. Su, and R. F. Boisvert, "High-Speed Quantum Key Distribution System for Optical Fiber Networks in Campus and Metro Areas," in Quantum Communications and Quantum Imaging VI (R. E. Meyers, Y. Shih, and K. S. Deacon, eds.), *Proceedings of the SPIE* **7092**, pp. 70920I-1 – 70920I-15 (2008).

Technical Reports

1. I. Beichl and R. Boisvert, "Mathematical Foundations of Measurement Science for Information Systems: Report of a Planning Workshop," NISTIR 7465, October 24, 2007.

Visualizations Published

1. H. K. Hung and S. G. Satterfield, "Atomic Scale Simulation of the Electron Distribution in an InAs/GaAs Quantum Dot", Cover Image, *Journal of Research of the National Institute of Standards and Technology* **113** (May-June 2008).
2. W. George, S. Satterfield, M. Olano, and J. Terrill, "Flow of Suspensions", Movie shown at the Argonne National Laboratory booth, Supercomputing 2008, Nov 15-21, 2008, Austin, Texas.

Accepted

1. I. Beichl and B. Cloteaux, "Measuring the Effectiveness of the S-metric to Produce Better Network Models," *Proceedings of the Winter Simulation Conference*, Miami, FL, December 7-10, 2008.
2. A. K. Gaigalas, L. Wang, and F. Hunt, "Interpretation of Measurements of Photochemical Reactions in Focused Laser Beams – Part 2," *Journal of Research of NIST*.
3. D. Gilsinn, B. Balachandran, and T. Kalmar-Nagy, eds., *Delay Differential Equations: Recent Advances and New Directions*, Springer.
4. D. E. Gilsinn, "On Algorithms for Estimating Computable Error Bounds for Approximate Periodic Solutions of an Autonomous Delay Differential Equation," *Communications in Nonlinear Science and Numerical Simulation*.
5. R. Kacker, A. Forbes, R. Kessel, and K.-D. Sommer, "Bayesian Posterior Predictive p-value of Statistical Consistency in Interlaboratory Evaluations," *Metrologia*.
6. P. E. King-Smith, B. A. Fink, J. J. Nichols, K. K. Nichols, R. J. Braun, and G. B. McFadden, "The Contribution of Lipid Layer Movement to Tear Film Thinning and Breakup," *Investigative Ophthalmology and Visual Science*.
7. V. Marbukh and D. Genin, "Towards Understanding Metastability in Large-scale Loss Network with Mobile Users: Emergence and Implications for Performance," in *Proceedings of IEEE Globecom,* New Orleans, December, 2008.
8. S. P. Mates, R. Rhorer, E. Whitenton, T. Burns, and D. Basak, "A Pulse-heated Kolsky Bar Technique for Measuring the Flow Stress of Metals at High Loading and Heating Rates," *Experimental Mechanics*.
9. C. S. Pande, K. P. Cooper, and G. B. McFadden, "Grain Size Distribution in Two Dimensions in the Long Time Limit," *Acta Materialia*.
10. R. Radebaugh, Y. Huang, A. O'Gallagher, and J. Gary, "Calculated Performance of Low-Porosity Regenerators at 4 K with He-4 and He-3," *Proceedings of the 15th International Cryocooler Conference*.
11. J. S. Sims, G. W. Bryant, W. L. George, T. G. Griffin, J. G. Hagedorn, J. T. Kelso, T. M. Olano, A. P. Peskin, S. G. Satterfield, and J. D. Terrill, "Computation and Visualization of Nanostructures and Nano-optics," *Computing and Visualization in Science*.
12. W. Wang, S. Sampath, Y. Lei, and R. Kacker, "An Interaction-based Test Sequence Generation Approach for Testing Web-applications," in *Proceedings of 11-th IEEE High Assurance Systems Engineering Symposium*, Nanjing, China, December 3-5, 2008.

Submitted

1. J. W. Bullard, C. F. Ferraris, E. J. Garboczi, N. S. Martys, P. E. Stutzman, and J. E. Terrill, "Virtual Cement and Concrete" *Innovations in Portland Cement Manufacturing*.
2. A. Carasso, "False Characteristic Functions and Other Pathologies in Variational Blind Deconvolution. A Method of Recovery," *SIAM Journal on Applied Mathematics*.
3. D. L. Cotrell and G. B. McFadden, "Boundary Slip Effects on the Linear Stability of Circular and Spiral Poiseuille Flow," *Journal of Fluid Mechanics*.
4. R. Datla, R. Kessel, A. Smith, R. Kacker, and D. Pollock, "Uncertainty Analysis of Remote Sensing Optical Sensor Data: Guiding Principles to Achieve Metrological Consistency," *International Journal of Remote Sensing*.
5. A. Dienstfrey, T. L. Oreskovic, L. T. Hudson, H. S. Bennett, "Analysis of ISCD-NIST Survey for Bone Health," *Journal of Clinical Densitometry*.
6. M. Elliott, C. Caves, and B. Eastin, "Graphical Description of Pauli Measurements on Stabilizer States," *Physical Review A*.
7. D. E. Gilsinn, "Bifurcations, Center Manifolds, and Periodic Solutions," *Delay Differential Equations: Recent Advances and New Directions*.
8. P. D. Hale, D. F. Williams, A. Dienstfrey, C. M. Wang, A. Lewandowski, T. S. Clement, and D. A. Keenan, "Calibration of Pulsed Waveforms," *IEEE Transactions in Instruments and Measurements*.
9. F. Hunt, A. K. Gaigalas and L. Wang, "Interpretation of Measurements of Photochemical Reactions in Focused Laser Beams-Part 2," *Conference on Cybernetics and Information Technologies, Systems and Applications: CITSA 2008*.
10. R. Kuhn, Y. Lei, and R. Kacker, "Advances in Combinatorial Testing," *IEEE Spectrum*
11. J. Lawrence, R. Kacker, Y. Lei, D.R. Kuhn, and M. Forbes, "Two-valued Covering Arrays," *Designs, Codes and Cryptography*.

