

High Performance Computing and Visualization

SSS: Screen Saver Science

William L. George

Jacob Scott (Univ. 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 computer would participate only when it would otherwise be idle, that is, when its screen saver would be running. This work is based on Jini, a Java-based open software architecture intended for the development of robust network services.

This project has several goals. First, we hope to utilize the idle processing power of the many PCs and workstations available at NIST to execute production scientific codes. The power of personal PCs and workstations continues to increase, and they are now becoming very capable of executing large-scale applications. Second, research on grid computing has been accelerating, and the SSS environment will provide a platform for local experimentation with highly parallel and distributed algorithms suitable for grid environments. Finally, the development of SSS applications will give us the opportunity to explore the use of Java on scientific applications, a topic of independent interest.

Up until recently, a project of this scope would have required a large amount of software development just to become minimally functional, and hence was not practical. But, with the introduction of Jini, and its 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 1980's by David Gelernter of Yale University. It provides the necessary infrastructure to support the development and use of remote services in applications.

The SSS project began in the summer of 2002. Substantial progress has been made in designing and implementing the basic SSS infrastructure. A generic compute server has been implemented, suitable for embedding in a screen saver, and a new Jini service, a remote file server, has been developed to provide SSS applications with remote file I/O capabilities. With the basic SSS infrastructure now 90% complete, we have begun to develop the first large application for SSS, a quantum Monte Carlo application in *ab-initio* chemistry. This application is currently running on our Linux cluster but does not require the use of the cluster's communications network since each task is completely independent. This makes it ideal for porting to the SSS environment.

W. George presented the preliminary design of SSS at the 2002 International Conference on Parallel Architectures and Compilation Techniques in September 2002.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management (virtual measurements).

Parallel Genetic Programming

Judith E. Devaney

John G. Hagedorn

Because the design and implementation of algorithms is highly labor-intensive, the number of such projects that can be undertaken is limited by the availability of people with appropriate expertise. The goal of this project is to create a system that will leverage human expertise and effort through parallel genetic programming. The human specifies the problem to be solved, provides the building blocks and a fitness function that measures success, and the system determines an algorithm that fits the building blocks together into a solution to the specified problem. We are implementing a generic Genetic Programming (GP) system with some unique features. These are intended to improve the operation of the system particularly for the types of real-world scientific problems to which we are applying the system at NIST. Genetic programming is also a meta-technique. That is, it can be used to solve any problem whose solution can be framed in terms of a set of operators and a fitness function. Thus it has applications in parameter search. NIST scientists have many special purpose codes that can be used as operators in this sense.

This year we continued development of our system and applied it to the study of hydrating plaster. We presented this work at two key conferences: the *Genetic and Evolutionary Computing Conference* (New York, July 2002) and the *Fifth International Conference on Discovery Science* (Lubeck, Germany, November 2002). Papers appear in the proceedings of both conferences.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management (intelligent interconnected systems).

Parallelization of Feff X-ray Absorption Code

James Sims

Howard Hung

Charles Bouldin (NIST MSEL)

John Rehr (University of Washington)

<http://math.nist.gov/mcsd/savg/parallel/xray/>

X-ray absorption spectroscopy (XAS) is used to study the atomic-scale structure of materials, and is employed by hundreds of research groups in a variety of fields, including ceramics, superconductors, semiconductors, catalysis, metallurgy and structural biology. Analysis of XAS relies heavily on ab-initio computer calculations to model x-ray absorption. These calculations are computationally intensive, taking days or weeks to complete in many cases. As XAS is more widely used in the design of new materials, particularly in combinatorial materials processing, it is crucial to speed up these calculations. One of the most commonly used codes for such analyses is FEFF. Developed at the University of Washington, FEFF performs ab initio multiple scattering calculations of x-ray absorption fine structure (XAFS) and x-ray absorption near-edge structure (XANES) spectra for clusters of atoms. The code yields scattering amplitudes and phases used in many modern XAFS analysis codes. FEFF has a user base of over 400 research groups, including many industrial users, such as Dow, DuPont, Boeing, Chevron, Kodak, and General Electric.

To achieve faster speeds in FEFF, James Sims and Howard Hung of MCSDD worked with Charles Bouldin of the MSEL Ceramics Division to develop a parallel version, FeffMPI. In modifying the code to run on the NIST parallel processing clusters using a message-passing approach, they gained a factor of 20-30 in improvement in speed over the single processor code. FeffMPI has

now been ported to a wide variety of machines including NTs and Macs. Timing results for an 87 atom GaN test case on various systems is given below.

