2003 - 2004

DEIXIS

THE DOE CSGF ANNUAL

2003 - 2004

2003 - 2004

DOE CSGF

DEIXIS – THE DOE CSGF ANNUAL

DEIXIS – THE DOE CSGF ANNUAL

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

Funded by:
Office of Science and
National Nuclear Security
Administration's Office of
Defense Programs

The Krell Institute
1609 Golden Aspen Drive, Suite 101
Ames, IA 50010
(515) 956-3696
www.krellinst.org/csgf

DOE CSGF


EXIT
Frederick A. Howes

THE DOE CSGF ANNUAL

DEIXIS

Department of Energy Computational Science Graduate Fellowship

TABLE OF CONTENTS

2003 - 2004

Editor
Shelly Olsan

Design
Juls Design, Inc.

Contributors
Jacob Berkowitz
William J. Cannon
Victor D. Chase
Peter Gwynne
Karen Hede
Lori Valigra

DEIXIS, The DOE CSGF Annual is published by the Krell Institute. The Krell Institute administers the Department of Energy Computational Science Graduate Fellowship program for the DOE under contract DE-FG02-97ER25308. The Krell Institute is a non-profit company that works in partnership with national science and technology research and education communities to help them achieve national priorities.

For additional information about the Department of Energy Computational Science Graduate Fellowship program, the Krell Institute, or topics covered in this publication, please contact:

Editor, *DEIXIS*
The Krell Institute
1609 Golden Aspen Drive, Suite 101
Ames, IA 50010
(515) 956-3696
www.krellinst.org/csgf

DOE CSGF

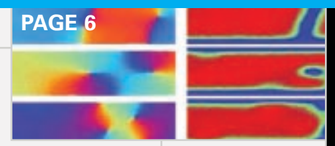

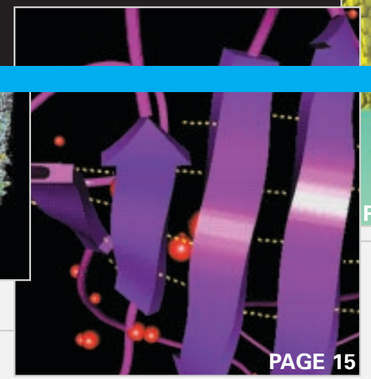

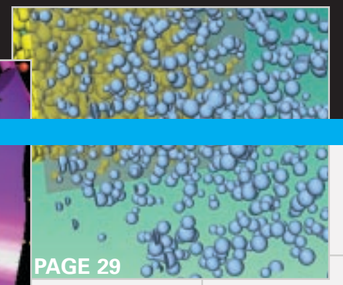

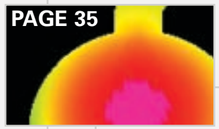

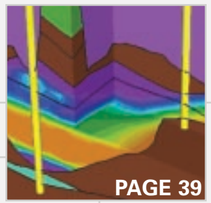

DEIXIS (ΔΕΙΞΙΣ) *transliterated from classical Greek into the Roman alphabet, (pronounced dāksis) means a display, mode or process of proof; the process of showing, proving or demonstrating.* DEIXIS *can also refer to the workings of an individual's keen intellect, or to the means by which such individuals, e.g. DOE CSGF fellows, are identified.*

DEIXIS *is an annual publication of the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program.* DEIXIS *illustrates work done at eight multi-program DOE laboratories and highlights the DOE CSGF fellows and alumni. The DOE CSGF is funded by the Office of Science and the National Nuclear Security Administration's Office of Defense Programs.*

High Performance Practica

FOR OVER 10 YEARS, the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program has been providing support to the nation's best new scientists. The DOE CSGF program's innovative requirement of cross-disciplinary study in science/engineering, applied mathematics, and computer science sets it apart from any other fellowship program in the country.

Mary Ann Leung, Randall McDermott, Ahmed Ismail, and Boyce Griffith at the 2003 DOE CSGF conference in Washington, DC.

ONE requirement of the fellowship program is that fellows take part in a three-month practicum at one of the Department of Energy's national laboratories. For fellows, the practicum experience is an opportunity to work with highly talented scientists in world-class facilities. After the experience, fellows leave the lab with increased skills, knowledge and aptitude as scientists. In fact, the four fellows profiled here found the practicum to be an invaluable experience that impacted their research in many positive ways.

A Calling Realized

MARY ANN LEUNG

University of Washington | Lawrence Berkeley National Laboratory | Story by Jacob Berkowitz

In 1997, Mary Ann Leung sold her house, quit her successful 16-year-long career as a trainer and instructional designer in the computer industry and, at age 38, returned to college to become a physician. Or so she thought. On her way to the medical sciences she discovered a long unrealized passion for math and the computational sciences. Now a Department of Energy Computational Science Graduate Fellow (DOE CSGF), she is using high performance computing to understand the role of quantum mechanics in applications ranging from nanoscale engineering to the emerging field of quantum computing.

Leung found her groove during a National Science Foundation-sponsored undergraduate summer stint in the lab of Dr. William (Bill) Reinhardt, a University of Washington professor of chemistry and adjunct professor of physics.

"I fell in love with the work I did," says Leung, now a third-year doctoral student in Dr. Reinhardt's lab at the University of Washington. Her job that first summer was to develop a computer program to visualize the time propagation of solitons in the Bose-Einstein Condensate (BEC). While planning a slide presentation of her work, she came across an image of a BEC generated by one of Dr. Reinhardt's computer programs. It was a moment of scientific and personal epiphany.

"The image had a strong likeness to a piece of artwork that I'd created 20 years earlier. I decided at that point that I'd been visualizing math much earlier in life, but I just hadn't realized it. So, seeing this image made me feel like I'd found my calling," recalls Leung.

Ironically, had she returned to school several years earlier she would never have seen her calling. In 1924, Albert Einstein and Satyendra Nath Bose hypothesized that noninteracting atoms chilled to supercold temperatures (about 190 nanoKelvin) would condense into a new form of gaseous matter, the BEC. What's remarkable about the BEC, they explained, is that in this supercold state the atoms would all go into the same quantum state: They would behave as a single super atom, a macroscopic quantum object large enough to be visible to the naked eye.

However, it wasn't until 1995 that three independent research teams were successful in creating the first BEC by supercooling rubidium-87, lithium-7, and sodium atoms. This breakthrough has set off a global wave of experimental work with BECs.

"They're a wonderful place for people to explore the fundamentals of quantum mechanics because you can see the quantum effects," says Leung, noting that BEC research has applications to fields ranging from superconductivity to nonlinear optics.

In the process of her doctoral work, initially focused on developing parallelized codes to understand and visualize the dynamics of solitons and vortices in BEC, Leung found her own kind of quantum coherence—a fellowship program that was just right for her.

"When I found the DOE CSGF, I thought: This is perfect for me. I love all of these things," recalls Leung. The Fellowship requires participants to take science, applied math, and computer science courses, and Leung was particularly inspired by an algorithms course.

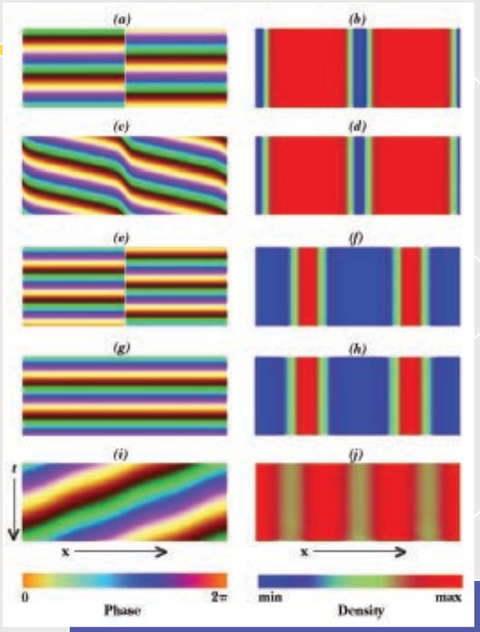

The five stationary solutions types in one dimension under periodic boundary conditions are propagated in time numerically for 100 natural time units. The phase and density are shown for (a) and (b) real solutions, repulsive case; (c) and (d) complex solutions, repulsive case; (e) and (f) antisymmetric real solutions, attractive case; (g) and (h) symmetric real solutions, attractive case; and (i) and (j) complex solutions, attractive case.

While planning a slide presentation of her work, she came across an image of a BEC generated by one of Dr. Reinhardt's computer programs. It was a moment of scientific and personal epiphany.

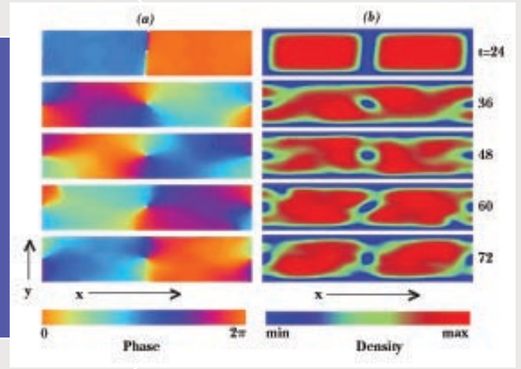

In 2D band solitons decay into pairs or opposites charged vortices. We added 0.01% stochastic noise to an initial two-soliton stationary solution in order to demonstrate this effect. Five slices of the xy-plane are shown at the midpoint of the z coordinate, equally spaced in time. The longitudinal direction has periodic boundary conditions. Note that the phase leads the instability.

But it was the Fellowship practicum that let her really test her mettle as an emerging computational scientist.

During a workshop she attended on the computational tools at DOE's Lawrence Berkeley National Lab, Leung met Dr. Andrew Canning, a staff scientist in Berkeley Lab's Computational Research Division. She was intrigued by his work: Although also exploring computational methods for quantum systems, Dr. Canning models solid state systems. It's a realm that offered a new challenge for Leung, one that she readily took on when offered the chance of a three-month practicum with Dr. Canning.

Leung's practicum research was part of an ongoing collaboration between, among others, Dr. Canning and Dr. Z.Q. Qiu, a solid state physicist who holds a joint appointment with the Lab's Materials Sciences Division and the Physics Department of the University of California at Berkeley. Dr. Qiu's group is a world leader in the creation of very pure, nano-scale metallic films. These thin metal films (sometimes only several atoms thick) are ideal for studying quantum mechanical effects and are at the frontier of a new realm of materials engineering.

"In the past, people have designed materials to have specific mechanical properties, such as strength," says Dr. Canning. "Now, because we can engineer at the atomic level, the idea is to design materials in which we can control the properties of individual electrons."

Using Berkeley Lab's IBM SP supercomputer, Leung modeled Dr. Qiu's experiments conducted at the Lab's Advanced Light Source, to understand the effect of the addition of a nanoscale nickel monolayer on quantum well states in copper. Quantum well states are an energy state in which an electron is sandwiched between two layers of atoms so that its motion is confined to a single dimension. They are thought to be responsible for the giant magneto resistance effect, the basis for the creation of very high density disk drives.

Leung and Canning's simulations have already provided a more detailed physical understanding of the nickel-copper quantum well state experimental results, as well as confirming some of the theoretical models used by Dr. Qiu's group.

Leung says that it was great to work closely with Dr. Canning and his post-doctoral students in "a collaborative environment that was really about pursuing interesting science." The experience provided Leung with a clearer sense of her potential.

"One of the important lessons that I learned from that experience was that you can apply yourself to a new field and within a relatively short amount of time, you can learn quite a lot and make progress," she says.

It a lesson she's taken back to her doctoral research at the University of Washington. Along with her supervisor Dr. Reinhardt, Leung is exploring the potential of the BEC and high performance computing for the study of quantum computing.

For over 10 years, the DOE CSGF program has encouraged the training of computational scientists by providing financial support to some of the most talented graduate students in the nation. Praise for the fellowship appeared in the National Science Foundation's Division of Mathematical Sciences publication *Mathematics and Science.*

But it was the Fellowship practicum that let her really test her mettle as an emerging computational scientist.

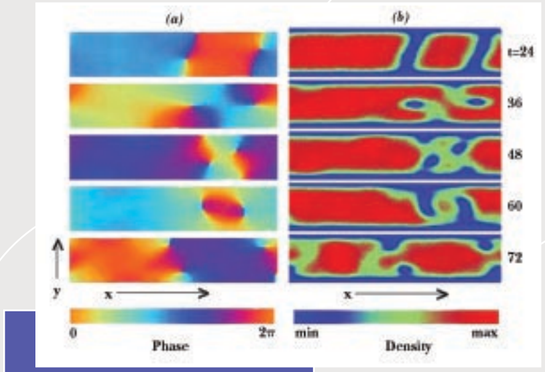

Collision of two initial band solitons that have evolved into two oppositely charged vortex pairs. We have added a 0.1% stochastic noise. As in the image on page 6, the relative velocity is induced by phase discontinuities across the band solutions, but the interaction is that of two vortex dipoles. This is an example of how to use the noise-induced instability of band solitons in 2D and 3D to study vortex interactions.

Dr. Reinhardt stresses that the work is currently at the very early stages of creating an original and challenging PhD research project. "It's an exciting, frustrating and turbulent process," he says.

But it's one that Leung has eagerly grasped.

"My vision is to create a computational model that will simulate the behavior of a quantum computer," says Leung. "So in my case, computational science not only means the application of high performance computing to a scientific problem, it also includes the investigation of new computational techniques using high performance computing."

Amidst the enormous international scientific buzz about quantum computing, Dr. Reinhardt is cautiously optimistic that BEC research offers a unique approach.

"Because BECs are big and you can control them with lasers, we're trying to figure out if we can use these very special macroscopic properties to make extra-stable pieces of quantum computers," he says, noting that as a large quantum object the BEC might be less prone to the quantum decoherence that disturbs single atom quantum computing systems.

At the same time as she is working towards numerically modeling a quantum computer, Leung is also helping others push scientific boundaries. She volunteers with organizations dedicated to encouraging girls to study science and engineering.

"In hindsight, I realize that I was always really good in math when I was a kid," says Leung. "But I never really received any encouragement to do anything with math. Quite the contrary, I was discouraged: 'That's not a girl thing to do'. Since I got such a late start in science, I really want to encourage young girls and women to go into science; it's a fascinating field."

PROGRAM REQUIREMENTS

Students selected for fellowships agree to undertake study and research in computational science. The program of study must provide background in a scientific or engineering discipline, computer science, and applied mathematics.

In order to be considered for the DOE CSGF, students must be U.S. citizens or permanent resident aliens and working toward a PhD at a United States university.

Students applying for fellowships must be undergraduate seniors or in their first or second year of graduate study.

Prior to the third year of the fellowship, fellows must complete a practicum assignment at a Department of Energy laboratory. Currently, approximately 20% of fellows who graduate from the DOE CSGF program work or have worked in a Department of Energy laboratory.

DISCIPLINES PURSUED

The fellows involved in the DOE CSGF study widely varying subjects. However, they all are using high performance computing towards their research goals. Fellows' disciplines include biophysics, chemistry, biochemistry, civil engineering, computer science, aerospace engineering, applied math, physics, bioengineering, aeronautical engineering, chemical engineering, robotics, computational chemistry, and computational mechanics.

Challenge of Simplification

AHMED ISMAIL

Massachusetts Institute of Technology | Oak Ridge National Laboratory |
Story by Jacob Berkowitz

Simplify. It's a way of being for some people. To rid themselves of clutter, they remove all that's non-essential so as to be left with the key elements of life. For Ahmed Ismail it's the key to his research. This computational minimalist spends his days trying to remove everything that's numerically extraneous, honing codes and algorithms until he's left with a gleaming gem of Truth.

The challenge is how to get to this place of computational enlightenment. While the goal is to simplify, the route there involves numerous complex steps. For Ismail, his journey to greater simplicity involved a southern sojourn, formally known as a practicum, at the DOE's Oak Ridge National Lab (ORNL) — a trip that opened this computational devotee's eyes to new ways to simplify.

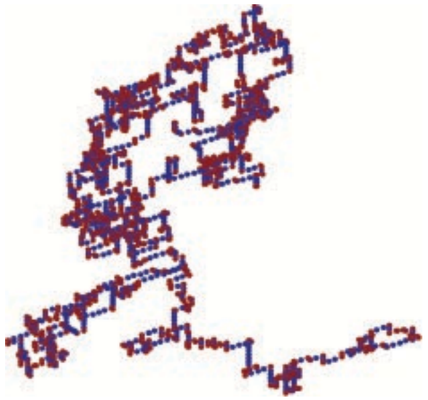

A 2048-bead polymer, modeled as a random walk on a three-dimensional cubic lattice. The red points indicate bends in the chain corresponding to stiffness penalties.

"My graduate and practicum research are focused on finding ways to systematically reduce the number of variables needed to computationally simulate a polymer system," says Ismail, a sixth-year chemical engineering graduate student at the Massachusetts Institute of Technology (MIT).

Known as coarse-graining, this process of simplification is crucial to the development of manageable, rapid and reliable simulations of complex processes in chemical engineering (and in other disciplines).

"Polymers are perhaps the largest and fastest growing area of materials for engineering applications," says Dr. Gregory Rutledge, Director of MIT's Program in Polymer Science and Technology, and one of Ismail's co-advisors. "The development of advanced, atomically detailed modeling methods is crucial to determining the molecular-level origin of a wide spectrum of polymer properties."

Ismail, and others in Dr. Rutledge's lab, are creating the algorithms for multiscale models. These models involve the study of the simultaneous behavior of polymer systems at multiple size levels, from electron-pair interactions at the level of quantum mechanics to the bulk properties of commercially produced materials. They also involve time steps ranging from femtoseconds to seconds. Integrating these multiscale processes necessitates finding ways to simplify the models. "With polymers, there can be hundreds of thousands of atoms per molecule. So, it isn't feasible to do any sort of direct calculation for a melt of these polymers," explains Ismail. "For example, if you have 100,000 of these polymers, and each one has 100,000 atoms, you'll never finish your simulation. However, for most of the properties that you want, you don't need to worry about the position of every single atom and electron. You just need to know the basic average properties of the system."

To accelerate the calculation of these average properties, Ismail uses a multiresolution technique called the wavelet transform. This mathematical tool acts as a filter through which data can be repeatedly processed until it merges thousands of variables into a single average value.

This process of simplification can follow one of the main types of coarse-graining: time-based or position-based. While both methods are designed to reduce the number of variables (or degrees of freedom) in a simulation, in practice they are significantly different.

Ismail's doctoral research has focused on the use of position-based coarse-graining to simplify a range of chemical engineering problems, beginning with theoretical simulation of Ising models, lattices of magnets.

"If you're in industry and you're looking at a new material and want an estimate of its phase behavior, the methods that I've been developing will very quickly give you an order-of-magnitude estimate for where this behavior is going to occur," says Ismail.

However, his contented path to simplicity took an abrupt turn last year when he got a call from Dr. William Shelton, a Senior Research Staff Member in the Computational Materials Science Group at ORNL inviting Ismail to do his practicum in Shelton's lab. A Connecticut native and a Yale graduate, Ismail had never spent time in the southern U.S. But his culture shock would involve more than just Knoxville's slower pace and a summer of ceaseless heat. Dr. Shelton's project was focused on time-based coarse-graining.

"A lot of graduate students tend to fall into this dynamic where they're looking just at their one research problem and lose perspective on what's happening in the whole field that they're working in. There's a larger world out there than the one little problem that you're trying to answer."

"From my point of view it was a different outlook on how to approach problems in my field," recalls Ismail. "To some extent I'd had blinders on for a little bit in the sense that I was looking at my tiny niche problem. A lot of graduate students tend to fall into this dynamic where they're looking just at their one research problem and lose perspective on what's happening in the whole field that they're working in. There's a larger world out there than the one little problem that you're trying to answer."

For Dr. Shelton, the project (prompted by anticipation of the DOE's new Computational Nanoscale Proposals) was a chance to tackle a key hurdle in the numerical simulation of molecular systems on longer time scales.

"A common problem in these systems is that the time step used is often on a femtosecond time scale. This makes it very difficult to get to the real physical dynamical process, which may occur at nano- or milliseconds," says Dr. Shelton.

"I was trying to find a way to effectively and rapidly model the fast degrees of freedom, such as bond bending, in a way that doesn't introduce too much error, but still allows you not to have to explicitly model these interactions. Because most of what slows down the simulation is the computation of the interactions between different parts of the system," says Ismail.

During the three-month practicum, Ismail, Dr. Shelton and Dr. Phani Nukala, a post-doctoral fellow in materials science at ORNL, were able to create algorithms to successfully demonstrate how to separate these degrees of freedom. The work will form the basis for a submission by Dr. Shelton as part of the Computational Nanoscale Proposals.

"The ideas work well in principle," says Ismail. "But implementing them with respect to currently existing codes was a bit of a challenge."

Modifying a code that has been gradually built up over more than a decade is like playing a game in which you have to remove a piece of structure without having it cause another piece to fall. If one code component is simplified, then other components that interact with it must also be altered.

For Ismail, his journey into numerical time scales involved "a lot more independence" than he was used to as a graduate student. At the same time, he says, "There was more cross-collaboration than I'd expected."

In late August of 2002, Ismail's DOE Fellowship journey took him back to MIT.

"When I returned to Boston the first thing that I thought when I stepped out of the airport was 'Wow, it's cold'," says Ismail.

It wasn't just his adaptation to a more southerly climate that had taken effect. He returned with a larger vision of coarse-graining. His doctoral work will now involve integrating time- and position-based coarse-graining approaches to polymer models.

As he looks to the future, this student of numerical simplicity sees the opportunity to apply his reductionist approach to the discipline itself.