12. J. Lawrence, R. Kacker, Y. Lei, D.R. Kuhn, and M. Forbes, "Binary Covering Arrays," *Designs, Codes and Cryptography*.
13. N. S. Martys, W. L. George, S. G. Satterfield, and M. Olano, "Spatial-Temporal Correlations in Start-up Flows of Colloidal Suspensions," *Physical Review Letters*.

Presentations

Invited Talks

1. R. Boisvert, "Succeeding in a Federal Lab," Future Practitioner Colloquium, INFORMS Annual Meeting, Washington, DC, October 11, 2008.
2. A. Carasso, "APEX Blind Deconvolution in Astronomy, Electron Microscopy, and Optical Engineering." 16th Annual Meeting of the Council on Ionizing Radiation Measurements and Standards, NIST, Gaithersburg, MD, October 22, 2007.
3. B. Cloteaux, "Generating Network Models," Center for Computing Sciences, Bowie, MD, August 28, 2008.
4. M. Donahue, "Micromagnetic Modeling of Spin Transfer Torque," University of Alabama, Tuscaloosa AL, November 28, 2007.
5. J. T. Fong, **"**Verification and Uncertainty Estimation of Simulations based on Finite Element Method," Department of Mechanical Engineering Seminar, Stanford University, Stanford, CA, May 7, 2008.
6. S. Glancy, "Diagnosis of Pulsed Squeezing in Multiple Temporal Modes," Quantum Estimation in Theory and Practice Workshop, Perimeter Institute, Waterloo, Canada, September 25-30, 2008.
7. F. Hunt, "Convergence and Sensitivity of a Sequence Alignment Algorithm," Dynamical Systems and Biology Conference, New York University, New York, NY, April 12, 2008.
8. R. Kacker, "Automated Combinatorial Testing for Software," George Washington University Department of Computer Science, Washington, DC, March 10, 2008.
9. D. Lozier, "Numerical Analysis in the NIST Digital Library of Mathematical Functions," International Conference of Numerical Analysis and Applied Mathematics, Kos, Greece, September 18, 2008.
10. W. F. Mitchell, "Strategies for *hp*-Adaptive Refinement," Sixth International Conference of Numerical Analysis and Applied Mathematics, Kos, Greece, September 16-20, 2008.
11. B. Rust, "Statistical Stabilization of Ill-Posed Problems," George Mason University, Dept. of Mathematical Sciences Colloquium, April 11, 2008.

Conference Presentations

1. G. W. Bryant, M. Zielinski, W. Jaskólski, J. Sims, and J. Aizpurua, "Squeezing the Light Out of Quantum Dots Using External Strain," 5th International Conference on Semiconductor Quantum Dots (QD2008), Gyeongju, Korea, May 11-16, 2008.
2. B. Cloteaux, "Generating Network Models Using the S-Metric," International Conference on Modeling, Simulation and Visualization Methods, Las Vegas, NV, July 15, 2008.
3. M. Donahue, "Precession Axis Modification to a Semi-analytical Landau-Lifshitz Solution Technique," 52nd Annual Conference on Magnetism and Magnetic Materials, Tampa FL, November 6, 2007.
4. J. T. Fong, **"**A Design-of-Experiment Approach to Verification and Uncertainty Estimation of Simulations based on Finite Element Method," 2008 Annual Conference of the American Society of Engineering Education, Pittsburgh, PA, June 25, 2008.
5. J. T. Fong, "An Intelligent Flaw Monitoring System: From Flaw Size Uncertainty to Fatigue Life Prediction with Confidence Bounds in 24 Hours," 8th World Congress on Computational Mechanics, Venice, Italy, June 30, 2008.
6. J. T. Fong, "Robust Engineering Design for Failure Prevention," ASME Pressure Vessels & Piping Conference, Chicago, IL, July 31, 2008.
7. J. T. Fong, "A Web-based Data Analysis Methodology for Estimating Reliability of Weld Flaw Detection, Location, and Sizing," ASME Pressure Vessels & Piping Conference, Chicago, IL, July 31, 2008.
8. J. T. Fong, "Structural Aging Monitoring via Web-based Nondestructive Evaluation Technology," ASME Pressure Vessels & Piping Conference, Chicago, IL, July 31, 2008.
9. W. George, J. Lancien, N. Martys, J. Terrill, and E. Garboczi, "Large Scale Simulations of Suspensions," NASA Booth, Supercomputing 2007, Reno, NV, November 10-16, 2007.