System	Time	n	Processor	Speed	Compiler
IBM SP (NERSC II)	00:27	32	nighthawk		xlf -O -qaxmem=9200
IBM SP NERSC	01:20	32	R6K		xlf -O -qaxmem=9200
SGI Origin2K	02:04	8	R10K	250	f77 -Ofast
Athalon 24cluster	03:09	4	athalon tbird	1000	pgf77 -fast
CRAY T3E NERSC	04:12	8	alpha EV-4		f90 -O2
Apple 8 node G4	04:14	8	G4	533	f90 -O
dual ALPHA	04:43	2	alpha EV-6	500	fort -fast lmpi
ALPHA	07:23	1	alpha EV-6	500	fort -fast
dual PIII linux	07:51	2	athalon smp	860	pgf77 -fast
SGI Octane	14:46	1	R10K	250	f77 -Ofast LNO:opt=0
dualPC686 linux	15:31	2	P3	450	pgf77 -fast
Mac G4 linux 86	22:40	1	G4	400	g77 -O2
dualPC686 linux	27:00	1	P3	450	pgf77 -fast

Time is given in minutes:seconds. n is the number of processors. Speed is given in MHz.

Jim Sims, Howard Hung, and Charles Bouldin used FeffMPI and a Continuous Random Network model to reproduce the main features in the X-ray absorption near edge structure of crystalline and amorphous germanium. This was published in the *NIST Journal of Research*. They also collaborated with J.J. Rehr and A. Ankudinov on a calculation of electron multiple scattering that was published in *Physical Review B*.

NIST Strategic Focus Areas. *Established Industries:* Materials processing. *Emerging Industries:* Information and Knowledge Management.

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 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 of less than 10 microhartrees (0.00001 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.

Author	N	Energy (a.u.)
Li ⁺		
TK [31]	308	-7.2799 1341 2669 3020
Drake [20]	2114	-7.2799 1341 2669 3059 0000
This work	4284	-7.2799 1341 2669 3059 6489
Be ⁺⁺		
TK [31]	308	-13.6555 6623 8423 5829
Drake [20]	2114	-13.6555 6623 8423 5867
This work	4648	-13.6555 6623 8423 5867 0206
B ⁺³		
TK [31]	308	-22.0309 7158 0242 7777
Drake [20]	2114	-22.0309 7158 0242 7815
This work	4648	-22.0309 7158 0242 7815 4163

Comparison with previously explicitly correlated calculations for 1¹S He-like ions.

In their early work Sims and Hagstrom used the method to compute not only energy levels, but also other atomic properties such as ionization potentials, electron affinities, electric polarizabilities, and transition probabilities of two, three, and four electron atoms and other members of their isoelectronic sequences. The Hy-CI Li⁺ result represented the best upper bound for the ion considered for more than 15 years. This result has been considerably improved this year, culminating in the best result so far. However, the nonrelativistic energy is only accurate to 6 decimal digits, clearly suggesting the need for higher precision calculations. Achieving such higher precision is the goal of this research.

In work published this year in the *International Journal of Quantum Chemistry*, Sims and Hagstrom have 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 extended precision. To obviate the high CPU and memory costs in a high precision calculation, the calculation was parallelized. It is the parallelization that opens up the possibility of Hy-CI method calculations for three, four (and hopefully more) electron atomic systems. Parallelization was accomplished by modifying their generalized eigenvalue problem solver (which is based on inverse iteration) to generate matrix elements in the appropriate block order. Then the Message Passing

Interface (MPI) standard was used to run the same program on multiple processors (on the same or different hosts), giving each host a block of the matrix, with no need to redistribute the matrices for the inverse iteration step. Hence, the matrix calculation runs in parallel, and the blocks of the matrix are spread across processors, solving both the memory and CPU speed problem. As far as we know, this is the first high precision calculation for few electron atomic systems to employ parallel computing.

The Hy-CI technique is still being used today, with the original 1970's work still being referenced in the peer-reviewed literature. In preprints that have been sent to Sims, Hy-CI has been extended to the relativistic domain, to the computation of atomic resonances, and to the computation of the best nonrelativistic energy ever for $2p$ -squared triplet P state of the hydrogen negative ion.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management (virtual measurements).

Cement and Concrete Projects

The NIST Building and Fire Research Laboratory (BFRL) does experimental and computational research in cement and concrete. Recently, MCS D 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. MCS D 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, Holcim, Dyckerhoff, Verein Deutscher Zementwerke (VDZ), 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 amount of 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. MCS D 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 four projects are included in this effort.

