Part II

Project Summaries

Applied Mathematics

Time-Domain Algorithms for Computational Electromagnetics

Bradley Alpert
Andrew Dienstfrey
Leslie Greengard (New York University)
Thomas Hagstrom (University of New Mexico)

http://www.boulder.nist.gov/itl/div891/alpert/td-em/index.html

Acoustic and electromagnetic waves, including radiation and scattering phenomena, are increasingly modeled using time-domain computational methods, due to their flexibility in handling wide-band signals, material inhomogeneities, and nonlinearities. For many applications, particularly those arising at NIST, the accuracy of the computed models is essential. Existing methods, however, typically permit only limited control over accuracy; high accuracy generally cannot be achieved for reasonable computational cost.

Applications that require modeling of electromagnetic (and acoustic) wave propagation are extremely broad, ranging over device design, for antennas and waveguides, microcircuits and transducers, and low-observable aircraft; nondestructive testing, for turbines, jet engines, and railroad wheels; and imaging, in geophysics, medicine, and target identification. At NIST, applications include the modeling of antennas (including those on integrated circuits), waveguides (microwave and photonic), transducers, and in nondestructive testing.

The objective of this project is to advance the state of the art in electromagnetic computations by eliminating three existing weaknesses with time-domain algorithms for computational electromagnetics to yield: (1) accurate nonreflecting boundary conditions (that reduce an infinite physical domain to a finite computational domain), (2) suitable geometric representation of scattering objects, and (3) high-order convergent, stable spatial and temporal discretizations for realistic scatterer geometries. The project is developing software to verify the accuracy of new algorithms and reporting these developments in publications and at professional conferences.

Riccati Equations and Inverse Scattering. Much work of the project this year has been devoted to understanding the scattering and inverse scattering implications of new representations for acoustic waves in unbounded domains. These representations, developed with Yu Chen (Courant Institute, New York University), by which scattered waves are represented in the frequency domain with nearly optimal efficiency on a line (in two dimensions) or plane (in three), are suggestive of a new approach to reconstruction--- or imaging---of acoustic scatterers, based on layer stripping. The current state of the art in acoustic inverse scattering is fully successful only for weak scattering, such as is used for fetal ultrasound imaging, a problem that is linearized by the Born approximation. For higher contrast media, such as those presented by human tissue with bones or by manufactured components, the problem is highly nonlinear and existing approaches do not scale well in either efficiency or accuracy. While practical inverse scattering in this general setting remains a highly speculative undertaking, it is hoped that the layer stripping approach will lead both to reasonable computational complexity, to manage costs, and to systematic regularization, to control errors, for large-scale problems.

This year a Riccati equation was developed to construct the scattering matrix, incrementally by layers, from a scatterer, and was presented in a seminar, "Representations of Waves in Unbounded

Domains, Scattering, and a Riccati Equation," at the DARPA Applied and Computational Mathematics Program, Virtual Electromagnetic Testrange program review, August 18-20, 2003, St. Louis, Missouri. Subsequently another Riccati equation was developed that is expected to more directly allow the incorporation of scattering causality constraints.

Periodic Structures and Novel Materials. Another continuing focus of the project is the solution of periodic acoustic and electromagnetic scattering problems. A suite of codes was developed this year for computing the optical scattering of an arbitrary 2D array of dielectric/paramagnetic spheres. The crux of this research lies in the computation of the doubly infinite sums of electromagnetic vector multipoles—the so-called "lattice sums." Vector multipoles appear, for example, in Mie's textbook solution for electromagnetic scattering by spheres. The research described here may be viewed as a periodization of that classical work.

Initial investigation into the computation of the requisite periodization operators proceeded from the plane wave expansion formulas for the vector multipole translation operators. However, further investigation revealed that this approach suffers from an instability that will ultimately limit the achievable accuracy. A new idea for circumventing this difficulty is to recompute the lattice sums using the ideas of Ewald in his seminal 1921 work on periodic electrostatics. This approach is currently being pursued.

An immediate motivation for the computations is that they may serve as an aid and guide to NIST researchers in the Electromagnetics Division who are attempting to model and fabricate thin-film meta-materials. This is a class of materials whose relative length-scales lie somewhere between the extremes of homogenization theory and that of photonic-crystals. In the former, the wavelength of light is sufficiently long relative to the characteristic size of the scatterers so that the precise micro-geometry may be replaced by a homogeneous material with averaged (effective) parameters. These, in turn, may be used to estimate reflection and transmission properties of the composite medium. In the "photonic-crystal" regime, the scatterers are at least comparable in size to a single wavelength of light. Consequently no homogenization theory is available at this time. It is felt that somewhere in between these two extremes, scatterers may continue to interact resonantly giving rise to some of the exotic properties usually relegated to photonic crystals. Nevertheless, extensions of ideas from homogenization theory may apply which could result in scattering properties that are more readily characterized and controlled. It is hoped that resolved computation of optical scattering for these geometries will validate and/or help refine this type of modeling effort.

Recent Publications. The paper, "Causal Characteristic Impedance of Planar Transmission Lines," Dylan F. Williams, Bradley K. Alpert, Uwe Arz, David K. Walker, and Hartmut Grabinski, *IEEE Transactions on Advanced Packaging* 26 (2), 165-171 (2003), which presents earlier NIST collaborative research, appeared. In addition, another paper presenting recent NIST collaborative research, "Near-Field, Spherical-Scanning Antenna Measurements With Nonideal Probe Locations," R. C. Wittmann, B. K. Alpert, and M. H. Francis, *IEEE Transactions on Antennas and Propagation*, was submitted for publication.

Impact. The work of the project is supported in part by the Defense Advanced Research Projects Agency (DARPA). The work has been recognized by researchers developing methods for computational electromagnetics (CEM) and has influenced work on these problems at Boeing and HRL (formerly Hughes Research Laboratories). It has also influenced researchers at Yale University and University of Illinois. In each of these cases, new research in time-domain CEM is exploiting discoveries of the project. In particular, some efforts for the new DARPA program on Virtual Electromagnetic Testrange (VET) are incorporating these developments. We expect that design tools for the microelectronics industry and photonics industry, which increasingly require accurate electromagnetics modeling, will also follow.

Solidification Modeling

Geoffrey McFadden
William Boettinger (NIST MSEL)
Sam Coriell (NIST MSEL)
Katherine Gurski (George Washington U.)
Daniel Josell (NIST MSEL)
Michael Miksis (Northwestern U.)
A. Roytbird (U. of Maryland)
Robert Sekerka (Carnegie Mellon U.)
K. Thornton (Northwestern U.)
Peter Voorhees (Northwestern U.)
James Warren (NIST MSEL)

In collaboration with J. Guyer, W. Boettinger, and J. Warren in the Metallurgy Division of MSEL, G. McFadden has developed a novel approached to the solution of boundary value problems in electrochemistry. The underlying model employs a diffuse interface formulation of the phase boundary that separates the electrode and the electrolyte in a multi-component electrochemical cell. The resulting phase-field model is able to incorporate a description of electrocapillarity of equilibrium interfaces, and also captures the correct kinetic behavior of a growing or corroding interface in the presence of currents. An adaptive spectral method was produced to help resolve the electrical double layer (see figure) that occurs at the interface between the electrolyte and electrode. The resulting numerical solutions were able to provide guidance in performing a matched asymptotic expansion of the solution in the limit of vanishing interface width. This expansion is useful in relating the new model to classical sharp-interface models of the double layer. The phase-field model exhibits an unusual expression for the surface energy of the electrode-electrolyte interface which is not positive definite, unlike other phase-field models for multi-component systems for neutral systems. The expression for the surface energy was checked by extending the code to compute equilibria in a cylindrical geometry, where the effects of capillarity can be compared with more traditional sharp-interface models. The results confirm that the expression for surface energy in the phase-field model reproduces the expected results obtained with sharp-interface methods, providing confidence in the phase-field approach. A two-part paper on the work has been accepted for publication in *Physical Review E*, and several invited and contributed talks on the subject were given during the past year.

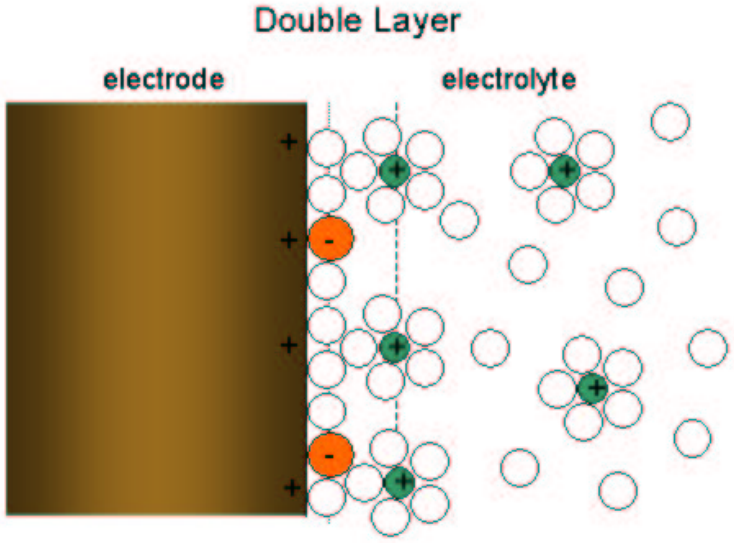

Illustration of an electrical double layer occurring at the interface between electrolyte and electrode

In related work, G. McFadden, in collaboration with S. Coriell, T. Moffat, D. Josell and colleagues in MSEL performed a linear stability for a model of electrodeposition. The model involves a three-component alloy in which the interface evolution is mediated by a surface catalyst, whose coverage of the interface is modeled explicitly. The linear stability results are consistent with some experimental observations of Josell and Moffat during the process of superconformal deposition. This process is important for the processing of interconnects in microelectronic devices. The research was published recently in the *Journal of the Electrochemical Society*. Currently this work is being extended by the numerical computation of finite amplitude solutions that complement the linear stability analysis. A manuscript describing the current work is now in preparation.

In collaboration with K. Gurski, formerly an NRC Postdoctoral Fellow in MCSD and now on the faculty at George Washington University, G. McFadden recently completed work on the instability of quantum wires due to anisotropic surface energy. This work is an extension of classical results due to Lord Rayleigh describing the instability of liquid jets with isotropic surface energies. A paper describing the work appeared this year in the *Proceedings of the Royal Society of London*. More recently, in collaboration with M. Miksis of Northwestern University, they worked on a model of the instability of nanoscale quantum wires that are in contact with an underlying substrate. The work extends the previous work by McFadden and Gurski on the instability of isolated wires, and also extends previous work by Miksis et al. on the effects of a substrate on the instability for isotropic surface energies. This work is in support of an NSF Nanotechnology and Interdisciplinary Research Team grant to NWU.

G. McFadden, in collaboration with K. Thornton and P. Voorhees, Northwestern University, A. Roitbird, University of Maryland, and J. Slutsker, J. Warren, and W. Boettinger of MSEL, has studied phase-field models of elasticity in a two-phase system. The models feature a diffuse interface structure that may prove useful in providing a useful model of surface stress, which becomes important in thin film deposition and other current applications in nanotechnology. A manuscript is in preparation on this work, which describes how the phase-field model includes the effects of elasticity on the surface energy in the system, and also provides a natural model of the surface stress in the system in the spirit of classical work by J.W. Gibbs. Gibbs developed a notion of surface excess quantities to describe the thermodynamics of surfaces, and in other contexts, phase-field models have been effective in modeling these concepts. A talk on this subject was given recently at the 2nd Annual Workshop on the Evolution and Self-Assembly of Quantum Dots held at Northwestern University.

In collaboration with S. Coriell, W. Boettinger, and J. Warren of MSEL, and R. Sekerka of Carnegie Mellon University, G. McFadden has been working on models of diffusion in solids that include the effects of vacancies and stress, which is an area of increasing interest in MSEL. Some progress in formulating models of these processes has been made this year, and the modeling has been applied to a description of the stress-induced expansion of diffusion couples in a binary alloy system. This work is motivated by published experimental measurements that measure the deformation during the diffusion process. A simplified model that we have been able to solve analytically seems to provide qualitatively reasonable description of the deformation during the diffusion process.

Materials Data and Metrology for Machining Simulation

Timothy Burns
Debasis Basak (NIST MSEL)
Gerald Blessing (NIST MEL)
Matthew Davies (U. of N. Carolina)
Brian Dutterer (NIST MEL)
Richard Fields (NIST MSEL)
Michael Kennedy (NIST MEL)
Lyle Levine (NIST MSEL)
Richard Rhorer (NIST MEL)
Eric Whitenton (NIST MEL)
Howard Yoon (NIST PL)

Among large U.S. manufacturing companies such as The Boeing Corporation, the use of finite-element software packages to simulate high-speed material removal processes is emerging as a potentially useful aid in the design and optimization of machining operations. Compared to older analytical approaches to the scientific study of metal-cutting operations, finite-element methods require fewer assumptions and simplifications. In addition, they provide high-resolution, and apparently highly realistic, visualizations of physical properties of materials in the vicinity of the tool-chip interface, such as stress, strain, and temperature fields. Unfortunately, there are insufficient material response constitutive data for use as input into these sophisticated simulation packages, and it is also uncertain whether or not the data that are currently available can be extrapolated reliably to the very high-strain-rate, high-temperature conditions that occur during a high-speed machining process. Furthermore, there are few experimental results that can be used to check the accuracy of finite-element predictions of the extreme conditions that are present during routine operations in modern automatically controlled machining centers. Also, metrology techniques and standards for developing such data are not well established.

To address some of these issues, MCSD has been involved in a long-term collaboration with the NIST Manufacturing Engineering Laboratory (MEL) on problems related to the modeling, simulation, and metrology of high-speed machining processes. Over the past three years, major funding has been provided by the NIST Intramural Advanced Technology Program for a multi-laboratory effort, involving ITL, MSEL, PL, and MEL researchers, to develop new experimental techniques for obtaining the material response data that are required for finite-element simulations of high-rate manufacturing operations such as milling and turning. The work has involved the development of a fairly standard high-strain-rate mechanical testing device called a split-Hopkinson, or Kolsky bar. What has made the device unique has been the adaptation of an existing NIST capability for controlled, very rapid preheating of the metal test specimens, together with some sophisticated non-contact thermometry for monitoring and controlling the test system. With this new experimental system, mechanical testing of material response can take place for the first time under conditions that approximate those that occur during high-speed machining, where a thin region of material is sheared so rapidly that it is heated from room temperature to near its melting point in a fraction of a second.

The results of our ATP-supported work to date confirm the hypothesis that we started with in this project: the rapid heating that occurs during high-speed machining of common steels of interest in manufacturing gives a harder stress-strain response than has been previously measured by dynamic testing, where the material has been pre-heated much more slowly. This result has important implications for finite-element simulations of machining processes. While portions of this work have been documented in papers and conference proceedings, we are in the process of preparing detailed reports on the new testing methods and material response data. Although the focus of this project has been on applications to high-speed machining, the new experimental testing capability will be broadly applicable to other areas of scientific interest, such as studies of dislocation dynamics.

Mathematical Problems in Construction Metrology

Javier Bernal
David Gilsinn
Christoph Witzgall
Geraldine Cheok (NIST BFRL)
William Stone (NIST BFRL)

During the past decade, laser-scanning technology has developed into a major vehicle for wide-ranging applications such as cartography, bathymetry, urban planning, object detection, and dredge volume determination, just to name a few. The NIST Building and Fire Research Laboratory is actively investigating the use of such technology for monitoring construction sites. Here laser scans taken from several vantage points are used to construct a surface model representing a particular scene. Another aspect of the project envisions that CAD-generated geometry sets will be transformed into a library of 3D construction site objects, which will serve multiple purposes during construction. The objects would be loaded into an augmented simulation system that tracks both equipment and resources based on real-time data from the construction site obtained from laser scans. With some additional enhancements, the end result will be a world model of the site, in which as-built conditions can be assessed, current construction processes can be viewed as they occur, planned sequences of processes can be tested, and object information can be retrieved on demand. A project could be viewed and managed remotely using this tool. Pick-and-place control of construction site objects is another application of the 3D construction site object library. With automation and robotics entering on the construction site scene, vision systems, such as LADAR (laser distance and ranging), will be incorporated for real time object identification, using the 3D library of templates. Once objects, such as I-beams, are located and their pose determined, robotic crane grippers can be manipulated to acquire the I-beam.

LADAR technology is currently being tested for locating equipment on construction sites. LADAR scans of I-beams in various poses have been generated. The LADAR scans generate a large

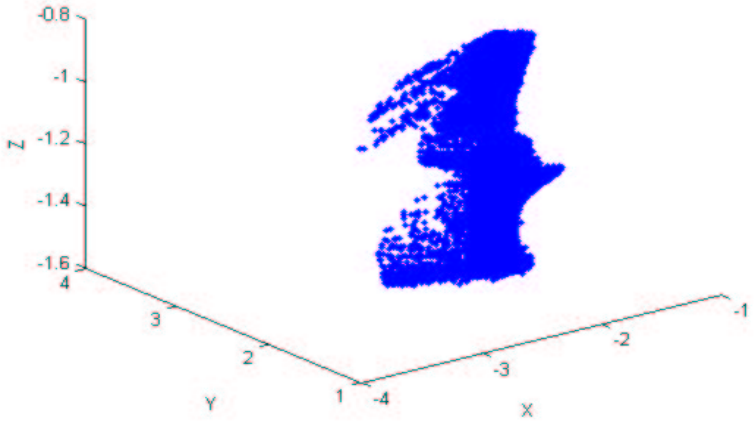

This figure shows the raw data acquired by a LADAR scanning from a position above an I-beam approximately along the diagonal from X = -1 to Y = 4. Notice the phantom pixels caused by scanner beam splitting at the top flange edges and the sides of the middle vertical web. Floor hits are shown in front and to the side of the I-beam. The view is from the end of the I-beam along the major axis.

number of noisy data values, however, arising from phenomena from floor hits to phantom pixels caused by beam splitting at sharp object edges. Preprocessing algorithms are currently being developed to clean the scanned data. (See the accompanying figures.) These include both Triangulated Irregular Networks (TIN) techniques for surface triangulation as well as using 3D binning as methods of reducing the large data set to manageable voxel elements. Joining center-of-gravity estimates with principle components analysis has shown potential for identifying the scanned target's location and pose. These provide the first steps needed to apply an iterated closest point (ICP) algorithm for matching a 3D object from the library against the scanned target image.

Objects in the 3D construction site library can be modeled by various means including multivariate splines and finite elements. Traditional least squares fitting algorithms can lead to ringing at sharp edges. A new minimum energy principle has been developed to fit Clough-Tocher finite elements to objects with sharp edges in such a way that ringing is suppressed. The minimization algorithm incorporates Gauss-Seidel relaxation along with a reweighting and regularization strategy.

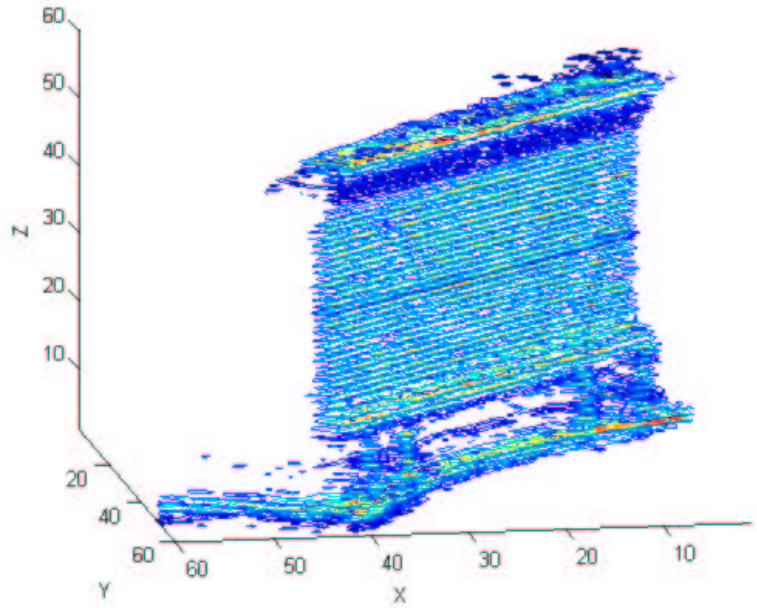

This is a contour plot at each z level of the raw data combined into a reduced number of 3D bins. The binned data has been partially filtered. Scanner beam splitting at sharp edges is still detectable as well a floor hits. Further filtering rules are being developed. Notice that the I-beam is sitting on four blocks that are partially distinguishable.

Micromagnetic Modeling

Michael Donahue
Donald Porter
Robert McMichael (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 MSEL, PL, and EEEL, to improve the state-of-the-art in micromagnetic modeling.

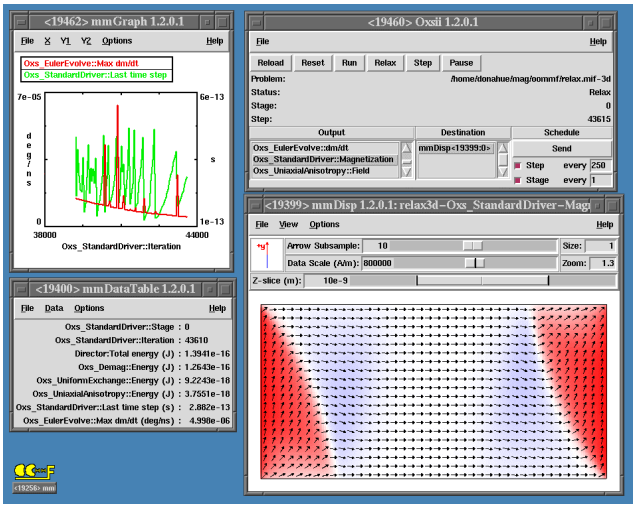

Michael Donahue and Donald Porter in MCSD have developed a widely used public domain computer code for doing computational micromagnetics, the Object-Oriented Micromagnetic Modeling Framework (OOMMF). OOMMF serves as an open, well-documented environment in which algorithms can be evaluated on benchmark problems. OOMMF also provides a fully functional micromagnetic modeling system, handling both two and three-dimensional problems, with sophisticated extensible input and output mechanisms. OOMMF has become an invaluable tool in the magnetics research community. In fiscal year 2003 alone, the software was downloaded more than two thousand times, and use of OOMMF was acknowledged in more than 40 peer-reviewed journal articles.

OOMMF is part of a larger activity, the Micromagnetic Modeling Activity Group (muMAG), formed in 1995 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. Discussion on a new problem, involving thermal effects, is planned for a muMAG workshop to be held at the Joint MMM/Intermag international magnetics conference in January 2004.

In large devices, random thermal effects tend to be self-canceling, but as device size decreases thermal effects grow in relevance. This is especially true in high sensitivity low field magnetic sensors, where performance is generally limited by thermal noise. A NIST Competence award was made this year to a four-laboratory project (EEEL, MSEL, PL and ITL) to design a new generation of such sensors. MCSD staff will be bringing thermal effects modeling into OOMMF as part of this project.

The MCSD micromagnetic project produced four journal papers and four conference presentations in the fiscal year 2003, and earned M. Donahue and D. Porter a NIST Bronze Medal in December 2002, an honor they shared with R. McMichael in MSEL.

A Direct Procedure for Classifying Image Smoothness

Alfred S. Carasso

Most commonly occurring images *f(x,y)* are not differentiable functions of the variables *x* and *y*. Rather, these images display edges, localized sharp features, and other fine-scale details or *texture*. Many digital image-processing tasks require prior specification of the correct mathematical function space in which the true image lies. If an image is incorrectly postulated to be too smooth, the processing algorithm may produce an overly smoothed version of the true image in which critical information has been lost. Examples of this problem might be the unsuspected smoothing out of small lesions in a denoised mammogram image or the unintended digital elimination of the heat-shield tiles in a deblurred image of the Columbia space shuttle.

During the last 10 years, a very considerable amount of image analysis research has been based on the assumption that most images belong to the space of functions of bounded variation. However, it has been subsequently discovered that such so-called total variation image processing sometimes results in unacceptable loss of fine-scale information. This phenomenon is now known as the *staircase effect*. In papers published in 2001, French researchers Gousseau, Morel, and Meyer, showed that most natural images are, in fact, not of bounded variation, and that total variation image-processing techniques must inevitably smooth out texture.