"There's a lot of work in coarse-graining that needs to be developed," says Ismail. "There is a lot out there on how to do coarse-graining, but right now a lot of it is on an ad hoc basis. People pick a particular model because it works, but we need to develop strong first principles approaches and a systematic method for choosing how to simplify them."

The same walk shown as coarse-grained chain of 64 beads, with each bead located at the center-of-mass of the corresponding 32-bead segment of the original 2048-bead chain. The colors (moving from blue to red) indicate the relative configurational energy of the bead, as determined by summing up the individual stiffness penalties of each segment.

SCOPE OF PROGRAM

Since its inception, the DOE CSGF program has supported 204 students in over 50 universities all over the U.S. Currently it supports nearly 75 students in 21 states: Arizona, California, Georgia, Illinois, Indiana, Iowa, Massachusetts, Michigan, Minnesota, New Jersey, New Mexico, New York, North Carolina, Oregon, Pennsylvania, Tennessee, Texas, Utah, Virginia, Washington, and Wisconsin.

Thinking Small to Predict Large-Scale Turbulence

RANDALL McDERMOTT

University of Utah | Sandia National Laboratories – California | Story by Lori Valigra

Randall McDermott's research may sound like *Gulliver's Travels:* the chemical engineer spends his days in a world of tiny flow structures, with the ultimate goal of capturing the large-scale effects of turbulence, one of the most complex — and computationally expensive — research areas of fluid mechanics.

McDermott, a PhD candidate in chemical engineering at the University of Utah, is tackling a new field of science — simulation science — which combines mathematics, computer science, and physics. McDermott is using this field to predict the chaotic behavior of turbulent flows. If he is successful, such simulations could help climate modeling and weather forecasting, have significant impact on developing cleaner power plants, and even influence astronomers who model stars and galaxies.

Examples of turbulence surround us in everyday life. Noisy faucets, puffy clouds, rumbling engines, and waves in the ocean are all evidence of fluid turbulence. The vapor rolling off a pot of boiling water is a particularly good visualization of turbulent flow. Turbulence can have positive and negative effects: If left to its own devices, it can stoke a fire wildly out of control, but if harnessed, it can allow engineers to design cars and airplanes with less aerodynamic drag, reducing wear and tear and improving fuel efficiency.

But what is turbulence really? The *American Heritage Dictionary* defines turbulence as, “the motion of a fluid having local velocities that fluctuate randomly.” McDermott’s description, however, is somewhat different. “What motion we consider turbulent depends on scale,” he says. “Think of a rough ocean. To the naked eye, the sea is very turbulent. But, if you could zoom in closer and closer, eventually you would reach a point where the water looked like a placid lake. Here the fluid is laminar, because at such a small scale, the forces of viscosity are stronger than other disturbances.”

The unpredictable nature of turbulence has made its mathematical description challenging. Physicists believe that all the complexities of turbulent flows can be described by the celebrated Navier-Stokes equations, which were derived in the mid-1800s, and come directly from Isaac Newton’s laws of motion. They describe fluid motion down to a microscopic scale. However, even today's most sophisticated computers aren't powerful enough to solve these equations for large-scale turbulent flows, and until they are, some engineers say turbulence remains an unsolved problem.

"Imagine watching a fire in your fireplace for a minute. It takes us months and tens of thousands of computer hours to generate that calculation,” McDermott said, “even with models to describe the unresolved turbulent eddies." To put that into the context of the typical Pentium-based home computer, it would take roughly 1,000 personal computers running for two days straight to produce about 10 seconds of information about a fire.

The term “eddy” is often used to refer to the whirling motion of a fluid, like the motion of water behind a rock in a river. The approach McDermott employs is called Large Eddy Simulation (LES), and it has shown great promise. LES is a method for approximately solving the Navier-Stokes equations. Models are required because, as McDermott puts it, “LES is essentially a blurry motion picture of the Navier-Stokes equations. The models are what refine the image.”

McDermott is developing and validating models for the small scales of turbulence, but this is not a small task. It involves creating simplified equations, or mathematical descriptions of the physical world, and turning them into software programs that can work on today's computers. Specifically, he is working to convert the One-Dimensional Turbulence model, developed by Alan Kerstein at Sandia National Laboratories in Livermore, California, into an LES.

McDermott explained that his type of modeling, known as synthetic field modeling, takes certain points of turbulence activity and then tries to fill in the blanks between the points. "It's like taking a low-resolution picture with a digital camera and running it through some kind of imaging software that tries to refine the picture. That's what happens with us in our work on turbulence. We're trying to fill in what's between the points.”

McDermott became interested in combustion and turbulence as an engineer at John Zink Co., a Tulsa, Oklahoma company that develops new combustion products, technologies and clean-air products. This was McDermott's job after graduating with a BS in chemical engineering from the University of Tulsa in 1994. He returned to school for his PhD at the University of Utah in 1999, and became a DOE Computational Science Graduate Fellow in 2001.

The stimulus behind his move for a DOE CSGF fellowship was a desire to apply his background in practical, applied engineering to theoretical problems. "That led me back to graduate school," said McDermott. “I realized while working at Zink that I needed a bigger tool box to tackle the types of problems encountered in combustion applications. Combustion is really fun. In one way or another it touches on every aspect of chemical engineering.”

His advisor at the University of Utah, Philip Smith, encouraged him to take the DOE CSGF. And, Professor Smith is on a roll: McDermott is the second of his graduate students to be honored with the prestigious fellowship. Diem-Phuong Nguyen, highlighted in *Deixis* last year, graduated from the program in 2003. Her fond experiences likewise motivated McDermott to pursue the DOE fellowship.

When it came time to choose a practicum site, Randy chose Sandia National Laboratories in California. “Choosing the practicum location was not easy at all. I had boiled it down to the Bay area. In the end, our group at Utah has close ties with Sandia. But, I really don’t think I could have made a bad choice.” For McDermott, the practicum took five months, and led him down a path that he and his advisors at Utah and Sandia say changed his research focus for the better.

McDermott says he entered the practicum with a formulation that, after months of effort, proved not to work. He credits the open environment at Sandia, and the interactions there with the nation's top scientists in various fields, with changing his focus during the practicum. He now has working simulations, and he will use this research as part of his PhD thesis. "It's better to be right than to be done," he said of not finishing his work during the typical three-month practicum period. "Now, my thesis

For McDermott, the practicum took five months, and led him down a path that he and his advisors at Utah and Sandia say changed his research focus for the better.

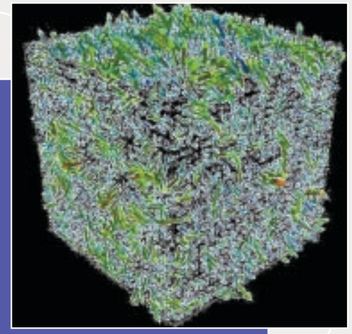

Velocity vectors colored by velocity magnitude for a 32^3 ("32 cubed") LES of isotropic turbulence decay in a periodic box. The point behind this simulation is to test subgrid models for proper energy dissipation.

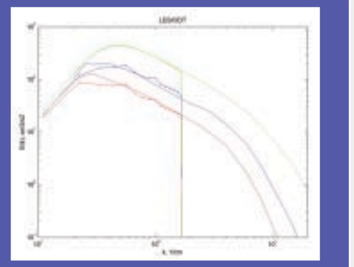

Three dimensional energy spectrum of the 1971 Comte-Bellot and Corrsin isotropic turbulence decay experiment. The solid lines without dots are the interpolated experimental data. The lines connected by dots are results of the large eddy simulation with ODT as the subgrid model. The colors represent three different points in time. The green is the initial condition and the energy decays from there.

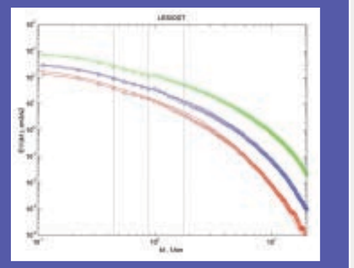

E11 energy spectrum of the combined LES and ODT fields for the 1971 Comte-Bellot and Corrsin experiment. The ODT is fully resolved with 512 points along one dimension. The vertical dotted line to the far right is the LES Nyquist limit. The one in the middle is the wavenumber at which the LES and ODT are joined. In other words, the data to the right of this line is from ODT and data to the left is from LES. The dotted line furthest left represents the wavenumber of the largest allowed eddy in the ODT model.

DOE CSGF HIGHLIGHTS

- ***Payment of tuition and required fees***
- ***Yearly stipend of $28,000***
- ***A $1,000 allowance for related expenses***
- ***Matching funds of up to $2,500 for a computer workstation***
- ***Opportunity to complete a practicum working with scientists and researchers at a DOE Laboratory***
- ***Yearly fellows conference with opportunities to meet other fellows and industry and government professionals***
- ***Renewable up to four years***

For more information: www.krellinst.org/csgf

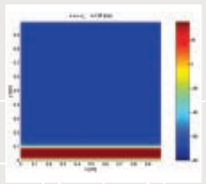

A

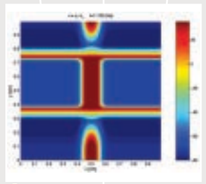

B

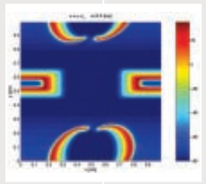

C

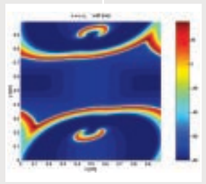

D

and practicum are converging upon one another." He said his interactions at Sandia were a highlight of the program. "Interacting with those scientists gets you a long way, seeing the way they think and approach a problem."

As Phil Smith, who is a professor and chair of the Department of Chemical and Fuels Engineering at the University of Utah, explained, "Randy is applying some very good models at smaller scales. He recognized that when we cannot resolve the turbulence, we need synthetic field models of what could be occurring, allowing us to simulate events at smaller scales than we can resolve in detail. He is seeing what information can be extracted from a large-scale event, like a fire caused by spilled jet fuel, to see what's going on at a smaller scale." These small-scale simulations one day might end up on supercomputers, simulating bigger events. Such simulations may eventually help fight fires better, lead to better containers for explosive chemicals and better building designs, and help rescuers know what to do in case of an accident.

McDermott's relentless efforts in the lab have impressed both of his advisors. "Randy is very focused on what he wants out of his research and what he will get out of an educational experience. He very much wants to see how theoretical science can be applied to real life," said Smith.

"Randy has special qualities. He takes ownership of the whole problem," said Alan Kerstein, McDermott's practicum advisor at Sandia. "He's not looking to narrow the focus to the smallest specialty as students often do to define a problem they can cope with. He is thinking broadly. He has chosen the most ambitious goal possible. He's both strong-minded and flexible, and is always pushing beyond what's known."

Kerstein said McDermott is very well organized, has a grasp of scientific problem-solving, and is willing to bounce ideas openly off colleagues. This is especially important in a field of science like simulation, which is very different from traditional science in that it uses an inductive approach, as opposed to a deductive approach, to explain experiments. "We are trying to find a set of equations for turbulence that will work in different applications," Kerstein said. "Randy has had a very valuable learning experience. He's built a new model from the ground up."

Simulating the Heart

BOYCE GRIFFITH

New York University | Lawrence Livermore National Laboratory | Story by Lori Valigra

When New York University mathematics student Boyce Griffith embarked on a cross-country car trip to California's Lawrence Livermore National Lab (LLNL) last year, he had only a hint of the ambitious research projects he ultimately would tackle during his DOE CSGF fellowship practicum.

Now preparing for his second practicum at LLNL, Griffith plans to build on his earlier work of developing software for large, parallel computers that aims to simulate how fluids and solid structures interact within large physiological systems, and get it to work on a three-dimensional model of the heart. The result might be a better understanding of heart problems, such as arrhythmias and sudden cardiac death, which cause about 300,000 deaths yearly in the United States alone.

Heart disease remains one of the leading causes of death in developed nations, and it can be difficult to obtain experimental measurements much beyond the surface of an intact heart. Griffith and his colleagues, who include Richard Hornung, his practicum advisor at LLNL, and Charles Peskin, his academic advisor at the Courant Institute of Mathematical Sciences at New York University, are hoping the numerical simulations Griffith is developing will work in a three-dimensional model of the heart, and thus reveal more about the inside mechanical, electrical and other dynamics of that vital organ.

We often use the expression, "in a heartbeat," but what those words mean when broken down into mathematical equations and computer language is quite a complex matter. "It takes 24 hours of dedicated access to a Cray supercomputer with 16 processors to simulate one-and-one-half heartbeats," said Griffith. The current software, developed by Peskin and David McQueen, also a researcher at NYU, was designed to take full advantage of shared memory parallel computers such as the Cray. However, the most powerful computers today are distributed memory machines, which presents additional programming and algorithmic difficulties. Griffith is redesigning the simulation software so that the heart model can be run on a distributed memory architecture. "We'll then hopefully be able to get answers faster and with more accuracy," he said.

Griffith's advisor at New York University, Peskin, along with McQueen, developed a computational 3-D mechanical heart model that enables

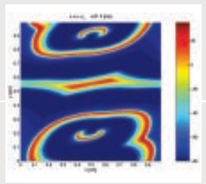

E

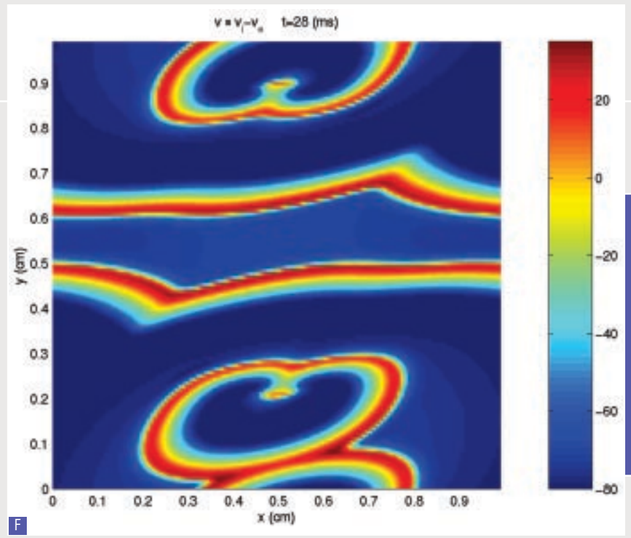

F

Images A-F show the formation of a spiral wave in cardiac tissue from a 2D simulation.

Images G and H are from a simulation of electrical activity in a 3D model of cardiac tissue.

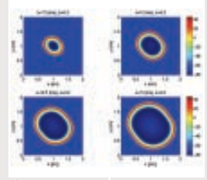

G

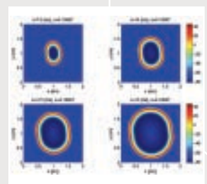

H

experiments in silico which would not be possible in vivo. To devise the model, he used mathematical modeling, high performance computing, fluid dynamics, physiology, engineering and other disparate disciplines to create the simulated heart. The ongoing collaborative project has already taken more than two decades.

Although several approaches can be taken to simulate the heart's activities, Peskin's is based on the immersed boundary method, a numerical method he invented for the study of the interaction of blood and heart valves. For example, a valve is elastic, so the fluid flowing through it will tend to deform it as it is passing through. Similarly, the movement of the valve will influence the flow of the blood.

One approach for describing this situation is to divide up space into two regions, one that contains the fluid and one that encloses the tissue. This kind of approach requires complicated descriptions for both the fluid and tissue regions, complicated further by the fact that these descriptions change over time. The immersed boundary method simplifies this situation by blending together the tissue and fluid models. This approach still allows for a complex description of the tissue geometry, but doesn't require reinitializing of either the fluid or the tissue regions as they deform over time. Essentially, in the immersed boundary method, the tissue is treated as if it is part of the fluid that exerts elastic forces both on itself and on the rest of the fluid.

The immersed boundary method, although used largely in bio-medical applications, can also be applied to model turbulent flows, such as the flow of wind by automobiles or trucks. But the method is typically applied to objects that are more flexible, such as the heart.

As part of his work, Griffith is developing new algorithms and developing new software to marry knowledge of the electrical to knowledge of the mechanical properties of the heart. He also is writing new software so that it can work on distributed memory supercomputers.

One reason Griffith chose LLNL for his practicum is that the lab, under the guidance of practicum advisor Richard Hornung, had developed software tools known as SAMRAI, which stands for Structured Adaptive Mesh Refinement Applications Infrastructure. SAMRAI is designed to simplify the development of complex simulations. Adaptive mesh refinement (AMR) makes it possible to reduce the amount of work spent on the parts of the simulation that do not need high resolution and to focus on only those regions that do need maximum resolution. This saves on computing time and costs. At least one region in which high resolution is needed is the area of tissue interaction with the fluid, and the immersed boundary method formulation makes this region especially easy to track. "At the time I was choosing practicum sites, I was more interested in the electrophysiology modeling/simulation end of the heart project, so AMR and the immersed boundary method were all new things to me," he said.

Griffith said he hopes that by the time he finishes his PhD research, with the help of his advisor he will have meshed the mechanical heart model that Peskin developed with an electrical model that works to trigger the contractions of the heart. As he continues his software development, Griffith said, he also is using AMR, and multigrid solvers, as well as investigating the use of inexact Newton methods to make the computations more efficient and accurate.

In an unusual move, Griffith will complete a second practicum (most students have only one). In the first practicum, he developed some AMR software for the immersed boundary method. He worked with Hornung to get a basic, non-AMR immersed boundary simulation working on some test problems, but did not get far enough along to have an entire heart model running on that software. He hopes to finish that work, with some additional electrical modeling, during his second practicum.

By Karyn Hede

Unfolding Proteins

PROTEIN FOLDING has often been likened to origami — the ancient Japanese art of paper folding. It's nearly impossible to discern by studying the pattern of creases in an unfolded two-dimensional sheet of paper that it can transform into a delicate swan or a scorpion poised to strike. A similarly vexing challenge faces those who attempt to see the swan in the ugly duckling — the tangled string of amino acids that form the starting material of all completed, neatly folded protein structures.

>> Bioengineering
>> Structural Dynamics
>> Fusion Energy Science
>> Energetic Materials
>> Quantum Chromodynamics
>> Process Engineering
>> Environmental Science
>> Computational Biology

Understanding how a string of amino acids, each with its own unique chemical and physical properties, folds up to form a fully functional protein has presented a formidable challenge for decades.

In fact, just as a whole branch of mathematics is dedicated to the study of origami forms, the challenge of protein folding has engendered an impressive array of mathematical models designed to predict protein folding behavior. And with the avalanche of DNA and protein sequence information being generated by genomics and its counterpart, proteomics, the rush is on to convert that data into information about fully functional folded proteins.

Solving protein structure requires a combination of knowledge from physics, biology, chemistry, and perhaps the most limiting factor: massive computational firepower.

"High performance computing has transformed our ability to take on the challenge of global energy minimization," says Teresa Head-Gordon, assistant professor of Bioengineering at the University of California at Berkeley, and faculty staff scientist at Lawrence Berkeley National Laboratory LBNL). "We couldn't have done this ten years ago."

Head-Gordon, together with Silvia Crivelli of the Visualization Group in Berkeley Lab's Computational Research Division and their collaborators, have developed a new approach to solving protein structure that borrows from knowledge of robotics and computer gaming and combines it with physical chemistry and global optimization algorithms that predict the most stable three-dimensional structure for a given protein.

The key to predicting protein structure, says Head-Gordon, is finding the global energy minimum. The principle goes back to a 1972 Nobel Prize-winning theory developed in the late 1960s by Chris Anfinsen of the National Institutes of Health, which stated that the three-dimensional structure of a protein is the one in which the free energy is at its lowest. The free energy takes into account the watery environment of the protein and all the various atomic interactions of the amino acids.

Since these initial experimental insights, scientists have discovered that proteins fold into three-dimensional shapes that include a limited number of preferred configurations that define their secondary structure. These shapes include alpha helices, beta pleated sheets and various types of coiled loops. The final folded protein may bring amino acids that are far apart from each other in the starting amino acid chain close together, just as, in the creation of an origami swan, opposite corners may end up folded together. Taking into account the limitations of movement imposed by the protein's chemical components still creates a mathematical origami that can tax the capacity of even the powerful parallel computers available at LBNL.

"We had to start with an initial structure in order to search the possible conformational states," says Crivelli. "We have to have secondary structure already formed to do global optimization. We were using local optimization with some constraints to form that secondary structure. It took hours or days of computational time to get those structures formed. I thought that was too inefficient, and as we started to look at larger proteins it became even more of a problem."

Crivelli's idea was to be able to manipulate long strings of amino acids without "breaking" them.