Parallelization of a Model of the Elastic Properties of Cement

Robert Bohn

Edward Garboczi (NIST BFRL)

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

Almost all real materials are multi-phase, whether deliberately, when formulating a composite, inadvertently, by introducing impurities into mono-phase material, or by the very nature of the material components, as in the case of cement-based materials. Predicting the elastic properties of such a material is dependent on two pieces of information for each phase: how each is arranged in the microstructure, and its elastic moduli. Cement paste is extraordinarily complex elastically, with many different chemically and elastically distinct phases (20+) and a complex microstructure. This complexity further increases in concrete, as aggregates (sand and gravel) are added. Elastic moduli are important in themselves, as mechanical properties (especially at early ages), and because they can be a good predictor of compressive strength.

A finite element package for computing the elastic moduli of composite materials has been written by staff of the NIST Building and Fire Research Laboratory and has been available for several years. This program has worked successfully for many different material microstructures. However, it is a single-processor code, which limits the systems that can reasonably be studied. Moreover, the predictions of the code for early age cement paste did not agree well with experiment. The initial goal of this project is to update and parallelize this code and use it to study early-age cement paste.

The main benefactor of this effort will be the cement and concrete industries that are members of the Virtual Cement and Concrete Testing Laboratory (VCCTL) consortium. These include most of the world's largest cement, aggregate, and chemical admixture (chemicals added to concrete to change properties) companies. Making this code parallel will allow much larger systems to be studied, which should help resolve the problem of obtaining accurate cement paste computations at early ages. Large systems are also required to study the elastic properties of random shapes, like aggregates found in concrete. This is needed to be able to extend the computations from the cement paste (cement + water) to the full concrete level.

This year Robert Bohn completed the parallelization of the five materials properties codes using MPI. He wrote a User Manual for use of these new programs. Edward Garboczi used the parallel code to compute the early age elastic properties of cement paste. The modified parallel code largely removed the discrepancies between the serial code and experiments for early age cement paste. Results were presented to the VCCTL companies in May 2002 (NIST) and in December 2002 (Dusseldorf, Germany).

New data on early age cement pastes is driving still further modifications in the code, and the parallel finite element code will continue to be used in this research. Examining the elastic properties of large digital representations of real aggregate will also utilize the MPI elastic code this year.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management (dynamic data infrastructure).

The Visible Cement Dataset

Steve Satterfield

Peter Ketcham

William George

Judith Devaney

James Graham (NIST CIO)

James Porterfield (NIST CIO)

Dale P. Bentz (NIST BFRL)

Symoane Mizell (NIST BFRL)

Daniel A. Quenard (CSTB)

Hebert Sallee (CSTB)

Franck Vallee (CSTB)

Jose Baruchel (ESRF)

Elodie Boller (ESRF)

Abdelmajid Elmoutaouakkil (ESRF)

Stefania Nuzzo (ESRF)

<http://visiblecement.nist.gov/>

To produce materials with acceptable or improved properties, adequate characterization of their microstructure is critical. While the microstructure can be viewed in two dimensions at a variety of resolutions (e.g., optical microscopy, scanning electron microscopy, and transmission electron microscopy), it is often the three-dimensional aspects of the microstructure that have the largest influence on material performance. Direct viewing of the three-dimensional microstructure is a difficult task for most materials. With advances in X-ray microtomography, it is now possible to obtain three-dimensional representations of a material's microstructure with a spatial resolution of better than one micrometer per voxel.

The Visible Cement Data Set represents a collection of three-dimensional data sets obtained using the European Synchrotron Radiation Facility (ESRF) in Grenoble, France in September of 2000

as part of an international collaboration between NIST, ESRF, and Centre Scientifique et Technique du Batiment (CSTB) also of Grenoble. Most of the images obtained are for hydrating Portland cement pastes, with a few data sets representing hydrating plaster of Paris and a common building brick. The goal of this project is to create a web site at NIST where all researchers could access these unique data sets. The web site includes a text-based description of each data set and computer programs to assist in processing and analyzing the data sets. In addition to the raw data files, the site contains both 2-D and 3-D images and visualizations of the microstructures.

This year all the data was processed and transferred to the web site. The total number of datasets available is 105 (91 cement, 10 plaster, and 4 brick). A web page was created for each of the datasets, with links to the raw volume data file, JPG images of several slices through the dataset, a 3D volume image, and a movie file. The web site is now publicly available. A paper describing the dataset was published in the *NIST Journal of Research*.