10. T. Gerrits, T. Clement, S. Glancy, S. W. Nam, R. Mirin, E. Knill, "Generation of Optical Cat States by squeezed photon subtraction," Southwest Quantum Information Technology Network Workshop, Santa Fe, New Mexico, February 14-17, 2008.
11. S. Glancy, "Calibration for Slightly Unbalanced Homodyne Detection," Southwest Quantum Information Technology Network Workshop, Santa Fe, New Mexico, February 14-17, 2008.
12. J. Kelso and S. G. Satterfield, "Plane Measurement", "DIVERSE Birds-of-a-Feather," ACM SIGGRAPH, Los Angeles, CA, August 11-15, 2008.
13. R. Kessel and R. Kacker, "Framework for Evaluation of Uncertainty with Test of Linearity Using Covering Arrays," Conference on Advanced Mathematical and Computational Tools in Metrology and Testing (AMCTM), Paris, France, June 23-25, 2008.
14. E. Knill, "Building Quantum Computers," International Symposium on Information Theory 2008, Toronto, Canada, July 6-11, 2008.
15. E. Knill, "Experimental Quantum Registers," Gordon Conference on Quantum Information Science, Big Sky, MT, August 31- September 5, 2008.
16. D.W. Lozier, "Asymptotics in the NIST Digital Library of Mathematical Functions," AMS-SIAM Special Session on Asymptotic Methods in Analysis with Applications, Joint Mathematics Meetings, San Diego, CA, January 6, 2008.
17. V. Marbukh and K. Mills, "Demand Pricing and Resource Allocation in Market-Based Compute Grids: A Model and Initial Results," The 7th International Conference on Networking, Cancun, Mexico, April 13-18, 2008.
18. V. Marbukh, "Can TCP Metastability Explain Cascading Failures and Justify Necessity of Flow Admission Control in the Internet?," 15th International Conference on Telecommunications, St. Petersburg, Russia, June 16-19, 2008.
19. V. Marbukh and K. Sayrafian, "A Framework for Joint Cross-Layer and Node Location Optimization in Mobile Sensor Networks," ADHOC-NOW, Nice, France, September 10-13, 2008.
20. V. Marbukh, "Towards Modeling Network Self-Organization/Evolution: Effects of Fairness, Risk Averseness, Competition, and Economic Pressures," 4th European Conference on Complex Systems (ECCS 2008), Jerusalem, Israel. September 14-19, 2008.
21. N. S. Martys, W. L. George, S. G. Satterfield, and M. Olano, "Spatial Temporal Correlations in Start-up Flows of Colloidal Suspensions," ACBM (Advanced Cement Based Materials Consortium) Semi-annual Technical Meeting, Northwestern Univ., Evanston, IL, March 26, 2008.
22. W. F. Mitchell, "A Parallel Multigrid Preconditioner for High-order and hp-Adaptive Finite Elements," 13th SIAM Conference on Parallel Processing for Scientific Computing, Atlanta, GA, March 12-14, 2008.
23. W. F. Mitchell, "Application of a Parallel Adaptive Finite Element Code to Confined Interacting Atoms," SIAM Annual Meeting, San Diego, CA, July 7-11, 2008.
24. M. Olano, "Shaders", "DIVERSE Birds-of-a-Feather," ACM SIGGRAPH, Los Angeles, CA, August 11-15, 2008.
25. B. Rust, "Comparing MSU T2 Time Series with Ground-based Measurements," NOAA-NIST Workshop on Calibration for Climate Quality Time Series, Camp Springs, MD, January 14, 2008.
26. B. Rust, "A Truncated Singular Component Method for Ill-Posed Problems," 15th Conference of the International Linear Algebra Society, Cancun, Mexico, June 16-20, 2008.
27. B. Rust, "Regularization by Residual Periodograms," SIAM Annual Meeting, San Diego, CA, July 7-11, 2008.
28. B. Rust, "Atmospheric Retention of Man-made Carbon Dioxide Emissions," SIAM Annual Meeting, San Diego, CA, July 7-11, 2008.
29. B. Saunders, "Tensor Product B-Spline Mesh Generation for Accurate Visualizations in the NIST Digital Library of Mathematical Functions," Seventh International Conference on Mathematical Methods for Curves and Surfaces, Tonsberg, Norway, June 26 - July 1, 2008.
30. J. Terrill, "Opportunities at NIST", Birds-of-a-Feather: Opportunities in Federal NITRD High End Computing Interagency Working Group (HEC-IWG) Agencies, Supercomputing 2008, Austin, Texas, November 18, 2008.