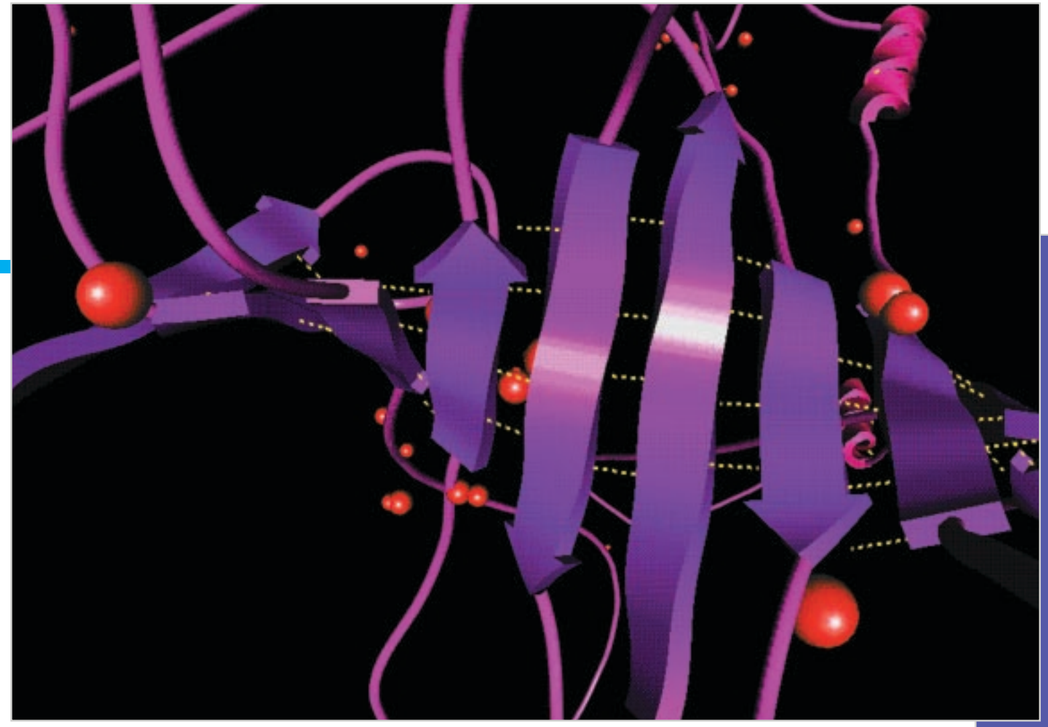

Atom collision and hydrogen bond rendering in ProteinShop.

To solve the problem, Crivelli sought a way to cut the time needed to reach a reasonable secondary structure, one that incorporated all the likely alpha helical regions and beta sheets, to an hour or two.

Rather than using time-consuming local optimization methods, she sought to create a tool that would allow a user to experiment among various potential conformations while simultaneously monitoring the energy profile of the protein. She enlisted the help of Nelson Max, a computer visualization expert at the University of California at Davis, and his graduate student Oliver Kreylos, along with LBNL colleague Wes Bethel.

Crivelli's idea was to be able to manipulate long strings of amino acids without "breaking" them. "We wanted to be able to move a secondary structure such as an alpha-helix or a beta-strand from point 'a' to point 'b' while maintaining the integrity of the protein. We wanted it to look natural." After trying several approaches, the group hit on the idea of borrowing from the field of robotics. They took advantage of a technique called "inverse kinematics," which calculates the movements of several articulated joints at once to reach a desired point in space. The method is used in both computer gaming and robotics to make a character's movements on screen as natural as possible. For example, it allows a virtual swan to bend its neck in a realistic way.

"In robotics," Crivelli explains, "when you move an arm, all the joints move. Our problem was similar except we had 80 or 90 joints."

The result is called "ProteinShop," an interactive computer visualization tool that allows a user to experiment with various protein conformations before settling on a "best guess," which can then be subjected to rigorous global minimization algorithms. Just as in the cell, amino acids are added to the structure one at a time. As the protein grows, the program searches through a database of known protein structures

COLLABORATORS

Teresa L. Head-Gordon is both a Professor of Bioengineering at the University of California at Berkeley and a Faculty Staff Scientist in the Physical Biosciences Division at Lawrence Berkeley National Laboratory (LBNL). She received her BS in chemistry from Case Western Reserve University and her PhD from Carnegie Mellon University in theoretical chemistry. Since joining LBNL in 2001, she has continued to hold the position of Department Head of the Computational Structure Group, Physical Biosciences.

Silvia Crivelli has spent over ten years developing software in high performance computing. She currently works as a Staff Scientist in the Computational Research Division at LBNL, and as a Visiting Associate Research Engineer at the California Institute for Quantitative Biomedical Research at the University of California at Berkeley. She received her BS in applied mathematics from the Universidad Nacional del Litoral, Argentina, and both her MS and PhD in computer science from the University of Colorado at Boulder.

Further Reading:

S. Brown & T. Head-Gordon (2003). Cool-Walking: A New Markov Chain Monte Carlo Sampling Method. *J. Comp. Chem. PAK Symposium* 24, 68-76.

T. Head-Gordon & S. Brown (2003). Minimalist models for protein folding and design. *Curr. Opin. Struct. Bio.*, in press.

S. Crivelli, O. Kreylos, N. Max, and W. Bethel (2003). *ProteinShop: A Tool for Interactive Protein Manipulation.* Submitted to: Journal of Molecular Graphics and Modeling.

S. Crivelli, B. Bader, R. Byrd, E. Eskow, V. Lamberti, R. Schnabel, and T. Head-Gordon (2002). *A Physical Approach to Protein Structure Prediction.* Biophysical Journal 82, 36-49.

Contact:

Teresa L. Head-Gordon
TLHead-Gordon@lbl.gov

Silvia Crivelli
SNCrivelli@lbl.gov

PRACTICUM COORDINATOR

Phillip Colella
pcolella@lbl.gov

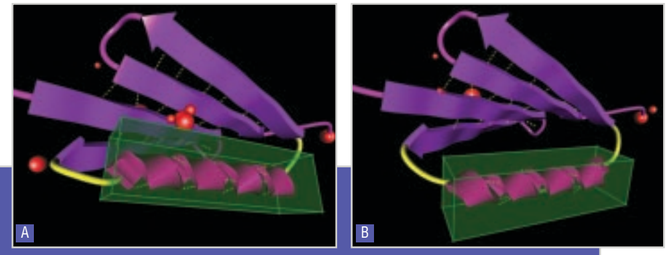

These figures illustrate the manipulation feature in ProteinShop. Figure A shows a user selecting an alpha-helix and the coil regions that will change with the manipulation (yellow lines). Figure B shows the effect of the manipulation: the alpha-helix has moved and the collisions that it was creating have disappeared.

to select the most likely secondary conformation. At any point, a user can manipulate sections of the protein to form new conformations while ProteinShop monitors the movement's effect on the energy of the system.

"The user can click on two amino acids to form a hydrogen bond and the program will automatically align them," says Crivelli. "We found that when a user tries to form a hydrogen bond in two dimensions, it's hard to see if those bonds are aligned, so we let them know when they are getting close. You can see dotted lines forming and that guides you on the two-dimensional screen. Similarly, if the user tries something that is going to bring two atoms into collision, we let them know there's something wrong. Orange spheres pop up on screen to warn them."

Once the initial conformation takes shape, Head-Gordon's expertise in biophysical chemistry and optimization algorithms takes over and produces a predicted global conformation. Predictions are determined by mathematically "trapping" a protein by use of an algorithm Head-Gordon has dubbed "antlion," after an insect that traps its prey by excavating a hole and waiting for an ant to fall in. The simulation program takes into account the forces on each atom and calculates each atom's position as forces are applied over time. In addition, the team uses a global optimization approach called stochastic perturbation that draws on secondary structure patterns to provide guidelines and hints for the optimization method. The process is repeated many thousands of times until a calculation settles on the most stable configuration.

Complicating the analysis is the difficulty of calculating the energy of a protein chain in the watery environment of a living cell. Head-Gordon's latest algorithm, which models proteins in water, takes into account the natural tendency of fatty (hydrophobic) amino acids to cluster together, just as oil droplets converge in a pot of water.

The team's approach was put to the test when they participated in a biennial competition that pits various modeling strategies against one another in a head-to-head competition to predict structures of proteins whose three-dimensional shape is unknown. Teams competing in "Critical Assessment of Techniques for Protein Structure Prediction," or CASP, are given only the amino acid sequence and a deadline for producing a predicted protein structure. Simultaneously, structural biologists are busy solving the protein structures experimentally using X-ray crystallography and other techniques that form a picture of the protein's actual shape. When the competition is over, the structures are revealed and competitors are rated on how well their predictions conform to the solved protein structure.

"It's like drinking from a fire hose," says Head-Gordon. "The sequences just keep coming and the deadlines are short, a matter of a few weeks. It really taxes your resources."

In the most recent competition, CASP5, held in the spring and summer of 2002, the LBNL team, which also includes Bobby Schnabel, Richard Byrd, and Betty Eskow from the University of Colorado, tackled 20 proteins for which no tertiary structure information was available. These proteins, called "new folds," are the most difficult to solve.

When the initial results were revealed in a December 2002 meeting held at Alisomar, California, Head-Gordon's team had placed between 13 and 15 out of nearly 200 participants.

"I was very proud of our effort," says Head-Gordon. "Predicting beta sheets is more complicated than alpha helices. In CASP5 we went all out and competed across the whole spectrum of proteins, including those with high beta content and very difficult topologies. We learned that our infrastructure is solid."

In particular, colleagues and competitors alike were uniformly enthusiastic about the potential of ProteinShop to become a universal tool that can be applied to many different protein folding methodologies.

Crivelli and her colleagues are busy modifying ProteinShop to allow researchers to adapt it to their various needs.

"As global optimization evolves, we would like ProteinShop to be able to talk directly to the parallel computing machine," says Crivelli. "The idea is that users could manipulate structures from a biological point of view and then put them back in the queue for more optimization. This would be a terrific project for a graduate student."

Meanwhile, Head-Gordon is looking toward being able to design materials never thought of by nature.

"If we had a basic understanding of the rules of self-assembly, we could change the sequence of a protein and create a new topology," she says. "This would give us the ability to create new enzymes that are inspired by nature but that nature never thought to make. … Or we could design new materials with the structural and functional properties we want."

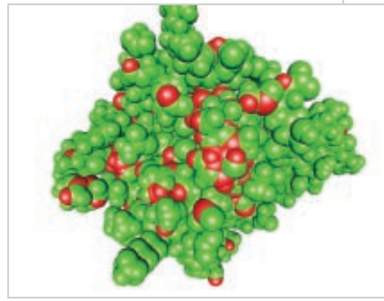

An energy visualization done in ProteinShop.

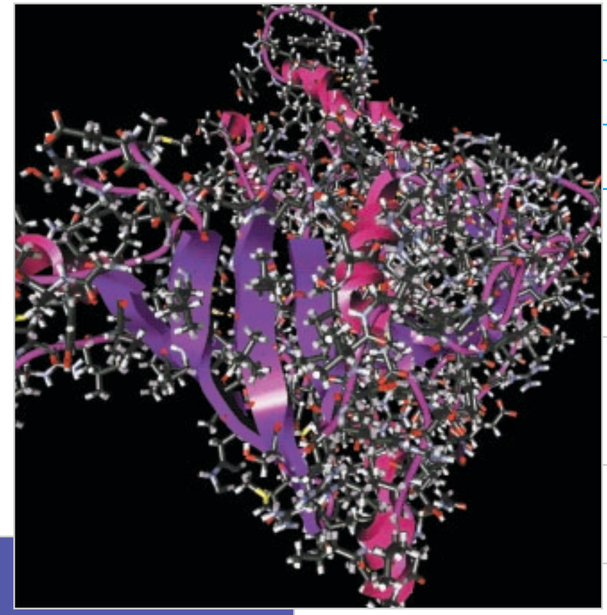

ProteinShop rendering of structure cartoons and side chains.

By Victor D. Chase

Studying Structures

QUESTION: **What do an aircraft carrier, a submarine, nuclear weapons, and a miniscule gear (many of which could fit on the head of a pin) have in common?**

ANSWER: Salinas.

"Salinas" is the Spanish word used as the name of the computer code that has been used to model and analyze the response of these vastly different structures to the various stresses they may endure.

Developed over the past seven years at Sandia National Laboratories, Albuquerque, New Mexico, Salinas falls into the category of finite element analysis programs, many of which are commercially available and used by engineers in designing all sorts of structures. But Salinas is different, in that it has significantly raised the bar on the speed and complexity of the analyses that can be run, which is what enables it to model structures as intricate as aircraft carriers and nuclear weapons with a high degree of accuracy.

"With Salinas, we can solve problems that no one else can solve," says Manoj Bhardwaj, a member of the Salinas team. That this is no idle boast is evidenced by the fact that the Salinas developers received a prestigious 2002 Gordon Bell Award at the Supercomputing 2002 conference. Considered to be the Super Bowl of computing, the Gordon Bell competition honors those who have developed computer applications that provide significant performance improvements.

Days to Minutes

The significance of the improvements made possible by Salinas becomes evident when one considers that for the Bell award, the Salinas team used 3,700 CPUs to solve some 100 million equations in six minutes, according to Bhardwaj. Typically, several days are required to solve far fewer equations.

Salinas achieves these dramatic results by being one of the first finite element analysis programs with the ability to run on massively parallel supercomputers. These machines are made up of thousands of CPUs that operate in parallel. This means that extremely complex equations can be divided into millions of parts, many of which are then solved simultaneously by the various CPUs and instantaneously reassembled to give the desired result. The supercomputers being used by the Salinas team are part of DOE's Advanced Simulation and Computing Initiative (ASCI), which encompasses several DOE sponsored laboratories. *(See Detecting Defects, on page 26, which discusses the use of an ASCI machine to model explosives at Lawrence Livermore National Laboratories).*

The ability to run on supercomputers is also what makes Salinas a one-code-fits-all program that can conduct analyses of devices ranging in size from aircraft carriers down to miniature gears.

"Parallel computing allows complicated geometries to be broken up into tiny pieces that are solved in conventional ways," explains Garth Reese, Salinas's principal investigator. "For example, an aircraft carrier is a plate-type structure, so we can break it up into little rectangles, solve the equations of stress and strain and motion of those rectangles without too much difficulty, and then tie each solution together. You end up with a large system of equations that tie together the solutions from each of these little finite elements."

Salinas's considerable capability allows it to be used to model numerous "what if" scenarios without having to conduct actual nuclear tests.

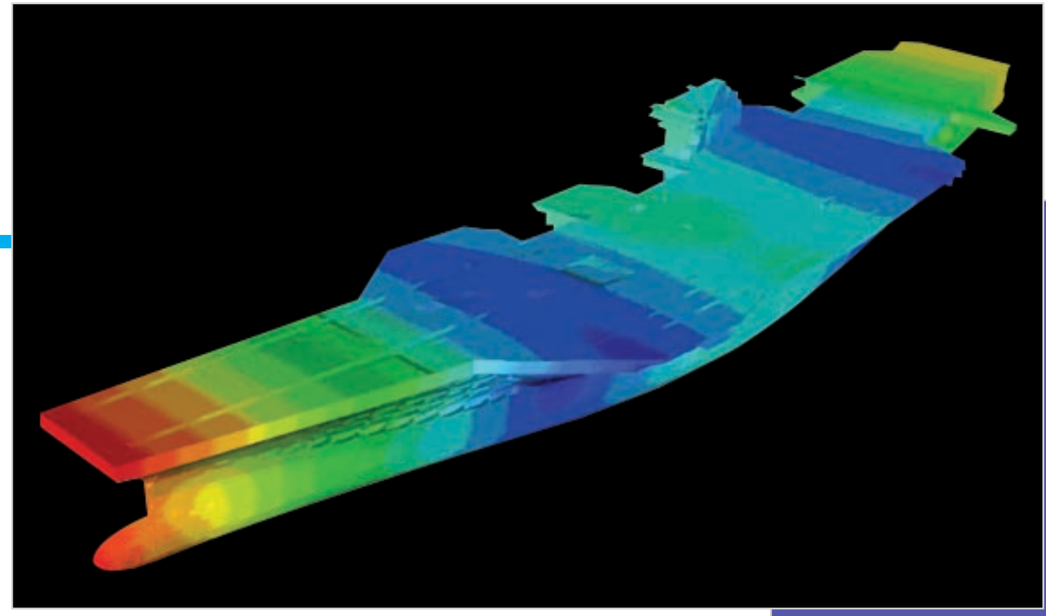

Finite element model of aircraft carrier being displaced in one of its natural frequencies. The colors represent the relative displacement of different parts of the model. The results were generated under a joint effort by the U.S. Navy (Naval Sea Systems Command), Sandia National Laboratories, and Northrop Grumman Newport News.

Bulletproof Bombs

Though it has a myriad of applications, Salinas was created primarily to model nuclear weapons and how they will react to hostile environments. This assignment stems from the fact that the DOE, which sponsors the Salinas program, is responsible for developing and maintaining the nation's nuclear weapons stockpile. Salinas plays a particularly important role in this program because the U.S., in compliance with the spirit of the Comprehensive Nuclear Test Ban Treaty, has not done any nuclear testing since 1992. Salinas's considerable capability allows it to be used to model numerous "what if" scenarios without having to conduct actual nuclear tests.

Additionally, many of the events that can be modeled by Salinas would be difficult if not impossible to create even if nuclear weapons tests were being conducted. And, even if such hostile environments could be created, accurate data extraction would be nigh on to impossible.

Salinas can, for example, computationally examine the results of a nuclear weapon being struck by the pressure blast from an

GARTH REESE

Garth Reese received his BS in physics from Brigham Young University in 1980. In 1984 he received his PhD in physics from Arizona State University. Immediately following graduation, he began work at Motorola GEG in Scottsdale, AZ, where he designed and analyzed surface acoustic wave filters and resonators. In 1986 he began working at Sandia National Laboratories, doing device development and modeling of piezoelectric components. Since 1991, he has worked in the structural dynamics and vibration controls department, where he has been involved in dynamics analysis, system identification and health monitoring. Garth is involved in development of structural system identification software for reconciliation of finite element structural dynamic models with test results. He is a principal author of Sandia's *Virtual Environment for Optimal Test.*

Manoj K. Bhardwaj received his BS, MS, and PhD in aerospace engineering from Virginia Tech. Immediately after receiving his PhD, Manoj started work in the Structural Dynamics group at Sandia. He has been part of the Salinas team for six years.

Contact:
Garth Reese
gmreese@sandia.gov

Manoj K. Bhardwaj
mkbhard@sandia.gov

PRACTICUM COORDINATORS

Albuquerque - Marcus Martin
mamart@sandia.gov

Livermore - Paul Nielan
pen@sandia.gov

anti-ballistic missile as it hurls toward its target, or the consequences of a fratricide event in which multiple nuclear weapons are launched and one is detonated near another.

"Simulation gives a far better understanding of what's happening in these situations and the data is far more complete than could be derived from real tests," says principal investigator Reese. This data can include information not only about whether a weapon will fail or not, but even more specifically about which parts are susceptible to failure. To do this, Salinas shows how the load generated by a strike on a weapon would be propagated from the impulse down to the various inner components. In fact, the program gets so specific, "I can say which particular component is going to break," notes Reese. "The intent is to learn how to design the structure so it will stand up to a certain level of attack."

Sea and Battle Worthy

Similarly, the goal of the aircraft carrier analyses is to aid designers in creating the most seaworthy and battle worthy ship possible while it is still on the drawing boards. This can eliminate a lot of trouble and expense, such as that which was incurred when an early carrier was found to ride too low in the water at the bow, after it was launched. Utilizing Salinas, "We will let designers have a picture of what's going on before the ship is built, so such retrofits will not be needed," said Reese.

One facet of the picture painted by Salinas shows a ship's natural vibration frequencies. "A large boat will twist and turn in the sea; you get metal bending everywhere," explains Reese. He cites, for example, the load experienced by a ship as it passes through two waves simultaneously, causing each end to be riding on a crest, while the middle is in a trough. Under these conditions some deformation is natural, and "It's important to know how it deforms naturally, what the natural modes of the vibration are," says Reese. Armed with this knowledge, engineers can design the ship so that it does not vibrate excessively and be damaged in rough seas. In this sense, aircraft carriers are no different from automobiles, albeit on a grander scale. If, for example, a car's tires are out of balance, the car will vibrate excessively at certain speeds, eventually causing damage to other parts of the car.

Once the natural vibration frequencies of a carrier are known, Salinas can examine its response to more hostile environments, such as what would occur if it were struck by a large conventional weapon.

Most of the damage done by weapons fire results not from shrapnel, but rather from the huge air blast created by the explosion. That is what happened, for example, to the guided missile destroyer the USS Cole during an October 2000 terrorist attack in the port of Aden, Yemen. When a small boat laden with explosives was detonated beside the ship, the resulting air blast blew a hole in the side of the ship, killing 17 sailors and injuring 39.

Using the results of Salinas's computational analyses, engineers can design ships to better withstand such massive blasts. In the case of aircraft carriers, for example, the size and placement of sponsons — the huge horizontal plates that protrude out over the hull and serve as part of the flight deck — can make a difference, since they also help absorb the effects of air blasts.

Stealth Submarines

Studying the effects of air blasts is part of the recent expansion of Salinas to include analyses of fluid action, in addition to the structural interactions it was originally designed to model. In this case fluids include gases (including air), and liquids (such as water). The Salinas group may apply this newly created capability to understanding what happens when a sound wave, or a ping, hits a submarine. Knowing how the structure responds, and how the ping reflects back out, are key to being able to design stealth submarines.

Fluid structure interaction analysis is also important to the fail-safe design of very small parts like the microelectromechanical (MEMS) gears, which may someday find their way into nuclear weapons. These gears can be built into silicon integrated circuit chips, enabling them to operate in concert with microelectronic systems.

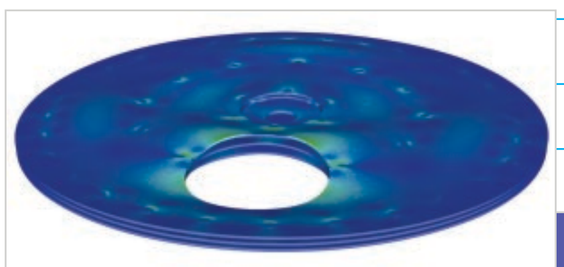

Finite element analysis of a three layer MEMS gear. The colors in the image represent the stresses in the gear.

Studying the effects of air blasts is part of the recent expansion of Salinas to include the analyses of fluid action, in addition to the structural interactions it was originally designed to model.