Image of cement

The Visible Cement Data Set web site will continue to be used by NIST and will serve as a valuable resource to both the construction materials and visualization research communities.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management (dynamic data infrastructure).



Nicos Martys (standing) and John Hagedorn review results of fluid flow simulations.

Parallelization, Visualization of Fluid Flow in Complex Geometries

John Hagedorn

Judith Devaney

Nicos Martys (NIST BFRL)

Jack Douglas (NIST MSEL)

<http://math.nist.gov/mcsd/savg/parallel/lb/>

<http://math.nist.gov/mcsd/savg/vis/fluid/>

The flow of fluids 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, and the service life of building materials. The latter two applications are of particular concern for our NIST collaborators. The detailed simulation of such transport phenomena in varying geometries and subject to varying 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, we developed a lattice Boltzmann (LB) algorithm that simulates multiple fluids, various forces, and wetting characteristics within arbitrary geometries. We parallelized the algorithm using MPI to enable the study of large systems.


Recently, 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. We have also modified code to solve the Brinkman equation, which is useful for modeling flow in multi-scale porous media.


A paper documenting our results appeared in the special volume *Materials Science of Concrete*. Papers have also been submitted to *Macromolecules* and *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 in our results.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management.



Breakup of a fluid thread due to Taylor Tomitika instability in a parallel plate geometry (left) and in a cylindrical geometry (right).



Comparison of model predictions (left) with experiment (right) at different times.

Computational Modeling of the Flow of Concrete


James Sims
Terence Griffin
Steve Satterfield

Peter Ketcham
Nicos Martys (NIST BFRL)
Edward Garboczi (NIST BFRL)


<http://math.nist.gov/mcsd/savg/parallel/dpd/>
<http://math.nist.gov/mcsd/savg/vis/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. We use a new computational method called dissipative particle dynamics (DPD), which has several advantages over traditional computational dynamics methods while naturally accommodating such 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). Parallelization of the algorithm is important in order to adequately model size distributions, and to have enough resolution to avoid finite size effects. This year J. Sims improved the earlier QDPD parallel code to better handle shear boundary conditions (such as the Lees-Edwards boundary conditions). He tested this version on SGIs, IBM, and Linux PC clusters.



Model Rheometer studied with QDPD code



More realistic computations with modeled aggregate

In addition to the parallel code, we are developing visualizations of actual aggregates modeled with spherical harmonics. Such modeled aggregates will enable more realistic concrete flow simulations. Concrete is a composite of mortar plus coarse aggregate. There is the potential for rheology to be greatly affected by aggregate shape, so aggregate shape characterization is needed to be able to successfully model concrete properties.

In the future, the QDPD code will be extended to handle a broader size distribution of aggregates. Visualization techniques will be developed to display the aggregates now used in the simulations.

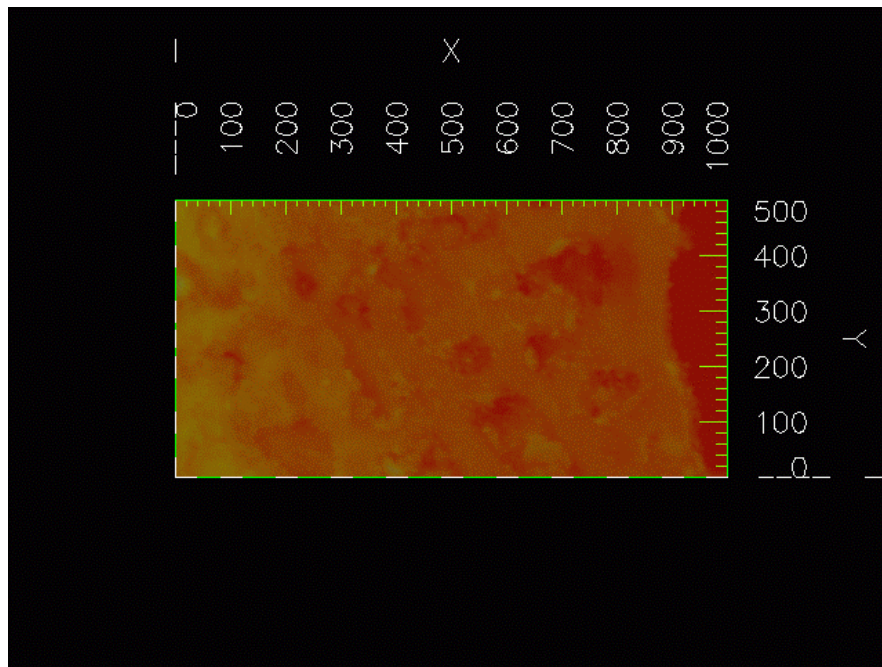
NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management.