Patents

1. Alfred S. Carasso, Singular Integral Image Deblurring Method, Patent Number 7,437,012, October 14, 2008.

Software Released

1. DiVisa[14]: a multi-dimensional visualization tool developed for researchers to understand the behavior of their data. Version 08.08.14 (August 8, 2008). – *C. Houard and J. Hagedorn.*
2. NGraph++[15]: a simplified C++ graph library for analysis of complex networks. Versions 0.9, 0.95. – *R. Pozo*.
3. OOF [16]: Object Oriented Finite Elements, modeling of materials with complex microstructures. Versions 2.0.5a1 (November 2007), 2.0.5a2, 2.0.5a3, 2.0.5a4, and 2.0.5a5, 2.0.5a6 (September 2, 2008). – *S. Langer.*
4. PHAML[17]: Parallel Hierarchical Adaptive Multi Level, general purpose software for 2D elliptic partial differential equations. Versions 1.2.0 (October 11, 2007), 1.3.0 (January 24, 2008), and 1.4.0 (April 25, 2008). – *W. Mitchell.*
5. Rendering Error Metrics[18]: methods for assessing geometric errors introduced by the rendering of three 3D geometric primitive forms: points, line segments, and triangles. April 10, 2008. – *J. Hagedorn, J. Terrill, A. Peskin, and J. Filliben.*
6. RENGEN[19]: cryocooler modeling. Version 3.3 (May 6, 2008). – *A. O'Gallagher.*
7. Tcl/Tk[20]: Tool Command Language and Toolkit. Versions 8.5b2 (October 26, 2007), 8.5b3 (November 19, 2007), 8.5.0 (December 20, 2007), 8.5.1 (February 5, 2008), 8.4.18 (February 8, 2008), 8.5.2 (March 28, 2008), 8.4.19 (April 18, 2008), 8.6a1 (June 25, 2008), 8.5.3 (June 30, 2008), 8.5.4 (August 15, 2008), 8.6a2 (August 22, 2008). – *D. Porter.*
8. TNT[21]: Template Numerical Toolkit, an interface for scientific computing in C++. Version 3.0.11. – *R. Pozo*.

[14] http://math.nist.gov/mcsd/savg/software/divisa
[15] http://math.nist.gov/~RPozo/ngraph/index.html
[16] http://www.ctcms.nist.gov/oof/
[17] http://math.nist.gov/phaml/
[18] http://math.nist.gov/mcsd/savg/vis/metrology/
[19] http://math.nist.gov/~AOGallagher/auxfiles/regenstart.html
[20] http://tcl.sourceforge.net/
[21] http://math.nist.gov/tnt/

Conferences, Minisymposia, Lecture Series, Shortcourses

MCSD Seminar Series

1. P. A. Lott (University of Maryland), "Efficient Numerical Simulation of Advection Diffusion Systems," December 19, 2007.
2. A. T. Benjamin (Harvey Mudd College), "Combinatorial Trigonometry (and a method to DIE for)," January 11, 2008.
3. R. Kessel (MCSD), "Evaluation of Uncertainty Associated with the Avogadro Constant," February 19, 2008.
4. A. Knyazev (University of Colorado Denver), "Block Locally Optimal Preconditioned Eigenvalue Xolvers (BLOPEX) Software Package," March 6, 2008.
5. J. Van Deun (Univ. of Antwerp), "Near Best Rational Approximation and Spectral Methods," May 20, 2008.
6. W. Walter (Technical Univ. of Dresden), "Some Thoughts on Modern Computer Arithmetic," August 14, 2008.

Local Events Organized

1. M. Donahue, Member, Local Organizing Committee, Hysteresis and Micromagnetic Modeling Symposium, NIST, Gaithersburg, MD, May 11-14, 2009.
2. D. Gilsinn, Instructor, Shortcourse on Numerical Computing with Matlab, NIST, Gaithersburg, MD, October 3, 4 and November 7, 8, 2007 and January 15, 16, 2008.
3. D. Porter, Secretary, Local Organizing Committee, Hysteresis and Micromagnetic Modeling Symposium, NIST, Gaithersburg, MD, May 11-14, 2009.

External Events Organization

1. R. Boisvert, Member, Program Committee, PSE Workshop08: Innovative and Collaborative Problem Solving Environments in Distributed Resources, at 4th IEEE International Conference on e-Science, Indianapolis, IN, December 7-12, 2008.
2. J. Fong, Organizer, Symposium on Failure Prevention via Robust Design and Continuous NDE Monitoring, at ASME 2008 Pressure Vessels and Piping Conference, Chicago, IL, July 27-31, 2008.

3. J. Fong, Organizer, Symposia on (1) Materials Property Databases, (2) PVP In-service Inspection (ISI) and Failure Event Databases. (NDE), (3) Field & Laboratory Non-destructive Evaluation (NDE) of Residual Stresses, (4) Reliability of Continuous NDE Monitoring and Uncertainty Estimation of Defect Detection, Location, and Sizing, at ASME Pressure Vessels and Piping (PVP) Division Conference, Prague, the Czech Republic, July 26-30, 2009.
4. F. Hunt, Member, Organizing Committee, Dynamical Systems in Biology, New York University, April 11-13, 2008.
5. W. Mitchell and R. Boisvert, Organizers, Minisymposium on PDE Software in Applications, SIAM Annual Meeting, San Diego, CA, July 7-11, 2008.
6. W. Mitchell, Member, Scientific Committee, International Conference of Numerical Analysis and Applied Mathematics (ICNAAM 2008), Kos, Greece, September 16-20, 2008.
7. W. Mitchell, Member, Program Committee, International Conference on High Performance Computing, Networking and Communication Systems (HPCNCS-08), Orlando, FL, July 7-10, 2008.
8. B. Rust and D. O'Leary, Organizers, Minisymposium on Regularization Parameter Selection for Ill-Posed Problems, SIAM Annual Meeting, San Diego, CA, July 7-11, 2008.
9. J. Kelso, S. Satterfield, and M. Olano, Organizers, DIVERSE -Flexible Open Source VE API, Birds-of-a-Feather, ACM SIGGRAPH, Los Angeles, CA, August 11-15, 2008.
10. J. Terrill, Organizer, NIST Exhibit and Demonstration Booth, Supercomputing 2007, Reno, NV, Nov 10-16, 2007.
11. Y. Parker and J. Terrill, Organizers, NIST Exhibit and Demonstration Booth, Supercomputing 2008, Austin, TX, Nov 15-21, 2008.