In the case of the MEMS gears, the fluid in question is air, which can have a deleterious effect on the gears as they rotate at 50,000 cycles per second. "A MEMS gear can be designed to run for a long time, but may not, because air sucks the energy out of the structure," explains Reese. Hence, "understanding the interaction between the air and the structure is both very important and quite difficult to model."

Salt of the Earth

Following a tradition among Sandia computer scientists, the Salinas code was named after a local pueblo. The namesake for this program is a pueblo on the east side of New Mexico's Sandia Mountains. Salinas is also the Spanish word for salt, and, says Reese, "While structural dynamics is not as flashy as is analyzing parts falling off a structure, or studying a large crash, it is crucial for the success of any structure. It is the bread and butter, or the salt. It is the kind of down-to-earth thing you have to have to get the analysis done for any system."

A WORLD-CLASS SOLVER

>> ***What makes Salinas different from other finite element analysis programs is its ability to run on massively parallel supercomputers (see accompanying article), and what gives it that ability is the "world-class linear solver" that serves as Salinas's backbone, says principal investigator Garth Reese.***

And what the solver does is make Salinas scalable. "This means that if we take a model with one million different pieces in it and run it on 500 processors, and then take a model that is twice as large and run it on 1,000 processors, we get the solution just as quickly in either case," says Reese.

On the surface it seems obvious that if a problem is made 50% more complex, and 50% more processors are used to solve it, the time required to achieve the solutions to both problems will be the same. But when it comes to scalability of computational equations on parallel computers, such seeming logic does not apply. In fact the reverse can be true — more processors can slow down the pace of calculation. This is so for two reasons: One is that the processors need time to communicate with each other, so the greater the number of processors, the greater the communication time required. Secondly, parts of some problems don't lend themselves to being solved by parallel computers; thus waiting for those slowpokes to be processed can delay the entire process.

The situation is not unlike the "too many cooks spoil the broth" proverb, or to put it in another context, two carpenters can reasonably be expected to get the construction of a house completed in half the time it would take one. But, if 100 carpenters were on the job, they would get in each other's way, causing a delay in the building process. The same is true for parallel computers.

"The coordination effort can often exceed the gain in work as the number of processors grows," says Reese. "Unless you are really careful, you can't scale more than two or three processors or you get to the point where each processor is waiting on another, or spending too much time in communication."

And therein lies the beauty of the linear solver — the result of a collaborative effort between Sandia and the University of Colorado at Boulder — that restores logic to the world of computing by enabling scalability. And it is that scalability that allows Salinas to run equations on thousands of processors simultaneously so it can quickly and accurately analyze aircraft carriers, miniscule gears, and anything in between.

super

By Peter Gwynne

Futuristic Fusion

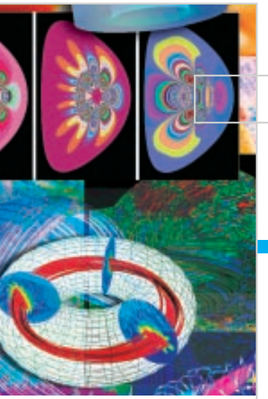

THE QUEST TO TAME the fusion power that makes the Sun shine relies on more than experimental devices. A group at the Oak Ridge National Laboratory is using computer simulations to design a futuristic fusion reactor.

Scientists have tried for half a century to tame the type of energy that makes the Sun shine. The process, called nuclear fusion, occurs when two atomic nuclei overcome the usual force of repulsion and slam together in a collision that creates a single, heavier nucleus and an excess of energy. The most common form of fusion involves isotopic forms of hydrogen, the simplest and lightest element in existence. When nuclei of the hydrogen isotopes deuterium and tritium collide, the impact creates a nucleus of helium, the next simplest element, and produces enough energy, in the case of the Sun, to light and heat the solar system and to support life on Earth.

On Earth itself, that same process creates the explosive power of the hydrogen bomb. However, just as a controlled version of the reaction responsible for the atom bomb provides the energy of nuclear power plants, fusion can theoretically be constrained in such a way as to produce a source of peaceful power far greater and longer lasting than any available from oil, gas, and nuclear stations.

Scientists recognize that commercial fusion reactors won't appear for several years. For the moment, scientific teams in the Department of Energy's national laboratories are working on the fundamental science and engineering that will lead to that goal.

Their research involves the creation of plasmas — gases at temperatures so high that their atoms' electrons and nuclei separate and move independently. Maintained long enough at a high enough temperature, a hydrogen plasma can stimulate enough fusion reactions to generate net energy. The teams have two main approaches: inertial confinement and magnetic confinement.

Inertial and Magnetic

Inertial confinement comes in several forms. The National Ignition Facility project at Lawrence Livermore National Laboratory provides a typical example. It will use laser beams to blast pellets that contain hydrogen isotopes symmetrically from all directions. The bombardment squeezes the atoms so close together that their nuclei fuse to produce helium nuclei and energy. The small, intense pulses of energy produced by igniting burning pellets will theoretically keep a fusion reactor producing power in much the same way that explosions of gasoline in an automobile engine keep the car moving.

The other method, magnetic confinement, more closely resembles the process of solar fusion. In the sun, the force of gravity holds together and compresses the plasma of hydrogen nuclei. In magnetic confinement, powerful magnetic fields perform the same task. The goal of this approach is to maintain the plasma at temperatures and pressures so high that fusion produces more energy than that originally used to generate the temperature and pressure.

Research on magnetic confinement relies in large measure on computer simulation. Why? "The problems are so difficult," explains Don Batchelor, section head of the theory group in the Fusion Energy Division at the Department of Energy's Oak Ridge National Laboratory (ORNL). "The equations are non-linear, you have spatial non-uniformity, and you encounter extreme anisotropy with factors up to a million to one."

In strictly layman's terms, computer simulation helps to compensate for a fundamental difficulty of research on plasmas: containing the material. No solid can hold them, because plasmas exist at temperatures far above the melting points of all known materials. That leaves magnetism as the only practical means of keeping plasmas in place. Creating the necessary magnetic fields is both complicated and expensive, however. Thus, the more research scientists can do with computers, the faster their projects will advance.

Batchelor's group has three main objectives. "First," he says, "we want to understand the physics of plasmas. Second, we are developing computational tools to ask questions, find the solutions and what they mean. Experiments give you a lot of data but don't tell you what the data mean; we use our computer codes to understand what the experiments have done. And third, we predict what new experiments are going to do; this gets more difficult and more important as the machines get larger and we understand more of the physics."

In strictly layman's terms, computer simulation helps to compensate for a fundamental difficulty of research on plasmas: containing the material.

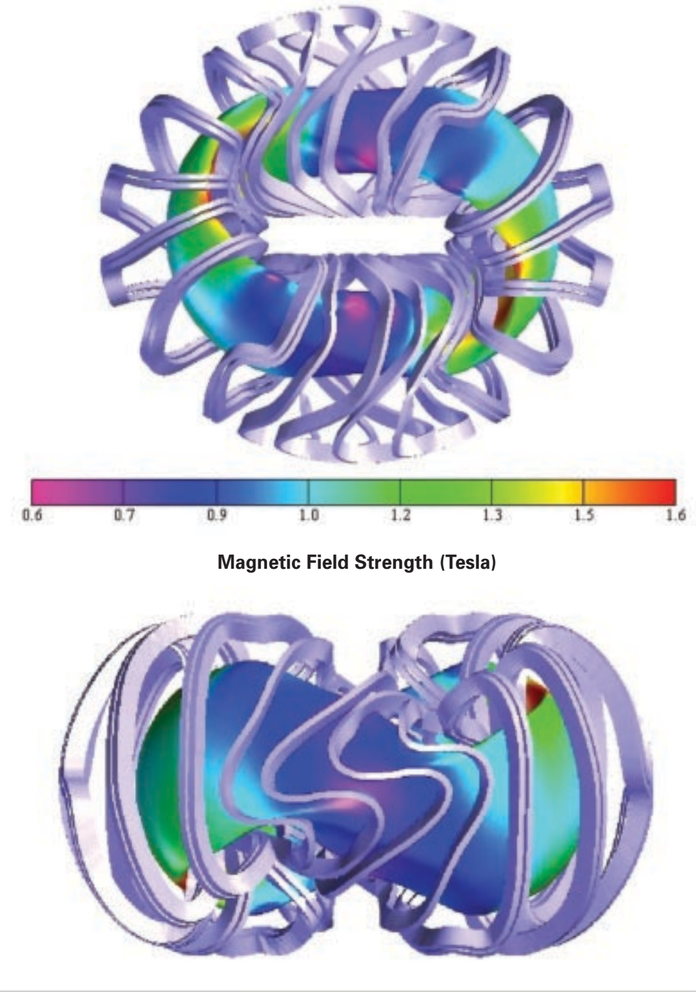

Physics/engineering based parallel optimizations used to synthesize innovative compact QPS stellarators.

DONALD B. BATCHELOR

Don Batchelor's arrival at the Oak Ridge National Laboratory 27 years ago to work on plasmas represented a return to his geographic and scientific roots. He was born at Oak Ridge during the Manhattan project, and he made a plasma jet as a high school science project. Having earned his Bachelor's degree in mathematics from MIT, he returned to the world of fusion, taking a PhD in plasma physics from the University of Maryland in 1976. At that point, he recalls, "I found a very large, very active theory group at ORNL doing very interesting things, including an oddball experiment called the ELMO bumpy torus, a device driven by plasma waves. So I came to ORNL. Since nobody else here was studying plasma waves more than qualitatively, I got into that work. I've been here ever since."

Further Reading:

Fusion Energy Science published by DOE, Jan. 2001, DOE/SC-0029

Fusion Science: Harnessing the Energy of the Stars published by a consortium including the APS, Division of Plasma Physics, Fusion Power Associates, General Atomics, MIT, PPPL, and University Fusion Associates.

Web sites

General Atomics - http://fusioned.gat.com/

PPPL - http://www.pppl.gov/fusion_basics/pages/fusion_basics.html

Coalition for Plasma Science - http://www.plasmacoalition.org

Contact:
Donald B. Batchelor
batchelordb@ornl.gov

PRACTICUM COORDINATOR

Jeffrey Nichols
nicholsja@ornl.gov

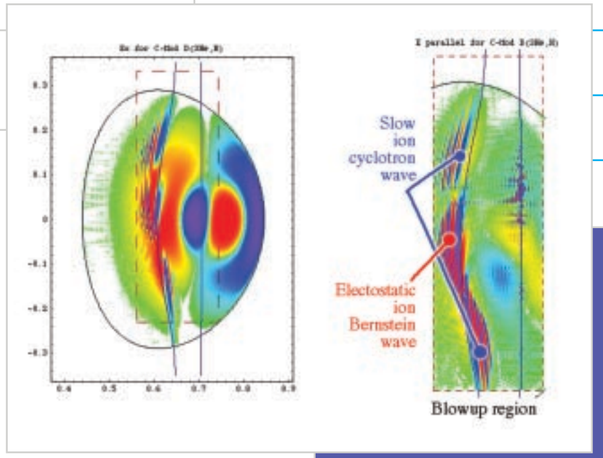

First fully resolved 2D calculations of an important wave process, conversion of fast waves to short wavelength modes, obtained on CCS machines. We find that fast, long wavelength electromagnetic waves launched from the right can be converted to slow electromagnetic ion-cyclotron waves, as well as the previously expected electrostatic ion Bernstein waves.

Tough Challenges

Understanding the basic physics of plasmas presents a series of tough challenges. "Here at ORNL we have two teams studying waves and heating in plasmas; the work is very complicated geometrically so that there's very little you can do with pencil and paper," says Batchelor. The most complex research involves understanding how turbulence within a plasma causes the system to lose energy, and hence perform less efficiently than it should. "Understanding turbulence is a very difficult nonlinear problem. It's difficult to work out even for normal liquids, let alone plasma hotter than the Sun," Batchelor continues. "We try to get a set of equations to describe the processes and then solve them. We start with models. We then put them on the computer, get solutions in simplified cases, and go through iterations to get closer to the real thing."

Success in the research projects obviously relies heavily on the choice of hardware and software. "We have a range of computer tools," Batchelor says. "Some of the physics teams spend a lot of their time working on very large software tools and models. We're always starting new efforts in software. But we have a lot of valuable resources out there to build on."

To run the software, ORNL's Center for Computational Science offers four massively parallel computers that can work on several different calculations simultaneously. Batchelor's team can also call on external computing power. "Our access to big computing here has been essential," he says. "But it doesn't matter where the computers are. All our studies have made a lot of use of local computing power at ORNL. We also use the increasingly powerful computing strength in our department. We have a new Cray vector machine and are expecting more within the next year." Batchelor wouldn't object to extra hardware. Asked how much computing power his group needs to carry out its research, he responds: "How much have you got?"

Two Types of Reactors

Physicists have devised two mainline approaches to fusion reactors. Both use magnetic fields to hold plasmas. They differ in the positioning of those fields. In tokamaks, invented in the Soviet Union in the late 1960s, the electric current that creates the magnetic fields runs inside the plasma, which is shaped like a doughnut. Stellarators, first developed in the United States a few years earlier, use external coils that surround plasma shaped like a cruller to generate the magnetic containment.

Stellarators created to date have suffered one significant drawback; they have had to be huge to be able to hold in the energy and particles that make up plasma. Batchelor's team has come up with a new stellarator concept that promises to overcome the problem. "We have developed an approach that can be much smaller but can still have good confinement and stability," he says. "It has the unusual property that, as the pressure goes up, the stability and the ability to hold the particles gets better."

To indicate the size of a stellarator, physicists use a criterion called the aspect ratio. That is the ratio of the larger diameter to the smaller in the asymmetrically shaped plasma. The best example of a traditional stellarator, a machine called the W7X now being built in Germany, has an aspect ratio of 11. "Ours has an aspect ratio of 3," says Batchelor. "We're looking at a tractor tire rather than a bicycle tire."

The group's scientists used their theoretical understanding of and experience with plasmas to create the new concept. And they extended the example of the W7X's designers by reversing the usual process of first designing a set of coils to produce the magnetic field and then working out the characteristics of the plasma that it would hold. "We started with the ideal target plasma that doesn't yet exist and then worked our way back to the coils that we would want to contain it," Batchelor explains. "We've put together a lot of different physics analyses in one computer model and have used non-linear optimization processes to improve the physics environment and simultaneously come up with a geometry that's feasible to build."

Finding Funding

As they continue to refine their simulations, the ORNL team has also involved itself in politics and economics. "We're trying to get funding to build our stellarator," Batchelor says. For that effort, outside help is available. The team is designing two stellarators — one for the Princeton Plasma Physics Laboratory and the other for ORNL.

The team also collaborates with the University of Montana and with stellarator teams in Germany and Japan. And the fusion program as a whole involves several other groups at ORNL, including physics, fusion materials, and engineering.

The program has hosted at least one DOE CSGF fellow. Eric Held worked in the group from 1995 to 1999. "We always have room for more," says Batchelor.

The program plays a double role in fulfilling the goals of the Department of Energy. "DOE has an energy and a science mission," explains Batchelor. "We're in the uncomfortable position of being in both. Until about five years ago, fusion appeared in an account called energy. Then we were moved to the science account. Our research is definitely a science program, but it's driven by an energy goal."

It's also driven by success. Last year the stellarator design team won ORNL's award as research team of the year. Continued success seems guaranteed as the team continues work on simulations that will point the way to a working fusion reactor.

Stellarators created to date have suffered one significant drawback; they have had to be huge in size to be able to hold in the energy and particles that make up plasma.

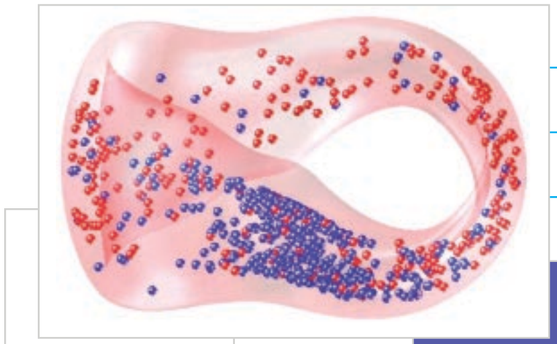

Massively parallel particle simulations for physics analysis.

THE PATH TO COMMERCIAL FUSION

>> In the 1970s, physicists working on fusion anticipated that, with an appropriate level of funding, their research could lead to working fusion reactors at any point from the year 2000 on according to the amount of funding their research received. Since then, funding has dwindled considerably, forcing researchers to push back the date at which they expect commercial reactors to make their debuts. In the interim, however, they have learned a great deal more about the science and technology of fusion reactors.

The path to commercial fusion contains two milestones. The first 'scientific breakeven' is a controlled fusion reaction that puts out at least as much power as it receives to start it off. Further along the route to commercialization is 'engineering breakeven.' A fusion reactor in that state would generate at least a small amount of power that can be tapped for use outside the reactor.

The first milestone is within reach. "We're within a factor of three of scientific breakeven," says Don Batchelor of Oak Ridge National Laboratory. "There's no doubt that we can achieve that." The International Thermonuclear Experimental Reactor (ITER), a project that brings together scientists from Europe, Japan, Russia, and the United States, is designed to produce about 700 megawatts, an amount of power that will take it well past scientific breakeven.

Progress beyond that will rely as much on economics as on physics and computer science. "We've made an incredible amount of scientific progress," says Batchelor. "Tokamaks and stellarators have much better performance than anyone thought possible. But more money is needed to move toward commercial feasibility. The U.S. can't do it alone." In fact, the Department of Energy is now formulating a plan that will see fusion power reaching the electrical grid in about 35 years. The scientific foundation is firmer than previous ones because scientists now understand the scientific and technical path to the goal in a way that earlier fusion researchers did not. But funding remains the critical and still uncertain component necessary to guarantee the use of fusion power.

By Victor D. Chase

Detecting Defects

WHEN SOMETHING YOU'VE PURCHASED IS DEFECTIVE, you usually send it back for a replacement, or throw it out and buy a new one. If you happen to be in the explosives business, however, the opposite holds true, for when it comes to creating explosions, defects are a good thing. In fact, detonation cannot take place without them.

But, as with virtually everything else in life, moderation is the key to desirable defects. Too many defects make an explosive material unstable, and therefore unsafe, while if there are too few, it won't detonate at all. But with just the right number of defects, the material will explode precisely when it is supposed to, and under just the right circumstances.

The defects in question are pockets, or voids, within and between the grains of chemical explosive materials. Once such a material starts to explode, the resulting pressure attempts to fill in those voids. As the material changes shape to flow in and around the defects, it becomes hotter than material away from them. The process of heating by changing shape also occurs when one bends a paper clip back and forth. The work of flexing releases heat.

In the case of explosives, as the voids are filled in and the nearby material becomes hotter, chemical reactions begin to occur at a faster pace. If the hot spots are big enough and hot enough, the organic molecules in the explosive begin to decompose, forming hot gas, and boom! An explosion takes place.

But determining the optimum number of defects to build into explosives is no easy task. Trial and error has been the answer to date. But, of course, when it comes to explosives, error can be especially undesirable. So it was to take the guesswork out of the equation, by bringing the precision of computational science methodology to the task, that the Grain-Scale Dynamics in Explosives (GSDE) project was created at Lawrence Livermore National Laboratory (LLNL), Livermore, California.

"What we are trying to do is harness the power of the Advanced Simulation and Computing Initiative (ASCI) at Lawrence Livermore to computationally look in detail at the behavior of explosives to help guide their design," says Jack Reaugh, who heads the project.

Under ASCI, a set of supercomputers has been created by tying together thousands of central processing units (CPUs) into one "massively parallel" machine. Complex problems are divided up among the CPUs, each of which solves some of the equations, thereby greatly speeding up the computing process. "The achievement of these machines is not so much that each CPU is much faster than any other, but that there are so many of them harnessed together," explains Reaugh. As a result, when it comes to supercomputers, the ASCI system "is as super as you get," he says.

ASCI was created to advance the state of computing hardware and software for all of DOE, and Reaugh's project is but one of many that make use of its computing capabilities. In fact, without ASCI it would be a practical impossibility to run the kinds of detailed computations necessary to analyze defects in explosives.

Nuclear Stockpile Safety

The GSDE project has been under way for four years, and though Reaugh hopes to expand its research to cover other types of explosives in the future, for the moment the group is using its computational prowess to focus on keeping the nation's nuclear stockpile safe, since chemical explosives are used as part of the nuclear firing sequence.

Safety is an issue for the chemical part of the stockpile because these explosives are organic, and as such they can change over time. If the defects change, the explosives may become more or less sensitive than they were originally intended to be. The goal of the computational group is to be able to predict whether such changes impact the effectiveness and safety of the explosives. "So far, we've accomplished some things, but we are not at the point of declaring victory," says Reaugh.