Correct characterization of the lack of smoothness of images is a fundamental problem in image processing. It turns out that so-called *Lipschitz spaces* are the appropriate framework for accommodating non-smooth images. The L^P Lipschitz exponent α for the given image, where $0 < \alpha < 1$, measures the fine-scale content of that image, provided the image is relatively noise free. Heavily textured imagery has low values for α, while large values of α indicate that the image is relatively smooth. Estimating an image's Lipschitz exponent is a delicate problem. Prior art has been based on the technique of lossy wavelet compression, which requires substantial expertise in wavelet theory, fast wavelet algorithms, and image compression. A. Carasso of MCSD has developed a new computational technology for estimating α that requires no knowledge of wavelet techniques or of image compression. It merely requires blurring the image by convolution with a specific kernel, and evaluating the discrete L^P norm of the difference between the blurred and original images. The rate at which this L^P norm tends to zero, as the kernel approaches the Dirac δ-function, is directly related to the Lipschitz exponent α. Thus, define the linear operator U^t on $L^P(R^2)$, $t > 0$, by

$$U^t f = \int \psi(u,v,t)\, f(x-u, y-v)\, du\, dv, \tag{1}$$

where $\psi(x,y,t)$ is the singular Poisson kernel. Define $\mu(t)$ by

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

Given a digitized image *f(x,y)* whose smoothness needs to be determined, the procedure consists in calculating $\mu(t)$ for a sequence of *t*-values tending to zero. Deep mathematical theorems on singular integral operators are used to relate the Lipschitz exponent α to the slope of the graph of log $\mu(t)$ versus log *t*, after excluding very small values of *t*. Using commonly available fast Fourier transform (FFT) algorithms, the required convolution in Eq. 1 can be accomplished by multiplying two functions. Very minimal computational effort is thus needed to implement this singular integral procedure. In addition, Carasso's approach has the advantage of allowing consideration of substantially wider Lipschitz spaces than is mathematically possible with existing procedures, thereby encompassing a much wider class of images.

Figure 1 is an example of the use of this procedure to estimate the Lipschitz exponent in the Mandrill image. The solid curve *A* is a plot of $\mu(t)$ versus *t* on a log-log scale. The dashed straight

line Γ asymptotes the behavior for very small values of t. This is the region that must be excluded from consideration. The dashed straight line Σ represents true behavior in the Mandrill image. The slope of that line yields the Lipschitz exponent α.

Figure 2 is especially instructive. The original Whirlpool galaxy image on the left has an L^1 exponent $\alpha = 0.530$. The APEX enhanced image on the right displays considerable fine-scale structure, and has a substantially smaller L^1 exponent $\alpha = 0.239$. The corresponding L^1 Poisson traces, solid lines A and B, have identical slopes for very small values of t, confirming the fact that behavior at very small values of t is artificial, is not indicative of image smoothness, and must be excluded from consideration. Rather, it is the dashed lines Σ_A and Σ_B, with their widely differing slopes that distinguish the different smoothness characteristics in these two images.

The nine images in Figure 3 and Table 1 form an interesting collection that includes natural as well as man made objects, exhibiting a wide range of sizes. The last row contains a nanoscale electron microscopy micrograph, a galactic scale object, and a cosmological scale structure. For this diverse collection, it is found that α lies in the range $0.2 < \alpha < 0.7$, in either L^1 or L^2. Notice also that the values of the constant C in Table 1 are confined to a narrow range.

Significant potential applications of this technology include the routine monitoring of image sharpness and imaging performance in electro-optical imaging systems, the performance evaluation of image reconstruction software, the detection of possibly abnormal fine-scale features in some medical imaging applications, and the monitoring of surface finish in industrial applications. In addition, as will be shown next, specifying the correct Lipschitz space wherein an image lies is vital for correctly solving the ill-posed image-deblurring problem.

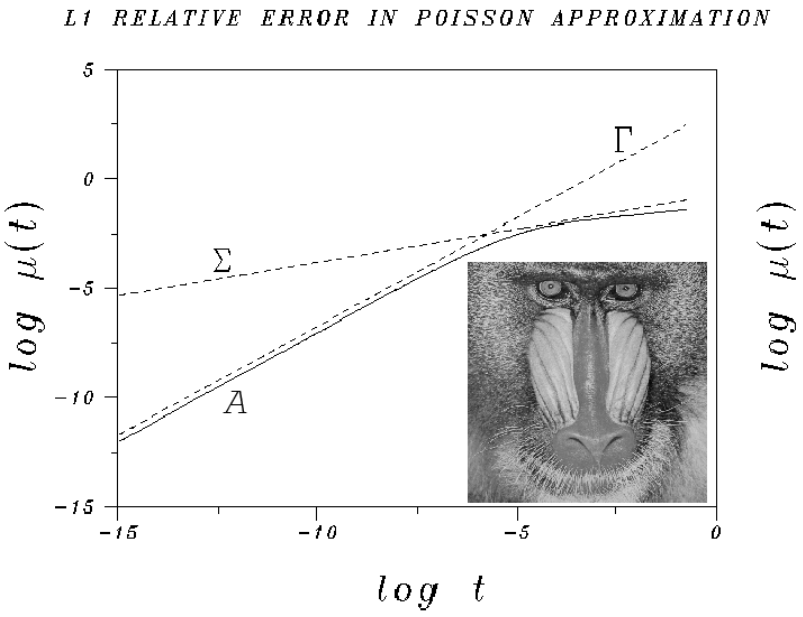

Figure 1. 512×512 Mandrill image is not of bounded variation. Solid curve A is a plot of $\mu(t) = \| U^t f - f \|_1 / \| f \|_1$ versus t on a log-log scale. Majorizing dashed straight line Γ, defined by $\log \mu(t) = 3.2 + 0.994 \log t$, accurately captures linear behavior for very small values of t, but is grossly inaccurate at larger values of t. Such linear behavior at very small t is a misleading finite dimensionality artifact, and is not indicative of true image smoothness. Majorizing dashed straight line Σ, defined by $\log \mu(t) = -0.75 + 0.306 \log t$, accurately reflects behavior for $-6 \le \log t \le -1$, while being grossly inaccurate at very small t. Behavior along Σ indicates true image smoothness. Slope of line Σ implies L^1 Lipschitz exponent $\alpha = 0.306$ for Mandrill image, while $\alpha = 1.0$ is required for bounded variation.

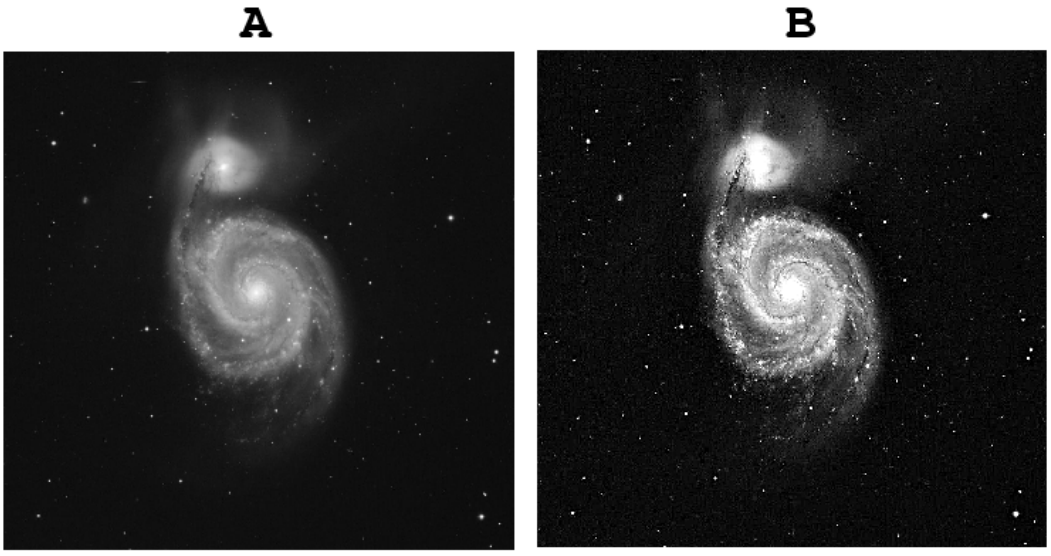

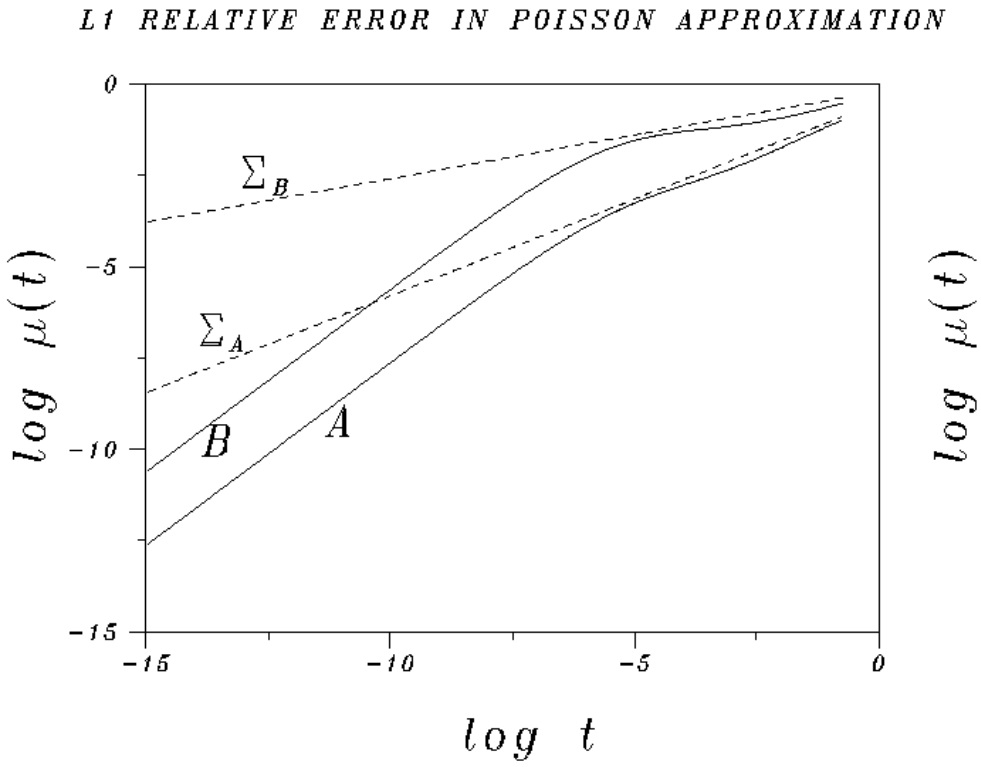

Figure 2. Whirlpool Nebula M51. Original and sharpened images have noticeably different L^1 Poisson traces $\mu(t) = \| U^t f - f \|_1 / \| f \|_1$ reflecting substantially different values of Lipschitz exponent α. (A) Original *1024×1024* image taken by T. Rector and M. Ramirez, National Optical Astronomy Observatory, (NOAO/AURA/NSF). Poisson approximation produces solid trace *A*, majorized by dashed straight line Σ_A defined by $\log \mu(t) = -0.5 + 0.530 \log t$. This implies that image (A) has an L^1 Lipschitz exponent $\alpha = 0.530$. (B) Blind deconvolution of (A) using APEX method brings out significant fine scale detail, and results in solid trace *B*, majorized by dashed straight line Σ_B defined by $\log \mu(t) = -0.2 + 0.239 \log t$. This indicates that deblurred image (B) has an L^1 Lipschitz exponent $\alpha = 0.239$. Image (B) strongly resembles M51 plate taken by famous astronomer Milton Humason using the *200* inch Mt Palomar telescope.

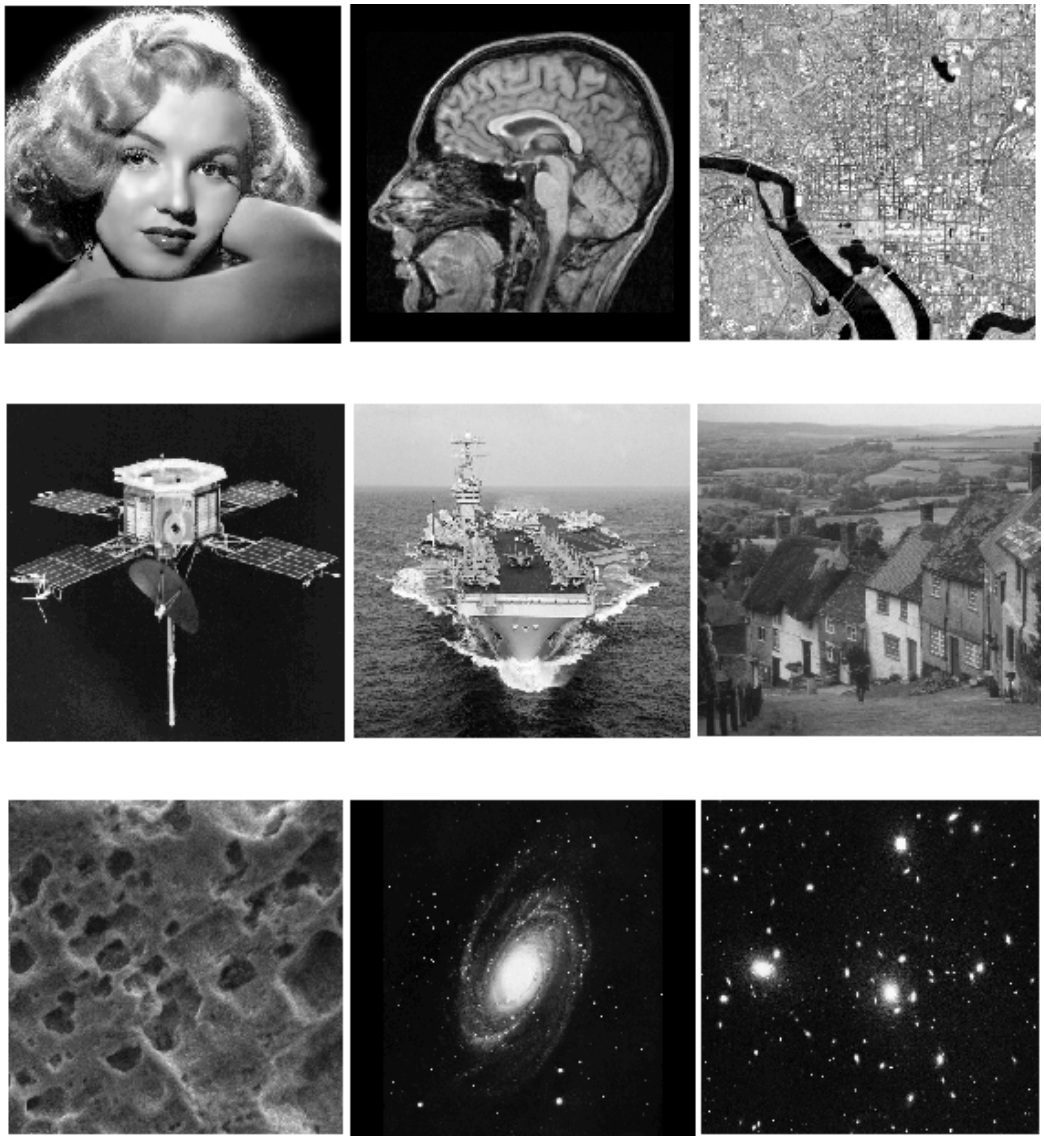

Figure 3. A significant class of high-resolution 8-bit images has Lipschitz exponents α in the range $0.2 < \alpha < 0.7$, in either L^1 or L^2, and are not of bounded variation.

Table 1. Values of (C, α) in $\| U^t f - f \|_p \le C \le \| f \|_p t^\alpha$, $0 < t \le 0.1$, $p = 1, 2$, for each image $f(x,y)$ in Figure 3, when U^t is Poisson operator in Eq. 1. Note narrow range of values in both C and α.

Image	**Size**	(C, α) **in** $L^1(R^2)$	(C, α) **in** $L^2(R^2)$
Marilyn Monroe	512^2	$C = 0.77, \alpha = 0.565$	$C = 0.68, \alpha = 0.474$
Sagittal brain MRI	512^2	$C = 1.28, \alpha = 0.590$	$C = 1.02, \alpha = 0.520$
Washington DC Landsat	512^2	$C = 0.45, \alpha = 0.341$	$C = 0.55, \alpha = 0.340$
Mariner 5 spacecraft	512^2	$C = 0.90, \alpha = 0.448$	$C = 0.99, \alpha = 0.417$
USS Eisenhower	512^2	$C = 0.47, \alpha = 0.420$	$C = 0.50, \alpha = 0.362$
English Village	512^2	$C = 0.49, \alpha = 0.472$	$C = 0.55, \alpha = 0.439$
Nanoscale micrograph	1024^2	$C = 0.45, \alpha = 0.415$	$C = 0.55, \alpha = 0.415$
Spiral galaxy	1024^2	$C = 0.68, \alpha = 0.365$	$C = 0.78, \alpha = 0.327$
Cluster of galaxies	1024^2	$C = 0.65, \alpha = 0.222$	$C = 0.97, \alpha = 0.216$

Singular Integral Image Deblurring Methods That Recover Texture

Alfred S. Carasso

Digital image acquisition plays an ever-increasing role in science, technology, and medicine, and, as a result, image deblurring has become an increasingly important activity. One important application at NIST involves the sharpening of Scanning Electron Microscope (SEM) imagery, in support of NIST's research programs in nanotechnology. In these applications, fast computational throughput for large size imagery is very desirable. Some deconvolution methods involve computationally intensive nonlinear iterative procedures that may require hours of CPU time. Direct (non-iterative) deconvolution methods, that can process images of size 1000×1000 pixels in less than a minute of CPU time, are considered *real-time* methods, and are highly sought after.

During the last ten years, there has been considerable interest in the use of nonlinear partial differential equations (PDE) methods in image deblurring. Such anisotropic diffusion equations typically originate from the use of steepest descent methods for minimizing a cost functional. Such functionals express the variational formulation of the ill-posed image deconvolution problem, and generally incorporate the regularizing constraints as an added penalty term. The Marquina-Osher finite difference scheme is one of the best known such PDE methods. It is based on the assumption that the unknown desired sharp image is of bounded variation. In that case, the penalty term involves the total variation (TV) norm of the unknown image. However, TV image deblurring is a computationally intensive nonlinear method that preserves edges, but often results in unacceptable loss of fine-scale information. Such loss of texture may be unsuspected, yet have very adverse consequences in medical imaging, nondestructive evaluation, or electron microscopy. The mathematical reasons behind this so-called staircase effect are profound, and have only recently been fully understood.

In fact, correct characterization and calibration of the lack of smoothness of images is crucial in image deblurring. As functions of x and y, most images are not differentiable. However, they are significantly better behaved than are the most general L^2 functions, while being significantly less smooth than are functions of bounded variation. For this reason, both the L^2-Tikhonov-Miller and the TV-Marquina-Osher methods are incorrectly formulated, and generally give poor results. Very recently, it has become apparent that so-called Lipschitz spaces are the proper framework for accommodating non-smooth images. A very significant breakthrough in that theory is provided by Carasso's use of singular integrals to estimate Lipschitz exponents, as discussed in the preceding article.

In a major new theoretical contribution to appear in the *SIAM Journal on Applied Mathematics*, Carasso reconsiders the image-deblurring problem in Lipschitz spaces. Using singular integrals, Carasso has created an entirely new cost functional for image deblurring. Rather than use the TV norm, Carasso's functional incorporates the L^2 Lipschitz space information $||U^t f - f \| \leq C \| f \|_{t^\alpha}$, $0 < t \leq t^*$, (such as might be found in Table 1 in the preceding article), into a new added penalty term. Most significantly, this new cost functional has the property that its minimum can be found explicitly in the Fourier transform domain. Therefore, using Fast Fourier Transform (FFT) algorithms, the deblurred image can be obtained directly, by inverse transforming this closed form Fourier expression. This provides an enormous computational advantage over TV deblurring for large size imagery. The Poisson operator U^t in the above theory leads to the so-called Poisson Singular Integral (PSI) deblurring method. However, U^t is only one of an infinite variety of singular integral operators that can be used. Gaussians and other radially symmetric Lévy stable laws may also be used.

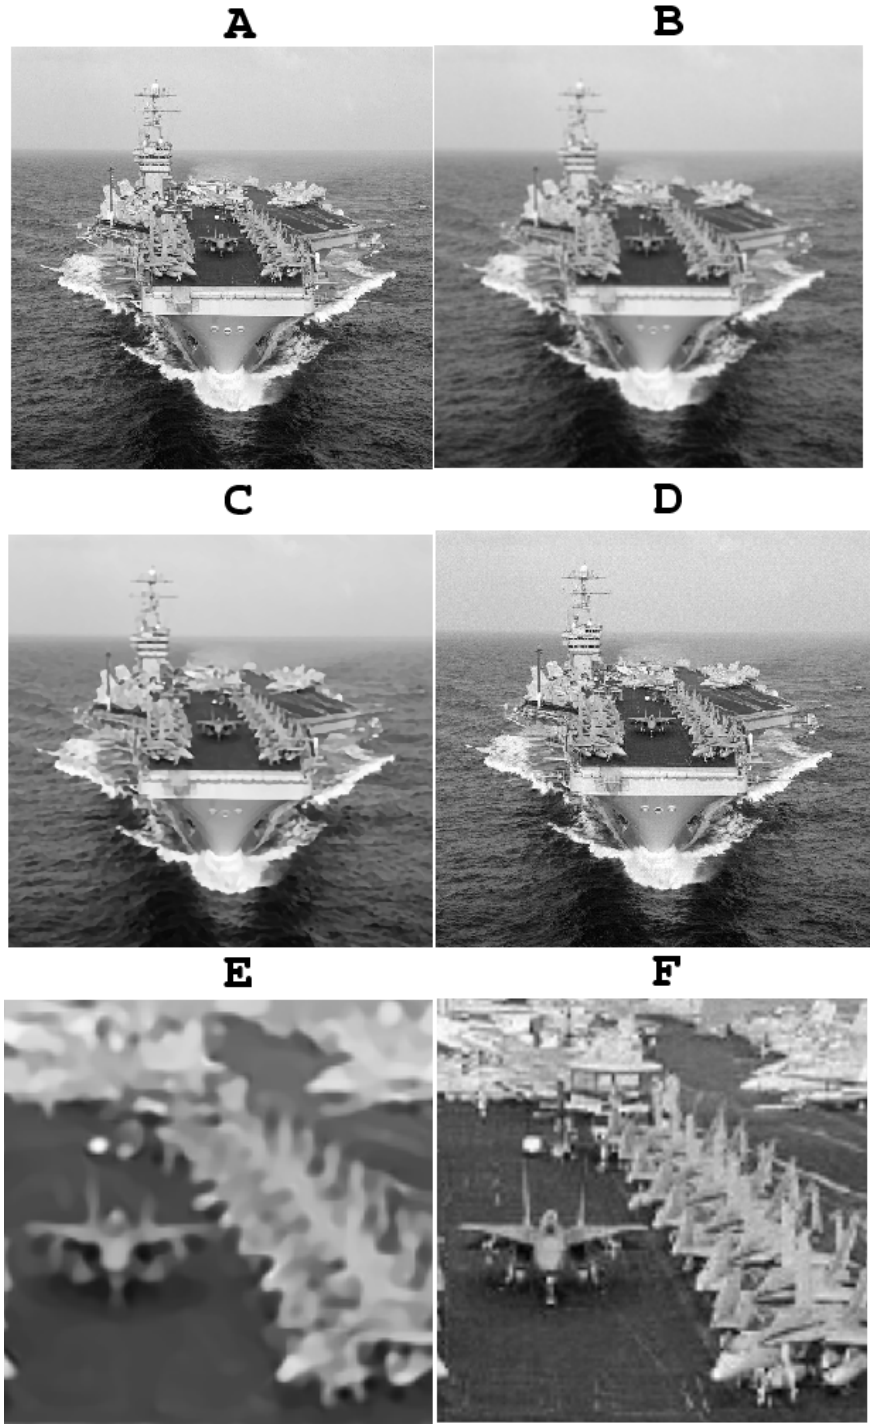

Figure 1. Comparison of total variation and PSI deblurring on mildly blurred image. Zooming on selected parts of the image enables meaningful comparisons between the two methods. (A) Original sharp USS Eisenhower image. (B) Mildly defocused image. (C) Total variation (TV) deblurring by applying PDE finite difference scheme. (D) PSI deblurring using Lipschitz space information, $\alpha = 0.362$, $C_t = 0.50$ in Table 1 in preceding article. (E) Zooming in TV deblurred image reveals significant loss of structural detail. (F) Zooming on same region in PSI image.