Multi-Modal Visualization

*Peter Ketcham
Steven Satterfield
Barbara am Ende
Judith Devaney*

*Joy Dunkers (NIST MSEL)
Marcus Cicerone (NIST MSEL)
Lyle Levine (NIST MSEL)
Gabrielle Long (NIST MSEL)*

All areas of materials science confront 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 collaboration with the NIST Materials Science and Engineering Lab, which is sponsored by the NIST ATP program, we are combining data from different 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.



Volume rendering of polymeric tissue engineering scaffold containing stained osteoblasts and bone matrix.

Our collaborators are gathering measured data using a variety of techniques, including optical coherence tomography and ultra small angle X-ray scattering imaging. These measurement techniques provide information on a sample's function, structure, chemical composition, dynamic properties, and polarization sensitivity. When combined in a manner that is visually apparent, these measurements can yield unprecedented insight towards the comprehension of complex relationships among large amounts of correlated data. These methods have applications to the characterization of biomaterials, the failure analysis of polymer composites, and the reliability of semiconductor devices.

Work during the first year of this three-year project has focused on the characterization of biomaterials, particularly in the area of tissue engineering. We have developed protocols and infrastructure for transferring the data from various measurement devices to a common, hierarchical storage format that is amenable for display and manipulation in our immersive visualization environment. We have produced static images of tissue scaffolds as a preliminary step to ensure the integrity and viability of our data transfer and storage system.

Over the next two years, we will continue to integrate measured data from several sources and will focus on the registration of correlated data sets. We will also establish techniques and infrastructure that enable us to view data over multiple length scales in our immersive visualization environment.

NIST Strategic Focus Areas. *Emerging Industries:* Nanotechnology; Information and Knowledge Management (dynamic data infrastructure).

Computation and Visualization of Nano-structures and Nano-optics

*Julien Franiatte
Judith Devaney
Steve Satterfield*

*James Sims
Howard Hung
Garnett Bryant (NIST PL)*

<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 modeling tools needed for the engineering of these systems. Applications include 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.

MCS D is parallelizing computational models for studying nanostructures in collaboration with the NIST Physics Lab. Parallel processing has enabled significant speedup in the existing sequential code. Codes that took 10 hours are now being completed in one hour using 16 processors. As the computational model is extended to handle more complex and larger systems by including not only

the nanocrystals but also the substrate and environment around them, parallel processing will become even more of a necessity. This year the code will be extended to study self-assembled quantum dots.


Scientific visualization also plays a key role in this work, not only in understanding the output of a simulation, but also in preparing input. We have built an immersive virtual reality application for this purpose, which we call nanobuilder. This utility allows users to interactively create nanostructures in an immersive environment. A menu of atoms is shown on one display wall. Using the wand, the scientist grabs an atom by pointing and clicking. Musical tones play when the grab is successful, thus providing immediate sensory feedback. Once the atom is captured, the scientist moves the atom to the desired location by moving the wand. The atom is deposited at the point by another click of the wand. Atoms already in place can also be repositioned. Grabbing and moving atoms is simplified as the environment is limited to the menu atoms and the atoms in the created structure; a virtual flashlight shows clearly which atom is being pointed to. A click of the wand writes the coordinates of the created nanostructure to a file. This file is input to our parallel code that computes the required orbitals.

Additional visualizations for exploring the output of nanostructures calculations have also been created. For example, a movie was created that showed a fly-through of a nanostructure with the two views, structures and bonds, side by side.

A paper on this work was presented at the Nanotechnology at the Interface of Information Technology Conference in New Orleans in February 2002.

The parallel code will be extended to handle more complex structures and studied for additional speedup potential. The visualization will be more closely tied to the computation. This will enable interactive exploration of nanostructures.

NIST Strategic Focus Areas. *Emerging Industries:* Nanotechnology; Information and Knowledge Management.



Atoms in a section of a nanostructure (top) with orbitals (bottom).