One of the first hurdles the group surmounted involved developing a method of computationally describing the makeup of the plastic-bonded explosives used in the stockpile. These explosives are produced by using a plastic binder to coat the grains, or crystals, comprising the material and then compressing them into a high-density mass. So that the grains fit together during the compression

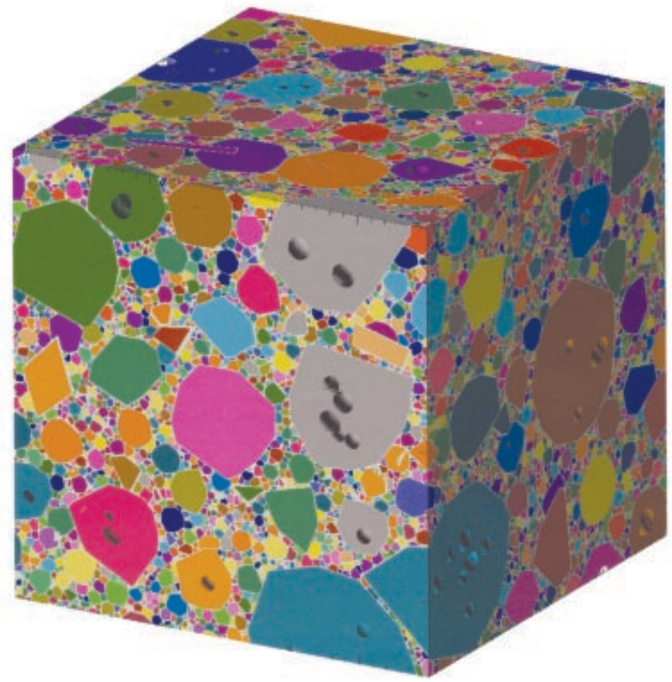

The computer-generated 0.3 mm cube contains approximately 0.1 million HMX grains with a bimodal grain-size distribution. Defects, shown as spherical voids, are present in the larger grains. Those that are intercepted by the surface of the cube appear as divots. Each color represents a distinct crystal orientation.

If the defects change, the explosives may become more or less sensitive than they were originally intended to be. The goal of the computational group is to be able to predict whether such changes impact the effectiveness and safety of the explosives.

JACK REAUGH

Jack Reaugh has worked at Lawrence Livermore National Laboratory (LLNL) since 1984. He is currently the leader of both the ASCI Grain-Scale Dynamics in Explosives project and the High Explosives and Organics Group, H-Division, Physics and Applied Technologies Directorate. Jack holds both a BS and MS in physics from the University of Illinois. Prior to his work at LLNL, he spent 18 years working in the private sector on computer simulations and response model development.

Further Reading/Viewing:
An animation of a 10GPa shock interacting with defects in an explosive can be viewed at http://www-cms.llnl.gov/s-t/hydrodynamic_mod_images/hydro.qt

An article describing the project appears in the March 2003 issue of *Science and Technology Review,* an LLNL publication. It may be viewed at http://www.llnl.gov/str/March03/Reaugh.html

Two related papers presented at the 12th International Detonation Symposium may be downloaded from the site http://www.sainc.com/onr/detsymp/

Click on: Technical Papers of 12th International Detonation Symposium (search for papers by S. Bastia, J.E. Reaugh and E.L. Lee)

Contact:
Jack Reaugh
jreaugh@llnl.gov

PRACTICUM COORDINATOR
Carol Woodward
cswoodward@llnl.org

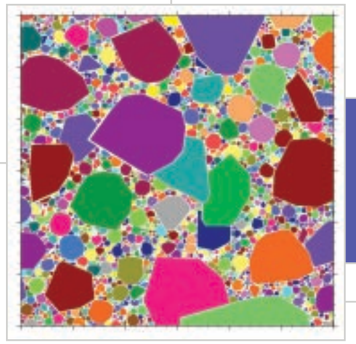

Synthetic photomicrograph of a computer-generated plastic-bonded explosive illustrating a bimodal grain-size distribution without defects. Each color represents a distinct crystal orientation. The plastic binder shows as a white outline around each grain. Slice is 0.3 x 0.3 mm.

process, they must be of widely varying sizes. This allows the smaller plastic coated pieces to fit into the spaces between the larger pieces.

Creating a computational mesh that accurately accounts for the particle size distribution is an extremely complex task. And, "If you don't get the particle size distribution right, you don't get the distribution of the defects right," notes Reaugh. Although the researchers have succeeded in developing a realistic computational array of particle size distribution, one of their ongoing tasks is to strive for ever better mathematical descriptions.

Microscopic Tricks

To derive the information from which they create computer models of explosives, the researchers borrow from empirical experiments and from their knowledge of basic physics. One way of deriving empirical information is "to use certain tricks of microscopy to look at crystals that go into the formation of these explosives and visualize what the defects look like within each of the particles," says Reaugh.

But microscopes are not of much help when information is needed about what actually happens in an explosive shockwave as the defects collapse and additional heat is released.

"One of the keys to knowing the extent to which an explosive will begin to detonate under a shock pressure is knowing something about the speed of the deflagration at high pressure," says Reaugh. But, he adds, "That was not known because it is difficult to perform experiments at pressures of a few hundred thousand atmospheres."

Yet such difficulty did not dissuade another group of LLNL scientists, whose help he enlisted to perform the high-pressure experiments. To do so, they are using a laboratory apparatus known as a diamond anvil cell. It includes two opposing diamonds, each about one-half carat, separated by a gasket containing a small sample of the explosive material being tested. The strength of the diamonds keeps the high pressure contained as a laser ignites the material, and the resulting deflagration is recorded on high-speed film.

A Cast of 30

The use of the expertise of researchers from outside the GSDE project to conduct these diamond anvil experiments is but one example of the collaboration among LLNL scientists from many groups to further this project. Even the funding for the diamond anvil cell experiments, which comes from the Munitions Technology Development Program, is the result of joint sponsorship by the Department of Defense and the DOE.

A related example of cross-pollination is found in the development of the software that Reaugh's group is using. "A cast of about 30, who are independent of my project, are developing physics and chemistry models and the numerical methods needed to solve the models in the computer simulation program that we use," says Reaugh.

The program they developed is dubbed ALE 3D, with the ALE standing for Arbitrary Lagrangian Eulerian — Lagrange being a French mathematician and Euler a Swiss mathematician. The hydrodynamic equations used in the program make use of two frames of reference, each associated with one of the mathematicians. The Lagrange frame provides a mathematical mesh that, in the simulation, moves with the explosive material. "It's as though a mesh is painted on the material, and as it moves about the mesh moves with it. The alternative is a Eulerian frame where the mesh is fixed in space and the material flows through it," explains Reaugh, a physicist himself.

By making use of this general software tool, which is available to many researchers using the ASCI computers, "hydrodynamic flow, heat transfer, and chemical reactions were already a part of the simulation program that we didn't have to develop for this project," says Reaugh.

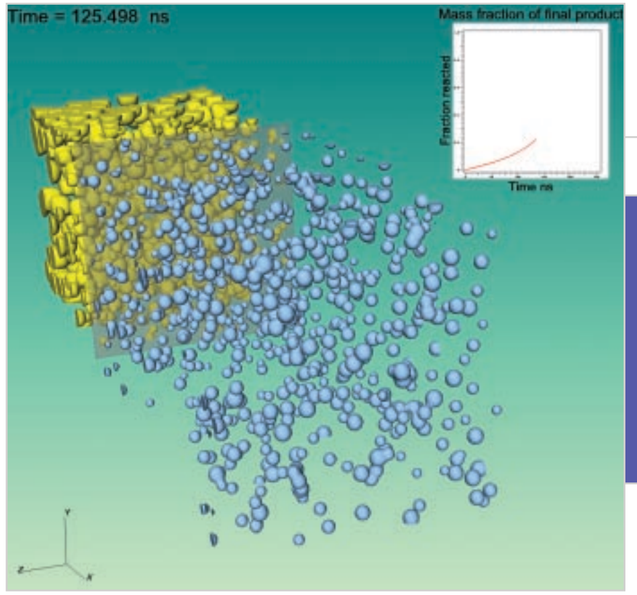

Snapshot from a simulation of a 10 GPa shock wave passing through a 1.2 x 0.3 x 0.3 mm brick of HMX with a spherical defect population. The shock (translucent gray surface) is progressing from back (left) to front (right). The voids (spherical bubbles) have collapsed behind the shock front, and the resulting hot spots have triggered flame fronts (yellow surfaces) that spread outward and link up in the reaction zone. The increased pressure in the hot gas behind the shock front drives the shock faster until a detonation is achieved within the 1.2 mm travel distance. Inset shows the mass fraction converted to gas product in the flames.

Utilizing input gleaned from the laboratory tests, ALE 3D, and their own substantial calculations, Reaugh's group has developed a model that can simulate a very small piece of explosive.

"Right now, the largest piece we can look at in computational detail is a cube a few tenths of a millimeter, and that's certainly not big enough to be practical," admits the group leader. And even at that size, the GSDE calculations have already pushed the ASCI supercomputer system to its limits. It took 400 CPUs about 3.5 hours to computationally assemble that miniscule cube of explosive material, and doing an actual calculation of the assembly then required over 300 CPUs running an additional 100 hours.

"We could, in principle, do the calculation for a large piece of explosive," he explains, "but it would require so much memory and so long a computational time that with the present ASCI machines we can't do it," he adds.

Which does not mean that the GSDE project's work is stymied. On the contrary, "We have plenty to do," says Reaugh. "We went through the process for the first time, and like anything you do for the first time, you learn how you should have done it, and what you should have done better.

"In doing these calculations we use models to describe the behavior of materials, which are, by their nature, approximate. In some cases they were accurate, and in some cases they weren't, so revising the models, describing the materials better, and looking in more detail at what the defect population really is, are all factors that can be examined computationally with the tools we have now."

Current planning calls for the GSDE project to continue its work for at least four more years. And when its goals are fully and safely accomplished, it will undoubtedly go out with a bang, albeit one that will contain no surprises.

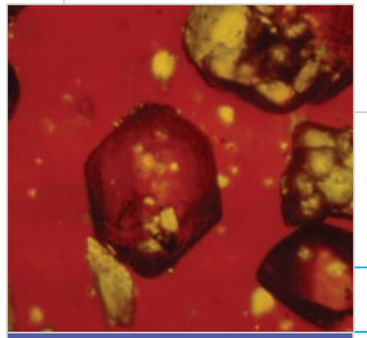

Photomicrograph of defects in crystals of the explosive HMX. The defects appear as yellow areas both inside and on the surface of the crystals. The central crystal is approximately 0.15 mm in diameter. Photo courtesy of Dr. Mark Hoffman, LLNL

By Jacob Berkowitz

Calculating Quarks

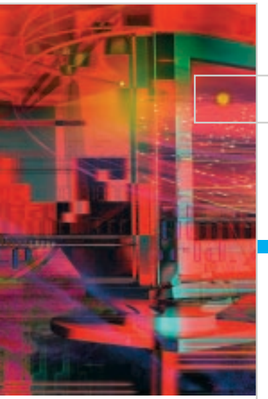

PARTICLE PHYSICIST DR. ROBERT (BOB) MAWHINNEY paints a picture of the subatomic world that makes you want to hit the replay button. It's dizzying. You lean closer, squint, as he describes his specialty, quantum chromodynamics (QCD), the study of the interactions among particles that have color charge — quarks and gluons.

OK, quarks aren't that foreign. The six flavors of quarks are the bare-bones constituents of all matter. We are mostly quarks. Combine two up-quarks and a down-quark and you have a proton. As for color charge (hence the "chromo" part of QCD) it's the quark-gluon equivalent of the electric charge in electromagnetism. Just as electrically charged particles interact by exchanging photons, in strong nuclear interactions color-charged particles interact by exchanging gluons.

When it gets messy, even for QCD devotees, is when you try to *calculate* the interactions between quarks and gluons. Sure, at high energies this world of strong quantum fluctuations is quite linear. It's at low energies, at the level of the proton, that things get messy. A quark turns into a quark plus a gluon. Then it interacts with another quark via a second gluon, while the first gluon turns into a quark/antiquark pair, then back to a gluon, and is reabsorbed by the first quark.

"You just have this enormously complicated, seething nonlinear sea of virtual particles. It's a perfectly posed computational problem," says Dr. Mawhinney, a professor of theoretical physics at Columbia University. "We have great faith in the underlying QCD equations because they're built on the principles of relativity and quantum mechanics, and at high energies, where we can calculate analytically, the results agree with experiment. But there are many phenomena predicted by the equations that we couldn't calculate with pencil and paper. So, it comes down to a question of computational strength to be able to calculate the physical consequences."

To do this, Mawhinney is part of a team from Columbia, the RIKEN Brookhaven Research Center (RBRC) at Brookhaven National Laboratory (BNL) and IBM that is designing and building the latest in a series of massively parallel computers to numerically probe the world of QCD. The 10-Tflops peak performance supercomputer, dubbed "QCD On a Chip" (QCDOC) is set to be booted-up at BNL in early 2004. QCDOC won't just push the bounds of quark-gluon physics — including understanding of the universe's early evolution — it's also a world-leading demonstration of the power and promise of topical computing.

Custom Made, Please

"There are only a few machines in the world that are built specifically for a topic," says Ed McFadden, who manages Brookhaven's Scientific Computing Facility, where QCDOC will be housed. But the 40-year BNL veteran (whose first computational color problem had nothing to do with quarks, but rather with solving the four-color map problem) notes that these purebred machines are attracting more interest than ever, including the interest of DOE's SciDAC (Scientific Discovery through Advanced Computing) program, which is partially funding QCDOC's design.

QCDOC won't just push the bounds of quark-gluon physics — including understanding of the universe's early evolution — it's also a world-leading demonstration of the power and promise of topical computing.

QCDOC will be the latest iteration of a computational vision first adopted by Dr. Norman Christ in the early 1980s. Since that time, the now senior Columbia University theoretical physicist has been leading the design of QCD-specific computers. When Dr. Mawhinney arrived at Columbia as a post-doc in 1990, Dr. Christ was modeling quarks and gluons on a custom-built 256-node, 16-Gflops machine. Within a couple of years, the pair — along with Dr. Al Gara, now at IBM, and a team of graduate students and post-docs — had set their sights on greater things: a 1-Tflops, 20,000 processor massively parallel computer, the QCD on Digital Signal Processor (QCDSP). The first QCDSP, funded by DOE, came online at Columbia in early 1998, and a second one, funded by RBRC, at BNL later that year.

The QCDSP was the perfect computational fit for the experimental work being done in BNL's Relativistic Heavy Ion Collider (RHIC).

"One of RHIC's primary goals is to detect a quark-gluon plasma, a phase of matter thought to exist in the early universe, in which quarks and gluons aren't bound into protons but are a gas or plasma. It's a very tricky thing to do experimentally, and nobody's been able to convincingly detect it yet. However, in the QCDSP we can calculate the properties of this quark-gluon plasma state, such as temperature and pressure, which can feed into the theoretical predictions for RHIC physics and our understanding of the early universe," says Dr. Mawhinney.

Still, why the years of time and effort to build a custom machine?

"Cost," says Dr. Mawhinney without a pause. He estimates that the approximately $4 million to build the QCDSP was about one-tenth the price tag of an equivalent commercial massively parallel machine at the time.

But for Mawhinney and the rest of the QCDSP and QCDOC team, the benefits of a dedicated machine extend far beyond the initial cost savings. These machines represent the chance to optimize the hardware and software to your science. In this case, the result has been the creation of compact, low-operating-cost machines of enormous reliability. *(See sidebar, Small is Beautiful on page 33).*

COLLABORATORS

Edward McFadden has worked at Brookhaven National Laboratory for 40 years and is a member of the Information and Technology Department. He is currently working as director of the Brookhaven Computing Facility, and he specializes in parallel hardware.

Robert Mawhinney received his PhD in theoretical physics from Harvard University in 1987 and is currently a professor of theoretical physics at Columbia University. His research focuses on theoretical particle physics, with primary emphasis on the nonperturbative structure of nonabelian gauge field theories. Currently he is involved in studying QCD with the Columbia 16-gigaflop parallel computer, built five years ago by Professor Norman Christ and his collaborators. Their simulations include the effects of two light quarks, which is about 100 times harder computationally than simulations without the light quarks.

Further Reading:
For more information on the QCDOC project, visit: http://www.bnl.gov/cdic/Sci_Projects/Basic_Energy/MolecularDynamics/MolecularDynamics.htm

Contact
Robert Mawhinney
rdm@physics.columbia.edu

Edward McFadden
emc@bnl.gov

PRACTICUM COORDINATOR

James Glimm
glimm@bnl.gov

Science on a Grid

The point of departure for the optimization of these QCD machines is lattice gauge theory. The theory enables the calculation of quantitative QCD equations by divvying up space and time into a regular grid of points. "It's a critical step in how we take the equations and map them into the computer," says Dr. Mawhinney.

What the QCDSP and QCDOC teams have done is to configure the computational architecture specifically for these grid-based lattice QCD calculations. With both machines, the processors are organized in a regular grid, four-dimensional for the QCDSP and six-dimensional for the QCDOC. Each processor handles all of the variables for one local volume of space-time. This design optimizes nearest neighbor communication, the critical issue in QCD since the forces between neighboring points are local.

"QCDSP and QCDOC use a simple mesh *in the hardware.* So, we're able to achieve very low latency for the communication between these neighboring nodes. This allows us to get very good efficiency."

It also means that these machines are able to apply a large number of processors to a fixed size problem. "This is one of the most difficult kinds of scaling for a parallel computer," notes Mawhinney.

He stresses that the software components are as important as the hardware.

"The advantage of targeting QCD is not only that we're looking at one application in terms of the hardware, but also that we have a narrower range of things to handle with our software," Mawhinney says. As a result, the custom operating system is "lean and small," with little code, in order to minimize the potential for bugs. The project is now also able to benefit from its age: the physics software written for QCDSP in C++ will be migrated to QCDOC.

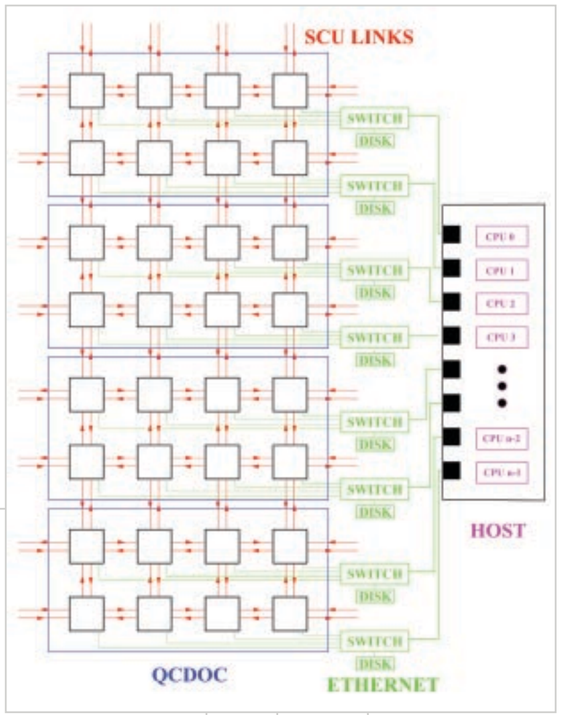

Gen-Next Hardware and Science

QCDOC's design is benefiting both from the experience gained with QCDSP and the widening, now international, lattice QCD computational collaboration. Two QCDOCs are presently being built, one for the RBRC at BNL and the second 10-Tflops machine for the United Kingdom QCD group based at the University of Edinburgh. In addition to DOE and RBRC support, this US-UK partnership has funneled money into the design project. The international collaboration has also added brainpower, primarily in the form of Dr. Peter Boyle, a post-doc from the University of Edinburgh who is based at Columbia, and Dr. Tilo Wettig, a Yale physics professor who has a joint appointment at the RBRC. With $5 million in funding from RIKEN (the Institute of Physical and Chemical Research in Japan) to complete the project in 2003, QCDOC will outpace its predecessor in three main ways: a faster processor, more local memory, and faster nearest neighbor communication.

A QCDOC daughterboard with two ASIGs and two DDRIMMs. The cubic blocks are connectors.

QCDOC's 10,000 nodes, each of which will be 20 times faster than those on the QCDSP, are custom-designed based on an IBM Power 440 ASIC template. The QCDOC ASIC — whose design has been led by the Columbia team in collaboration with Dr. Gara — will have 4 Mbytes of memory on the silicon wafer and a high-speed interface to this memory, designed by Dr. Gara and collaborators at IBM. The additional memory increases QCDOC's potential to be used for other more memory-intensive theoretical physics grid-based applications, a possibility that's now being considered in terms of broadening the machine's use.

The increased computing power QCDOC will provide is just what Dr. Mawhinney and others need to clear the QCD fog of non-linearity. To date, he says, researchers have used an approximation, called the quenched truncation, to navigate around a key part of this non-linearity.

"The real qualitative step that everyone is looking for as we go from a 1-Tflops machine to a 10-Tflops one is to remove this approximation so that we will actually be handling all the non-linearity," says Dr. Mawhinney. QCDOC will also allow the improved study of quarks when in a bound state as color-neutral hadrons.

However, the impact of these custom machines is extending beyond QCD science, sending ripples throughout the high performance computing community.

"Many other high performance users are extremely impressed with QCDSP," says BNL's McFadden. "But on the other hand, we get into this whole issue of whether it's better to buy commodity computers. There's a huge effort with QCDSP and now QCDOC in terms of developing the hardware, and then writing the operating system and maintaining this whole thing. That scares the hell out of a lot of people. But I think the success with QCDSP will have people giving it a second look, especially in terms of cost."

Indeed, the QCD-computing program has already had a notable computational spin-off: IBM's planned Blue Gene/L supercomputer. This 360-Tflops peak performance machine's main hardware architect is Dr. Gara of QCDSP and QCDOC fame. Bound for DOE's Lawrence Livermore National Laboratory in 2005, Blue Gene/L represents an extension of the grid-based application hardware structure pioneered in the QCD machines.

While others watch and assess the performance of these QCD machines, for Mawhinney the excitement over the design of a new machine is mutating into the expectation of what it will help tell us about the world of quantum chromodynamics.

"There is a tremendous advantage to being able to put theory and experiment to the test, and QCDOC will enable us, ever more accurately, to test our theories against the experiments. And where the discrepancies will show up, where we'll see the next hint of new physics, you don't really know. But it's the tool that lets you ask the questions."