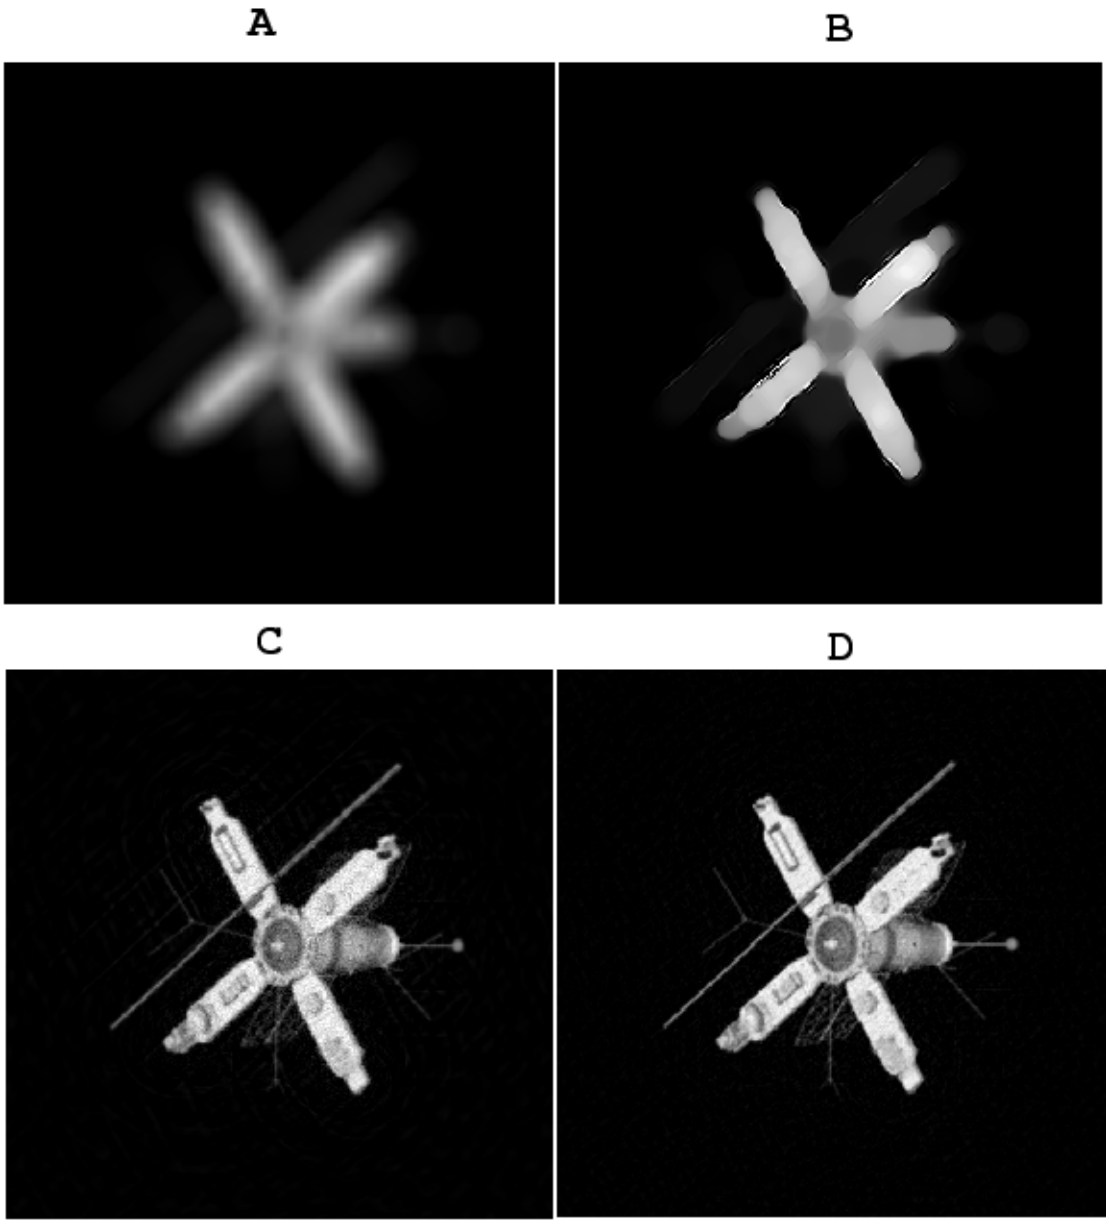

Figure 2. Use of substitute information. Robust PSI method produces remarkably good reconstruction using Lipschitz space data corresponding to image of "similar" object. (A) Strongly defocused USAF satellite image with multiplicative noise. (B) Total variation (TV) deblurring using PDE finite difference scheme results in severe loss of structural detail. (C) PSI method using substitute prior Lipschitz data $\alpha = 0.417$, $C_t = 0.99$ obtained from Mariner 5 image in Table 1 in preceding article. (D) True Wiener filtering with exact power spectra $|\hat{n}(\xi,\eta)|$, $|\hat{f}_e(\xi,\eta)|$. Realizable fast PSI deblurring, using substitute data, closely matches unrealizable true Wiener filtering.

We have stressed the fact that most images are not smooth. The PSI method is predicated on locating the unknown image $f(x,y)$ in the correct Lipschitz space, while TV deblurring assumes $f(x,y)$ to be of bounded variation. It is sometimes argued that such refined smoothness measures are primarily applicable to $f^O(x,y)$, the original intensity field that gave rise to the digitized finite dimensional image $f(x,y)$, but are not meaningful for $f(x,y)$ itself. Indeed, since all norms are equivalent in finite dimensional space, it remains to be seen whether such abstruse function space notions are ultimately of any computational significance.

This is illustrated in Figure 1. The original sharp USS Eisenhower image shown in (A) was synthetically defocused, and multiplicative noise was then added to form the blurred image (B). The nonlinear TV PDE finite difference scheme applied to image (B) produced the deblurred image in (C). (D) is the PSI deblurred image, using exact L^2 prior Lipschitz space information $\alpha = 0.362$, $C_{t^*} = 0.50$, from Table 1. Zooming on selected parts of these images, in (E) and (F), shows significant loss of structural detail in the TV image as compared with PSI deblurring. In this example, the bounded variation assumption underlying the PDE scheme clearly eliminates valuable information.

In general, the prior Lipschitz space information that is required in the PSI method may not be known. However, inspection of Table 1 suggests that the values of (C,α) are confined to a fairly narrow range. Indeed, a plausible guess for (C,α) might be $(0.5, 0.5)$ in many cases. In other situations, a sharp image of a similar object may provide valuable estimates for (C,α). The PSI method is sufficiently robust as to provide useful reconstructions, even with inexact Lipschitz data. Such initial reconstructions can then be interactively refined through *fast* simultaneous computation and display of multiple trial PSI images, corresponding to neighboring (C,α) values. The next example shows how good such initial reconstructions can be.

The strongly defocused USAF satellite image in Figure 2(A) was blurred synthetically, and multiplicative noise was then added. As may be expected in such a severely blurred image, TV PDE deblurring can only provide a very sketchy reconstruction as shown in Figure 2(B). However, the Mariner 5 image, shown in Figure 3 in the preceding article, may be considered a "similar" object, and the corresponding L^2 information in Table 1, $C = 0.99$, $\alpha = 0.417$, may be used in the PSI method. Remarkably, this produces the reconstruction shown in Figure 2(C). For comparison, Figure 2(D) is the *true* Wiener filtered image, using the *exact* power spectra of each of the noise and unknown sharp image. This much prior knowledge is almost never available, and true Wiener filtering is practically unrealizable. However, the initial PSI image appears to be an excellent approximation to this optimal Wiener image.

Modeling and Computational Techniques for Bioinformatics Based Data Mining

Fern Y. Hunt
Anthony J. Kearsley
Agnes O'Gallagher

The purpose of this project is to develop generic algorithms that can utilize large-scale and fast computational methods to perform the comparison and alignment of a large number of long biological sequences, a critical problem in bioinformatics. Almost all methods in current use are based on an optimization problem that is solved using the methods of dynamic programming. Starting from a Markov decision theory approach, MCSD mathematicians instead solve a linear programming problem. The goal is to avoid the exponential increase in computation as the number of sequences increases and to make use of the high performance algorithms available for solving linear programming problem. Potential applications for this work are found in drug discovery, where large numbers of sequence sites must be evaluated for use as drug targets, and in the development of

pharmacogenomic therapies, i.e. therapies that are tailored to the genetic makeup of the patient. Multiple sequence alignment also plays an important role in the characterization of biological threats.

During the year 2003, we continued work on creating a software environment for aligning biological sequences given user supplied cost functions and training sequences. We used the package to undertake an extensive study of the quality of the alignments suggested by our method as a function of cost function parameters. For a set of cytochrome p450 sequences we characterized the parameter values that produced alignments that were most consistent with those obtained by the widely used alignment method CLUSTALW. The results of this work was reported at the at an invited presentation at the Biological Language Conference held at the University of Pittsburgh in November 2003 and received a favorable reception at a presentation before the Computational Biology Division at the NIH. Presentations were also given at the IEEE Computational Systems Bioinformatics Conference in August and the Bioinformatics Institute of the University of Maryland, Baltimore County. Written accounts of our work appear in a short article in *Applied Mathematics Letters* and in will appear in conference proceedings. We also prepared two papers, which are about to be submitted to journals. Theoretical and computational progress was made in work with alignments based on the dual linear programming problem. Hunt obtained sufficient results that imply that the optimal solution generates costs that are optimal in a stronger sense than just optimal "on average". Optimal costs are lower than other costs with high probability. Working with protein data supplied by Arlin Stolzfuss of CARB, Hunt and SURF student Shirin Mehraban constructed alignments of these sequences using the solution of the dual linear programming problem.

The parameter study suggested changes that we should make in the linear programming problem that would provide more flexibility in the class of alignments we produce. We will also be experimenting with the use of simulated training data coming from a protein evolution model to see if we can obtain comparable or better results.

Linear Stability of Spiral Poiseuille Flow with a Radial Temperature Gradient

David Cotrell
Geoffrey McFadden

We consider fluid flow in an annulus driven by the combination of an axial pressure gradient, rotation of the inner and outer cylinders, and a radial temperature gradient (i.e., spiral Poiseuille flow with a radial temperature gradient). D. Cotrell and G. McFadden are investigating the transition from the simplest flow possible, steady flow with two nonzero velocity components and a radial temperature gradient that varies only with radius) to the next simplest flow possible, steady flow with three nonzero velocity components and a temperature profile that vary in the radial and axial directions). This work is motivated by electrochemical processes in rotating cylinder electrodes, heat transfer in rotating machinery, flow-amplified electrophoretic separations, and vortex flow reactors for shear-sensitive biological systems. It extends the case of isothermal spiral Poiseuille flow for which D. Cotrell and A. Pearlstein (University of Illinois at Urbana-Champaign) recently computed complete linear stability boundaries for several values of the radius ratio and rotation rate ratio, and shows how the centrifugally driven instability, beginning with steady or azimuthally-traveling-wave bifurcation of circular Couette flow (flow driven solely by wall motion), connects to a non-axisymmetric Tollmien-Schlichting-like instability of nonrotating annular Poiseuille flow (flow driven solely by an axial pressure gradient). Results for the non-isothermal case show that the stability boundary shifts either up or down depending on the sign of the temperature difference. For the isothermal case, it is also known that in many instances there is no instability for small enough axial flow rates. For the non-isothermal case, however, Cotrell and McFadden have shown that for any nonzero temperature

difference between the inner and outer radii, a new non-isothermal mode of instability causes the base state to be destabilized under these conditions.

Monte Carlo Methods for Combinatorial Counting Problems

Isabel Beichl
Francis Sullivan (IDA Center for Computing Sciences)

I. Beichl and F. Sullivan have developed a new method using stratified sampling to estimate the number of independent sets in a graph. An independent set is a subset of the set of nodes of a graph where none of the elements of the subset has any connections with any other elements of the subset. This problem is important because of applications in wide ranging applications from thermodynamics to data communications. It is dual to the problem of counting the number of cliques in a graph and thus gives graph connectivity information. The new technique is similar to a method based on importance sampling that developed by Beichl and Sullivan in previous years. In stratified sampling, the space to be sampled is divided into strata and it is required that each stratum be sampled. In our method one complete sample is made from many subsamples, and each stratum must be represented by a subsample if possible. In doing this, a more complete sampling is made and so variance is reduced. He new importance sampling method has solved problems that could not have been done by any existing method. There have been recent improved results in computing the variance of this method. It was proved that stratified sampling always has lower variance than the importance sampling method. Beichl's implementation used innovative techniques dealing with extended precision. This work was recently accepted as a talk at the SIAM Workshop on Combinatorial Scientific Computing (CSC04).

Measures of Approximation in Delay Differential Equations

David Gilsinn

Delay differential equations have arisen in the study of machine tool chatter. Machine tool chatter is a self-excited oscillation of a cutting tool against a machined workpiece. Acoustically it can be detected as a high frequency squeal in a machine shop. Modeling chatter and other machining instabilities has been a problem of long-term interest to the NIST Manufacturing Engineering Laboratory. Mathematically, chatter is characterized as a limit cycle for a nonlinear delay differential equation.

D. Gilsinn of MCSD is developing a precise measure of how near an approximate periodic solution is to the "real" limit cycle. The measure developed is more precise than a standard estimate of order-of-magnitude error. Whereas in ordinary differential equations there are well-developed Green's function methods of determining the initial conditions for periodic solutions as well as known properties of fundamental solutions, in delay differential equations there are no such easily constructed relations. Developing characteristic multipliers and initial conditions for periodic solutions of delay equations involve solving Fredholm-type integral equations. These solutions are critical to developing precise measures of approximation. Matlab programs have been developed to compute the characteristic multipliers for the linear variational periodic delay equations associated with the nonlinear Van der Pol equation with delay. Research continues on the development of error and convergence estimates for the algorithms developed to estimate the characteristic multipliers.

Separating Signal from Noise in Measured Time Series

Bert W. Rust

The goal in time series modeling is a mathematical representation of the signal that reduces the residuals to white noise. Smoothing splines or regression splines are often used to model the signal when the dynamics of the process are not understood. The former use a knot at every data point and depend on the choice of a smoothing parameter, and the latter depend on the number and locations of the knots. We have developed diagnostic software for evaluating the performance of these procedures. As an example, consider the time series, plotted in Fig. 1. These pressure differences drive the equatorial trade winds and profoundly influence global weather patterns. The mathematical model, with 13 adjustable parameters, is

$$\Delta P(t) = C_0 + C_1 t + A_{12}\sin\left[\frac{2\pi}{12}(t+\phi_{12})\right] + A_{27}\sin\left[\frac{2\pi}{\tau_{27}}(t+\phi_{27})\right] + A_{33}\sin\left[\frac{2\pi}{\tau_{33}}(t+\phi_{33})\right] + A_{44}\sin\left[\frac{2\pi}{\tau_{44}}(t+\phi_{44})\right]$$

The dominant feature is a 12-month oscillation corresponding to the annual cycle of insolation. The smaller oscillations have periods of approximately 27, 33 and 44 months. The 44 and 27-month oscillations produce the El Nino, Southern Oscillation. The residuals are acceptably normal and only 3 of the 83 autocorrelation estimates fall outside the 95% band for white noise. A Fourier analysis, shown in Fig. 2, confirms the white noise character of the residuals.

Cubic smoothing splines are obtained by minimizing a weighted linear combination of the sum of squared residuals (SSR) and the integral of the squared second derivative. The SSR is given weight 1, and the integral a variable weight, called the smoothing constant, which is chosen to optimize the partition between signal and noise. If the noise variance is unknown, that weight is chosen to minimize the generalized cross validation, and the fit is called the "optimal" smoothing spline. For the present problem it looks like the fit in Fig. 1, but its SSR is much smaller. The Fourier analysis given in Fig. 3 indicates that it captured noise that should have gone into the residuals, especially at the lower frequencies. Also, a small remnant of the variance in the yearly cycle was left in the residuals. Choosing larger smoothing constants reduces the high/low frequency imbalance by increasing the 12-month remnant. Smaller values eliminate the remnant but exacerbate the imbalance. There does not appear to be a choice that will produce white noise residuals.

For regression splines the smoothing is determined by the number and locations of the knots. An effective strategy has been developed by Thijsse, *et al*. [Thijsse, B. J., Hollanders, M. A., and Hendrikse, J., *Computers in Physics*, Vol. **12**, No. 4 (Jul/Aug 1998) pp.393-399] who used a generalized Durbin-Watson statistic as a guiding parameter. Their program Spline2 chose 29 equally spaced knots to give a fit similar to the one in Fig. 1 and an SSR reduction about half that for the smoothing spline. The Fourier analysis in Fig. 4 rules out the white noise hypothesis for the residuals, but the frequency imbalance is not so bad as that for the smoothing spline, and the fit captured all of the variance in the yearly cycle.

It was mostly at lower frequencies that the splines captured noise as well as signal. Hopefully this does not indicate a fundamental limitation of piecewise polynomial approximations. I am currently developing a smoothing spline strategy that uses the squared length of the cumulative periodogram rather than the integrated squared second derivative to obtain a smoothing constraint. It is possible to express this length in terms of the spline coefficients, but it is messy to do, and the resulting minimization problem may be more difficult than the one currently being solved.

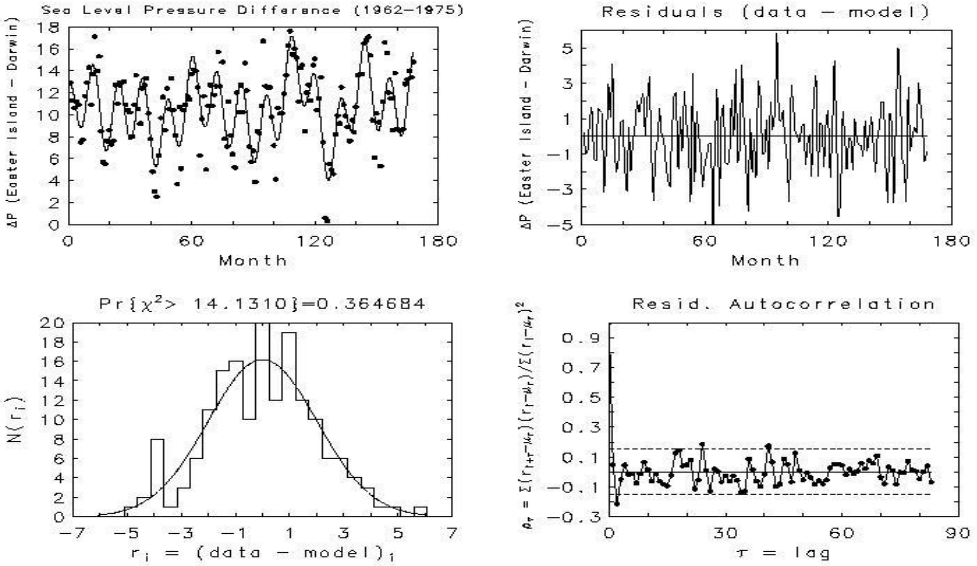

Figure 1. Upper left: Monthly average sea level atmospheric pressure differences, measured in millibars [mb], between Easter Island and Darwin, Australia for the years 1962-1975. The smooth curve is a nonlinear least squares fit of the mathematical model given below. Upper right: Time series of residuals for the fit. Lower left: Histogram distribution of the residuals. The smooth curve is a transformed standard normal distribution with the same standard deviation and area as the histogram. Lower right: Statistical autocorrelation function for the residuals. The dashed lines define 95% bounds for white noise.

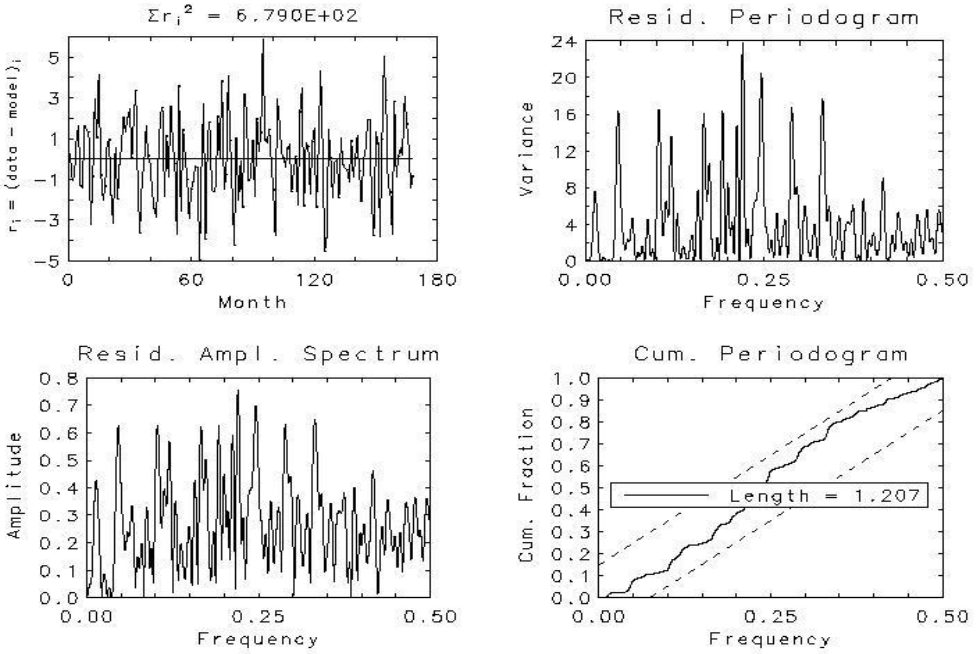

Figure 2. Upper left: The residual time series. Upper right: Fourier variance spectrum of the residuals. The frequency units are inverse months. Lower left: Fourier amplitude spectrum of the residuals. Lower right: Cumulative variance spectrum. The dashed lines define a 95% confidence band for white noise.

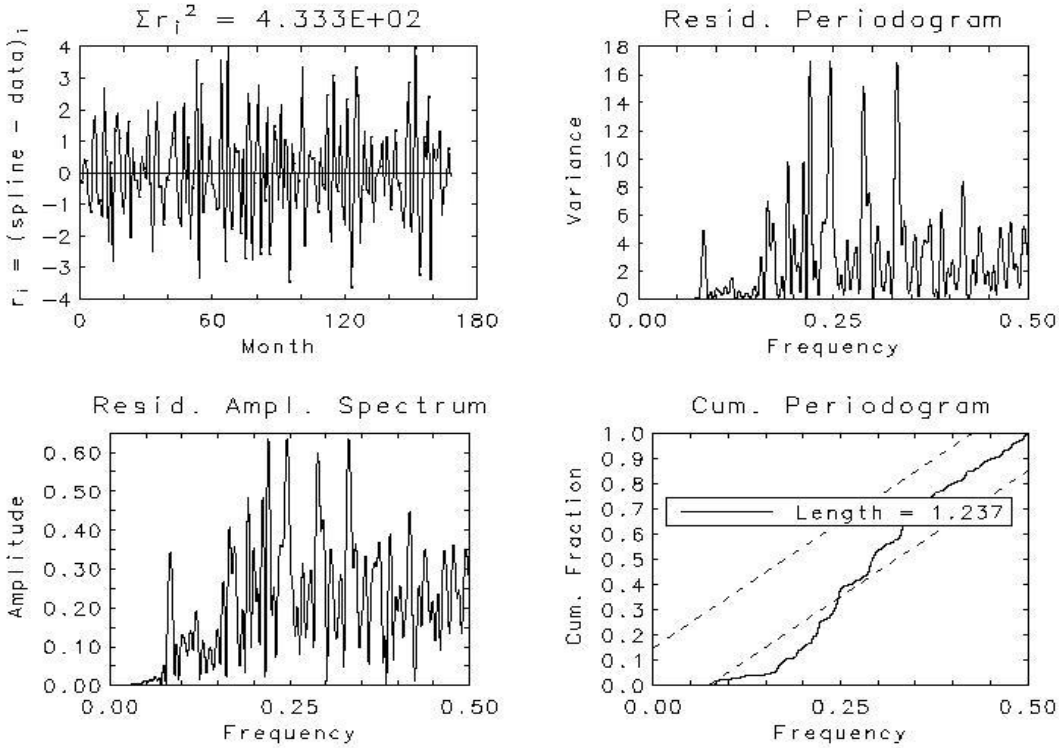

Figure 3. Fourier spectral analysis of the optimal smoothing spline residuals.