Modeling and Visualization of Dendritic Growth in Metallic Alloys

William George

Steve Satterfield

James Warren (NIST MSEL)

<http://math.nist.gov/mcsd/savg/parallel/pfm/>

<http://math.nist.gov/mcsd/savg/vis/dendrite/>

Snowflake-like structures known as dendrites develop within metal alloys during casting. A better understanding of the process of dendritic growth during the solidification will help guide the design of new alloys and the casting process used to produce them. MCS D mathematicians (e.g., G. McFadden, B. Murray, D. Anderson, R. Braun) have worked with MSEL scientists (e.g., W. Boettinger, R. Sekerka) for some time to develop phase field models of dendritic growth. Such diffuse-interface approaches are much more computationally attractive than traditional sharp-interface models. Computations in two dimensions are now routinely accomplished. Extending this to three dimensions presents scaling problems for both the computations and the subsequent rendering of the results for visualization. This is due to the $O(n^4)$ execution time of the algorithm as well as the $O(n^3)$ space requirements for the field parameters. Additionally, rendering the output of the three dimensional simulation also stresses the available software and hardware when the simulations extend over finite-difference grids of size 1000x1000x1000.

We have developed a parallel 3D dendritic growth simulator that runs efficiently on both distributed-memory and shared-memory machines. This simulator can also run efficiently on heterogeneous clusters of machines due to the dynamic load-balancing support provided by our MPI-based C-DParLib library. This library simplifies the coding of data-parallel style algorithms in C by managing the distribution of arrays and providing for many common operations on arrays such as shifting, elemental operations, reductions, and the exchanging of array slices between neighboring processing nodes as is needed in parallel finite-difference algorithms.

A paper on this work was published in the *Journal of Computational Physics* this year.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management.

Tools for Scientific Visualization

Steven G. Satterfield

Peter Ketcham

John Hagedorn

Adele Peskin

John Kelso (Virginia Tech)

Harry W. Bullen IV (Montgomery Blair H. S.)

Jessica S. Chang (Montgomery Blair H. S.)

Alex V. Harn (Montgomery Blair H. S.)

Sean P. Kelly (Montgomery Blair H. S.)

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 MCS D IV environment is an open source system named DIVERSE (Device Independent Virtual Environments -- Reconfigurable, Scalable, Extensible). The DIVERSE API (Application Programming Interface) 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 desktop systems to multi-wall stereographic displays with head tracking.

Included with DIVERSE is an extensible application called Diversifly that allows various techniques for navigation through user data loaded from a variety of external data formats. On top of this DIVERSE/Diversifly 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 Diversifly file loaders;
- developing application-specific files that are easily transformed into existing or newly created graphics file formats; and
- developing filters to connect data transformation pipelines.

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. For example, a sequence (.seq) file loader DSO has been created. The .seq file format provides a very simple input file to specify a "flip-book" style animation sequence. Another Diversifly extension that we have developed is a "turn-table" DSO. This DSO allows any visualized object, such as a concrete particle to be automatically rotated as if on a turntable for viewing. Since the DSO is a general implementation, simply loading the DSO at run-time easily creates this animation, and thus no specialized programming is needed.

Many other individual DSOs have been created to facilitate the simple use of the MCSD immersive visualization environment. Many of these have been released for public use. See <http://math.nist.gov/mcsd/savg/software>. A paper describing this work was presented at the Virtual Reality 2002 Conference held in March 2002 in Orlando.




A Diversifly DSO has been developed to switch between alternate views. In this example we see a cloud visualization on the left. A click of the wand turns on streamlines.



On the left is the output of a Diversifly DSO that provides a flashlight for the virtual environment. On the right is a DSO which monitors the physical position of the wand and head tracked glasses. If either device comes close to physically touching the screen, a large bright red Stop Sign is displayed.

Glyph Toolbox. The purpose of the Glyph Toolbox Project is to construct a collection of tools, i.e. individual UNIX style command line programs that can be used to build a polygon-based virtual environment or used as glyphs for data representation. The tools (command and filters) produce an ASCII-based file that is machine and rendering independent. The actual display of the ASCII files is handled by converting the output polygon file into a format suitable for display by a viewing program, such as DIVERSE/Diversifly, VRML, Open Inventor, etc. Many glyphs and glyph manipulation tools were designed. More than 50 such commands were created. Some are meta-commands that arrange other glyphs. Examples are illustrated below. The construction of these tools was mainly undertaken by a group of four student volunteers from Montgomery Blair High School in Silver Spring, MD.

NIST Strategic Focus Areas. *Emerging Industries:* Information and Knowledge Management (dynamic data infrastructure).



Examples from the Glyph ToolBox. gtb-confetti (left) displays arbitrary objects, with individual properties, at user-specified locations. gtb-smily (right) creates emoticon glyphs using the values of nearly three-dozen variables.