SMALL IS BEAUTIFUL

>> ***For a supercomputer, QCDOC is a mighty mouse. "It has a very small footprint," says Ed McFadden, manager of Brookhaven National Laboratory's Scientific Computing Facility. With a 10-Tflops peak performance, QCDOC will occupy 60 square feet. By comparison, the 3-Tera-Operations ASCI Blue Mountain at DOE's Los Alamos National Laboratory covers 11,000 square feet.***

This diminutive size has broad repercussions — in terms of cost, performance and reliability — for the quantum chromodynamics science that QCDOC will perform.

QCDOC's designers chose to use processors of mid-range speed, thereby optimizing overall performance and drastically reducing the energy consumption and heat output per node. Each QCDOC node will use less than two watts, allowing the nodes to be tightly packed in this water-cooled machine, and making it 50% more energy efficient than Japan's Earth Simulator.

"The science is very directly impacted by the low electrical power and the small footprint of these machines," says Columbia University particle physicist Dr. Robert (Bob) Mawhinney, one of QCDOC's principal architects. "Because the electricity and the room to house the machine don't cost as much, we can actually get enough computing power to really push the science ahead without sacrificing every post-doc, graduate student and travel slot in our entire budget."

QCDOC's compact nature also means fewer components and cables, and thus less chance of mechanical breakdown. This is already the case with QCDOC's predecessor, QCDSP.

"One of the most impressive things about QCDSP, and we expect about QCDOC, is that is has a 98.37% uptime during the past three years," says McFadden. "It's extremely reliable."

By Peter Gwynne

Modeling Metals

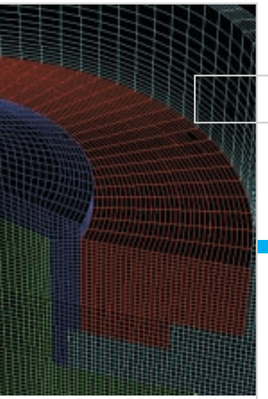

A METHOD OF SIMULATING the casting of liquid metals into solid shapes offers significant saving of time and money for engineers who process plutonium, uranium, and other substances critical for nuclear weaponry.

In 1994, nuclear engineer Doug Kothe requested internal research funds from the Los Alamos National Laboratory for a small scientific project. His objective was simple. He wanted to extend to three dimensions a two-dimensional computer model of free surface flow — the changing interface between liquids, such as liquid hydrogen and water, and gases, such as air, as they flowed past each other — that he had completed for the National Aeronautics and Space Administration. With financing from Los Alamos's Laboratory Directed Research and Development (LDRD) project, Kothe not only applied his computing code to an extra dimension; he also broadened it so that it could be used to model free surface flow between any gas and liquid.

Since then, the project has blossomed far beyond the single-scientist study that Kothe envisioned. Called Telluride, it helps to simulate the methods used for processing uranium, plutonium, and other metals critical to the country's stockpile of nuclear weapons. "The overall aim is to provide computer simulation of materials processing for the nuclear weapons complex," explains Telluride's present project leader Jim Sicilian. As an indication of Telluride's importance, responsibility for financing it has moved from LDRD to the National Nuclear Security Agency's Advanced Simulation and Computing (ASCI) program, a project that spends $700 million each year to give government scientists and engineers the technical capabilities they need to maintain a credible nuclear deterrent without underground nuclear testing.

Telluride has potential applications beyond simulating the processing of nuclear materials. "We collaborate with universities and other government labs to develop new models," says Sicilian. "We provide source code and exchange scientists with them, and we often have graduate students and postdoctoral fellows come to study with us. And some outside researchers have developed extensions of our work to apply to their own problems that we might ourselves pick up for the future. We have about a dozen collaborative agreements of that type."

Hopson did more than determine the vision for Kothe's untitled project. He decided that it would benefit from having a name. "As an avid skier, he named it after Telluride, Colorado," says Kothe. "We tend to name Los Alamos projects after southwestern geography." Following that example, the team later gave the name Truchas — Spanish for trout, and the name of a peak in the Sangre de Cristo mountain chain in New Mexico — to the first Telluride software that it released outside the Los Alamos National Laboratory. The idea of using simulation processing has about 20 years of history behind it. Sicilian himself devoted 15 years to working for a small company developing a commercial program for metal casting analysis. "The Telluride project continues to improve the breadth of physical phenomena modeled by the software and to validate it for application to manufacturing processes," Sicilian says. "No commercial products have all the features we need for simulation of our processes. That's the reason Telluride exists."

Called Telluride, it helps to simulate the methods used for processing uranium, plutonium, and other metals critical to the country's stockpile of nuclear weapons.

Opportunity for Application

The creation of ASCI in 1996 provided the opportunity to apply Kothe's original project to real and significant security needs. As it happened, Los Alamos had recently taken over the task of processing metals such as plutonium and uranium into parts for nuclear weapons, a task previously performed by a Department of Energy facility in Rocky Flats, Colorado. "My group leader, John Hopson, told me that my code could be used to model mold filling for our foundries," Kothe recalls. "Surprisingly, liquid metals behave much like water."

Simulation offers obvious advantages in the development of methods for processing difficult-to-handle metals such as plutonium and uranium. "It's a lot cheaper and faster to do it computationally rather than carry out the real thing, especially as the materials are expensive and dangerous," says Sicilian. "The payoff in the weapons projects is larger than it is in other areas because the experiments are exceedingly costly and time-consuming. To meet timelines by doing many experiments is not an acceptable approach."

Computer simulation offers metal processors a more fundamental advantage. "We can learn things that experimenters are unable to measure without disturbing the system," Sicilian continues. "For example, when people here do experiments, they measure the temperatures in the mold material, not in the metal itself. Recent simulations of ours have shown that there can be significant differences depending on the thermal history of the metal."

Two Problems

Almost as soon as he agreed to the concept of applying his research to real-life metal processing, Kothe realized that he faced two problems. "First," he says, "engineers at our two foundries — one for plutonium and the other for other metals — hadn't asked for models. And as a nuclear engineer by training and a computational fluid dynamicist by practice, I had no idea of the complexity of the physics involved."

The complexity stems in part from the fact that the free surface flow characteristics involved in pouring liquid metal into a mold cavity represent only one component of the processing that needs to be modeled. Even before the liquid metal is poured into the mold, the software

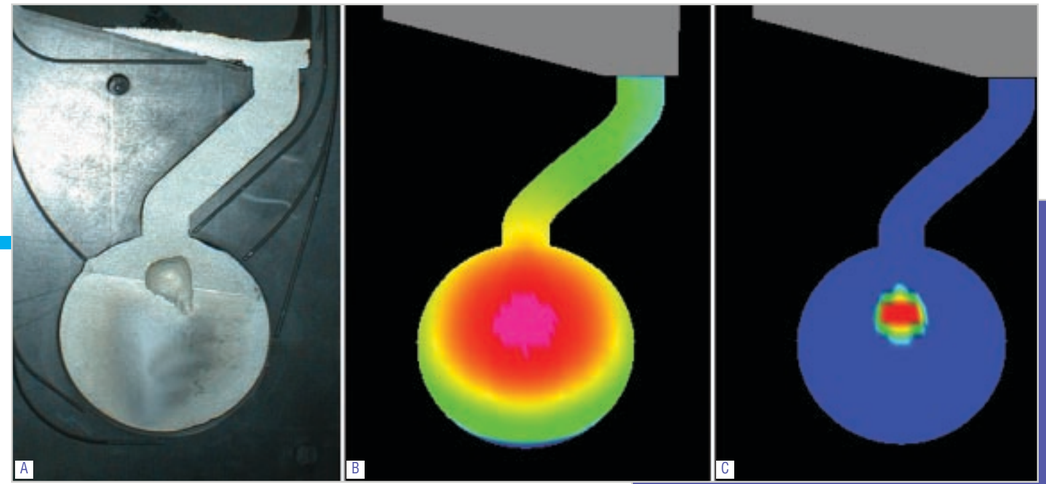

ALUMINUM BALL CASTING *A ball is a difficult shape to cast since the last place to solidify will tend to be toward the center of the ball and will typically exhibit shrinkage porosity as shown above. Aluminum is a challenging metal to cast as it has 7% shrinkage on solidification.*

Figure A is a photo of an aluminum ball casting showing porosity.
Figure B represents computed temperature distribution during solidification.
Figure C shows a computed volume fraction distribution; red is last to solidify.

Experiments and calculations: Deniece Korzekwa

COLLABORATORS

Douglas B. Kothe's home for most of the last 18 years has been the Los Alamos National Laboratory. He arrived at the lab in 1985 as a Purdue University graduate student in nuclear engineering and wrote his doctoral dissertation there. "It was a wonderful opportunity," he recalls. On receiving his PhD, Kothe took a staff position at the Lawrence Livermore National Laboratory. But a year later he returned to Los Alamos, where he has stayed ever since. The weapons laboratory also introduced Kothe to his specialty. "I got into computer simulation when I started my graduate research studying a hydrodynamic problem," he explains. "There was only so much you could do with pencil and paper. I was fortunate enough to work with computational scientists at Los Alamos." As a result of that experience, Kothe possessed the necessary intellectual tools to lead the Telluride project when it expanded from his original idea. He continued that role when he was promoted to the middle management position of group leader in October 2000, but handed over the leadership to Jim Sicilian two years later. Nevertheless, he remains an active member of the Telluride team, spending 20% of his time working on technical details of fluid flow algorithms and models.

James M. Sicilian returned to the lab that he had left more than two decades earlier to join the Telluride team in October 2001. Arriving at the Los Alamos National Laboratory in the mid-1970s with a newly minted PhD in nuclear engineering from Stanford University, he started work on nuclear safety issues that involved fluid flow. In 1980 he left for the private sector, joining local software development company Flow Science. Fifteen years later he set out on his own as a consultant. Then came the invitation to work on Telluride. "I was interested because of the challenges in the project and the opportunity to work with the computing resources at Los Alamos and to develop models in this environment," he says. He succeeded so well that he became project leader just a year later.

Further Reading:
For more information regarding the Telluride project, visit http://www.lanl.gov/telluride/workshop-2003/Presentations/index.shtml

Contacts:
Douglas B. Kothe
dbk@lanl.gov

James M. Sicilian
sicilian@lanl.gov

PRACTICUM COORDINATOR

Aric Hagberg
aric@lanl.gov

One of the beauties of Telluride's software is its ability to run on a broad range of computers.

must predict the distribution of temperature in the mold material induced by electromagnetic heating. Also, the alloy cools and solidifies in the mold. Further cooling of the solid metal to room temperature completes the casting process. Next, the alloy part undergoes further heat treatment to homogenize it. Finally, machining prepares the part for integration into the system for which it is designed.

That sequence of events means that Telluride must provide realistic models for solidification of the alloy, small and large segregation effects involving alloy components, free surface incompressible flow, heat transfer, complex three-dimensional geometric effects, microstructural nucleation and evolution, solid state transition effects, and the residual build-up of stress and strain that develops in response to those effects. So while the application of Telluride had the initial goal of helping foundry workers to understand and optimize their casting processes, it should also facilitate the design of new casting processes for alloy components that possess specific properties on the microstructural scale. "Our overall aim is to predict the properties of the metal after it's been cast," Sicilian says. "We're also involved in trying to develop new models of crystallization itself."

Plutonium and Uranium

Hopson and Kothe expected that Telluride would prove its value for plutonium processing first. After all, Kothe says, "plutonium is not your typical industrial metal." For example, its coefficient of linear expansion varies markedly and in stepwise fashion between 100° and 700° Celsius as it undergoes phase changes in its crystal structure in that temperature range. But as it turned out, he continues, "The people very interested in the modeling were those in the foundry that dealt with all metals except plutonium." That team, headed by Deniece Korzekwa, faced its own difficulties in casting uranium alloys. Uranium alloyed with niobium presents particular problems. "We have to look at how the concentration of niobium varies throughout a cast part from place to place," explains Sicilian.

Korzekwa's team started by comparing experimental reality with Telluride simulations of water poured into a plastic mold. Having satisfied themselves that the software accurately predicted reality, team members then used it to predict the cooling rate and expected grain size of real metallic castings. That work soon achieved two critical objectives: it verified the software and it validated Telluride's capability to model real events in metal and alloy processing by checking its predictions against experimental findings.

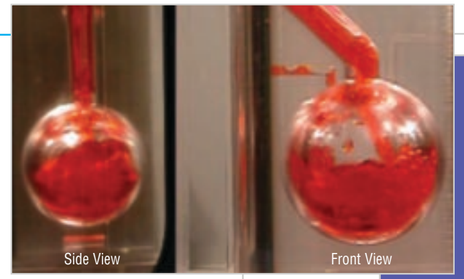

Mold filling experiment with water pouring into a lexan mold.

Experiments:
Deniece Korzekwa

Calculations:
Markus Bussmann
Deniece Korzekwa
Kin Lam

Once the foundry teams showed the benefits that Telluride modeling could provide, other teams at Los Alamos realized that the approach could help them. In 1999, the group responsible for welding transuranic metals and their alloys approached the Telluride team for assistance. "Welding had been something of a weak link," Kothe explains. More recently, a group involved in developing the foam parts used to cushion components in the complex structures of nuclear weapons signed up for the technology. "In this case," says Sicilian, "we use Telluride to model chemical reactions that occur during the curing process for the foam." The process resembles the transformation of bread from a wet material to a dry material in the baking oven.

Computers and Collaborations

One of the beauties of Telluride's software is its ability to run on a broad range of computers. "We can execute the program on a very wide variety of platforms, from personal computers to the latest ASCI computers here at Los Alamos," Sicilian points out. Ultimately, however, the project will benefit from the computing power available at Los Alamos. "Thus far we've been able to do development and testing work on fairly easily available machines," he continues. "But once we get into really serious calculations, we'll move up by two orders of magnitude to the big ASCI computers that are unique in the world."

Los Alamos offers the Telluride team another benefit: the variety and quality of its scientists. "The lab was enabling for us because it has so many people with diverse backgrounds," says Kothe. "We've had physicists, applied mathematicians, engineers, computer scientists, and others collaborate with us. We've had regular intimate contact with the experimentalists. A large research institution is the only place at which something like this can be done. That's what makes a national lab attractive."

Telluride's developers also seek collaboration beyond Los Alamos. In January 2003, they convened a workshop to share ideas about fluid mechanics and heat transfer and to discuss what Telluride can offer researchers interested in those areas. Over two days of working sessions, the team highlighted the kinds of phenomena that Telluride can model, the algorithms used in its software, and the problems that it is currently used to solve.

For the moment, the workshops and the technology will remain restricted to scientists in the government and academic sectors. "We have no desire to commercialize the software," says Kothe. "It's a research code; it's not really user-friendly." Nevertheless, the Telluride approach promises advantages for more general computer modeling of metal processing. "As the physical models and numerical methods are transferred into more commercial lines," says Sicilian, "there are international competitive reasons why it would enhance our security, by making processing more economical."

Over the medium term, Telluride-style modeling holds promise for basic research in environmental science. Scientists expect to be able to use the techniques to gain better understanding of such effects as the impact of meteors on Earth, the spread of wildfires, and the growth of tsunamis.

PUCK CASTINGS *Local solidification time (time in mushy zone) and the cooling rate through the epsilon phase have been correlated to grain size and coring. These simulations show variations between pucks as well as within each puck. This information can be compared to stockpile material and be used in subsequent simulation codes.*

Calculations: Deniece Korzekwa

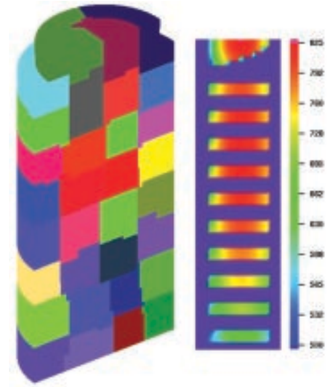

PHYSICS FOR THE FUTURE

>> ***Despite its impressive record of achievements since it started in 1994, Telluride remains a work in progress. So far, the software can model such physics phenomena as viscous incompressible Navier-Stokes fluid flow; heat transfer by conduction, convection, and radiation; phase changes from liquid to solid and between solid crystal structures in pure metals and alloys; flow of free surfaces with surface tension; the evolution of microstructures; and fluid turbulence. Goals for the near future include developing the ability to model compressible Navier-Stokes fluid flow; the evolution of stress in solid metals and alloys; nucleation; and the response of metals and alloys to small deformations.***

To ensure smooth progress in its development, the Telluride team has set out a series of specific objectives for coming months and years. Goals for the end of 2003, for example, include validation by an independent team of Telluride's Truchas software for analysis of a casting process, certification by an application expert of a Truchas analysis of a complete, integrated simulation of a casting process, and certification of a Truchas analysis of an integrated foam-curing process by the Energy Department's Kansas City plant. By six months later, the Telluride group anticipates that a Department of Energy manufacturing engineer will have used Truchas to improve manufacturing. It also expects to have five refereed scientific papers published or scheduled for publication. By the end of September 2004 the plan envisions more ambitious use of the software and hopes that the project will receive an achievement award from the Energy Department and the Los Alamos National Laboratory. And ultimately the group projects that Telluride technology will find application in modeling processes beyond casting, welding, and foam curing.

By William J. Cannon

STOMPing Ground

IF YOU LIKE A GOOD MISNOMER, you need look no further than the Scenic Highway site outside Baton Rouge, Louisiana. There, spread over 17 acres of graded dirt, are dozens of recovery wells. Down the road a mile and a half, at a place called Brooklawn, are more wells on 60 acres — 165 wells at the two sites, about a mile from the Mississippi River. Take in the view, but don't drink the water.

Scenic and Brooklawn comprise the EPA's Petro Processors Superfund Site where, in the 1960s and '70s, for 13 and 11 years respectively, a "disposal company used to pour organic waste from petroleum companies into large ponds and trenches," says Mark White, senior research engineer at the Department of Energy's Pacific Northwest National Laboratory.

White is code author and custodian of STOMP, which stands for "subsurface transport over multiple phases." STOMP parcels up a piece of ground into three-dimensional cubes, then simulates underground flow and transport of very bad things, environmental nightmares both happening and waiting to happen. White and his STOMP team at PNNL's Hydrology Group try to get on the case before the nightmare has happened; their intervention can lead to action that can prevent a given pollutant from reaching groundwater.

At Scenic and Brooklawn, the nightmare was happening. A site-restoration company had cleaned up the surface and, as common sense might dictate, drilled wells from which to pump out the nasty stuff that had seeped downward. White's word for that stuff sounds like "Dean Apples "— DNAPLs, denser-than-water nonaqueous (oily) phase liquids (NAPL), among the many forms of pollutants covered in the transport equations of STOMP.

In this context, the important thing about a DNAPL is that it sinks in water. Mart Oostrom, a senior research engineer in White's group, ran a STOMP simulation for Scenic only to discover that the army of pumps, which had cost a small fortune to install and to operate and were intended to suspend the organic pollutants, actually made things worse. The model, also applied for Brooklawn, showed "that as millions of gallons of water flowed upward, the water table dropped," White says. "The DNAPLs migrated farther down into the system and toward the drinking-water aquifer. They shut down the perimeter wells."

As a result, recovered groundwater continues to be treated to remove hazardous liquids and other contaminants. Unlike its namesake, the Scenic simulation and others turn out to be rather elegant. STOMP solves a series of equations that describe the physical properties in the hydrology system under investigation. The program can be tailored to track the migration of materials through water, brine, ice, oily liquid pollutants, through just about any fluids that might fill underground pores, thanks to a so-called variable source code that enables anyone using it to dial in the desired governing equations to be solved: water mass, air mass, dissolved-oil mass, oil mass, salt mass and heat. Equations solve problems specific to transport of chemicals in solution, radioactive decay and chemical reactions. At the end of this process is a time-lapse picture of underground plumes of pollutants expanding and receding, depending on how the invasive chemicals interact with the soils and elements in the ground, under any conditions the modeler can conceive of. "We have put the simulator through a rigorous verification procedure against analytical solutions, laboratory experiments and field demonstrations," White says. "I want to emphasize that the nature of this work is collaborative. Mart and I ponder why the numerical simulations don't agree with the experimental observations. Numerical simulation solves a collection of mathematical equations in an attempt to describe physical processes.

"Typically these equations do not completely describe the modeled system, as assumptions were taken or processes were ignored to generate the equation sets. This is where the art of numerical modeling comes in — deciding how to mathematically describe observed multi-fluid flow processes, and deciding which equations to include in the collection. If one

STOMP parcels up a piece of ground into three-dimensional cubes, then simulates underground flow and transport of very bad things, environmental nightmares both happening and waiting to happen.