Other Professional Activities

Internal

1. I. Beichl, Member, NIST Research Advisory Committee.
2. R. Boisvert, Member, NIST Scientific Computing Steering Group.
3. W. Mitchell, Member, ITL Awards Committee.
4. A. Peskin, Member, NIST Research Advisory Committee.
5. B. Rust, Chair, MCSD Seminar Series.
6. B. Rust, Member, ITL Awards Committee.
7. Staff members reviewed multiple manuscripts for the Washington Editorial Review Board (WERB) and the Boulder Editorial Review Board (BERB), as well as proposals for the NIST TIP, SBIR, and Construction Grants programs.

External

Editorial

1. I. Beichl, Editor-in-Chief, *IEEE Computing in Science & Engineering*.
2. R. Boisvert, Associate Editor, *ACM Transactions on Mathematical Software*.
3. R. Boisvert, Area Editor, (Numerical Analysis, Mathematical Software, and Computational Engineering, Finance, and Science), Computing Research Repository (CoRR), www.arXiv.org.
4. D. Gilsinn, Associate Editor, *ASME Journal of Computational and Nonlinear Dynamics*.
5. D. Lozier, Associate Editor, *Mathematics of Computation*.
6. G. McFadden, Associate Editor, *Journal of Crystal Growth*.
7. G. McFadden, Associate Editor, *Interfaces and Free Boundaries*.
8. W. Mitchell, Associate Editor, *International Journal of Applied Mathematics and Computational Science.*
9. W. Mitchell, Associate Editor, *Journal of Numerical Analysis*, *Industrial and Applied Mathematics.*

Boards and Committees

1. B. Alpert, Thesis Committee, Christopher Kurcz, University of Colorado, "Fast convolutions with Helmholtz Green's functions and radially symmetric bandlimited kernels."
2. R. Boisvert, Co-chair, Publication Board, Association for Computing Machinery (ACM).
3. R. Boisvert, Ex-Officio Member, ACM Council.
4. R. Boisvert, Chair, Working Group 2.5 (Numerical Software), International Federation for Information Processing (IFIP).

5. R. Boisvert, Member, ACM Awards Committee.
6. S. Glancy, Member, Science Advisory Council, University of Evansville, IN.
7. D. Lozier, Vice Chair, SIAM Activity Group on Orthogonal Polynomials and Special Functions.
8. B. Miller, Member, Math Working Group, World Wide Web Consortium (W3C).
9. D. Porter, Member, Tcl Core Team.
10. B. Saunders, Member, Selection Committee, Etta Z. Falconer Lecture, Mathfest 2008.
11. J. Terrill, Member, Federal High End Computing Implementation Task Force.
12. J. Terrill, Member, Federal High End Computing Research and Development, and Infrastructure Interagency Working Group.

Reviewing

1. Division staff members referee manuscripts for a wide variety of journals including *ACM Transactions on Mathematical Software*, *Communications in Numerical Methods in Engineering*, *Computer Standards & Interfaces*, *Computing in Science and Engineering*, *Contributions in Nonlinear Science and Numerical Simulations, Environmental Geology, High Performance Computing, Networking and Communication Systems, IEEE Transactions on Information Theory*, *IEEE Transactions on Instrumentation and Measurement*, *IEEE Transactions on Magnetics, Inverse Problems, Journal of Applied Physics, Journal of Crystal Growth, Journal of Physics A: Mathematical and Theoretical*, *Journal of Physics B: Atomic, Molecular, and Optical Physics*, *Journal of Scientific Computing*, *Journal of Vibration and Control*, *Mathematical Reviews*, *Nonlinear Dynamics*, *Physical Review A*, *Physical Review B*, *Physical Review E*, *Physical Review Letters*, *Physics of Fluids*, *Proceedings of the National Academy of Sciences*, *Proceedings of the Royal Society A, Quantum Information & Computation*, *SIAM Journal on Scientific Computing.*
2. Staff members review proposals for the following research programs: DARPA, DoD, DoE, NSF, and the National Security Science and Engineering Faculty Fellowship Program.

External Contacts

MCSD staff members make contact with a wide variety of organizations in the course of their work. Examples of these follow.