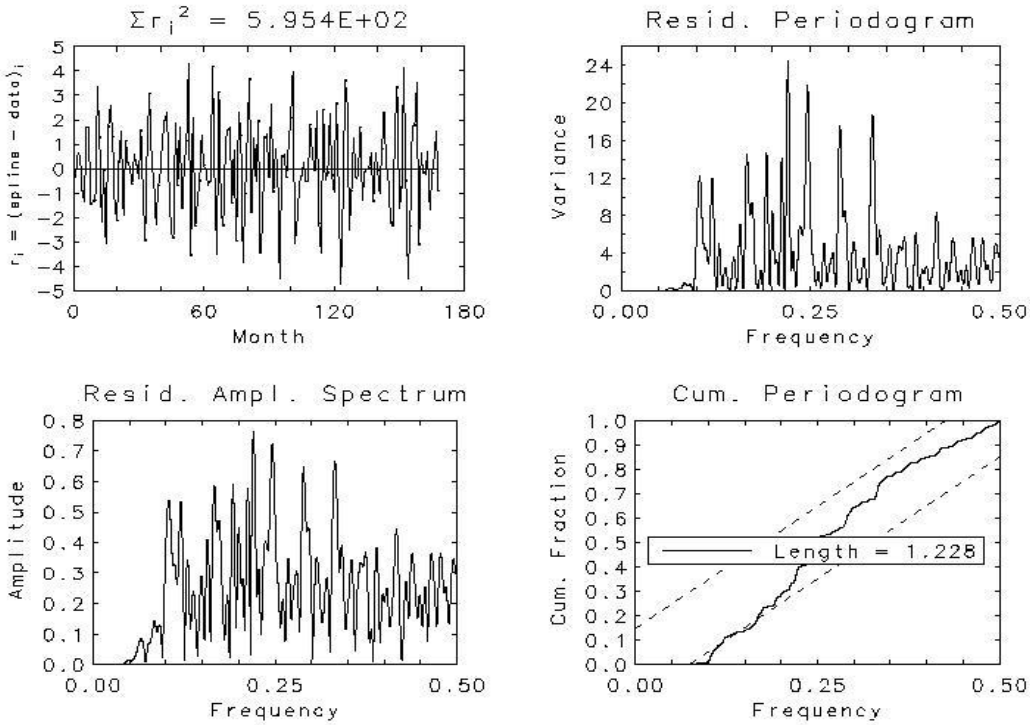

Figure 4. Fourier spectral analysis of the Spline2 residuals.

Systems Identification and Parameter Estimation

Bert W. Rust

An important modeling problem is to identify a system of ordinary differential equations (ODEs) describing the dynamical relationships between a set of measured time series variables. The ODEs usually depend on several unknown constants and initial values, so a related problem is to estimate those unknown parameters by fitting the solutions to the measurements. A timely example involves global fossil fuel emissions of carbon dioxide, the accumulation of carbon dioxide in the atmosphere, and global temperature variations. Good measured time series records for these three variables are readily available on the Web.

Figure 1 gives a plot of annual global total fossil fuel carbon dioxide emissions, measured in millions of metric tons of carbon [MtC]. The dashed curve was obtained by fitting a mathematical model consisting of a superposition of an exponential baseline and a sinusoid, with a period approximately 65 years, whose amplitude increases with the same exponential rate as the baseline. It was used as an initial approximation in determining the mathematical relationships between the emissions and the other two variables.

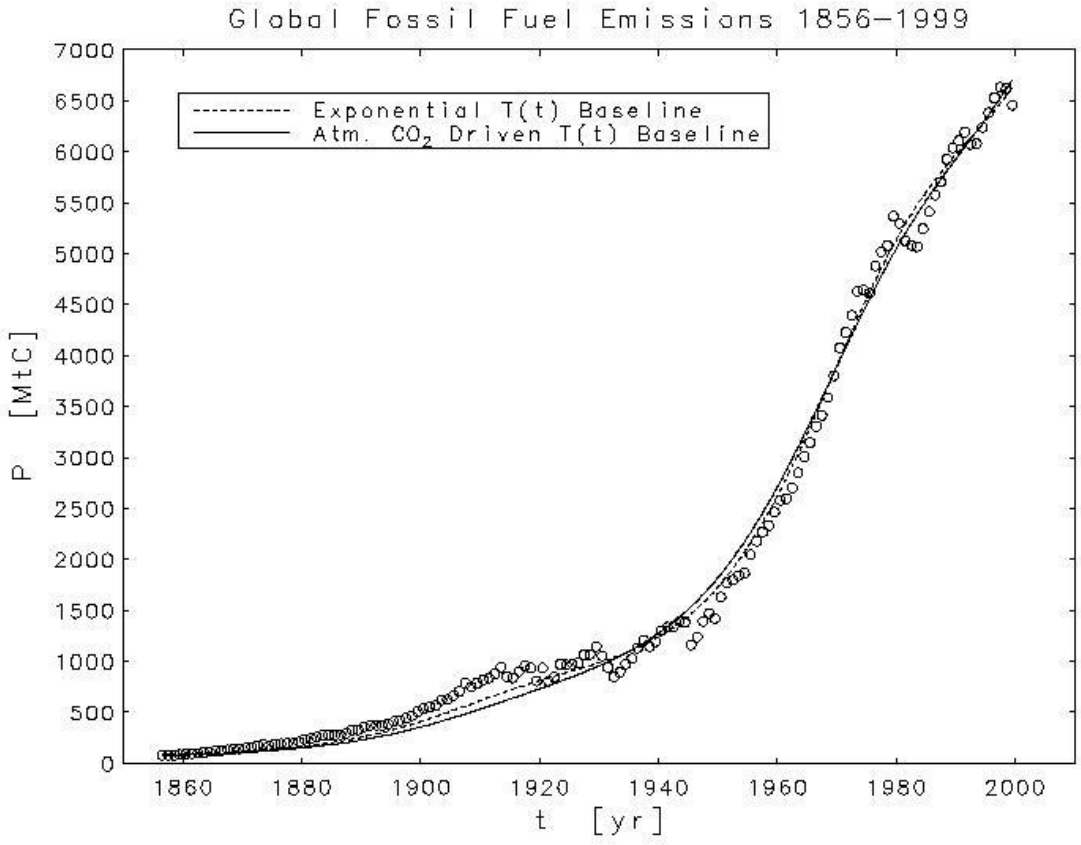

Figure 1. Annual global total fossil fuel carbon dioxide emissions. The dashed curve is the fit of a purely mathematical model, and the solid curve is the fit of a temperature-driven model.

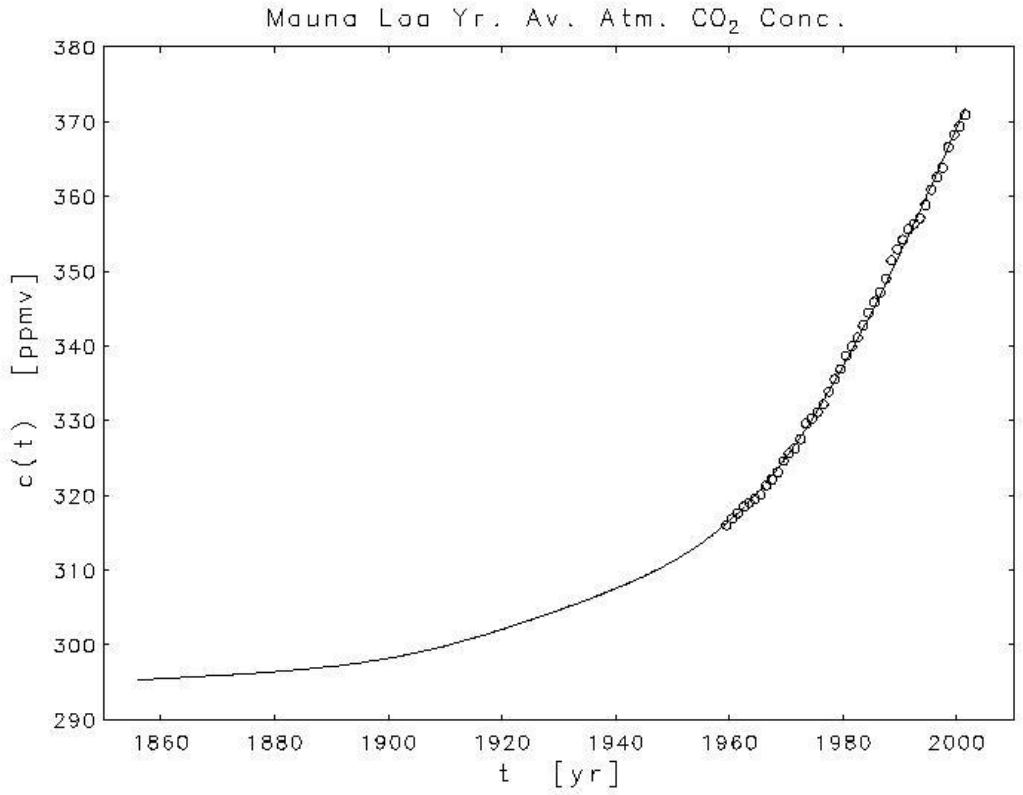

Figure 2. The circles are the yearly average atmospheric carbon dioxide concentrations, measured in units of parts per million by volume, for the years 1959-2001. The curve is the linear least squares fit of a two-parameter model based on the cumulative integral of the dashed curve in Fig. 1.

The connection between the emissions and temperature variations is established by the carbon dioxide in the atmosphere. Figure 2 gives a plot of yearly average carbon dioxide concentrations in the troposphere. Roughly one half of each year's emissions remain in the atmosphere, so the concentration is well modeled by a simple two-parameter model comprising a sum of a constant and a fraction of the cumulative integral of the emissions. Using the dashed curve in Fig. 1 to approximate the emissions gives the fit plotted in Fig. 2. The data go back only to 1959, but the fit was extended to 1856 in order to use it for estimating the temperature variations. The 65-year cycle was completely smoothed out, so its presence in the temperature record cannot be caused by its presence in the emissions record.

The data plotted in Fig. 3 are annual global average tropospheric temperature anomalies measured relative to the average temperature for the years 1961-1990. The dashed curve is a linear least squares fit of a mathematical model consisting of an exponential baseline, with rate constant exactly one half of the one for the dashed curve in Fig. 1, plus a sinusoid with period exactly the same as the sinusoid there and phase constant set to make the oscillation exactly one half cycle ahead (or behind) the one there. This inverse correlation was first noted by Rust and Kirk [B. W. Rust and B. L. Kirk, *Environment International*, Vol. 7 (1982) 419-422] who interpreted it as a feedback by which warming slows the rate of fossil fuel production.

Antarctic ice core reconstructions of atmospheric temperatures and carbon dioxide concentrations show that the two quantities have varied proportionately for the past 160,000 years. This suggests a simple two-parameter, linear relationship between them. Adding the sinusoid and fitting the resulting model gave the solid curve in Figure 3.

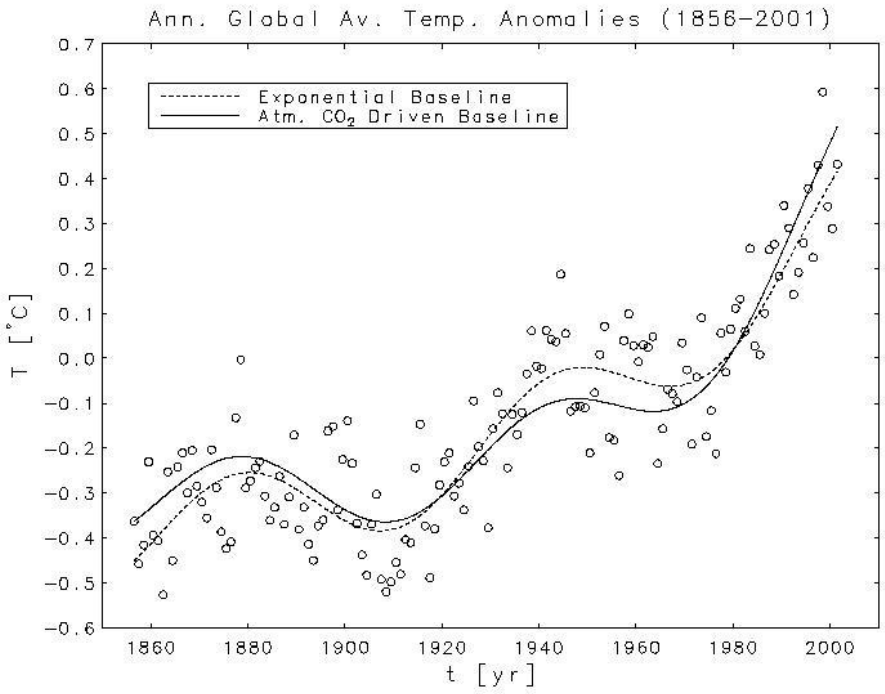

Figure 3. The circles are the observed temperature anomalies. The dashed curve is the fit of the mathematical model, and the solid curve is the fit of the atmospheric carbon dioxide driven model.

Since the 65-year cycle in the temperatures was not caused by fossil fuel emissions, its presence, with sign reversed, in the emissions suggests a temperature driven, negative feedback. Rust and Kirk successfully modeled this feedback with the ODE

$$\frac{dP}{dt} = \left(\alpha - \beta \frac{dT}{dt} \right) P , \qquad P(t_0) = P_0$$

where $P(t)$ and $T(t)$ are emissions and temperature at time t and α, β, and P_0 are free parameters. If the function represented by the solid curve in Figure 3 is substituted into this model and the ODE is fit to the emissions data, the result is the solid curve in Figure 1. Its similarity to the dashed curve initial fit is evidence for the correctness of the chain of assumptions used in its construction. It is probably possible to refine some or all of the models in the above chain of fits, but that may not be needed to identify the ODEs governing the three variables. Analysis thus far suggests the equations

$$\frac{dP}{dt} = \alpha P - \beta \left\{ \eta P + A \cos\left[\frac{2\pi}{\tau} (t + \phi) \right] \right\} P , \qquad P(t_0) = P_0$$

$$\frac{dc}{dt} = \gamma P , \qquad c(t_0) = c_0$$

$$\frac{dT}{dt} = \eta P + A \cos\left[\frac{2\pi}{\tau} (t + \phi) \right] , \qquad T(t_0) = T_0$$

where $c(t)$ is the atmospheric concentration of carbon dioxide and α, β, η, A, T, ϕ, P_0, γ, c_0, and T_0 are adjustable parameters. Current efforts are focused on fitting these three equations simultaneously to all three data sets.

Numerical Algorithms for Advanced Mass Spectrometry

Javier Bernal
Anthony Kearsley
Charles Guttman (NIST MSEL)
William Wallace (NIST MSEL)

Mass spectrometry (MSpec) is an important technique for obtaining information about materials in a wide variety of applications. Loosely speaking, a mass spectrometer (see Figure 1) detects the presence of chemicals inside a sample suspended in a matrix by bombarding it with ions (with the help of a laser), and then counting the ions as they "bounce" off the prepared sample. Each chemical causes ions to bounce off in a different way, thus leaving a signature. If one is interested in testing for the presence of a particular chemical inside a sample material, an MSpec experiment can be designed to result in large peaks being produced when the chemical in question is detected. For example, one must identify, separate, and sum the matrix peaks present in the Bradykinin MSpec output in Figure 2. Because of its accuracy, this method for determining the presence of chemicals has been employed with more and more frequency by scientists.

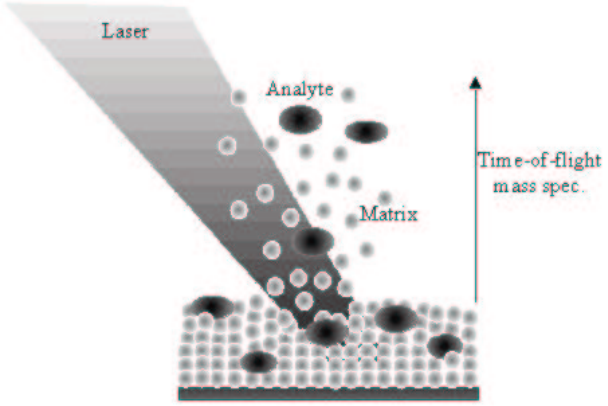

Figure 1: Illustration of MALDI Time-of-Flight Mass Spectrometer

While very detailed chemical procedures have been developed for MSpec sample preparation, very little work has been done on the development of numerical methods for the analysis of MSpec output. The objective of this project is to develop a reliable, and robust algorithm for *automatically* analyzing raw output from mass spectrometers and to demonstrate the algorithm by building a from-everywhere accessible public-domain web-based tool, making a high quality implementation of the algorithm widely available. The webtool is being constructed to directly analyze noisy mass spectrometer output, generating unbiased machine-independent analysis of the noisy data in a timely fashion, greatly reducing (or eliminating) the need for laborious and painstaking manual analysis.

Modern mass spectrometers produce astonishingly high-quality data. Scientists are regularly faced with the difficult and time-consuming task of sorting through enormous noisy data sets in search of structures (e.g., a "peak" or a "trough") corresponding to the presence of chemical in question. Because of unavoidable measurement errors, test data can vary considerably when identical samples are analyzed with different equipment or in different laboratories. This expensive and time-consuming task is usually done with the aid of packaged computational algorithms, many of which are 'black-box' codes, sold by software vendors or by companies manufacturing mass spectrometer hardware. In practice, these codes are often massaged into performing correctly by changing large numbers of algorithmic input parameters and/or testing subsets of the output data independently before analyzing the entire mass spectrometer output. Numerical behavior of currently available software can vary drastically as a function of numerous extraneous factors, including parameter selection, choice of peak shape for example, machine precision, and even the computer platform on which the software is being run. Simply put, this process is notoriously difficult.

We have succeeded in developing, implementing and testing an algorithm that, given raw MSpec data, can efficiently identify peak and trough structure without assuming any a priori size or

shape structure. Preliminary results are promising, and the NIST Polymers Division is now using the algorithm regularly. A prototype implementation of the algorithm has been made available as a web tool called MassSpectator; see http://www.ctcms.nist.gov/%7Ewallace/MassSpectator_intro.html.

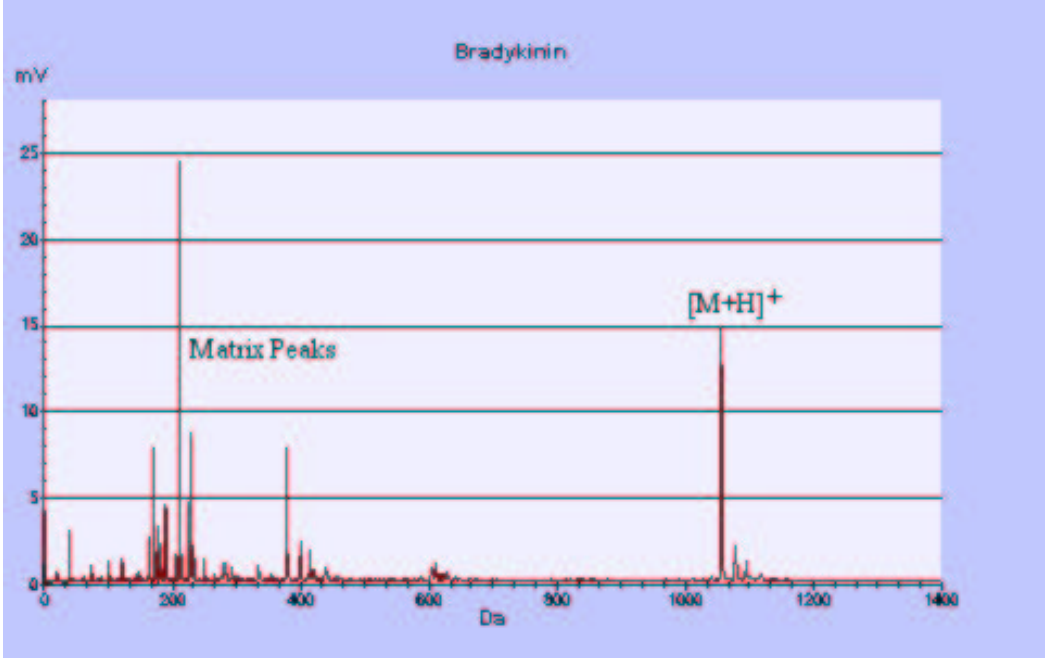

Figure 2: Bradykinin Mspec Output

Mathematical Software

Parallel Adaptive Refinement and Multigrid Finite Element Methods

William F. Mitchell
Eite Tiesinga (NIST PL)

http://math.nist.gov/phaml

Finite element methods using adaptive refinement and multigrid techniques have been shown to be very efficient for solving partial differential equations on sequential computers. Adaptive refinement reduces the number of grid points by concentrating the grid in the areas where the action is, and multigrid methods solve the resulting linear systems in an optimal number of operations. W. Mitchell has been developing a code, PHAML, to extend these methods for use on parallel computers. The expertise and software developed in this project are useful for many NIST laboratory programs, including material design, semiconductor device simulation, and the quantum physics of matter.

This year the effort was focused on two areas: enhancement and release of the PHAML software, and application of PHAML to solve Schrödinger's Equation for eigenvalues and eigenstates relevant to optical traps for neutral atoms, in collaboration with the E. Tiesinga of the Quantum Processes group in NIST's Atomic Physics Division. Understanding the interactions of adjacent atoms corresponding to qubits of a quantum gate require the computation of multiple eigenvalues in the middle of the spectrum, with eigenstates that have sharp gradients, a very challenging computation. The major accomplishments for FY 2003 are the following.

- Released PHAML 0.9.x and announced it on the major numerical analysis newsgroups and lists. Additional releases were made as enhancements warranted.
- Modified the software's structure so that PHAML is compiled into a library to be linked with an application.
- Added new capabilities including a more general initial grid on rectangular domains, a query function to obtain information about the grid and solution, new preconditioners, SPMD examples, and enhancements to the graphics.
- Added links into more third party libraries including LAPACK, MUMPS, SuperLU, PETSc, and hypre.
- Wrote an application on top of PHAML for solving the Schrödinger equation model of two interacting atoms in a trap.
- Experimented with different spectral transformations, iterative methods, and preconditioners for computing interior eigenvalues of the Schrödinger equation.
- Solved a model of interacting Cesium atoms in a trap with a grid of 1.5 million vertices on eight processors.

In addition, four talks were given at conferences on this topic, two of which were invited plenary lectures, and four papers were written and submitted to refereed journals.

Future work will include (a) continuing to enhance PHAML with additional capabilities and robustness, (b) seeking an effective iterative solver for the indefinite linear systems that arise when computing interior eigenvalues, (c) improving the robustness of the Schrödinger equation application code, and (d) extending the application to a multi-channel model with time-dependent systems of equations.