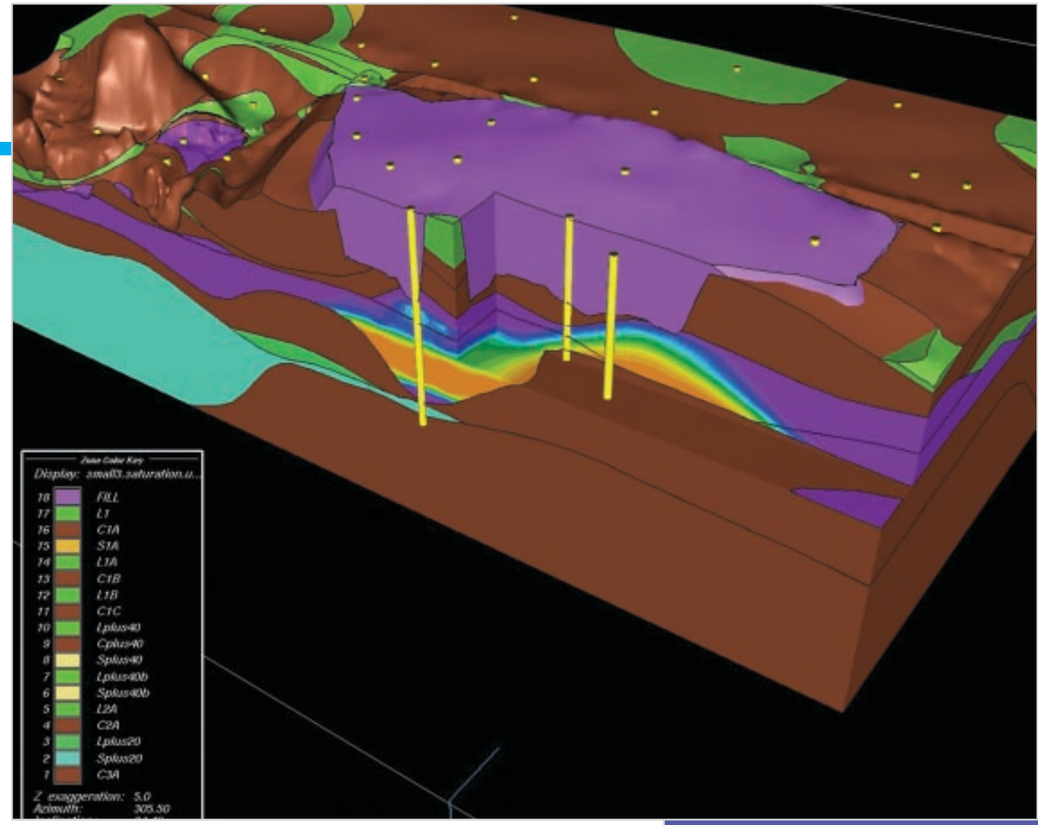

Simulated DNAPL saturations at year 2000 superimposed on alternative geology for the Scenic Site Operable Unit near Baton Rouge, LA

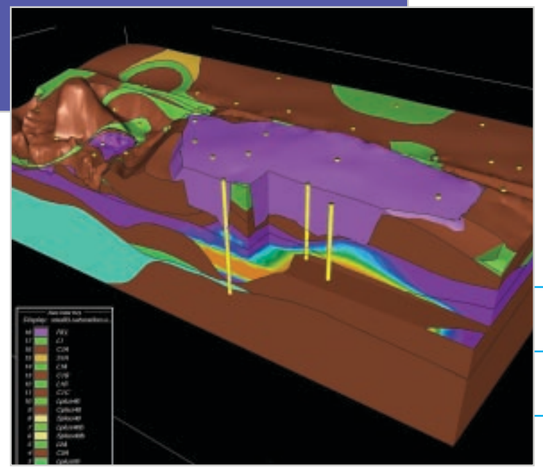

Simulated DNAPL saturations at year 2100 superimposed on alternative geology for the Scenic Site Operable Unit near Baton Rouge, LA.

includes too many equations and processes the execution speed suffers; conversely, if one excludes a critical equation or process, simulation results don't accurately represent the experimental observations. This is where the art comes in, where we can dream up some beautiful stuff that all fits together."

A given simulation can cover days, months and, in some cases, decades. At the Hanford Site, a former weapons-production area in PNNL's eastern Washington-state back yard, Oostrom and White turned STOMP toward predicting the movement of a toxic underground carbon tetrachloride plume. Between 1955 and 1973, Hanford workers dumped as much as 580,000 liters of liquid carbon tetrachloride, an organic cleaning solvent. The plume is so massive and unapproachable that a simulation is the only way to assay its threat to the groundwater layer below and to test likely methods for halting its movement and, ultimately, cleaning it up. It appears to have spread to a depth of 60 meters into the vadose zone, the area above the groundwater 65 meters below the surface.

"Preliminary results from ongoing STOMP simulations indicate that vapor transport of the carbon tetrachloride was significantly more important than its migration as a liquid," White says.

More recently, STOMP has been put to the task of recreating the history of buckling million-gallon, single-shelled storage tanks from the seams of which has leaked self-boiling radioactive waste, a byproduct of the plutonium production at Hanford that continued throughout the Cold War until 1988. The sludge in the 75-foot-diameter vessels at what is called the SX Tank Farm has heated to up to 140 degrees centigrade, eating away at the tanks' steel linings.

The filling and draining of the tanks with radioactive waste from 1955 to 1975 set up a cycle of heating and cooling that has altered the hydrology and affected the movement of leaking sludge. STOMP can provide information that would inform decisions on methods of containment or even on whether to attempt to actually transport material from the tanks.

"We're really interested in knowing two things," White says. "First, how the dynamic thermal and hydrologic conditions in the subsurface outside these tanks affected the migration of leaked contaminants, and second, how the current thermal and hydrologic conditions affect the future migration of these contaminants and future leaked contaminants from tank closure operations, such as sluicing."

STOMP simulations predict temperature and moisture "saturation profiles in the soil surrounding the tanks over time that included the surface recharge (moisture replenished by rain) and sludge leaking into the soil," White says. The simulation shows that when temperatures in the porous soil beneath the tanks reach more than 100 degrees centigrade, "it creates a steam plume around the shell that dries out the soil," White says. "With no liquid water, the capillary pressure wicks water back toward the tanks," where it evaporates.

"The water that evaporates migrates away from the tank as water vapor by advection and diffusion through the gas-filled pore spaces," White explains. "This water vapor then condenses where the underground temperature falls below 100 centigrade, forming liquid water. It's this liquid water that wicks back toward the tanks to be re-evaporated." This cycle of evaporation, migration, condensation and wicking is called "countercurrent heat pipe flow."

The heat pipe flow during the 1960s and '70s concentrated salts and other nonvolatile contaminants near the heated-tank surface "because those contaminants are left behind during the evaporation process, but the water wicked toward the tank surfaces can contain dissolved quantities of the contaminants." At first, the concentration of the sludge remained nearer the surface inside this hot zone beneath the leaking tanks. Through the years, though, as the temperature in the tanks has waned, so has the heat in the soil. The hot zone beneath the shells "collapses," White says, "and moisture (replenished by rain) returns." The ambient soil moisture has diluted the sludge that saturates deeper toward the groundwater layer, while the sludge's dissolved constituents migrate toward the groundwater at various rates.

Over the past decade, STOMP has become a well-traveled simulator, moving far beyond PNNL. Besides the Petro Processors Site, STOMP made some early predictions of the potential effects of using Yucca Mountain, Nevada, as the nation's high-level nuclear waste repository and, lately, STOMP has been proving its value and its versatility on an international stage. Or, more accurately, beneath one. Like many nations, Norway is pumping the industrially produced greenhouse gas carbon dioxide into a salt-water aquifer to sequester it, dissolved in the brine, from the atmosphere. Norway and others considering a similar plan need to know whether the gas will stay put.

To recreate those conditions in the laboratory, White and Oostrom turned to PNNL colleague Pete McGrail, a staff engineer in the Applied Geology and Geochemistry Group. McGrail fashioned a pressure cell a half-meter long and a third of a meter high, filled with well-graded sand and brine, then injected "supercritical" carbon dioxide, a between-phase form of the substance that "is not quite gas, not quite liquid and lighter than brine," White explains. At high pressure, carbon dioxide dissolves in the brine and increases the density of the fluid.

A simulation showed that after three hours the injected carbon dioxide "forms a distinct gaseous phase from the brine," White says. "The CO_2 rises along the left-hand wall and then starts to migrate across the top of the cell. During this process CO_2 dissolves into the aqueous brine, yielding a CO_2-saturated brine that is denser than the CO_2-unsaturated brine."

This dense CO_2-saturated brine overlying less dense and unsaturated brine is unstable and results in "fingering" of the CO_2-rich brine into unsaturated brine. The fingering mechanism promotes the dissolution of carbon dioxide. "We try to quantify through these simulations the dissolution-enhancement that occurs through fingering, to allow for better injection strategies, designs, and protocols," White says.

Gas, toxic oils, viscous radioactive sludge — as long as people consider the earth beneath their feet as a place to put what they don't want sharing the surface with them, STOMP will be there to tell them whether or not it's a good idea.

STOMP

>> ***STOMP author Mark White and collaborator Mart Oostrom are part of the Hydrology Group in the Environmental Technology Division of Pacific Northwest National Laboratory. White says that while each national lab has a program for modeling the migration of underground pollutants, there is no particular mandate to give those capabilities macho-sounding acronyms like STOMP. Still, Lawrence Livermore has its NUFT, and, not to be outdone, Lawrence Berkeley has TOUGH.***

By any name and any objective measurement, STOMP would have to be ranked among the top of those contenders as one of the most widely tested systems. The technology — Battelle, the nonprofit corporation that manages PNNL and other national labs, holds the copyright to the code — has transferred well. Besides the national labs and government agencies, STOMP has been enlisted by private companies and universities alike and has been run on a great variety of hardware, ranging from personal workstations to Beowulf clusters and teraflop-performance parallel supercomputers.

White, a mechanical engineer by training, continues to adapt the STOMP code to new applications and to ramp up its performance. STOMP's trials and successes are well documented in the professional literature, as well as in more accessible places: White's numerical framework is spelled out in a theory guide available on the Web, along with an application guide and a user's guide. Oostrom, whose background is in soil physics, is co-author of the guides and keeps the code honest, checking it against experimental results.

COLLABORATORS

Mark White is the principal author of STOMP and is a Senior Research Engineer in the Hydrology Group / Environmental Technology Division at Pacific Northwest National Laboratory (PNNL). He has been at PNNL since 1986 and his current research interests include developing mathematical descriptions of experimentally observed multiple-phase subsurface flow and transport processes and applying these descriptions to the effective remediation of contaminated subsurface environments through field-scale technology design, development, or improvement. He received his BS in biophysics from Pennsylvania State University, and his MS and PhD in mechanical engineering from Colorado State University.

Mart Oostrom is a Senior Research Engineer in the Hydrology Group/Environmental Technology Division at Pacific Northwest National Laboratory (PNNL), where he is responsible for projects related to multifluid flow and contaminant transport model development. Current projects include the development of a theory for relative permeability-saturation-pressure relations in mixed-wet porous media, the investigation of the surfactant-enhanced aquifer remediation technique for clean-up of aquifers contaminated with chlorinated solvents, and the development of algorithms for flow and transport. A native of The Netherlands, Dr. Oostrom received both his BSc (soil science and geology) and MSc (soil physics) from the Wageningen Agricultural University in The Netherlands. He received his PhD in soil physics from Auburn University.

Further Reading:
For more information on STOMP, White, Oostrom and their team, please see their Website at http://www.pnl.gov/etd/stomp.

Contact:
Mark White
mark.white@pnl.gov

Mart Oostrom
mart.oostrom@pnl.gov

PRACTICUM COORDINATOR

Janet Jones-Oliveira
jjo@pnl.gov

By Karyn Hede

Modeling Blood Flow

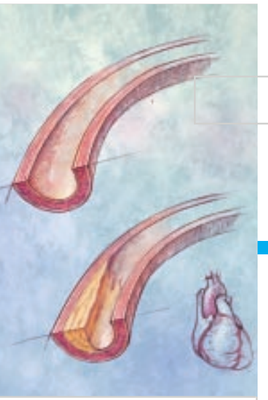

LITTLE DID PAUL FISCHER KNOW that a convivial hallway chat outside his office at Argonne National Laboratory would lead him on a "fantastic voyage" into the turbulent world of an artery on the verge of causing a stroke. Like the characters in the 1966 sci-fi film classic, a team of scientists has ventured into the bloodstream to see what happens inside a blocked artery. But this voyage is a virtual trip, combining computational prowess with advanced medical imaging and fluid dynamics — all with an eye to preventing a devastating stroke.

It was in the summer of 1998 that Fischer struck up what he calls a 'serendipitous conversation' with visiting scientist Frank Loth, a specialist in fluid mechanics from the University of Illinois at Chicago. Loth was working on computer simulation of blood flow inside arteries. The two began to chat about the challenges of modeling blood flow, particularly when blood rounds a bend in an artery where mechanical forces can cause tiny eddies and whirlpools to form.

"Frank mentioned that he had a couple of cases where the blood flow transitioned to turbulence and I said 'We can do that,'" recalls Fischer. It turned out that Fischer specializes in creating computer programs that model turbulent flow in complex domains, just what Loth needed to understand the effects of turbulent behavior on the formation of potentially deadly atherosclerotic plaque.

Trying to understand the physical forces that govern blood flow is not new; scientists have been publishing treatises on the subject for decades. Even the equations that Fischer uses to describe fluid flow have been in use since the nineteenth century. But only in the past few years has a confluence of technologies made it possible to create a realistic and useful computer model of turbulent blood flow.

"We want to understand what is particular about one person that causes atherosclerotic disease versus other people who don't have disease," says Fischer. "It turns out that geometry is very important in determining wall shear stress, which is the action of blood on vessel wall. Now we want to learn, is it the shape of the vessel? Is it the flow conditions? What are the important factors in formation and stability of plaque?"

Loth and his longtime collaborator at the University of Chicago, Hisham Bassiouny, a vascular surgeon, are specifically interested in the carotid artery, the main artery that supplies blood to the brain. It turns out that the carotid is a prime location for fatty plaque to build up, narrowing the vessel diameter and forcing the heart to work harder to get blood into the brain. Over time, this plaque build-up can become unstable. When small pieces break off, they travel to the brain and can become lodged in smaller vessels, blocking blood flow and causing brain cells to be starved for oxygen so that they die within minutes. The phenomenon is so common that since you started reading this article someone in the U.S. has had a stroke.

The Chicago-based scientists began the collaboration with Fischer with one question in mind: Who is most at risk for a stroke? And particularly, what are the conditions that are most likely to lead to stroke? Being able to predict with some certainty who is at greater risk would help tremendously in deciding who should undergo carotid endarterectomy, an operation in which surgeons split open the carotid artery to remove potentially deadly atherosclerotic plaque. In addition, says Fischer, the simulations could help determine if grafted arteries and veins are prone to failure due to mechanical forces.

But only in the past few years has a confluence of technologies made it possible to create a realistic and useful computer model of turbulent blood flow.

"Frank Loth is really the driver of the science behind this project," says Fischer. "As a biomechanical engineer, he understands the biology and he understands the fluid mechanics. He really holds it all together."

But to create realistic models they needed some way to model turbulent conditions inside the carotid. And it isn't enough to do the calculations on an idealized approximation of an artery, because each person's carotid artery is a different shape, changing dramatically the conditions inside the vessel walls.

To develop a realistic simulation, the team first had to develop a way to convert data collected from medical imaging devices such as a CT (computed tomography) or MRI (Magnetic Resonance Imaging) scans into a three-dimensional representation of the vessel wall. Seung Lee, who at the time was a student at the University of Illinois at Chicago, spearheaded that part of the project. Lee, now a graduate student at MIT and a DOE CSGF fellow, developed software to convert two-dimensional slice-based data into a three-dimensional mesh described by grid points and organized computationally into "bricks." Within each brick the physical forces of velocity and pressure are described mathematically in terms of a polynomial function. And all this must be calculated before the simulation even begins.

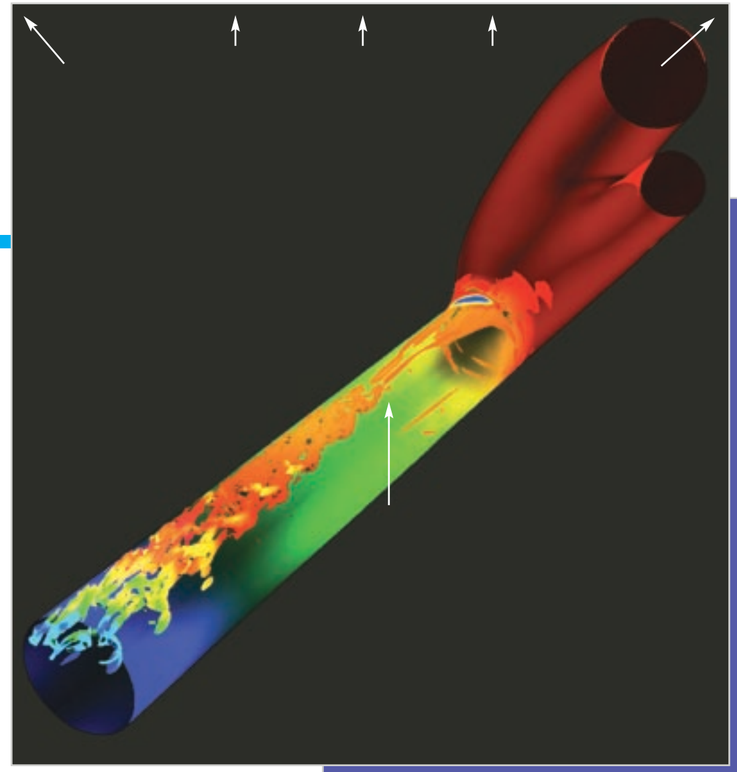

The transition process: hairpin vortices initiate and grow downstream of an arteriovenous graft due to the high flow rates in the graft.

At this point, Fischer and his Argonne colleague Henry Tufo entered the picture, providing the computational means to describe the turbulent flow.

"All the vortices associated with turbulence dramatically increase the number of grid points required," says Fischer. The reason is that in a smooth laminar flow the force is mostly uniform and in one direction, as you see in water running from a faucet that is turned on slightly.

PAUL FISCHER

Paul Fischer is a scientist in the Mathematics and Computer Science Division at Argonne National Laboratory (ANL) where he has worked since 1998. In addition to his work at ANL, he is also a senior fellow at the Computation Institute at the University of Chicago. Dr. Fischer holds a BS in mechanical engineering from Cornell University, an MS in mechanical engineering from Stanford University and a PhD in mechanical engineering from MIT. He took part in the development of a parallel spectral element code, Nek5000, which is currently available and supported at over a dozen institutions worldwide and was honored with the 1999 Gordon Bell Prize for sustaining 380 GFLOPS with 4096 processors of Intel's ASCI Red.

Further Reading

F. Loth, N. Arslan, P. F. Fischer, C. D. Bertram, S. E. Lee, T. J. Royston, R. H. Song, W. E. Shaalan, and H. S. Bassiouny, *Transitional Flow at the Venous Anastomosis of an Arteriovenous Graft: Potential Relationship with Activation of the ERK1/2 Mechanotransduction Pathway*, J. Biomechanical Eng. (in press).

P.F. Fischer, G.W. Kruse, and F. Loth *Spectral element methods for transitional flows in complex geometries, J. of Sci. Comput.* 17 1, pp. 81-98 (2002)

P.F. Fischer and J.S. Mullen, *Filter-Based Stabilization of Spectral Element Methods* Comptes Rendus de l'Académie des sciences Paris, t. 332, Série I - Analyse numérique, p. 265-270 (2001).

Contact:
Paul Fischer
fischer@mcs.anl.gov

PRACTICUM COORDINATOR

William Gropp
gropp@mcs.anl.gov

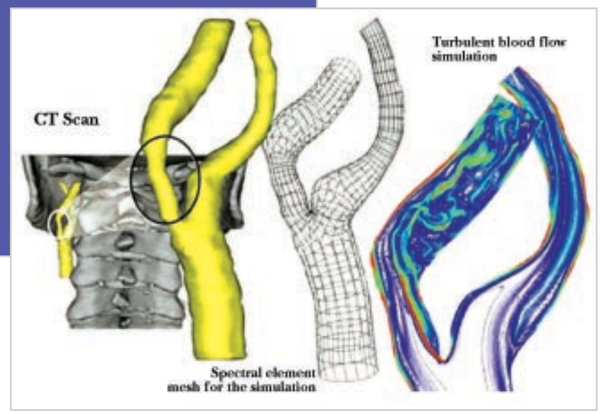

Patient specific simulation scenario: CT scan image of a diseased carotid artery is acquired (left), converted into a computational mesh (center), and turbulent blood flow is simulated to predict hemodynamic stresses downstream of the stenosis (right).

Hairpin vortices in an arteriovenous graft model.

If you suddenly turn the faucet on full force, you get jets and bubbles. "Imagine trying to describe that computationally and you get some idea what we are trying to do," he says.

"And on the time scale," he adds, "if we want to model the force exerted during one heart beat we suddenly go from needing a minimum of 30 time points per heartbeat to thousands of time points per heartbeat."

The behavior of blood flowing inside a vessel is governed by the Navier-Stokes equations. It is relatively easy to model smooth, laminar blood flow, and in fact there are probably at least 30 laboratories worldwide that are doing so, says Fischer. But when you introduce turbulence into the equation, the level of computational complexity increases tremendously.

To make the calculation possible, the team employs a computational tool called the spectral element method, which allows them to divide the calculation into discrete computational blocks that are ideal for numerically solving the Navier-Stokes equations on a massively parallel computer system. The spectral element method, explains Fischer, also excels at calculating flows that are just near the onset of turbulence, as is the case for blood flow.

Once the team had put all the elements in place it was time to put their algorithm to the test, to enter the world of a blocked artery and see what is happening inside. With Bassiouny providing real-world image data of a patient whose artery was severely blocked, the team ran their algorithm on Pittsburgh Supercomputing Center's Terascale System, which has a storage capacity 100,000 times that of most desktop PCs and a thousand times the computational capability. The goal: to arrive at a realistic simulation within 24 hours, quickly enough to generate data in a time frame where it can eventually be used to help direct clinical decisions.

The first simulation, done in mid-2000, passed with flying colors. Loth's group continues to conduct controlled experiments to compare with the simulations to provide quality control, and the agreement between the numerical and experimental models has been impressive.

From the computational standpoint, the algorithm can now calculate a single cardiac cycle in 10-20 hours of CPU time on 256 processors. In fact, the performance of the algorithm was so impressive that it garnered the coveted Gordon Bell Prize, a major computing award, for Fischer and Tufo.