Industrial Labs

ActiveState
Aerospace Corp.
Akustik Technologie Gottingen (Germany)
Altair
Ansoft Corporation
Aptech Engineering Services, Inc.
CIC NanoGUNE (Spain)
Core Minutes
Cray, Inc.
Dreamworks
Fakespace
Freescale Semiconductor
Geometry Factory
HydroMeteorlogical Service of Emilia Romagna
JMH Consulting
Maplesoft
Mentor Graphics
Model Integration
National Instruments
OpenTech Inc.
Orbital Sciences Corporation
Quantum Design
Rowan Williams Davies & Irwin Inc.
Setterhold Consulting
Siemens
Thousand Fee Consulting, LLC
Visual Numerics, Inc.
WL/Delft Hydraulics
Zebra Imaging

Government/Non-profit Organizations

Arnold Air Force Base
Association for Computing Machinery
Association for Women in Mathematics
Centrum voor Wiskunde en Informatica (Netherlands)
Ecole Supérieure d'Informatique et Applications de Lorraine (ESIAL) (France)
European Molecular Biology Laboratory (Germany)
French Embassy
IDA Center for Computing Sciences
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
National Center for Atmospheric Research
National Institute of Information and Communications Technology (Japan)
National Reconnaissance Office (NRO)
Naval Research Laboratory
National Science Foundation

Naval Surface Warfare Center
Nuclear Regulatory Commission
Oak Ridge National Laboratory
Sandia national Laboratory
Society for Industrial and Applied Mathematics
Swedish Royal Institute of Technology
US Department of Defense
US Department of Energy

Universities

Aachen University (Germany)
American University of Sharjah (UAE)
Beloit College
Cairo University (Egypt)
California Institute of Technology
Carnegie Mellon University
Case Western Reserve University
City University of Hong Kong
Cornell University
Dalhousie University (Canada)
ESIAL (France)
George Mason University
George Washington University
Hong Kong University of Science & Technology
Indiana University
Iowa State University
Institut Supérieur d'Informatique de Modélisation et de leurs Applications (France)
Jagiellonian University (Poland)
Jacobs University, Bremen (Germany)
Johannes Kepler University (Austria)
Massachusetts Institute of Technology
Macalester College
National Technical University of Athens (Greece)
New York University
Ningbo University (China)
North Carolina State University
Notre Dame University
Pennsylvania State University
Queen Mary University of London
San Diego State University
Stanford Linear Accelerator Center
Stanford University
Technische Universiteit Delft (The Nethelands)
Technische Universitat Munchen (Germany)
Universidade Federal do Rio Grande do Sul (Brazil)
Universiteit van Amsterdam (The Nethelands)
University of Antwerp (Belgium)
University of Bristol (UK)
University of California – Davis
University of Cologne
University of Colorado at Boulder
University of Delaware
University of Duisberg-Essen (Germany)
University of Edinburgh (UK)
University of Erlangen (Germany)
University of Illinois at Urbana-Champaign
University of Karlsruhe (Germany)
University of Kent (UK)
University of Maryland Baltimore County
University of Maryland College Park
University of Memphis
University of Minnesota
University of South Carolina
University of Southampton (UK)
University of St. Andrews (UK)
University of Toronto
University of Washington
University of Wisconsin, Madison
University of Wisconsin, Milwaukee
Utah State University
Washington University in St. Louis
Xi'an University of Science and Technology (China)

Part V

Appendices

Staff

MCSD consists of full time permanent staff located at NIST laboratories in Gaithersburg, MD and Boulder, CO. This is supplemented with a variety of faculty appointments, guest researchers, postdoctoral appointments, and student appointments. The following list reflects all appointments held during any portion of the reporting period.

Division Staff

Ronald Boisvert, *Chief*, Ph.D. (Computer Science), Purdue University, 1979
Robin Bickel, *Secretary*
Jeffrey Fong, Ph. D. (Applied Mechanics and Mathematics), Stanford University, 1966
Roldan Pozo, Ph.D. (Computer Science), University of Colorado at Boulder, 1991
Christopher Schanzle, B.S. (Computer Science), University of Maryland – Baltimore County, 1989

Guest Researchers
Barry Bernstein / Illinois Institute of Technology
Grant Erdmann, Commerce Science Fellow / Air Force Research Laboratory

Mathematical Modeling Group

Geoffrey McFadden, *Leader* , Ph.D. (Mathematics), New York University, 1979
Bradley Alpert (Boulder), Ph.D. (Computer Science), Yale University, 1990
Timothy Burns, Ph.D. (Mathematics), University of New Mexico, 1977
Alfred Carasso, Ph.D. (Mathematics), University of Wisconsin, 1968
Andrew Dienstfrey (Boulder), Ph.D. (Mathematics), New York University, 1998
Michael Donahue, Ph.D. (Mathematics), The Ohio State University, 1991
Fern Hunt, Ph.D. (Mathematics), New York University,
Raghu Kacker, Ph.D. (Statistics), Iowa State University, 1979
Anthony Kearsley, Ph.D. (Computational and Applied Mathematics), Rice University, 1996
Peter Ketcham. M.S. (Mathematics), University of Minnesota, 1997
Stephen Langer, Ph.D. (Physics), Cornell University, 1989
Agnes O'Gallagher (Boulder), M.S. (Applied Math), University of Colorado at Boulder, 1991
Donald Porter, Ph.D. (Electrical Engineering), Washington University, 1996