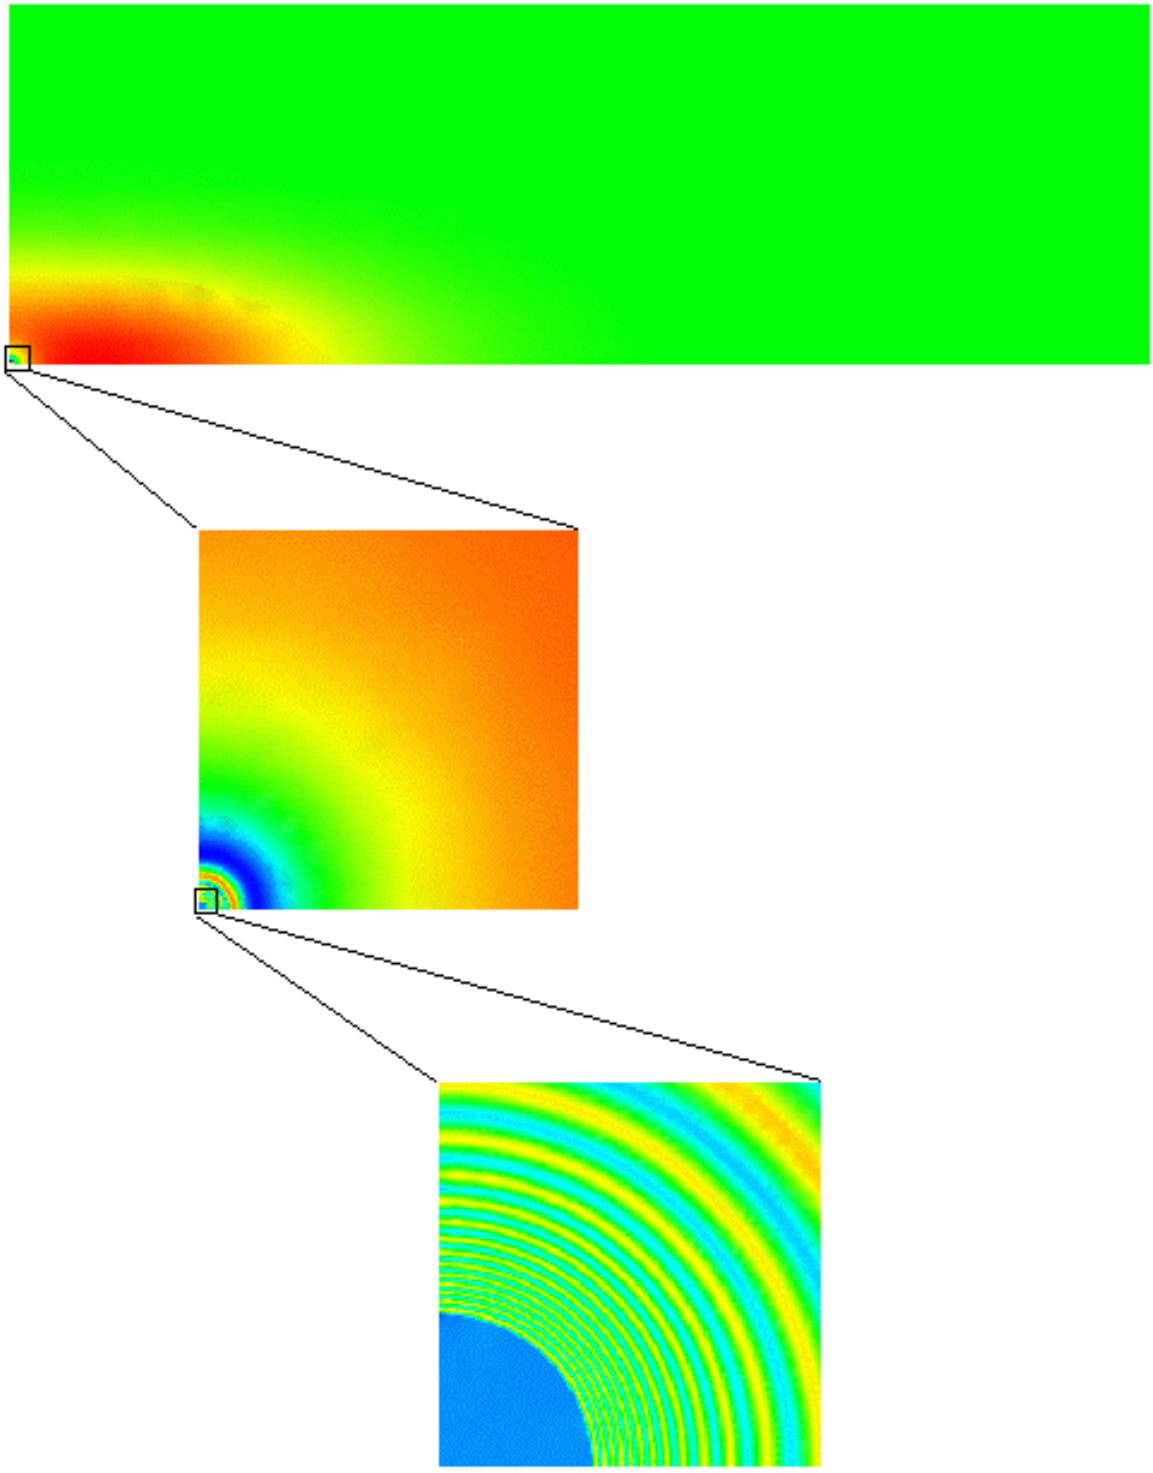

The wave function for the first trapping state of two interacting Cesium atoms in a cigar-shaped trap, shown at three zoom levels.

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. SciMark includes 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 megaflop rates for the numerical kernels, as well as an aggregate score for the complete benchmark. The current results database lists submissions form more than 1,800 computational platforms, ranging from palm devices to high-end servers, and contains reports from 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.

SciMark and its kernel components have become a pseudo-standard in industry and academia. They were adopted by the Java Grande Benchmark Forum, and Sun Microsystems used SciMark 2.0 to demonstrate the floating-point improvements to their JVM 1.4.2 (see http://java.sun.com/j2se/1.4.2/1.4.2_whitepaper.html).

As of December 2003, the record for SciMark is 550 Mflops, a 44% improvement over the best score reported one year ago (380 Mflops).

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 played a leading role in the latest 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 comprising fundamental matrix/vector operations common to most scientific computing applications. If the interface to such kernels is standardized, 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 both portability and performance. 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. Recently there has been renewed interest in extending this success to other areas. The BLAS Technical Forum (BLAST) recently completed the development of a new BLAS standard that addressed important new computational interfaces, and extended functionality of the previous proposals. BLAST is an international consortium of industry, academia, and government institutions, including Intel, IBM, Sun, HP/Compaq/Digital, SGI/Cray, Lucent, Visual Numerics, and NAG.

Among the most significant components of the new 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 the 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.

This year saw the integration of the Sparse BLAS interface in the NIST Template Numerical Toolkit (see below), and the continuing development of an ANSI C Sparse BLAS reference implementation to accompany the existing Fortran95 version. The TNT interface provides an object-oriented framework to support fundamental sparse matrix operations and memory management in application codes, while the reference implementation provides stand-alone code that serves as a testing and validation module for platform-dependent versions of the Sparse BLAS library.

TNT: Object Oriented Numerical Programming

Roldan Pozo

http://math.nist.gov/tnt

MCSD maintains an active research program in the design of object oriented mathematical software libraries. This work has led to some of the most highly used object oriented linear algebra packages, including Lapack++, Iterative Methods Library (IML++), Sparse Matrix Library (SparseLib++), Matrix/Vector Library (MV++), and the Template Numerical Toolkit (TNT). The latter package was downloaded more than 10,000 times during the past 12 months, for example, and is currently in use in a variety of industrial and commercial applications. This year saw several software releases and a major redesign in the latest version.

TNT incorporates many of the ideas explored by R. Pozo and colleagues with previous designs and includes new techniques that utilize the latest ANSI C++ capabilities. The package includes support for both C and Fortran multidimensional array layouts, arrays sections, and application modules, such as linear algebra, which includes fundamental algorithms for matrix decompositions (LU, Cholesky, SVD, QR), eigenvalue problems, and sparse matrix support such as vector scatter/gather operations, matrix multiplication with dense arrays, and the solution of sparse triangular systems, which can utilize the Sparse BLAS (see above).

With the latest TNT developments, R. Pozo has designed a separate 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 templates, or instrumented versions for debugging session. Under development are new extensions to provide support for iterative methods used in the solution of large-scale linear systems, including commonly used nonstationary methods such as Conjugate Gradient (CG), Generalized Minimal Residual (GMRES), Quasi-Minimal Residual (QMR), and related variants. 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.

OOF: Finite Element Analysis of Material Microstructures

Stephen Langer
Andrew Reid (Drexel University)
Seung-Ill Haan (U. of MD, Baltimore Co.)
Edwin Garcia (Penn State University)
Eric Ma (Montgomery Blair High School)
Kyle Stemen (Kent State University)
Edwin Fuller (NIST MSEL)
Craig Carter (MIT)

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

The OOF Project, a collaboration between MCSD, MSEL's Ceramics Division and Center for Theoretical and Computational Materials Science, and MIT, 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.

OOF is undergoing a transformation from the current version, OOF1, to a completely new version, OOF2. OOF1, which is freely available on the web, solves elasticity and thermal conductivity problems, while OOF2 will be easily extendable to a much wider variety of physics. OOF1 is therefore in a maintenance mode, where new features are not being added. During FY 2003 a few minor bugs were corrected and the code was upgraded to satisfy the requirements of new compilers, notably those used on Macintosh OS X.

OOF1 continues to be downloaded from the web at roughly the same rate as previous years. During FY 2003, there were approximately 1,600 downloads of the compiled OOF programs, and 700 downloads of the source code. The developers were asked to contribute four review articles to journals and books. OOF is being used in an educational context at North Carolina Agricultural and Technical State University in Greensboro, NC, and at the University of Washington in Seattle. Members of the OOF team traveled to Greensboro to give an introduction and tutorial on OOF to members of the NCAT Materials Science Department.

OOF2 is being written in a mix of C++ and Python, an object-oriented scripting language. OOF2 development during FY 2003 has been addressing the large number of tasks that must be completed before an initial release of the code. The emphases have been on preserving extensibility and maintainability through proper object-oriented design, preserving generality by making only minimal assumptions about the problems being solved, and preserving usability by designing a clear and intuitive user interface. Major accomplishments this year included:

- The first solution of a problem, from initial micrograph to stress analysis, completely within the OOF2 graphical user interface.
- New finite element mesh modification routines, including snapping mesh points to the boundaries between materials in the microstructure, and automatically removing badly shaped elements.
- Independent undo and redo operations for mesh modifications and mesh selections. Many OOF2 operations work on the currently selected mesh elements (or nodes, or segments). Those operations also modify the mesh, so the selection state must be propagated from one version of the mesh to another. Usability dictates that the user be able to independently undo and redo both the mesh selection and the mesh modification.
- A unified system for saving data files in any of three formats: Python scripts, which are editable and flexible, but insecure; non-Python ASCII, which is editable and secure, but inflexible; and binary, which is not editable or flexible, but is secure and compact. Security is

important because the developers eventually hope to establish an on-line library of OOF2 material definitions and microstructures, and needs to guarantee that the data files do not contain arbitrary executable python code.

- The first steps toward parallelization. OOF2 can now use the PETSc library of parallel matrix solvers, and the core of the program has been modified to support separate threads of control.
- Ports to new computer systems, including Macintosh OS X and Microsoft Windows. The Macintosh is now a primary development platform (along with Linux and SGI IRIX). The Windows port was a feasibility demonstration.

Computational Geometry Using Exact Arithmetic

Javier Bernal

The computation of three-dimensional Delaunay or regular triangulations requires two geometric tests. The first decides on which side of a plane spanned by three points a fourth point lies. The second decides on which side of the surface of a sphere spanned by four points a fifth point lies. Both tests require the determination of the sign of certain determinants. However, due to round-off error, the evaluation of such determinants in floating-point arithmetic may result in small non-zero values being obtained for determinants whose values are actually zero. A so-called "zero tolerance" is usually used for deciding whether a non-zero number is small enough to be considered zero. However, such an approach can lead to wrong decisions being made which in turn can lead to erroneous geometrical results.

In order to avoid round-off problem, J. Bernal has developed a scheme for carrying out additions, subtractions and multiplications using exact arithmetic. The scheme is based on operations using both double precision and integer arithmetic. Under this scheme input coordinates of points can have as many as 14 significant figures with as many as 9 significant figures to either the left or the right of the decimal point. Since exact arithmetic is used only when necessary the execution speed of the programs is only affected slightly. Currently this scheme is being incorporated into computational geometry software available at MCSD Web site.

Information Services for Computational Science

Ronald Boisvert *Bruce Miller*
Joyce Conlon *Roldan Pozo*
Marjorie McClain

http://math.nist.gov/
http://gams.nist.gov/
http://math.nist.gov/MatrixMarket/

MCSD continues to provide Web-based information resources to the computational science research community. The first of these is the Guide to Available Mathematical Software (GAMS), a cross-index and virtual repository of some 9,000 mathematical and statistical software components of use in science and engineering research. Both public domain and commercial software supported for use on NIST central computers is cataloged, as well as software assets distributed by *netlib*. GAMS users locate software via several search mechanisms. The most popular is the GAMS Problem Classification System, which provides a tree-structured taxonomy of standard mathematical problems solvable by extant software. Major math software library vendors have also adopted it.

A second resource provided by MCSD is the Matrix Market, a visual repository of matrix data used in the comparative study of algorithms and software for numerical linear algebra. The Matrix Market database contains more than 400 sparse matrices from a variety of applications, along with software to compute test matrices in various forms. A convenient system for searching for matrices with particular attributes is provided. The web page for each matrix provides background information and visualizations.

Web resources developed by MCSD continue to be among the most popular at NIST. The MCSD Web server at math.nist.gov has serviced more than 64 million Web hits since its inception in 1994 (12 million of which have occurred in the past year!). The Division server regularly handles more than 10,000 requests for pages each day, serving more than 463,000 distinct hosts on a yearly basis. Altavista has identified more than 8,000 external links to the Division server. Seven MCSD sites are listed in ITL's top web sites:

1. NIST Math Portal, http://math.nist.gov/
2. Matrix Market, http://math.nist.gov/MatrixMarket/
3. Guide to Available Mathematical Software: http://gams.nist.gov/
4. Division home page: http://math.nist.gov/mcsd/
5. ACM Transactions on Mathematical Software: http://math.nist.gov/toms/
6. Digital Library of Mathematical Functions: http://dlmf.nist.gov/
7. Template Numerical Toolkit: http://math.nist.gov/tnt/

The GAMS home page is downloaded more than 34,000 times per month by some 152,000 hostnames. The Matrix Market has distributed more than 14 Gbytes of matrix data, including nearly 160,000 matrices, since its inception.

High Performance Computing and Visualization

Virtual Cement and Concrete Testing Laboratory

Robert Bohn
Judith Devaney
William. George
Terence Griffin
John Hagedorn
Howard Hung
Peter Ketcham
Steve Satterfield
James Sims
Jack Douglas (NIST MSEL)
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 been collaborating with BFRL in the parallelization of their codes and in creating visualizations of their data. In January 2001 the Virtual Cement and Concrete Testing Laboratory (VCCTL) Consortium was formed. MCSD assisted in this effort through presentations of our work with BFRL and demonstrations of visualizations in our immersive environment. The consortium consists of NIST (BFRL and ITL) and nine industrial members: Cemex Trademarks Worldwide, Ltd., Holcim (US) Inc., Master Builders Technologies, National Ready Mixed Concrete Association, Association Technique l'Industrie des Liant Hydrauliques (ATILH), International Center for Aggregate Research (ICAR), W.R. Grace, Sika Technology AG, and Portland Cement Association. The overall goals of the consortium are to develop a virtual testing system to reduce the need for physical concrete testing, and to expedite the research and development process. This will result in substantial time and cost savings to the concrete construction industry as a whole. MCSD continues to contribute to the VCCTL through collaborative projects involving parallelizing and running codes, creating visualizations, as well as presentations to the VCCTL current and prospective members. The following projects are included in this effort.

Computational Modeling of the Flow of Concrete. Understanding the flow properties of complex fluids like suspensions (e.g., colloids, ceramic slurries and concrete) is of technological importance and presents a significant theoretical challenge. The computational modeling of such systems is also a great challenge because it is difficult to track boundaries between different fluid/fluid and fluid/solid phases. N. Martys and E. Garboczi of BFRL are utilizing a new computational method called dissipative particle dynamics (DPD), which has several advantages over traditional computational dynamics methods while naturally accommodating necessary boundary conditions. In DPD, the interparticle interactions are chosen to allow for much larger time steps so that physical behavior, on time scales many orders of magnitude greater than that possible with molecular dynamics, may be studied. Our algorithm (QDPD) is a modification of DPD, which uses a velocity Verlet algorithm to update the positions of both the free particles and the solid inclusion. In addition, the rigid body motion is determined from the quaternion-based scheme of Omelayan (hence the Q in QDPD). J. Sims has developed a parallelized version of the algorithm, which is needed in order to adequately model size distributions, and to have enough resolution to avoid finite size effects. T. Griffin and S. Satterfield have developed visualizations of model output.

A paper on simulations of sheared suspensions has been submitted to *Computer Physics Communications*. Modeled aggregate (see below) have been added to the code since the rheology can be greatly affected by aggregate shape. The code has also been modified to incorporate fibrous materials. In the future, the QDPD code will be modified to include lubrication forces so that very

dense suspensions can be modeled. During the next year we will be adapting the code to run on Federally supported supercomputers. A supercomputer access proposal entitled "Modeling the Rheological Properties of Concrete" was submitted to the INCITE project at the National Energy Research Scientific Computing Center (NERSC), Lawrence Berkeley National Laboratory by W. George, J. Hagedorn, J. Sims, and J. Devaney in collaboration with N. Martys of BFRL.

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

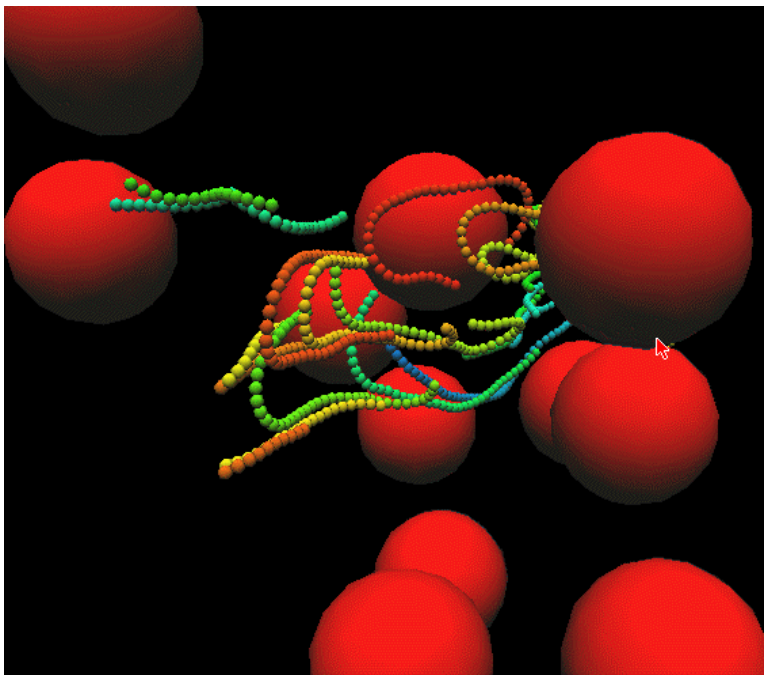

Influence of fibers on the rheology of cement-based materials.

Parallelization, Visualization of Fluid Flow in Complex Geometries. Fluid flow in complex geometries plays an important role in many environmental and technological processes. Examples include oil recovery, the spread of hazardous wastes in soils, the processing of polymer blends, droplet breakup, phase separation and chemical analysis in confined geometries, and the service life of building materials. The latter two applications are of particular concern for our NIST collaborators in BFRL and MSEL. The detailed simulation of such transport phenomena subject to varying geometries, environmental conditions or saturation, is a great challenge because of the difficulty of modeling fluid flow in random pore geometries and the proper accounting of the interfacial boundary conditions.

In order to model realistic systems, N. Martys of BFRL has developed a lattice Boltzmann (LB) algorithm that simulates multiple fluids, various forces, and wetting characteristics within arbitrary geometries. J. Hagedorn and J. Devaney parallelized the algorithm using MPI to enable the study of large systems. We have run a series of simulations of multiple fluid systems in a variety of confined geometries such as between parallel plates and within tubes. These simulations were designed to investigate the effects of confinement, lattice spacing, and discretization. The simulations have yielded results that agree well with experimental results while advancing our understanding of onset time and evolution of capillary driven instabilities in confined geometries. Hagedorn and Devaney also modified the code to solve the Brinkman equation, which is useful for modeling flow in

multi-scale porous media. In a recent paper, published in *Materials and Structures*, we reviewed our work concerning the lattice Boltzmann methods. The paper "Suppression of Capillary Instability of a Polymeric Thread via Parallel Plate Confinement" was published in the July 29, 2003 issue of *Macromolecules*. A paper entitled "Breakup of a Fluid Thread in a Confined Geometry: Droplet-Plug Transition, Perturbation Sensitivity, and Kinetic Stabilization with Confinement" reporting joint work with J. Douglas of MSEL has been accepted by *Physical Review E*.

In the future, we plan to extend the algorithm to accommodate a wider range of realistic material characteristics. We also plan to implement modifications such as adaptive grid techniques and varying time steps in order to enhance accuracy of results. Flow in tomographic-based images of paper and microfluidic devices will be studied.

http://math.nist.gov/mcsd/savg/parallel/lb/
http://math.nist.gov/mcsd/savg/vis/fluid/

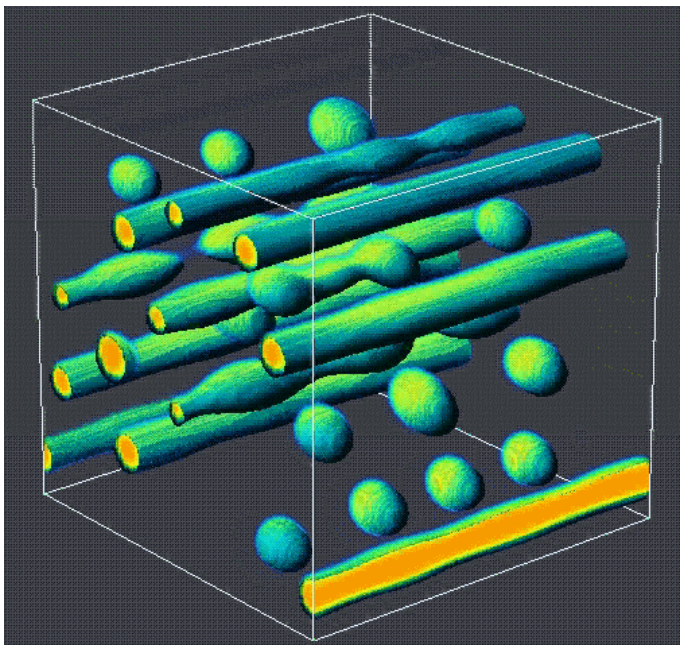

Lattice Boltzmann simulation of phase separation under steady shear.

Parallelization of a Model of the Elastic Properties of Cement. R. Bohn has parallelized computer codes originally developed by E. Garboczi of BFRL, that compute the linear electric and elastic properties of random materials based on 3D images of actual materials. Part of the intrinsic error that comes about with the use of 3D digital images to represent microstructure is digital resolution error. This can be quantified and eliminated by investigating the same problem at several resolutions, proceeding to large enough systems to see the asymptotic behavior. The parallel implementation of the existing serial elastic code enables this to be done. One application of interest to the VCCTL is the early age elastic properties of cement paste. To properly resolve the small necks of material that are holding the solid backbone together requires the fine digital resolution made possible with the parallel code. These codes calculate electrical and thermal responses via finite element techniques and responses to AC/DC fields using finite difference techniques. Documentation of the parallelization was published as NISTIR 6967 on June 19, 2003, entitled, "User Manual for Finite Element and Finite Difference Programs. A Parallel Version of NIST IR 6269."

http://math.nist.gov/mcsd/savg/parallel/epc/

Immersive Visualization of Flows of Concrete Aggregate from Scanned Images. In support of the VCCTL, T. Griffin has worked with N. Martys and E. Garboczi of BFRL to use X-ray tomography to create a database of realistic aggregates spanning about four decades in size. Examples include cement particles, sand and, even, rocks. These realistic aggregate shapes can be incorporated into codes used to model the rheological properties of cement based materials. The purpose of this project is to develop techniques to display flows of multiple types of aggregate in an immersive visualization environment.

http://math.nist.gov/mcsd/savg/parallel/dpd/

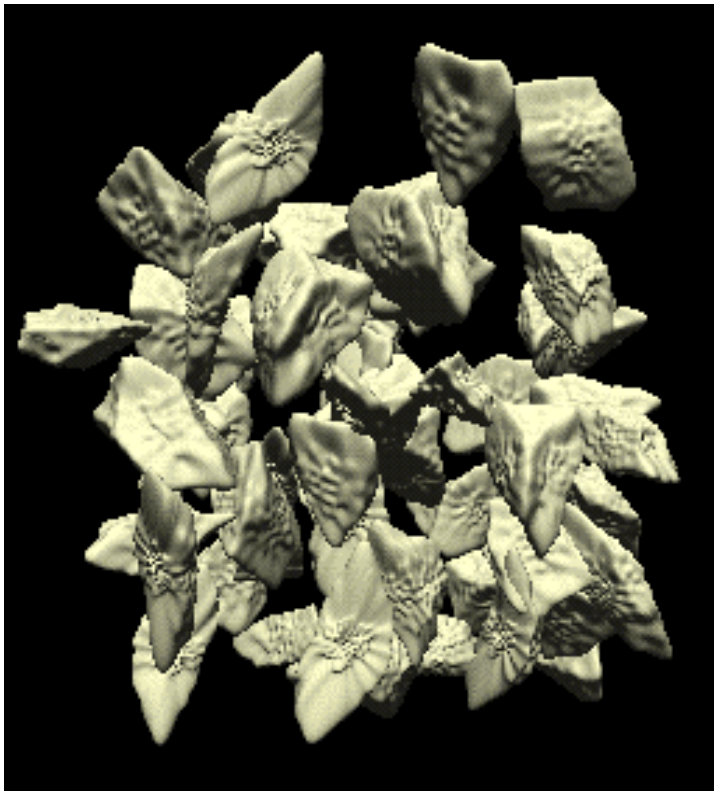

Image of modeled actual aggregate in a shear flow.