NRC Postdoctoral Associates
Valerie Coffman, Ph.D. (Physics), Cornell University, 2006
Aaron Lott, Ph.D. (Mathematics), University of Maryland, 2008

Contractors
Andrew C.E. Reid, Ph.D. (Physics), Queen's University, Kingston, Ontario, 1994

Faculty Appointees
Daniel Anderson / George Mason University
Dianne O'Leary / University of Maryland College Park
Florian Potra / University of Maryland Baltimore County

Guest Researchers
Mirit Aladjem / National Institutes of Health
Richard Braun / University of Delaware
David Cotrell / Lawrence Livermore National Laboratory
Gunay Dogan / Cornell University
Michael Forbes / MIT

John Gary / Boulder
Katharine Gurski / George Washington University
Seung-Ill Haan / University of Maryland Baltimore County
Sohyoung Kim / National Institutes of Health
Yu (Jeff) Lei / University of Texas at Arlington

Students
Kevin Dela Rosa / University of Texas at Arlington
Gillian Haemer / University of Southern California
Olga Kuznetsova / University of Maryland College Park

Mathematical Software Group

Daniel Lozier, *Leader,* Ph.D. (Applied Mathematics), University of Maryland, 1979
Marjorie McClain, M.S. (Mathematics), University of Maryland College Park, 1984
Bruce Miller, Ph.D. (Physics), University of Texas at Austin, 1983
William Mitchell, Ph.D. (Computer Science), University of Illinois at Urbana-Champaign, 1988
Bert Rust, Ph.D. (Astronomy), University of Illinois at Urbana-Champaign, 1974
Bonita Saunders, PhD (Mathematics), Old Dominion University, 1985

Faculty Appointees
Frank Olver / University of Maryland College Park
G.W. Stewart / University of Maryland College Park
Abdou Youssef / George Washington University

Guest Researchers
Leonard Maximon / George Washington University

Students
Liuyuan Chen / Montgomery Blair High School, Maryland

Discrete Mathematical Analysis Group

Ronald Boisvert, *Acting Leader*
Isabel Beichl, Ph.D. (Mathematics), Cornell University, 1981
Javier Bernal, Ph.D. (Mathematics), Catholic University, 1980
Brian Cloteaux, Ph.D. (Computer Science), New Mexico State University, 2007
David Gilsinn, Ph.D. (Mathematics), Georgetown University, 1969
Scott Glancy (Boulder), Ph.D. (Physics), University of Notre Dame, 2003
Emanuel (Manny) Knill (Boulder), Ph.D., (Mathematics), University of Colorado at Boulder, 1991
Vladimir Marbukh, Ph.D. (Mathematics) Leningrad Polytechnic University, 1986

NRC Postdoctoral Associates
Bryan Eastin, Ph.D. (Physics), University of New Mexico, 2007

Faculty Appointees
Saul Gass / University of Maryland College Park
James Lawrence / George Mason University

Guest Researchers
Stephen Bullock / IDA Center for Computing Sciences
Sita Ramamurti / Trinity University, DC
David Song / Korea Institute for Advanced Study
Francis Sullivan / IDA Center for Computing Sciences
Christoph Witzgall / NIST Scientist Emeritus

Students

Adam Meier / University of Colorado
Amelia Tebbe / St. Mary's College
Yanbao Zhang / University of Colorado

Scientific Applications and Visualization Group

Judith Devaney, *Leader*, Ph.D. (Information Technology), George Mason University, 1998
Yolanda Parker, *Office Manager*
Robert Bohn, Ph.D. (Physical Chemistry), University of Virginia, 1991
William George, Ph.D. (Computer/Computational Science), Clemson University, 1995
Terence Griffin, B.S. (Mathematics), St. Mary's College of Maryland, 1987
John Hagedorn, M.S. (Mathematics), Rutgers University, 1980
John Kelso, M.S. (Computer Science), George Washington University, 1984
Adele Peskin (Boulder), Ph.D. (Chemical Engineering), University of Colorado at Boulder, 1985
Steven Satterfield, M.S. (Computer Science), North Carolina State University, 1975
James Sims, Ph.D. (Chemical Physics), Indiana University, 1969

Faculty Appointees

Marc Olano / University of Maryland Baltimore County

Guest Researchers

Aurelei Canale / France
Dong Yeon Cho / Seoul National University, Korea
Edith Enjolras / Université Blaise Pascal, France
Cedric Houard / Université Blaise Pascal, France
Julien Lancien / France
Audrey Lemoussu / France
Alexandre Marie / France

Students

William Hess / Montgomery Blair High School
Christopher Leedy / University of Kentucky
Kevin Rawlings / Montgomery Blair High School
Axel Rivera Rodriguez / University of Puerto Rico at Humacao