Database on Sulfate Attack. NIST has a large amount of cement data including mixture design, cement characteristics, macroscopic measurements (expansion and modulus of elasticity), and microscopic observations (X-Ray, SEM). H. Hung has been working with C. Ferraris of BFRL to organize this data for use by the VCCTL. An Access database is being constructed data related to sulfate attack. Eventually, it will be expanded to organize all data linking macroscopic and microscopic measurements. Finally, the database will be linked to the VCCTL and will be accessible for search from outside NIST. Where possible, data will be collected from other researchers once the database is made public. This data will ultimately be used as the basis for model validation for the prediction of the cement performance. Currently SAVG has created a general interface and search engine to the Access Database to simplify interaction with, and display of results from, the database.

Visualization of Smart Gels

Steve Satterfield
Carlos Gonzalez (NIST CSTL)

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

A smart gel is a material that gels in response to a specific physical stimulus. For example, it may gel at a specific temperature or pressure. The mechanisms that create a gel in response to a given environment are not well understood. Developing this understanding is key to being able to engineer materials that gel at designated points. The potential for applications of smart gels is enormous. Since smart gels expand or contract in response to external stimuli, they could be useful in applications such as an artificial pancreas that releases insulin inside the body in response to high sugar levels. Smart gels might someday be used to make exotic foods, cosmetics, medicines, and sensors.

The NIST team is studying a subclass of these materials called shake gels. Through some as yet unknown process, these watery mixtures of clays and polymers firm up into gels when shaken, and then relax again to the liquid phase after some time has passed. A shake gel might be used, for example, in shock absorbers for cars. The material would generally be a liquid but would form a gel when the car drove over a pothole; the gel thickness adjusting automatically to the weight of the car and the size of the pothole. A more esoteric application might be the formation of gelled areas within a liquid where holograms could be created using a laser.

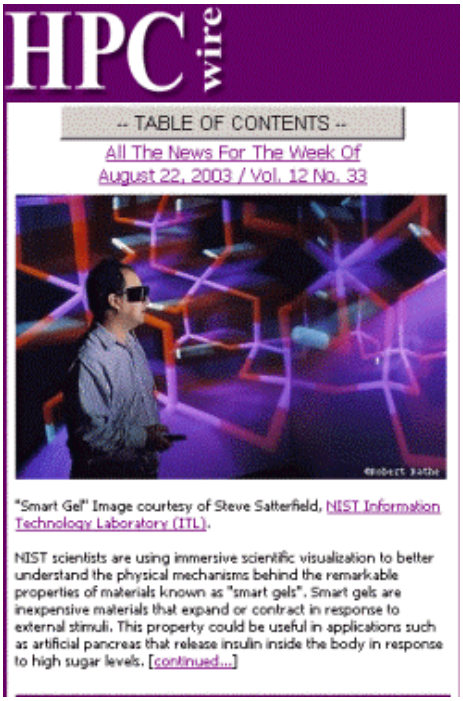

S. Satterfield collaborated with C. Gonzalez of CSTL to develop immersive visualizations of the output of numerical simulations of shake gels for molecular display. These enabled NIST chemists to answer questions that had defied attempts at solution. The 3D visualization helped the scientists see that for the shake gel it is the water's oxygen atoms, instead of the hydrogen atoms as previously thought, that attach to the clay. Future work will include visualization of the polymer in addition to the water and clay molecules.

The Shake Gel visualization work has received the following publicity:

- S. Satterfield worked with Gail Porter, Laura Ost and Ronald Meininger, NIST Public and Business Affairs Division, to create a Quicktime video presentation on the collaborative work of Satterfield and Gonzalez. The video titled "Recipe for a Shake Gel" was featured on the NIST home page in August 2003. It is permanently available at http://www.nist.gov/public_affairs/newsfromnist_smartgels.htm.
- S. Satterfield and C. Gonzalez put together a submission to the web based technical news publication HPCwire (www.hpcwire.com) on the shake gel collaboration. The submission, including an image, was used on the cover of HPCwire for the week of August 22, 2003, Volume 12, No. 33.
- A Quicktime video of the smart gel immersive visualization was submitted to SC2003 for use in the conference video.

- Reporters from CNN visited the MCSD immersive visualization lab in December 2003, interviewing J. Devaney and S. Satterfield for their Saturday science feature NEXT@CNN. The feature aired in January 2004.

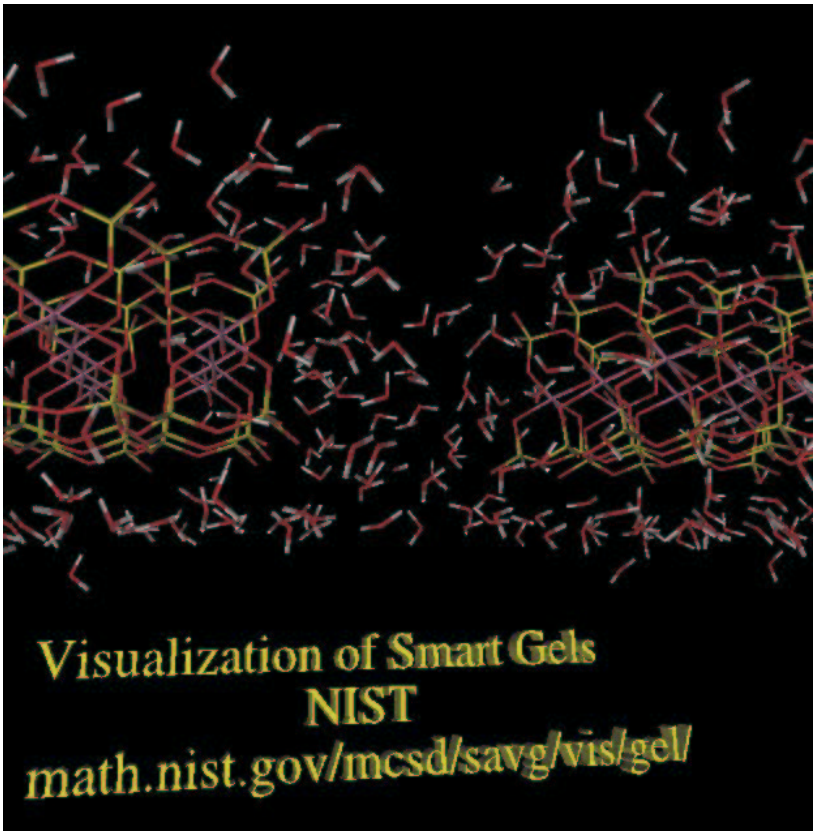

Computational Models of Microelectronic and Optoelectronic Devices

Howard Hung
Herbert S. Bennett (NIST EEEL)

The majority electron density as a function of the Fermi energy has been calculated in zinc Blende, n-type GaSb for donor densities between 10^{16} cm^{-3} and 10^{19} cm^{-3}. These calculations solve the charge neutrality equation self-consistently for a four-band model (three conduction sub-bands at Γ, *L*, and *X* and one equivalent valence band at Γ of GaSb. The results are important for interpreting optical measurements such as Raman measurements that are proposed as a non-destructive method for semiconductor wafer acceptance tests. Raman measurements give values of the Fermi energy. Their correct interpretations need physically reasonable physics models for electron density as a function of the experimental Fermi energy. The results were transferred to CSTL and to others in industry and universities. A joint paper entitled "Dependence of Electron Density on Fermi Energy in N-Type Gallium Antimonide" was published in the May-June 2003 issue of the NIST *Journal of Research*.

Multi-Modal Imaging and Visualization

Peter Ketcham *Joy Dunkers (NIST MSEL)*
Steven Satterfield *Marcus Cicerone (NIST MSEL)*
John Hagedorn *Lyle Levine (NIST MSEL)*
John Kelso *Gabrielle Long (NIST MSEL)*
Judith Devaney

Modern materials science research must address real systems and processes. In many cases, we can no longer advance science simply by studying model systems that are idealized in dimension and function. We must comprehend realistic, complex, three-dimensional systems in terms of their structure, function, and dynamics over a broad scale from nanometers to millimeters. In this ATP-sponsored collaboration with the NIST MSEL we are combining data from different physical measurement techniques that reflect both functional and structural information. The combined data from a single sample is then visualized in our interactive, immersive, virtual reality environment in order to gain new insights into the physics and materials science of complex systems.

Our collaborators are gathering measured data using a variety of techniques, including optical coherence tomography (OCM) and confocal fluorescence (CFM) imaging. These provide different types of information on a sample. When combined in a manner that is visually apparent, this can yield unprecedented insight into the complex relationships among large quantities of correlated data. These methods have applications to the characterization of biomaterials, the failure analysis of polymer composites, and the reliability of semiconductor devices.

We applied three different rendering techniques in order to transform from the scientific space to the visualization space to create graphical representations: volume rendering, surface rendering, and their combination. Each is described in more detail below.

Volume-Rendering Techniques. We applied volume-rendering techniques to the OCM/CFM data. Volume-rendering techniques utilize an algorithm in which every point along a line of sight both emits and absorbs (virtual) light to an arbitrary extent.

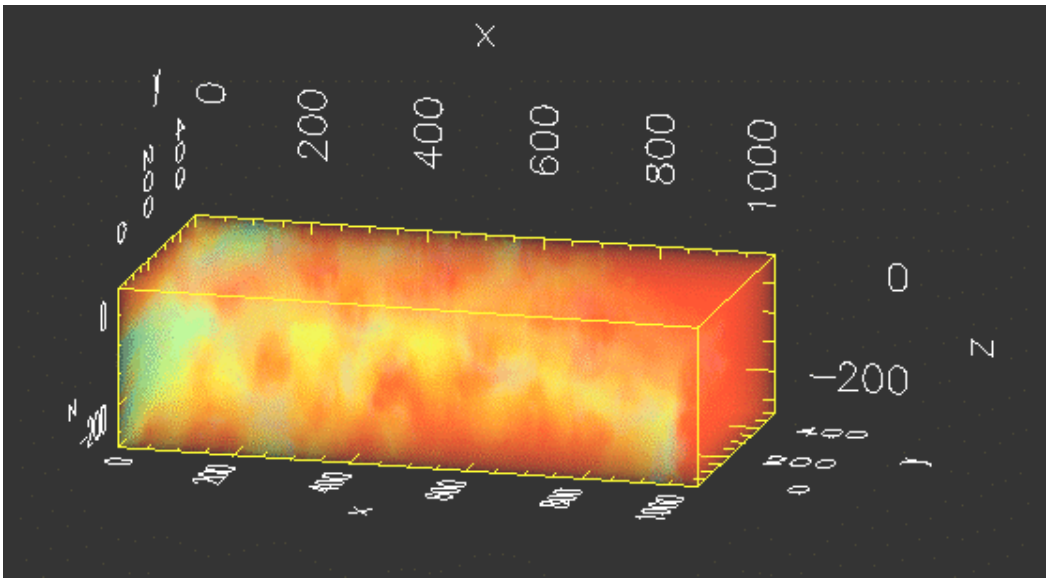

Volume rendering of OCM data.

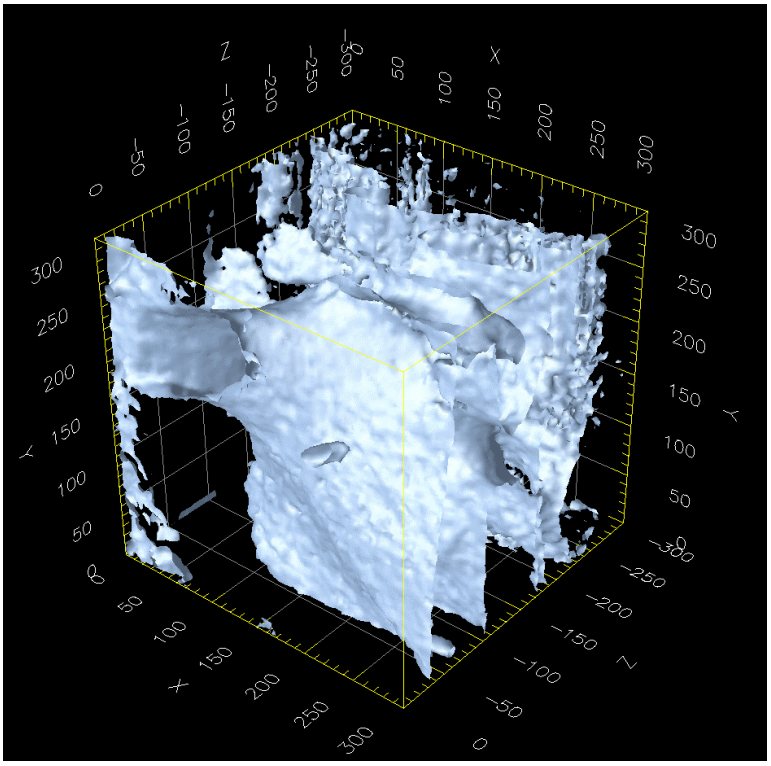

Surface rendering of OCM data at an isosurface at an isovalue of 2.0. This is viewable in the immersive visualization environment.

Surface-Rendering Techniques. For our next approach, we applied surface-rendering techniques to the OCM data. With surface-rendering techniques, a geometric representation is constructed from the data and (virtual) light is reflected off the surfaces of this representation. The following figure shows a surface-rendered image of OCM data with and without the application of a three-dimensional median filter. The purpose of the three-dimensional median filter is to reduce noise in the measured OCM data. We have investigated constructing geometric representations through a variety of means. We have implemented automated software programs to process the OCM/CFM data sets and convert the data to a geometric representation that can be displayed interactively. We are currently working on translating this geometry into a form suitable for display in our immersive environment.

Combined Volume/Surface-Rendering Techniques. We have also experimented with combined volume/surface-rendering techniques. In this method, the goal is to extract a scientifically interesting feature from an MMIV data set and to highlight it in context. The extraction can be done as an isosurface, through a segmentation process, or through the application of a transfer function to bring out detail in the volume data that might otherwise be hidden. Multiple views of the data can be accessed through the use of wand buttons that turn different features on and off to interactively highlight interesting aspects of the volume data. By combining the volume/surface-rendering techniques, interesting aspects of the data can be seen *in situ*. This is important in this application where understanding the factors that affect cell growth is important.

For the future we intend to focus on methods for interactivity, transparency techniques, fast methods for moving surface visualizations to the immersive environment, and combined surface-volume visualizations in the immersive environment.

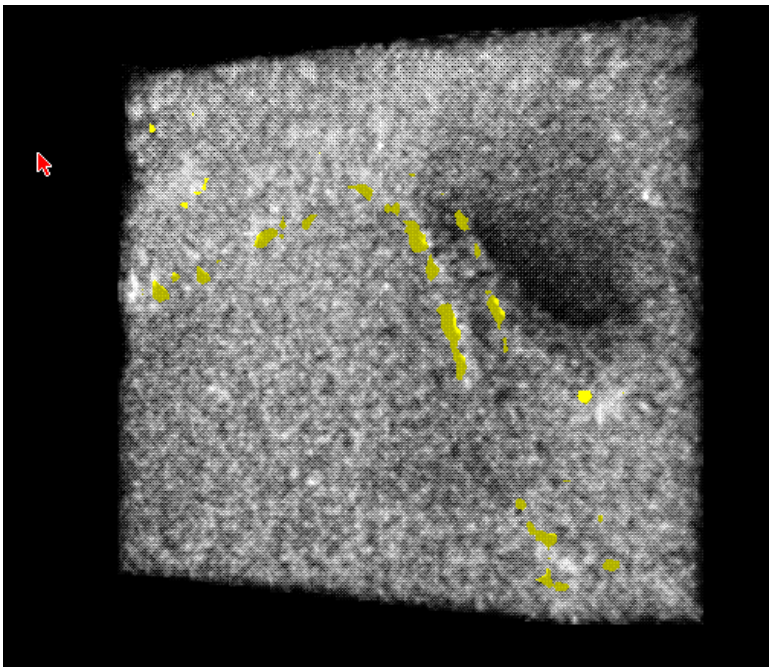

Combined volume/surface-rendering of OCM data. The yellow isosurface is drawn as polygon data and while embedded in the volume, clearly shows its structure. This is viewable in the immersive visualization environment.

Computation and Visualization of Nanostructures and Nano-optics

James Sims *Adele Peskin*
Howard Hung *Garnett Bryant (NIST PL)*
Steve Satterfield

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

Accurate atomic-scale quantum theory of nanostructures and nanosystems fabricated from nanostructures enables precision metrology of these nanosystems and provides the predictive precision modeling tools needed for engineering these systems for applications including advanced semiconductor lasers and detectors, single photon sources and detectors, biosensors, and nanoarchitectures for quantum coherent technologies such as quantum computing. Theory and modeling of nanoscale and near-field optics is essential for the realization and exploitation of nanoscale resolution in near-field optical microscopy and for the development of nanotechnologies that utilize optics on the size-scale of the system. Applications include quantum dot arrays and quantum computers. Atomic-scale theory and modeling of quantum nanostructures, including quantum dots, quantum wires, quantum-dot arrays, biomolecules, and molecular electronics, is being used to understand the electronic and optical properties of quantum nanostructures and nanosystems fabricated from component nanostructures. Theory and numerical modeling is being used to understand optics on the nanoscale and in the near field with applications including near-field

microscopy, single-molecule spectroscopy, optics and quantum optics of nanosystems, and atom optics in optical nanostructures.

Parallelization. This year MCSD's parallelization of G. Bryant's nanostructures code has been extended. Benchmarking runs on NIST's Linux cluster show that the new parallel code can handle 200,000 atoms with approximately 10 CPUs. Further parallelization resulted in a 41 processor working parallel code with a speedup of a factor of 27. These runs represent a significant achievement. For the future a mixed-model multi-level parallelization is being developed for further speedups.

Visualizations. An interactive immersive object manipulation application that can be used to view and setup nanostructures is being developed for this project. However, it is generic and can be used in other applications. A preliminary experiment in displaying the data as a volume visualization has also been conducted. More work is needed to determine if this idea will be useful.

Proposal Submission. A supercomputing proposal entitled "Complex Nanobiohybrid Structures for Nanoelectronics and Nanobiotechnology" was submitted to the INCITE project at National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory by J. Sims, W. George, H. Hung, and J. Devaney in collaboration with G. Bryant of the NIST PL.

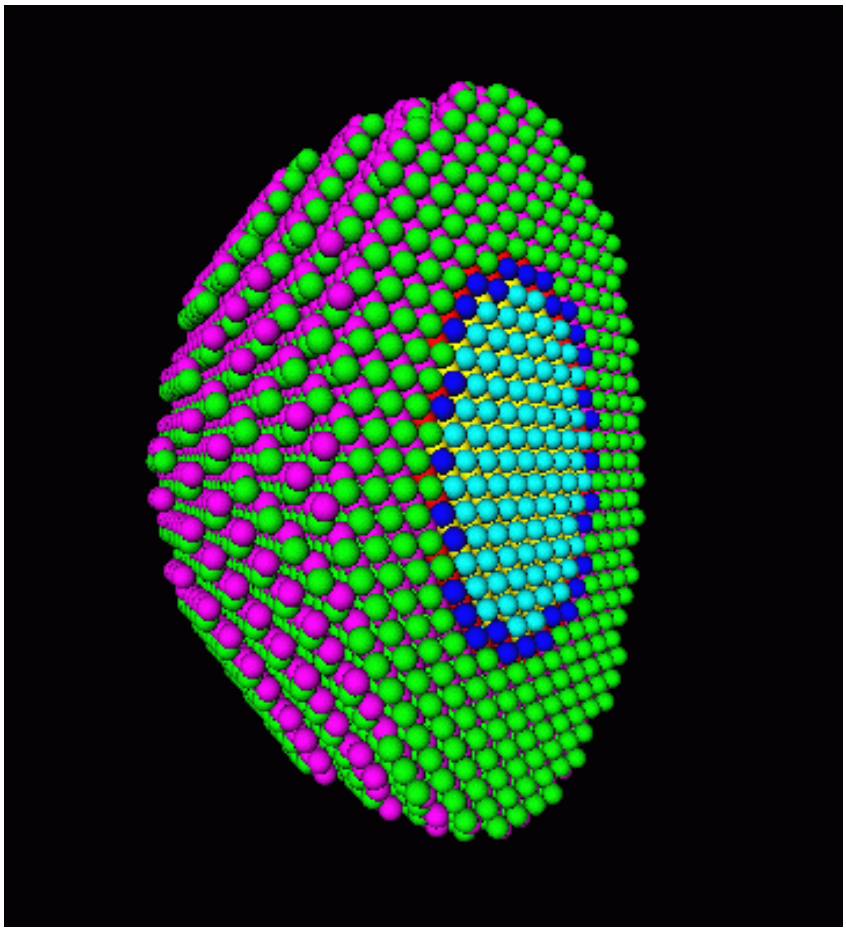

This visualization of concentric spheres was developed to test the MCSD system for visualization of nanostructures.

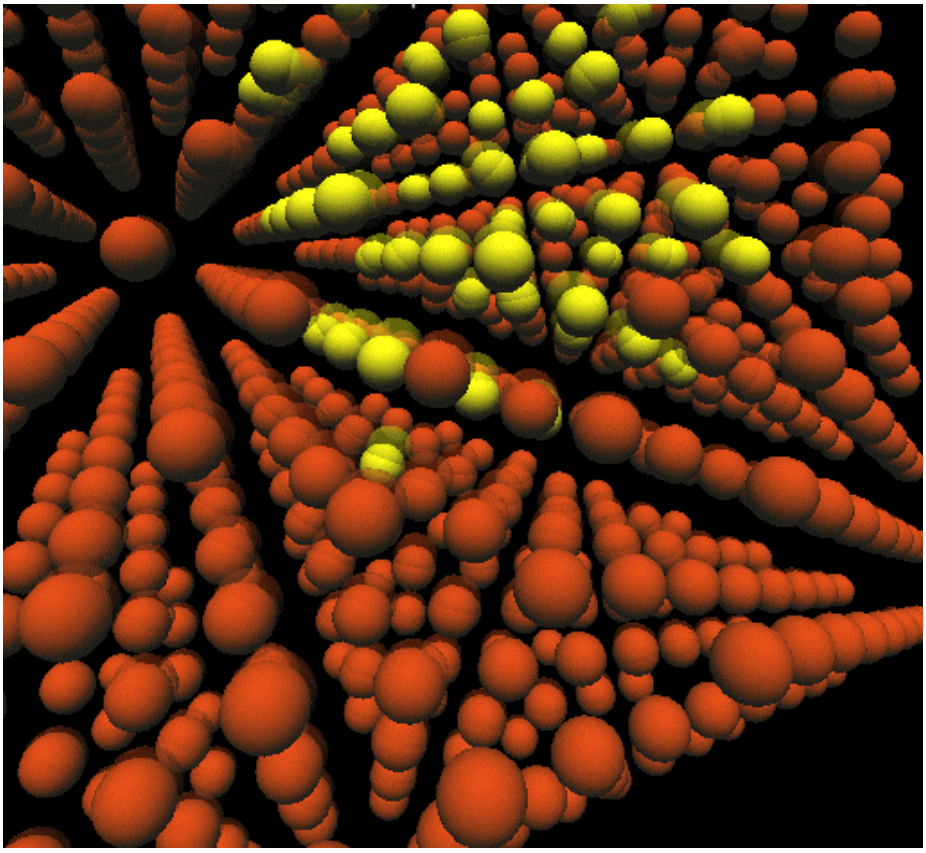

A visualization of a copper-gold nanostructure interface.

Visualization of Nanostructures for Material Reliability

Adele Peskin
Vinod Tewary (NIST MSEL)
Dave Reed (NIST MSEL)

MSEL's Materials Reliability Division (MRD) has developed an efficient algorithm for multiscale modeling of nanostructures consisting of a million atoms. The physical consequences of interactions among such a large number of atoms are not easy to interpret, even if the basic physics is well understood. Scientific visualization is an important tool for understanding the physical processes in such systems by allowing detailed examination of atomic positions and displacements. MCSD is helping MRD researchers by providing software tools for immersive visualization of nanostructures that can be eventually integrated with the mathematical model developed by the MRD.

A. Peskin created a demonstration of 17,000 atoms in two different spatial configurations as well as a subset of 1099 atoms. Spheres were used to represent the atoms; transparent spheres were used to represent the shifted positions of the atoms. She developed a preliminary system for the visualization of three-dimensional atomistic configurations of a gold nanoisland in copper. These were demonstrated in a September 2003 presentation for MSEL Director Les Smith and MRD Chief Alan Clark on the visualization of the copper-gold atomic structure.

Visualization of Macromolecules

Terence Griffin
Jack Douglas (NIST MSEL)
Marc Mansfield (Stevens Institute of Technology)

As part of a demonstration of a new path integration technique for the computation of both the translational diffusivity and the intrinsic viscosity of macromolecules, T. Griffin created several visualizations of large protein structure with correctly scaled atoms for J. Douglas of NIST MSEL and Marc Mansfield of the Stevens Institute of Technology. The images were designed to reveal internal structure. They appear in the paper "Numerical Path Integration Technique for the Calculation of Transport Properties of Proteins," to appear in *Physical Review E*.

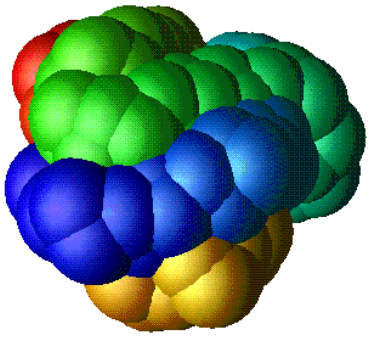

Visualization of the protein myoglobin. This representation corresponds to placing a sphere of 0.5 nm placed at the alpha carbon of each amino acid residue.

Interoperable MPI

William George
John Hagedorn
Judith Devaney

http://impi.nist.gov

The Interoperable MPI (IMPI) project supports the vendors of MPI, as they implement the IMPI 0.0 protocols, by maintaining the NIST IMPI conformance tester, managing the IMPI mailing list (interop@nist.gov), maintaining the IMPI specification document and its errata, and promoting the IMPI implementations. The following are accomplishments of the past year.

- W. George consulted with engineers from MPI Software Technology, Inc. as they worked on implementing the IMPI protocols. Several questions about specific IMPI protocols were answered and several small bugs were uncovered in the collective communications algorithms as presented in the IMPI specification. Corrections were added to the IMPI Errata and made available on-line.
- MPI Software Technology has made extensive use of the on-line NIST IMPI tester in the last year. They completed the implementation of IMPI for their commercial MPI/Pro library and demonstrated it at the SC 2003 conference in November in Phoenix.
- The High Performance Computing Center of the University of Stuttgart (Germany), has been using the NIST IMPI tester extensively as they develop their MPI library, PACX-MPI, which includes IMPI. (See http://www.hlrs.de/organization/pds/projects/pacx-mpi.)
- W. George corresponded briefly with Bernhard Klein of LogicaCMD, http://www.logicacmg.com, an IT company in Great Britain that is beginning to use IMPI with LAM/MPI in an internal project.
- Collaboration was initiated with Andrew Lumsdaine and Jeff Squyres of Indiana University to produce a journal article on the use of IMPI in scientific applications. The intended journal is the *IEEE Transactions on Parallel and Distributed Systems*. A short article "Parallel Programming with Interoperable MPI" has been published as NIST Interagency Report 7066. A version of this article will also appear in the February 2004 issue of the magazine *Dr. Dobb's Journal*.

Screen Saver Science

William L. George
Samuel Small (Johns Hopkins U.)
Jacob Scott (U. of California, Berkeley)

http://math.nist.gov/mcsd/savg/parallel/screen/

The Screen Saver Science (SSS) project aims to develop a computing resource composed of a heterogeneous set of PCs, scientific workstations, and other available computers, that can be easily used by scientists to execute large highly distributed, compute-intensive applications. Each individual computer in this system would make itself available for participating in a computation only when it would otherwise be idle, such as when its screen saver would be running. This project is based on Jini, an open software architecture built in Java and intended for the development of robust network services.

SSS has several goals. First, we hope to utilize the idle processing power of the many PCs and workstations we have available here at NIST to execute scientific production codes. The compute power of personal PCs and workstations continues to increase, becoming more capable of executing large compute intensive applications due to faster processors and larger main memories. Second, the research on grid computing has been accelerating and the SSS environment will allow us to develop and experiment with new highly parallel and distributed algorithms more suitable for emerging grid environments. Finally, the use of Java for scientific applications is of interest in general and so the development of applications for SSS will give us the opportunity to explore this topic on actual production quality applications.

Up until recently, this type of project would have required a large investment in software development just to become minimally functional and so was not practical, especially for a small team of programmers. However, with the introduction of Jini, and more specifically the Jini based network service called Javaspaces, the most difficult parts of this project have now become trivial. Javaspaces is a portable, machine independent, shared memory system that expands upon the tuple-space concepts developed in the 1980s by David Gelernter of Yale University.

The SSS project began in the summer of 2002. Substantial progress has been made in designing, implementing, and testing the basic SSS infrastructure. A generic compute server has been implemented and a new Jini service, a remote file server, has been developed to provide SSS applications with basic file I/O capabilities. A backup-restore facility has also been implemented to allow the system to be temporarily shutdown and moved or modified without loosing any statistical information, such as the number of tasks completed and computation time for these tasks, as well as user preferences that are stored in the SSS system.

W. George collaborated with S. Small, a 2003 SURF student, to design a job scheduler for SSS. This job scheduler is an independent Jini service that can be used to guarantee that no job submitted to SSS will be delayed from executing indefinitely (starvation-free scheduling). This design also implements a simple job preference scheme to allow some control over the execution of the available jobs, giving some jobs priority over others. Small presented this work at the SURF Colloquium in August 2003.

The control scripts needed for automatically starting and stopping an SSS Server on Unix systems have been completed. Existing hooks into the X Display Manager (XDM) as well as hooks into the popular open source X11 screen saver "xscreensaver" are used to control the operation of the SSS server. Testing has been performed on SGI IRIX systems but no SGI specific techniques were used so these scripts will also work on other Unix systems such as HP UX, IBM AIX, Sun Solaris, and Linux systems. A similar technique is under investigation for Windows based systems. These scripts enable the running of SSS tasks under two separate conditions: when no user is logged into the system, and when a user is logged in but is inactive (determined via the screen saver hooks).

The first demonstration application for SSS, a simple Monte Carlo algorithm for solving Laplace's equation in two dimensions has been completed. A companion program has also been written to monitor and display the progress of this computation graphically, showing the current state of each grid point as either: computed, submitted to SSS, running, or not yet submitted. A set of 12 SGI workstations has been configured to run SSS tasks for these tests. The computation was submitted from a separate Sun workstation. A version of this application was also developed that enables real-time display of the computation on the local immersive visualization system.

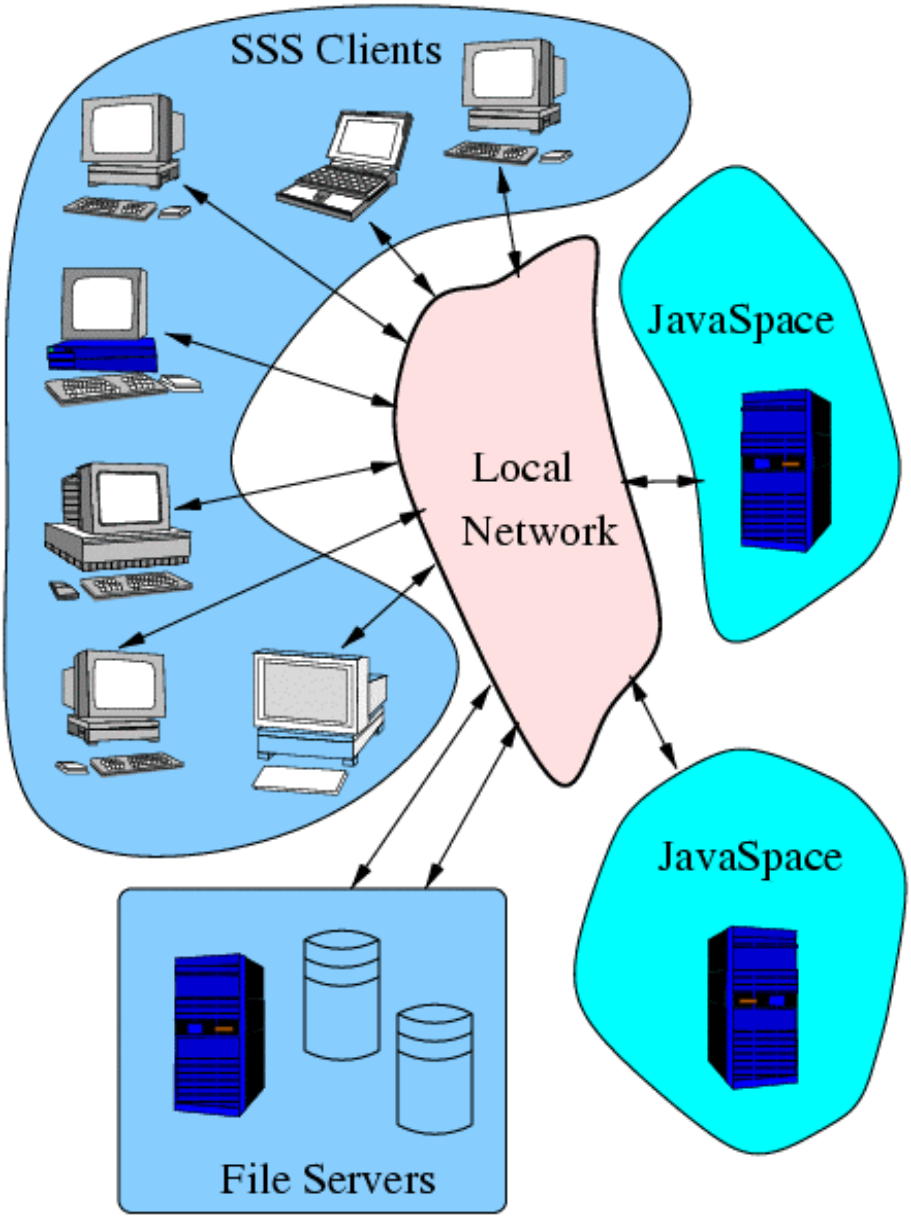

Schematic of the Screen Saver Science system.

Computation of Atomic Properties with the Hy-CI Method

James Sims
Stanley Hagstrom (Indiana University)

http://math.nist.gov/mcsd/savg/parallel/atomic/

Impressive advances have been made in the study of atomic structure, at both the experimental and theoretical levels. For atomic hydrogen and other equivalent two-body systems, exact analytical solutions to the nonrelativistic Schrödinger equation are known. It is now possible to calculate essentially exact nonrelativistic energies for helium (He) and other three-body (two-electron) systems as well. Even for properties other than the nonrelativistic energy, the precision of the calculation has been referred to as "essentially exact for all practical purposes", i.e., the precision goes well beyond what can be achieved experimentally. These high-precision results for two-electron systems have been produced using wave functions that include interelectronic coordinates, a trademark of the classic Hylleraas (Hy) calculations done in the 1920s.

The challenge for computational scientists is to extend the phenomenal He accomplishments (i.e., the ability to compute, from first principles alone, any property of any two electron atom or its ion to arbitrary accuracy) to three, four, and more electron atomic systems. Where three electron atomic systems (i.e., lithium (Li) and other members of its isoelectronic series) have been treated essentially as accurately as He-like systems, demand on computer resources has increased by 6,000 fold. Because of these computational difficulties, already in the four-electron case (i.e., beryllium (Be) and other members of its isoelectronic series) there are no calculations of the ground or excited states with an error less than 10 microhartrees (10^{-5} a.u.).

This is where a technique developed by Sims and Hagstrom in a series of papers from 1971 to 1976 becomes important. They developed the Hy-CI method, which includes interelectronic coordinates in the wave function to mimic the high precision of Hy methods, but also includes configurational terms that are the trademark of the conventional Configuration-Interaction (CI) methods employed in calculating energies for many-electron atomic (and molecular) systems. Because of this, the Hy-CI method has been called a hybrid method. This is the power of the method, because the use of configurations wherever possible leads to less difficult integrals than in a purely Hy method, and if one restricts the wave function to at most a single interelectronic coordinate to the first power, then the most difficult integrals are already dealt with at the four electron level and the calculation retains the precision of Hy techniques, but is greatly simplified.

In work published last year in the *International Journal of Quantum Chemistry*, Sims and Hagstrom computed the most accurate nonrelativistic energies for several states of Helium-like ions using the Hy-CI Method for two electrons, using extended precision arithmetic, large basis sets, and parallel computing. They used this technique 30 years ago to determine energy levels of Li (3 electrons), Li^- and Be (four electrons) much less accurately with supercomputers of the day.

In a paper about to be submitted to the *Journal of Physics B: Atomic, Molecular, and Optical Physics*, Sims and Hagstrom discuss changes they have made to their methodology to most effectively use modern day computers to further increase the size (number of terms) and accuracy of the calculations. They also discuss how to efficiently evaluate the only difficult integral arising when using this technique in the case of number of electrons $N \geq 3$, the three-electron triangle integral. They focus on recursive techniques at both the double precision and quadruple precision level of accuracy while trying to minimize the use of higher precision arithmetic. Also, they investigate the use of series acceleration to overcome problems of slow convergence of certain integrals defined by infinite series. They find that a direct + tail Levin u-transformation convergence acceleration overcomes problems that arise when using other convergence acceleration techniques, and is the best method for overcoming the slow convergence of the triangle integral. The techniques that are discussed are important in any attempt to extend high precision calculations to Be and beyond.

It is also important to know how *good* extrapolated results really are. In their quest for rapid numerical evaluation of these integrals, Sims and Hagstrom were led to seek a way to judge how to evaluate the accuracy of computed one-center three-electron correlated integrals. This quest led them to a paper by Remiddi giving an analytical formula for the simplest of these integrals. In the process of coding up Remiddi's Equation 42, they found two typos in Remiddi's paper, which led to a paper by Sims and Hagstrom published in *Physical Review A* this year in which they point out the two misprints in the paper and provide verification of the corrections with a short table comparing 30 significant digit results using Remiddi's formula (with their corrections) and the output of their recently developed triangle integral package.

The three-electron triangle integrals discussed by Sims and Hagstrom in these two papers have been the real bottleneck to highly accurate Hy-CI calculations. Now that this bottleneck has been removed, prospects for doing really accurate calculations on atoms with $N \geq 4$ becomes a real possibility. A calculation on Be (a 4 electron atom) is underway to test this hypothesis.

Tools for Scientific Visualization

Steven G. Satterfield *John Hagedorn*
Adele Peskin *Terence Griffin*
John Kelso *Judith Devaney*

http://math.nist.gov/mcsd/savg/software/applications/
http://math.nist.gov/mcsd/savg/software/dsos/
http://math.nist.gov/mcsd/savg/software/gtb/

Large and complex data sets are becoming more commonplace at NIST, as high performance parallel computing is used to develop high fidelity simulations, and combinatorial experimental techniques are used in the laboratory. Immersive visualization (IV) is significantly different from traditional desktop visualization and significantly more effective at illuminating such data sets. However, the benefits of IV can only be gained when scientists use it. The key ingredient to making IV accessible to scientists is to provide the ability to simply and quickly move their data into the immersive environment.

The primary software controlling the MCSD IV environment is an open source system called DIVERSE (Device Independent Virtual Environments – Reconfigurable, Scalable, Extensible). The DIVERSE Application Programming Interface (API) facilitates the creation of IV environments and asynchronous distributed simulations by handling many of the necessary details. The software runs on a variety of display devices from multi-wall stereographic displays with head tracking to desktop systems, even low-end Linux systems.

Included with DIVERSE is an extensible application called Diversify that allows various techniques for navigation through user data loaded from a variety of external data formats. On top of this DIVERSE/Diversify infrastructure, MCSD has developed additional tools and techniques for quickly moving research data into the IV environment, often with little or no special-purpose graphics programming. The approach used is to create small and reusable tools that fit into the immersive environment and which lend themselves to combination in a variety of ways to perform useful tasks, such as moving the results of a numerical simulation into the immersive environment. Based on this philosophy, MCSD is providing the key ingredient to making IV accessible to NIST scientists by developing

- simple, reusable graphics file formats as Diversify file loaders,
- application-specific files that are easily transformed into existing or newly created graphics file formats, and

- filters to connect data transformation pipelines through shared memory.

Diversifly is extensible. Distributed Shared Objects (DSOs) can be written to implement a new input graphics data format. These DSOs are much like subroutines (C/C++ functions). However, they are linked into the executable code at run time and the main Diversifly application need not be modified or recompiled. This year our group created the beginning of a useful set of tools for visualization with DIVERSE. These include a set of new DSOs to display discrete sets of data, methods to hide and display objects as needed, methods to communicate with outside programs, and tools for sound and lighting. The following are examples of our DSOs.

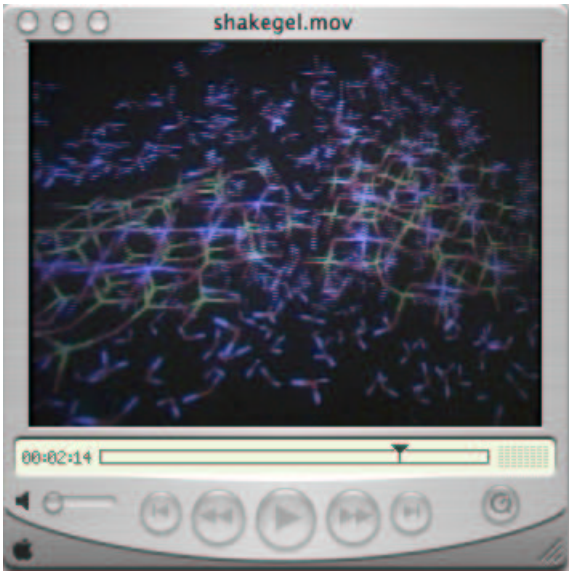

DIVERSE movieFromPathfile DSO

- AbswitchShell: displays pairs of objects viewed with the Diversifly program, and executes accompanying shell commands associated with the objects.
- AbgestureShell: a variation of ABswitchShell with a different object selection method.
- LinkedList: displays linked lists of objects one at a time.
- ObjectMover: displays sets of objects according to user defined positions, and then allows the user to move the objects around during the visualization.
- Slider, region, and isolate: methods for selecting objects for viewing.
- Sound: allows sounds to accompany object selection.
- WandPointer: adds a visual pointer to the wand to make object selection easier.

Several complete DIVERSE applications have also been developed. For example, Xwand is a program that enables the use of desktop GUIs in immersive systems such as a RAVE or CAVE. TIM (Things In Motion) is a set of related programs whose goal is to facilitate the visualization of time-series positional data.

In TIM the data to be displayed can be either read from an existing file, or it can be generated by a simulation process and asynchronously displayed by a visualization process as the data are being generated. The visualization and simulation processes do not need to run on the same system, allowing the simulation process to run on a large remote number-cruncher, while the visualization can be done on a local system. TIM supports systems ranging from a laptop to a multi-screen immersive display such as a CAVE(tm).

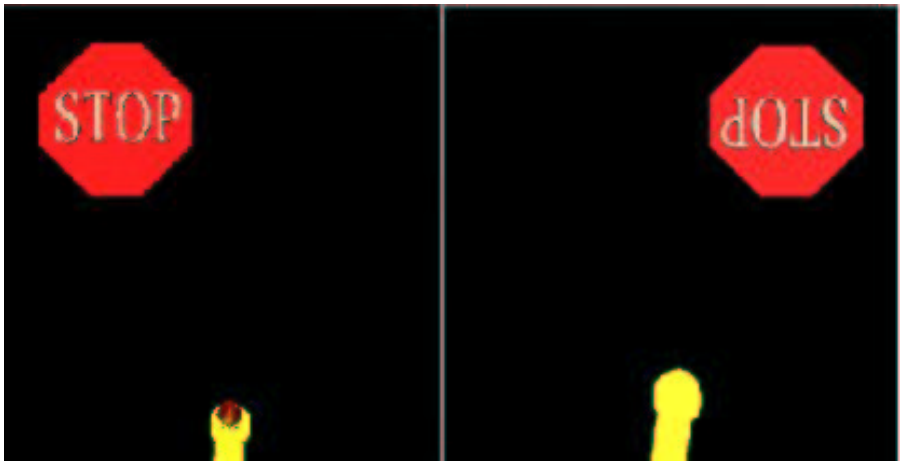

DIVERSE ABgestureShell DSO: Object in position A, then in Position B.

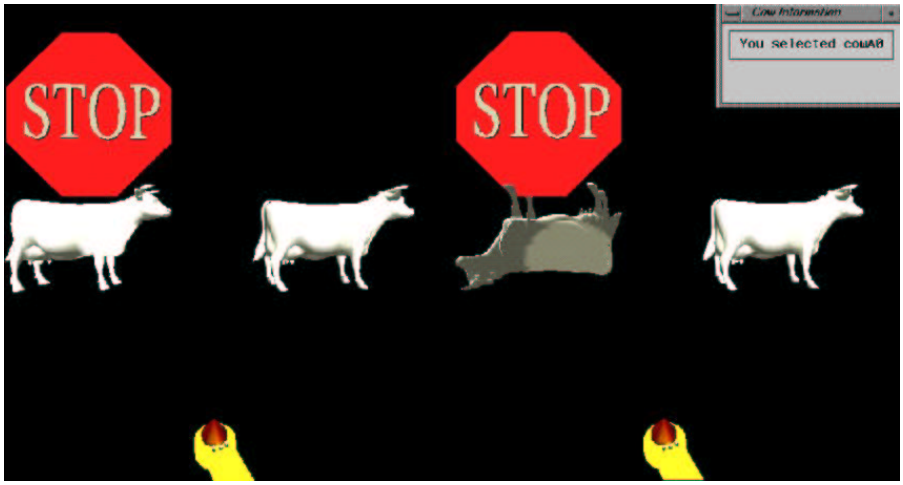

DIVERSE ABswitchShell DSO: Object in position A, then in Position B with command A executed.

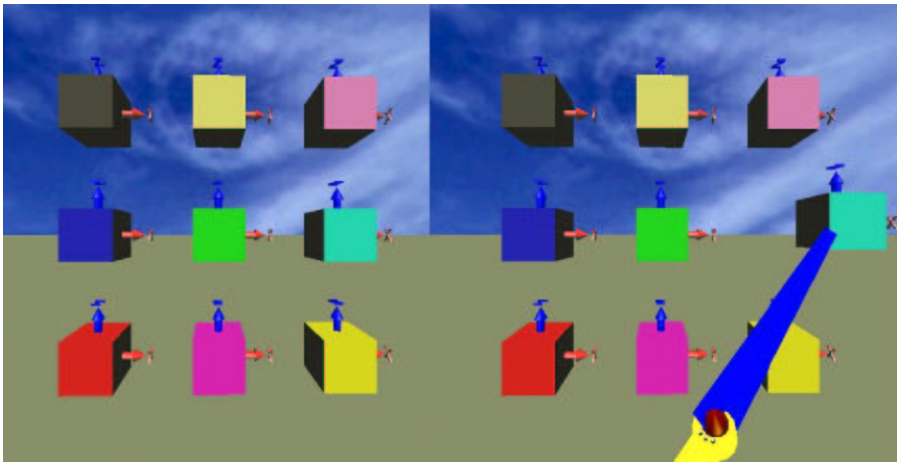

DIVERSE objectMover DSO: Object in position A, then in Position B.