Glossary of Acronyms

ACI	American Competitiveness Initiative
ACM	Association for Computing Machinery
ANSI	American National Standards Institute
API	application programmer's interface
AQM	active queue management
AS	autonomous system
ASME	American Society of Mechanical Engineers
AWM	Association for Women in Mathematics
BFRL	NIST Building and Fire Research Laboratory
BLAS	Basic Linear Algebra Subprograms
CCG	DoD Coordinated Calibration Group
CiSE	Computing in Science & Engineering
CMM	coordinate measuring machine
CMU	Carnegie Mellon University
CNST	NIST Center for Nanoscale Science and Technology
CODATA	Committee on Data for Science and technology
CPU	central processing unit
CRU	Climate Research Unit (UK)
CSTL	NIST Chemical Science and Technology Laboratory
CT	computed tomography
CWI	Centrum voor Wiskunde en Informatica (Amsterdam)
DARPA	DOD Defense Advanced Research Projects Agency
DIVERSE	Device Independent Virtual Environments — Reconfigurable, Scalable, Extensible (visualization software)
DLMF	Digital Library of Mathematical Functions
DOD	U.S. Department of Defense
DOE	U.S. Department of Energy
DOE	design of experiments
DPD	dissipative particle dynamics
DXA	dual-energy x-ray absorptiometry
ECN	explicit congestion notifications
EEEL	NIST Electronics and Electrical Engineering Laboratory
EM	electromagnetic
EM	expectation maximization
EMC	electromagnetic compatibility
ESCMSE	European Society of Computational Methods in Science and Engineering
FDA	Food and Drug Administration
FDS	Fire Dynamics Simulator
FEM	finite element method
FFT	fast Fourier transform
FY	fiscal year
GAMS	Guide to Available Mathematical Software
GCV	generalized cross validation
GPU	graphics processing unit
GUI	graphical user interface
HEC	high-end computing
HPC	high performance computing
HTML	hypertext markup language
Hy-CI	Hylleraas-Configuration Interaction technique
ICIAM	International Congress on Industrial and Applied Mathematics
IDA	Institute for Defense Analysis
IEEE	Institute of Electronics and Electrical Engineers
Iid	idenpendent and identically distributed
IFIP	International Federation for Information Processing

IML++	Iterative Methods Library
IMS	Innovations in Measurement Science
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
IQR	inter-quartile range
ISIMA	Institut Supérieur d'Informatique de Modélisation et de leurs Applications
IT	information technology
ITL	NIST Information Technology Laboratory
JAMA	Java Matrix package
MALDI-TOF	matrix-assisted laser desorption/ionization time-of-flight
MCMC	Markov chain Monte Carlo
MCSD	ITL Mathematical and Computational Sciences Division
MEL	NIST Manufacturing Engineering Laboratory
MICS	medical implant communication system
MIT	Massachusetts Institute of Technology
MKM	mathematical knowledge management
MMM	magnetism and magnetic materials
MPI	Message Passing Interface
MRAM	magneto-resistive random access memory
MRI	magnetic resonance imaging
MSEL	NIST Materials Science and Engineering Laboratory
MSU	microwave sounding unit
MV++	Matrix/Vector Library
μmag	Micromagnetics Activity Group
NASA	National Aeronautics and Space Administration
NBS	National Bureau of Standards
NIH	National Institutes of Health
NIST	National Institute of Standards and Technology
NISTIR	NIST Internal Report
NITRD	Networking and Information Technology Research and Development
NMR	nuclear magnetic resonance
NRC	National Research Council
NSF	NOAA's National Science Foundation
NUM	network utility maximization
NWS	National Weather Service
OAR	NOAA's Oceanic and Atmospheric Research Office
OASCR	DOE Office of Advanced Scientific Computing Research
ODE	ordinary differential equation
OLES	NIST Office of Law Enforcement Standards
OOF	Object-Oriented Finite Elements
OOMMF	Object-Oriented Micromagnetic Modeling Framework
OSPF	open shortest path first
PDE	partial differential equation
PET	positron emission tomography
PHAML	Parallel Hierarchical Adaptive Multi Level (software)
PCAST	President's Council of Advisors on Science and Technology
PITAC	President's Information Technology Advisory Committee
PL	NIST Physics Laboratory
PREP	Professional Research Experience Program
QDPD	quarternion-based dissipative particle dynamics
QKD	quantum key distribution
RAVE	Reconfigurable Automatic Virtual Environment
R&D	research and development
RF	radio frequency
RSS	received signal strength
SAVG	MCSD Scientific Applications and Visualization Group
SciDAC	Scientific Discovery Through Advanced Computing
SECB	slow evolution from the continuation boundary

SEM	scanning electron microscope
SIAM	Society for Industrial and Applied Mathematics
SIGGRAPH	ACM Special Interest Group on Graphics
SIMA	NIST Systems Integration for Manufacturing Applications Program
SPEC	System Performance Evaluation Corporation
SPIE	International Society for Optical Engineering
SRM	standard reference material
SSD	sum of squared deviations
SURF	Student Undergraduate Research Fellowship
SVD	singular value decomposition
TCP	transmission control protocol
TNT	Template Numerical Toolkit
UCLA	University of California at Los Angeles
UNESCO	United Nations Educational, Scientific, and Cultural Organization
UMBC	University of Maryland Baltimore County
UMCP	University of Maryland College Park
URL	universal resource locator
VCCTL	Virtual Cement and Concrete Testing Laboratory
VRML	virtual reality modeling language
W3C	World Wide Web Consortium
XML	Extensible Markup Language