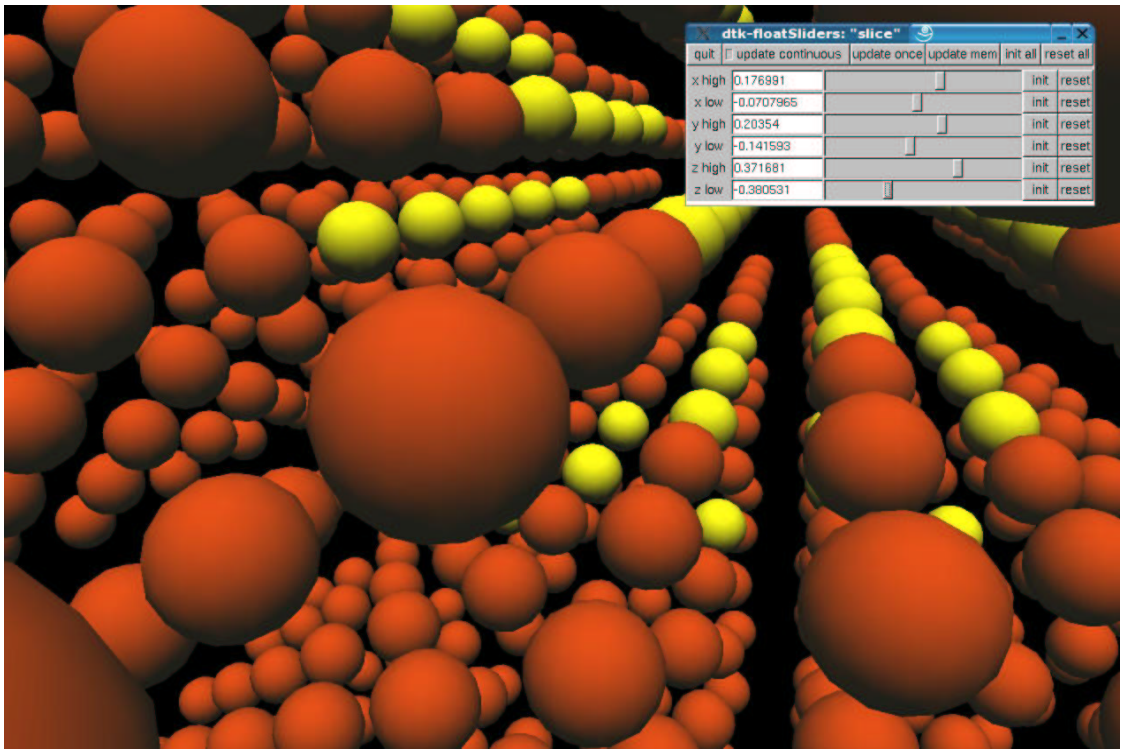

Xwand lets one use a desktop GUI in an immersive environment.
Here the wand manipulates the sliders in the GUI directly.

Special Projects

Digital Library of Mathematical Functions

Daniel Lozier
Ronald Boisvert
Joyce Conlon
Marjorie McClain
Bruce Fabijonas
Raghu Kacker
Bruce Miller
F.W.J. Olver
Bonita Saunders
Abdou Youssef
Charles Clark (NIST PL)
Charles Hagwood (NIST ITL)
Nell Sedransk (NIST ITL)
Qiming Wang (NIST ITL)
Sheehan Olver (student)
Elaine Kim (student)
Richard Askey (U. of Wisconsin, Madison)
Michael Berry (U. Bristol, UK)
Leonard Maximon (George Washington)
Morris Newman (U. California, Santa Barbara)
Ingram Olkin (Stanford)
Peter Paule (J. Kipler U., Linz, Austria)
Bill Reinhardt (U. Washington, Seattle)
Nico Temme (CWI, Amsterdam)

30 authors under contract

http://dlmf.nist.gov/

This project is striving to provide applicable mathematical knowledge to the general scientific community. In the past such knowledge has been provided in published handbooks. This method is still extremely valuable but important improvements in the dissemination of technical information can now be envisioned by making full use of current and anticipated capabilities of the World Wide Web.

The mathematical domain we are addressing is *special functions*. This domain includes those mathematical functions that arise naturally in the description of natural phenomena, as well as within mathematics itself. Among the many properties that are used in everyday scientific work are definitions, normalizations, graphs, derivatives, differential equations, integrals, integral representations, integral transforms, recurrence relations, power series, convergent and divergent expansions in terms of simpler functions, and interrelationships among the functions.

More specifically, we are constructing the Digital Library of Mathematical Functions (DLMF), a Web-based information resources, along with an associated 1000-page published handbook. The model for this endeavor is the immensely successful *Handbook of Mathematical Functions*, edited by Abramowitz and Stegun and published in 1964. The new handbook will replace the old one, approximately doubling the number of formulas contained. The Web site will offer the same information along with some additional formulas, lists of available software, and extended metadata. In addition, the Web site will provide this information in a form that can provide many advantages: downloading formulas into word processors and computer algebra systems, generating tables and graphs on demand, visualizing functions in three-dimensional animations, searching for formulas that contain arbitrary mathematical sub-expressions, linking to mathematical software documentation and programs, and linking to online references. To construct the DLMF we have assembled a NIST board of editors, an external board of associate editors, an external group of authors and validators, and a NIST staff.

The beneficiaries of the DLMF project will be, potentially, everyone who uses special functions. They will be able to go to the Web site to locate formulas, software, visualizations, and references where more information can be found, and they will be able to incorporate material from the DLMF Web site into their own documents and programs. And they will be able to rely on the

information they find. NIST has a tradition of providing reference knowledge that is well chosen, comprehensive, and rigorously characterized. Indeed, the engineering and scientific communities look to NIST for handbooks, and more recently Web sites, that can be used to carry out analysis, modeling and computation. These resources save time that would otherwise be spent poring through technical articles and monographs, resolving differences in notation, filling in gaps, and detecting and correcting errors. The DLMF project has established a careful process of topic selection, editorial control, and validation to assure that the results released to the public are faithful to the NIST tradition.

Major accomplishments are the result of team efforts. In FY2003 these include the following:

- **Content Development.** Outside authors are writing most chapters, with the NIST editors exercising strict control over style and content. The authors are being compensated in two stages with funds granted by the National Science Foundation. The editors have required substantial revisions from every author. By the end of November 2003, satisfactory initial drafts of most chapters had been received, marking the end of the first contractual stage, and the authors provided with their first increment of compensation. A few chapters have been troublesome, contributing to some delay in the project schedule. The editors believe that all core chapters are now on track for satisfactory completion. Non-core chapters may be deferred for inclusion in a later release of the Web site and a second edition of the handbook.

- **Authoring Tools.** The documents are authored using LaTeX, which we have extended with significant additional markup. These extensions clarify the document structure, provide useful metadata and reinforce the semantic content, particularly of the mathematics. The markup design, as well as general-purpose tools to convert LaTeX documents into XML, are under continuous development and refinement. In its current form, the tools are able to process the documents, store the components as XML in a database, and construct the Web site from these components. We have reached the stage where most of the mathematical content can be sufficiently parsed to produce MathML.

- **Graphics.** Great attention is being paid to graphics and visualizations. VRML is being used to construct 3D surfaces of functions that can be rotated and zoomed with a readily available browser plug-ins. Color is being used to represent height or, for complex functions, phase. The basic VRML capabilities are being extended in various ways to offer the user options for labeling, scaling, and viewing 2D slices. Production of these visualizations is a large and technically demanding task. Students have been employed to do some of this work. The DLMF will have approximately 500 figures, including 2D graphs as well as surfaces. We have prototypes for almost all of them, and nearly production-ready versions for more than half. Among the specific achievements of the past year are improvements in the legibility of labels and VRML controls, the addition of a color option that allows users to choose coloring of surfaces by height, quadrant or full-spectrum color circle, the addition of a triad feature that allows users to display 3D arrows to identify coordinate directions, the creation of five standard viewpoints, one of which is a view from above that reduces the surface to a density plot when scaled down to zero in the z-direction, and the implementation of movable labels that change with the standard viewpoints so that the labels are always visible.

- **Search.** Indexes are important for books but a mathematics-aware search engine is needed for the Web site. The problem is that search in the DLMF needs to work with mathematical queries, and all current engines are text-based. Good progress has been made in adapting one of these engines to locate equations through examination of the LaTeX encoding. The query language was designed to be as natural as possible for users who are familiar with LaTeX. The current system displays all equations that satisfy a mathematical query, each with a link that can be followed to go to its occurrence in the DLMF.

- **References.** The bibliography is an important part of the DLMF, both for identifying sources of proofs and for referring to extensions of material included in the DLMF. NIST staff is constructing and checking the bibliography with the aid of Math Reviews (MR) and other bibliographic resources. MR and Digital Object Identifier (DOI) links to reviews and full texts of articles are being included whenever they exist.
- **Validation.** Outside validators, compensated under contract to NIST with funds from the NSF grant, will check each chapter for accuracy. The editors have identified validators for most chapters and written a detailed validation guide. The validation process will begin in January 2004.
- **Administration.** The biggest administrative burden is management of contracts and grants. For example, in the past year the Mathematics Editor's original 5-year contract expired and was renewed. Because of increased high-level government review that required substantial added paperwork, two interim contracts were required to cover the period before the new contract went into effect.

The DLMF project has received considerable external recognition because of the importance of its subject matter and because it is the first large effort to construct a widely accessible body of technical mathematics on the Web. For example, a two-day workshop at Simon Fraser University in January 2003, called *Special Functions in the Digital Age*, was built around the DLMF project. A similar event, the *Second North American Workshop on Mathematical Knowledge Management*, has been planned at the Joint Mathematics Meetings in January 2004. While NIST staff participated in each of these events, external groups not formally associated with the DLMF project organized them. During the last year DLMF staff have given major addresses at conferences and colloquia in Southampton (UK), Strobl (Austria), Toronto, Montreal, Vancouver, Phoenix, Norfolk, and Washington. Two articles were published in *Annals of Mathematics and Artificial Intelligence* **38**(1-3) May 2003. Both had been presented at the First International Workshop on Mathematical Knowledge Management in 2001.

Members of the NIST staff on the DLMF project are participating in external research and development activities that are related to special functions and the presentation, semantic-based interchange, and traceability to a standard definition of mathematical knowledge in the Web environment. External groups in which we are participating are

- OpenMath (www.openmath.org).
- MathML (www.w3.org/Math).
- Java Grande (www.javagrande.org).
- OMDoc (www.mathweb.org/omdoc).
- CEIC (www.ceic.math.ca).
- SIAM-OPSF (www.siam.org/siags/siagos.htm).

Likely outgrowths of the DLMF project, in addition to maintenance tasks such as errata and Web improvements, include digital libraries in another fields of mathematics, R&D in issues related to mathematics and the Web, and R&D in issues related to formalization of mathematics using computers.

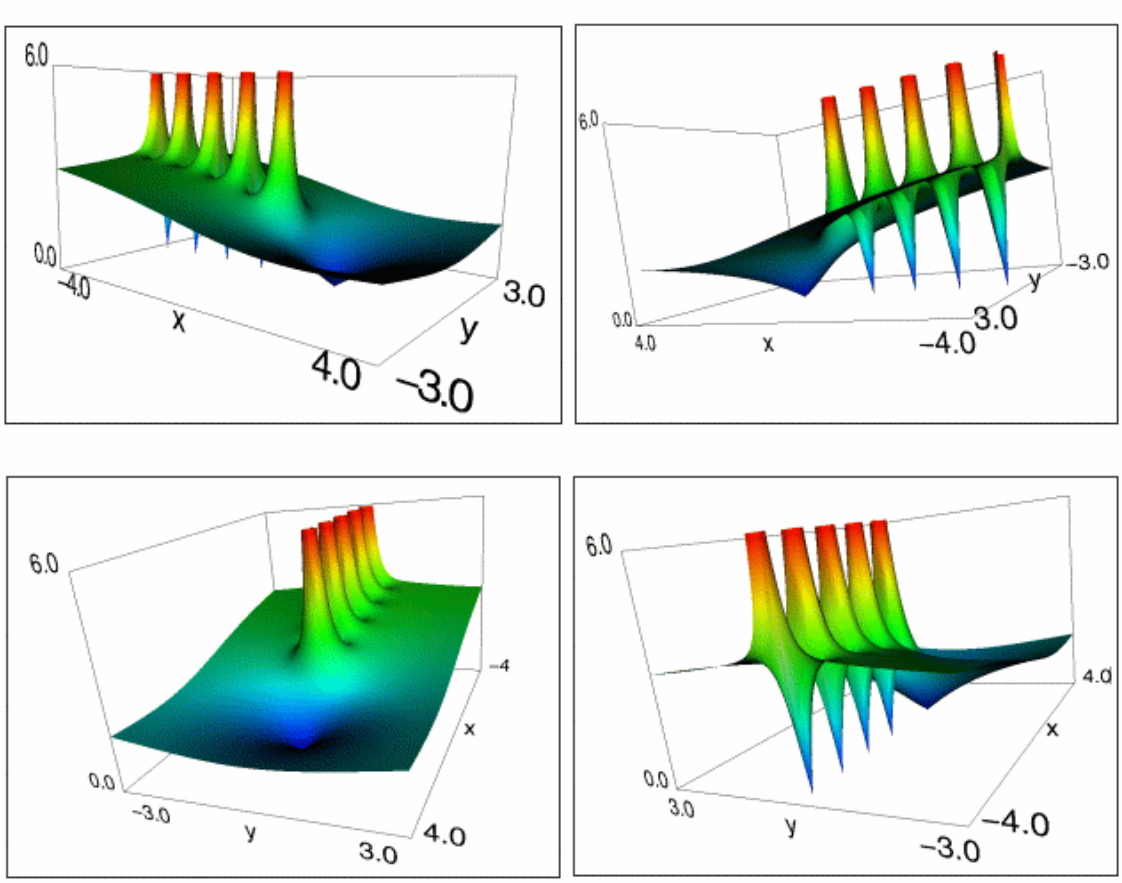

Four views of the digamma function taken from the VRML visualization in the DLMF Web site. Also known as the psi function, it is the logarithmic derivative of Euler's gamma function. It is a complex-valued function of one complex argument. The point on the surface above x+iy is the modulus (absolute value) of the function value. Coloring is by height. The surface depicts 5 zeros at one positive and 4 negative real values of the argument, and 5 poles at x = 0, -1, -2, -3, -4.

Quantum Information

Isabel Beichl
Stephen Bullock
Emanuel Knill
Dianne O'Leary
David Song
Gavin Brennan (NIST PL)
Igor Markov (U. of Michigan)
John Martinis (NIST EEEL)
Vivek Shende (U. of Michigan)
Francis Sullivan (IDA Center for Computing Sciences)
Carl Williams (NIST PL)
David Wineland (NIST PL)

Since the mid-1960s the dimensions of transistors on microchips have halved approximately every 18 months. It is estimated that logic capacity and speed are doubling, and memory capacity is quadrupling, every three years. These rapid developments have fueled the explosive growth in information technology we have witnessed in the last 20 years. Projecting this progression into the future, a transistor will be charged with a single electron in about 15-20 years. At that point, the laws of quantum rather than classical physics will begin to dominate, requiring fundamental changes in the physical basis for computer technology if increases in computing power are to be sustained. An intriguing idea is the use of quantum states of atoms, ions, or photons themselves to store, process,

and communicate information. Quantum mechanics is well known for its counterintuitive properties, and its foundations are still debated. However, due to this "quantum weirdness," a number of new capabilities conjectured to be impossible classically have been discovered. Among the potential benefits of quantum information technology are (a) the solution in polynomial time for problems considered exponentially difficult, examples being factoring and discrete logarithms, the bases of common public-key encryption systems, and (b) communication channels secure against eavesdropping.

The NIST Physics Laboratory (PL) has a mature experimental program in quantum computing, and ITL and PL have a maturing experimental program in quantum communication. MCSD is developing a program in quantum information theory to complement these efforts. Initially supported by the DARPA Quantum Information Science and Technology (QuIST) program, these efforts were considerably enhanced this year with the award of NIST Competence funding for a project entitled *Quantum Information Theory and Practice*. Work on this project will center on the development and analysis of architectures and algorithms for quantum information processing. Joining the program at the end of FY 2003 were Emanuel (Manny) Knill, a senior quantum information theorist formerly of Los Alamos National Labs, and Stephen Bullock, an NRC Postdoctoral Fellow. The research being undertaken by MCSD staff is summarized below.

Quantum Computer Benchmarks. E. Knill is working to help realize quantum information processors by developing the software and architecture needed to handle even high error rates in realistic physical devices. Different physical devices may require specialized tools for implementing the logical qubits that constitute a quantum information processor. To determine which tools to use requires characterizing and benchmarking the physical devices. One proposal for quantum information processing (due to Knill, Laflamme and Milburn) is based on using single photons and linear optics. It is characterized by having a high-detected error-rate by design. Surprisingly, detected errors are not a significant obstacle to quantum information processing. E. Knill has now formalized this observation, showing that scalable quantum computation is possible with detected error rates well above 20%. This is described in a manuscript entitled "Scalable Quantum Computation in the Presence of Large Detected-Error Rates" which has been submitted to *Physics Review A*.

E. Knill is actively participating in the ion-trap quantum computation projects led by D. Wineland of PL. He observed that measurement error can be overcome by using multiple read-out qubits. The ion trap group has just implemented this idea as a demonstration experiment using two ions (as yet unpublished).

Quantum Computer Architecture. The development of quantum processor and memory technology remains in its early infancy. So far, there have been very few serious attempts to develop architectural models for complete quantum computers. Such models are necessary to enable the real engineering of quantum computers. Because many physical implementation schemes for quantum computing rely on nearest-neighbor interactions, there is a hidden quantum communication overhead to connect distant nodes (qubits) of the computer. D. Song, along with PL colleagues G. Brennen and C. Williams, have proposed a physical solution to this problem which provides a pathway to a scalable quantum architecture using nonlocal interactions. The solution involved the concept of a "quantum bus" that uses entanglement as a resource to connect distant memory nodes. This work, which has appeared in *Physical Review A*, received online media coverage (see e.g., http://www.newsfactor.com/ perl/story/20485.html).

Several proposed quantum computer models include measurement processes in order to implement nonlocal gates and create necessary entanglement resources during the computation. Measurement can be much more expensive than other quantum operations, and hence it is important to understand what flexibility that compilers have in rearranging operations to optimize use of

resources. D. Song has studied schemes in which the measurements can be delayed for two- and three-qubit nonlocal gates.

Since we have an extremely reliable classical information-processing machine, it is expected that the first practical quantum computer will operate as a special-purpose unit attached to a classical system. D. Song is currently working on designing a semiclassical quantum computer architecture in order to maximize the classical part of a quantum computation while achieving the same quantum speed-up. This issue is important since quantum operations may be extremely difficult to control due to decoherence.

Quantum Cryptography. Cryptography has been one of the most fruitful applications of quantum information theory. Among the most fruitful applications of quantum mechanics to cryptography are distribution protocols for cryptographic keys. Quantum key distribution (QKD) aims to establish a pair of identical sequences of random binary digits at two separated locations without allowing a third party to learn these digits. Such shared secrets can be used as a cryptographic key for secure communication. D. Song has proposed a new QKD scheme based upon entanglement swapping. Using two Bell states, two bits of secret key can be shared between two distant parties. Song is working to collaborate with the NIST Quantum Communications Testbed experimental group in developing plans for implementing entanglement based cryptography protocols of this type.

Optimal Entanglement Swapping. D. Song has also studied methods for optimal entanglement swapping. Such techniques form the basis for nonlocal operations in quantum computers as well as certain QKD protocols. Entanglement swapping schemes consider bipartite states shared pairwise between two or more parties in a chain. The intermediate parties perform Bell measurements with the result that the two end parties will share an entangled state that can then be converted into maximally entangled states. In two partially entangled states, entanglement swapping by Bell measurement will yield the weaker entanglement of the two. This scheme is optimal because the average entanglement cannot increase under local operation and classical communication. However, for more than two states, this scheme does not always yield the weakest link. D. Song has considered projective measurements other than Bell-type measurement and showed, numerically, that while Bell measurement may not be unique, they are indeed optimal among these projective measurements.

Quantum Logic Circuit Design. A quantum logic circuit describes how a large quantum computation may be broken down into elementary operations on one or two quantum bits, just as a classical logic circuit decomposes a complicated computation on bit strings into a sequence of *AND*, *OR*, and *NOT* operations on pairs of classical bits. In each instance, obtaining good logic circuits is often a first step towards the design of efficient physical computers.

Recent work of S. Bullock, joint with the University of Michigan quantum circuits group, has identified generically optimal logic circuits in two quantum bits. This work has recently appeared in *Physical Review A.* They have also constructed smallest-possible logic circuits for diagonal quantum computations in any number of quantum bits, work which was recently accepted for publication in *Quantum Information and Computation.* The latter is the only known optimal circuit design algorithm in n quantum bits. Recent work in two qubits has also identified which two-quantum-bit computations are structured, i.e., which admit much smaller circuits than a generic two-quantum-bit computation. This newest result fits into a broad literature that generally seeks to classify which computations might be efficiently realized by quantum computers.

S. Bullock has developed and recently released a program, 3cnot, which produces a quantum logic circuit for any two-qubit computation which contains several one-qubit operations but at most three interactions between both qubits (see http://math.nist.gov/~SBullock).

Matrix Decompositions. A quantum computation is mathematically modeled by the application of a unitary matrix to a complex vector which models an initial quantum data state. A matrix factorization, $U = U_1U_2 \cdots U_n$, in this context describes how to realize a computation U by first applying U_n, then U_{n-1}, etc. until U_1. A matrix decomposition is an algorithm that produces a matrix factorization on any input U, thus providing an organized way of dividing any quantum computation into a sequence of simpler computations. A current project of S. Bullock, joint with Gavin Brennen of PL, seeks to build specialized unitary matrix decompositions with applications to quantum computing and to leverage the existing literature in pure and applied mathematics on matrix analysis to study these decompositions. There are two highlights of this project so far.

- A matrix decomposition of 4×4 unitaries, useful in the design of two qubit logic circuits, has been generalized to an arbitrary number of qubits. Applications to circuit design are as yet to be realized, but the decomposition has found immediate application to the physical theory of quantum entanglement. D. O'Leary has suggested efficient numerical algorithms for computing this new matrix decomposition, which in certain cases may require extending recently developed symplectic-Jacobi techniques of Fussbender-Mackey-Mackey.
- The standard *QR* matrix factorization has been adapted and optimized to computational settings in which data is encoded in the hyperfine energy levels of electrons in neutral atoms. Only certain such levels are paired, and this makes certain Givens rotations in the *QR* decomposition more easily realized physically than others.

Limits of Quantum Computation. Contextuality states that what one gets for an answer depends on the context in which one measures the property. In the process of investigating the idea of contextuality in quantum computing, I. Beichl and F. Sullivan have devised a new simplified proof of the Kochen-Specker theorem. This theorem says that there is no constrained coloring of the two-sphere. The coloring is by two colors so that in any orthogonal triple of vectors, exactly two vectors have the same color. In quantum computing, the colors are related to the spins of bosons. The rational 2-sphere, however does admit such a coloring. Beichl and Sullivan have devised a simplified proof of this fact and why it does not negate the Kochen-Specker theorem. The existence of the coloring is presumed when it is claimed that quantum computing gives 2^N parallelism for N qubits. Contextuality and the Kochen Specker theorem must be better understood before such parallelism can be claimed.

Quantum Information Theory and Practice Seminar. An informal seminar was begun in the Fall of 2003 to foster interactions among participants in the NIST competence project on Quantum Information Theory and Practice. Participants are theoreticians from the NIST Information Technology and Physics Labs in Gaithersburg and Boulder. Faculty and graduate students from the University of Maryland Baltimore County are also participating. The discussion topics and leaders of recent sessions follow.

1. Stephen Bullock, "Matrix Decompositions," September 23, 2003.
2. Anastase Nakassis, "Bit String Reconciliation for QKD," October 10, 2003.
3. Sam Lomonaco, "Continuous Quantum Algorithms," October 30, 2003.
4. David Song, "A QKD Protocol Based on Entanglement Swapping," November 6, 2003.
5. Gavin Brennen, "Universal computation with qudits with applications to multi-level atoms", December 4, 2003.
6. Stephen Bullock, "Concurrence canonical decomposition," December 18, 2003.