The team's spectral element code, Nek5000, employs a parallel iterative procedure known as domain decomposition. "What we've done," says Fischer, "is to break the calculation into pieces so that each piece can be solved independently and recombined to give an approximate solution. This procedure is repeated until the solution converges for the given time step. The simulation is then advanced in time and the iteration is repeated."

The initial calculations were, in fact, so promising that the team has stepped up efforts to push the process even further.

In order to have efficient iterative methods for systems involving millions of gridpoints, you need to have multilevel algorithms that can quickly improve the solution at the fine and coarse scales. To scale this approach to hundreds of processors, it is important to have a fast coarse-grid solver, since that's typically a communication intensive process. "One of the things we've done is develop a very fast parallel coarse-grid solver, but we are trying to further improve the speed of the overall process by introducing additional intermediate grids, which will result in a multigrid algorithm," says Fischer.

For vascular surgeons like Bassiouny, the real payoff will come when it is possible to gather raw data from medical imaging devices and feed the information directly into the modeling software. At that point, it would be possible to conduct a series of clinical trials in which the computer simulations are correlated to patient outcomes and it becomes feasible to predict who is most at risk for stroke.

"Our experience," explains Bassiouny, "is that many people can live for years with fairly extensive plaque and never have a stroke, while someone else may have less plaque, but if it's unstable the risk is much greater."

Fischer emphasizes that in addition to its medical applications, the algorithm can be used for studying heat diffusion and other types of fluid flow, whether liquid or gas.

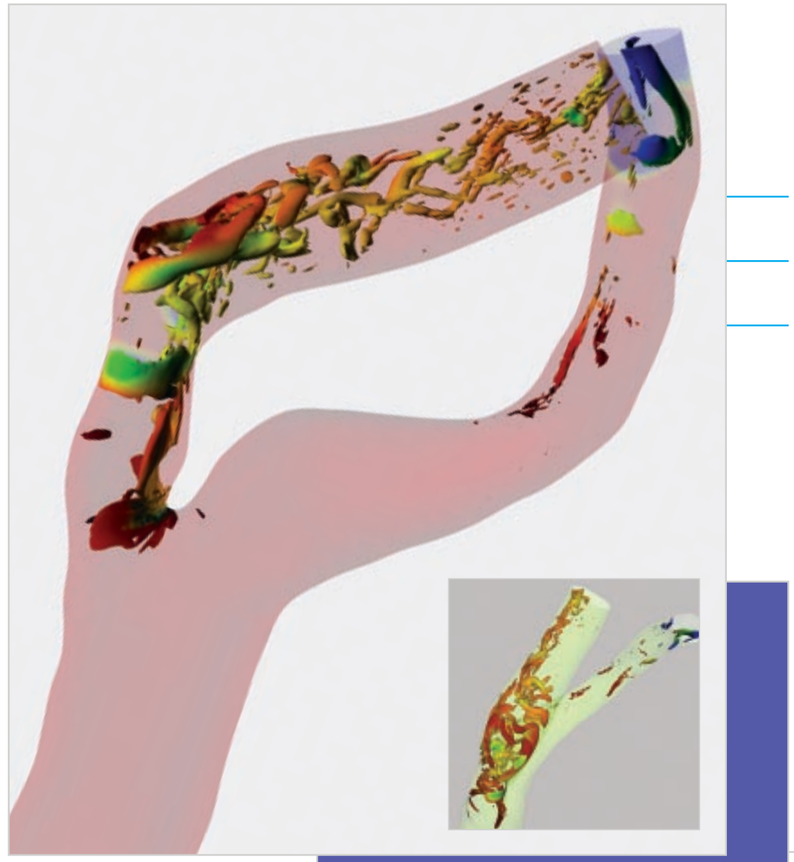

Vortices identified by the spectral element simulation indicate the nature and scale of eddies in the flow just prior to onset of full turbulence, which occurs at the peak of the cardiac cycle.

"We have a dozen groups around the world using the code for different applications," says Fischer. "The Nek5000 code provides researchers with direct access to advanced parallel algorithms."

In the words of Dr. Peter Duval, the surgeon in *Fantastic Voyage*, "We stand in the middle of infinity between outer and inner space, and there's no limit to either."

Tom Epperly

IN THE BIBLICAL STORY OF BABEL, people become disorganized and scatter across the land when they suddenly find everyone around them is speaking a foreign language. They surely could have used a guy like Tom Epperly.

As a project leader of the components group at Lawrence Livermore National Laboratory, Epperly creates code that allows computer languages to talk to one another. The group is part of a larger multi-institutional partnership called the Common Component Architecture Forum that is developing component technology for scientific computing.

"The main thing we are after here is code reuse," says Epperly. "We are providing technology that enables people to reuse other people's work, and in the process we are saving money and making sure the 'best in class' solution is available to other people." The goal of component technology, says Epperly, is to create a software library that allows programmers to pick and choose the applications best suited for their problem and then allows all the components to work together seamlessly. Component technology, says Epperly, is what makes it possible for users of Microsoft products to cut and paste an Excel spreadsheet into a Word document and then put the whole thing into a Frontpage Web site.

Of course, the programmers at Microsoft had interoperability in mind when they wrote their code. Epperly's job is much trickier, since he is working with languages created long before anyone envisioned today's high performance computing environment.

"Fortran is still a major language in scientific computing," says Epperly. "In contrast, the mainstream computer science world is dominated by languages such as Perl, Java, C++ and Visual Basic. My mission is to upgrade the state-of-the-art for software development in scientific computing."

Today, he says, most programmers are developing code in an object-oriented programming style and his group adapts code to work with these programming styles.

Epperly's main customers are programmers at the national labs and in academia who are developing high performance programs that run on large parallel computers. For example, his group developed a product they call "Babel" that, appropriately, allows programmers to mix the languages C, C++, Fortran77, Fortran90, Java and Python in a single application. Scientists at Pacific Northwest National Laboratory (PNNL) are using Babel to allow various computational chemistry components to talk to each other during a molecular simulation.

Epperly started out as a chemical engineer, obtaining his PhD in 1995 at the University of Wisconsin at Madison. His postdoctoral work at Imperial College in London focused on design under uncertainty: how to build a chemical plant when demand for the product is unknown. His designs optimized return on investment under conditions of uncertainty.

After a short stint working for a company that supplies software to the chemical industry, he decided to move on to new challenges.

"I wanted to be on the cutting edge of technology and the national lab presented a good balance of basic research and application," he says. "At the national lab there is an emphasis on both, and both are considered valuable. The work I do here has an impact on real life decisions and situations."

He says having a background in physical systems has helped him in his current work and has often made him an interpreter between the scientists and the computer programmers.

"It is important to understand the application area, which is modeling physical systems, in doing this work," he says. "There's often a gap between computer science types and the people who know the underlying physics. Having a joint background in engineering helps me bridge that gap."

"My mission is to upgrade the state-of-the-art for software development in scientific computing."

Marc Serre

IN THE AFTERMATH of the World Trade Center destruction many scientists mobilized to aid in understanding what happened. Engineers analyzed twisted girders. Forensic scientists helped identify victims. But as much of the initial frenzy of activity slows down and the event itself recedes into history, Marc Serre's work is just beginning.

Serre, an emerging expert in a specialty known as spatiotemporal geostatistics, is modeling the movement of toxic gases that emanated from the World Trade Center disaster site. Using a method he developed with former UNC mentor George Christakos, an expert in modeling of environmental systems, Serre is trying to reconstruct the path of airborne pollution released by the collapse of the burning towers.

"We have the best method to integrate all available knowledge," says Serre, assistant professor in the department of environmental sciences and engineering at the University of North Carolina at Chapel Hill. "We have satellite data, weather information and expert knowledge. When you are able to combine all that, you have a better assessment of what people were exposed to after the World Trade Center collapse."

Unlike previously used methods that could only model movement of toxic substances over space or over time, Serre's method does both simultaneously, and is able to incorporate flow transport equations and so-called "fuzzy knowledge" or "expert knowledge."

"This is a completely novel idea of how to incorporate that knowledge," says Serre. "The framework we have developed is more rigorous and more flexible than the classical approach and we are able to incorporate all the knowledge bases to come up with a better assessment of exposure to toxic agents."

"Professor Christakos had the idea for several years of using the Bayesian maximum entropy method to estimate the distribution across space and time of environmental agents," he says. "It came at a very interesting point, where my background was just ideal to implement it. There was a kind of a spark. I provided the numerical implementation of his idea, and it all came together."

Since the two developed this new way of tracking complex movement of substances during Serre's graduate student days, the two have made the resulting program available to anyone on their web site. Called BMElib (available for download at www.unc.edu/depts/case/BMELIB/), the program can be used for a variety of applications. Serre himself has tracked the movement of things as diverse as arsenic in Bangladesh and sexually transmitted diseases in inner cities. He says his program is now in use worldwide to study things such as groundwater purity in Brazil and the movement of bats in Spanish caves.

His current work focuses on improvement of the conceptual framework and the code performing the numerical calculation, which used to run mostly on high-end computers. However, the latest version can run on a PC, making it available to many more researchers. Serre has co-authored a book showing several applications of the software and is teaching it to graduate students.

"What we have learned through our [DOE CSGF] fellowship is to combine mathematical science with smart high-level numerical implementation," he says. "That's exactly what we have done here."

Serre is trying to reconstruct the path of airborne pollution released by the collapse of the burning towers.

Scott Zoldi

COMPUTATIONAL SCIENCE might not seem like the kind of training that would lead to a career as a private eye, but for Scott Zoldi a penchant for spotting changes in behavior has led to a career as a high-tech computational sleuth.

Zoldi's training in computational and theoretical physics turned out to be just what they were looking for at Fair Isaac Corporation, the largest U.S. company dedicated to detecting and reducing consumer fraud. As an analytic science manager, Zoldi leads a team that creates computer programs to analyze an enormous quantity of data and ferret out customers who are committing telecommunications fraud. This includes using call records to detect cloned phones — actual physical copies of people's cell phones that contain all information needed to make thousands of dollars of fraudulent calls in a very short time frame.

> "We apply our computational physics training to create statistical models that mitigate fraud risk."

"There is so much data out there that companies don't understand what to do with it," says Zoldi. "What you need is someone with a strong analytic and computational background to go and analyze terabytes and terabytes of information and reduce it down to predictable variables. What we do is squeeze out the good stuff and throw away a whole bunch of pulp."

The goal is to help companies identify and defuse fraudulent behavior as quickly as possible.

"Since we operate in a business environment, we have to focus very much on return on investment," says Zoldi. "We apply our computational physics training to create statistical models that mitigate fraud risk. We are measured on how well we reduce their fraud and credit losses."

For example, if a company is losing $50 million per year due to fraud and our models detect 25-50% of the fraudulent activity, that's a savings of tens of millions of dollars, says Zoldi.

Zoldi sees his current work as a direct extension of his graduate work.

"My research was kind of broad; it centered around complex problems with an abundance of data and looking for relationships in that information," says Zoldi.

In graduate school at Duke University, Zoldi used his computational skills to create algorithms that helped doctors determine if antidepressant medicine was working for the patient. The key was converting multiple brain wave readings into a useful pattern that could be compared before and after treatment.

Zoldi's training in computational and theoretical physics has served him well at Fair Isaac, where efficient use of computational time is essential.

"Within the company we have a large number of UNIX and mainframe computers, and the company builds machines from scratch specifically for training our neural networks," says Zoldi. "We build our own analytic tools and computational algorithms because when we are processing huge volumes of data, and we have such tight deadlines, it's really important that everything is done as efficiently as possible. If we write a bad code or don't think very carefully about tying our architecture to our program, it could take six years versus six weeks to solve a problem."

Zoldi says he enjoys the challenge of working in industry.

"There is a lot of innovation here," he says. "Our company employs well over 100 PhDs in modeling. Personally, I find it a great challenge. It's an environment where I can show the value and accuracy of my theories and my ideas. It is quite challenging and intellectually stimulating."

Howes Scholars

THE FREDERICK A. HOWES SCHOLAR

in Computational Science award was established to honor the late Frederick Anthony Howes, who managed the Applied Mathematical Science Program in the U.S. Department of Energy during the 1990s. Dr. Howes was highly respected and admired for his energy, dedication and personal integrity.

One of Howes' responsibilities was to oversee the Department of Energy's Computational Science Graduate Fellowship (DOE CSGF) program. He was extremely committed to this program, which supports graduate students in computational science. In fact, without his support, the program might not have survived.

To honor his memory and his dedication to the DOE CSGF program after his untimely death, one or two DOE CSGF fellows are chosen each calendar year as a Howes Scholar. Candidates are chosen based on their academic excellence, leadership and character and are nominated by their academic advisors. The honor provides the recipients with a substantial cash award, a Tiffany crystal paperweight, and the distinction of being named a Howes Scholar.

2003 Scholars

Two candidates were selected as Howes Scholars in 2003, Dr. Jon Wilkening of Courant Institute of Mathematical Sciences and Dr. Oliver Fringer of Stanford University. Both Dr. Wilkening and Dr. Fringer were DOE CSGF fellows between 1997 and 2001. Dr. Wilkening graduated from the University of California at Berkeley with a PhD in Mathematics, and Dr. Fringer graduated from Stanford University with a PhD in Civil and Environmental Engineering.

Dr. Fringer received his award at the annual DOE CSGF Fellows Conference held in July in Washington, D.C., where he also presented his research. Dr. Wilkening was in Australia at the time of the conference and was presented his award in October at the Courant Institute.

Margaret Wright of the Courant Institute (center) presents Oliver Fringer with his award at the 2003 DOE CSGF Fellows Conference. Also pictured are Barbara Helland (left) and James Corones (right) of the Krell Institute.

Peter Lax presents Jon Willkening with his Howes Scholar award at a reception at the Courant Institute.

For More Information

Contact Barbara Helland at helland@krellinst.org for more information regarding this award or the DOE CSGF program.

Alumni Directory

A

Asohan Amarasingham
Brown University
Cognitive Science
Fellowship Years: 1998-2002
Current Status: Student, Brown University

B

Allison Baker
University of Colorado
Applied Mathematics
Fellowship Years: 1999-2003
Current Status: Student, University of Colorado

Edward Barragy
University of Texas
Engineering Mechanics
Fellowship Years: 1991-1993
Current Status: Intel

William Barry
Carnegie Mellon University
Structural & Computational Engineering
Fellowship Years: 1994-1998
Current Status: Faculty, Asian Institute of Technology

Martin Bazant
Harvard University
Physics
Fellowship Years: 1992-1996
Current Status: Faculty, Massachusetts Institute of Technology

Edwin Blosch
University of Florida
Aerospace Engineering
Fellowship Years: 1991-1994
Current Status: CFD Research Corp

Dean Brederson
University of Utah
Computer Science
Fellowship Years: 1996-1998
Current Status: Student, University of Utah

Paul Bunch
Purdue University
Chemical Engineering
Fellowship Years: 1994-1997
Current Status: Eli Lilly & Company

Jeffery Butera
North Carolina State University
Mathematics
Fellowship Years: 1993-1997
Current Status: Staff, Hampshire College

C

Brandoch Calef
University of California – Berkeley
Applied Mathematics
Fellowship Years: 1996-2000
Current Status: Boeing

Patrick Canupp
Stanford University
Aerospace Engineering
Fellowship Years: 1991-1995
Current Status: Robert Yates Racing

Kent Carlson
Florida State University
Mechanical Engineering
Fellowship Years: 1991-1995
Current Status: Staff, University of Iowa

Bonnie Carpenter Cozad
University of Illinois
Mechanical Engineering
Fellowship Years: 1991-1995

Edward Chao
Princeton University
Plasma Physics
Fellowship Years: 1992-1995
Current Status: GE Medical Systems

Jarrod Chapman
University of California – Berkeley
Computational Biology
Fellowship Years: 1999-2003
Current Status: Student, University of California – Berkeley

Eric Charlton
University of Michigan
Aerospace Engineering
Fellowship Years: 1992-1996
Current Status: Lockheed Martin

Michael Chiu
Massachusetts Institute of Technology
Mechanical Engineering
Fellowship Years: 1992-1996
Current Status: Teradyne

Joshua Coe
University of Illinois
Chemical Physics
Fellowship Years: 2001-2002
Current Status: Student, University of Illinois

Ken Comer
North Carolina State University
Mechanical Engineering
Fellowship Years: 1991-1995
Current Status: Procter & Gamble

John Costello
University of Arizona
Applied Mathematics
Fellowship Years: 1998-2002
Current Status: Student, University of Arizona

Nathan Crane
University of Illinois
Civil Engineering
Fellowship Years: 1999-2002
Current Status: Sandia National Laboratories — New Mexico

Stephen Cronen-Townsend
Cornell University
Computational Materials Physics
Fellowship Years: 1991-1995
Current Status: Staff, University of Massachusetts

Robert Cruise
Indiana University
Physics
Fellowship Years: 1997-2001
Current Status: Staff, Indiana University

Joseph Czyzyk
Northwestern University
Industrial Engineering
Fellowship Years: 1991-1994

D

William Daughton
Massachusetts Institute of Technology
Plasma Physics
Fellowship Years: 1992-1996
Current Status: Los Alamos National Laboratory

Mark DiBattista
Columbia University
Computational Fluid Dynamics
Fellowship Years: 1992-1994

John Dolbow
Northwestern University
Theoretical and Applied Mechanics
Fellowship Years: 1997-1999
Current Status: Faculty, Duke University

Brian Dumont
University of Michigan
Aerospace Engineering
Fellowship Years: 1994
Current Status: Airflow Sciences Corporation

Amanda W. Duncan
University of Illinois
Electrical Engineering
Fellowship Years: 1991-1995
Current Status: Intel

Lewis Jonathan Dursi
University of Chicago
Astrophysics
Fellowship Years: 1999-2003
Current Status: Staff, University of Chicago

E

Thomas Epperly
University of Wisconsin – Madison
Chemical Engineering
Fellowship Years: 1991-1995
Current Status: Lawrence Livermore National Laboratory

F

Matthew Fago
California Institute of Technology
Aeronautical Engineering
Fellowship Years: 2000-2003
Current Status: Northrop Grumman

Michael Falk
University of California – Santa Barbara
Physics
Fellowship Years: 1995-1998
Current Status: Faculty, University of Michigan

Matthew Farthing
University of North Carolina
Environmental Science & Engineering
Fellowship Years: 1997-2001
Current Status: Staff, University of North Carolina

Michael Feldmann
California Institute of Technology
Computational Chemistry
Fellowship Years: 1999-2002
Current Status: Staff, California Institute of Technology

Stephen Fink
University of California – San Diego
Computer Science
Fellowship Years: 1994-1998
Current Status: IBM

Robert Fischer
Harvard University
Computer Science
Fellowship Years: 1994-1998
Current Status: Brigham & Women's Hospital

Gregory Ford
University of Illinois
Chemical Engineering
Fellowship Years: 1993-1995

Oliver Fringer
Stanford University
Environmental Fluid Mechanics
Fellowship Years: 1997-2001
Current Status: Faculty, Stanford University

G

Kenneth Gage
University of Pittsburgh
Chemical Engineering
Fellowship Years: 1998-2002
Current Status: Student, University of Pittsburgh

Nouvelle Gebhart
University of New Mexico
Chemistry
Fellowship Years: 2001-2003

Charles Gerlach
Northwestern University
Mechanical Engineering
Fellowship Years: 1995-1999
Current Status: Network Computing Services

Timothy Germann
Harvard University
Physical Chemistry
Fellowship Years: 1992-1995
Current Status: Los Alamos National Laboratory

Christopher Gesh
Texas A&M University
Nuclear Engineering
Fellowship Years: 1993-1997
Current Status: Pacific Northwest National Laboratory

Matthew Giamporcaro
Boston University
Cognitive and Neural Systems
Fellowship Years: 1998-2000

Kevin Glass
University of Oregon
Computer Science
Fellowship Years: 1996-2000
Current Status: Staff, University of Oregon

Larisa Goldmints
Carnegie Mellon University
Structural Mechanics
Fellowship Years: 1997-2001

William Gooding
Purdue University
Chemical Engineering
Fellowship Years: 1991-1994

Corey Graves
North Carolina State University
Computer Engineering
Fellowship Years: 1996-1999
Current Status: Faculty, North Carolina Agricultural & Technical State University

Noel Gres
University of Illinois
Electrical Engineering
Fellowship Years: 1999-2001

Eric Grimme
University of Illinois
Electrical Engineering
Fellowship Years: 1994-1997
Current Status: Intel

John Guidi
University of Maryland
Computer Science
Fellowship Years: 1994-1997
Current Status: Math High School Teacher

H

Aric Hagberg
University of Arizona
Applied Mathematics
Fellowship Years: 1992-1994
Current Status: Los Alamos National Laboratory

Glenn Hammond
University of Illinois
Environmental Engineering & Science
Fellowship Years: 1999-2003
Current Status: Sandia National Laboratories – New Mexico

Jeffrey Haney
Texas A&M University
Physical Oceanography
Fellowship Years: 1993-1996
Current Status: Dynacon, Inc.

Rellen Hardtke
University of Wisconsin – Madison
Physics
Fellowship Years: 1998-2002
Current Status: Faculty, California State Polytechnic University

Eric Held
University of Wisconsin – Madison
Engineering Physics
Fellowship Years: 1995-1999
Current Status: Faculty, Utah State University

Judith Hill
Carnegie Mellon University
Mechanics, Algorithms & Computing
Fellowship Years: 1999-2003
Current Status: Student, Carnegie Mellon University

Charles Hindman
University of Colorado
Aerospace Engineering
Fellowship Years: 1999-2003
Current Status: Air Force Research Laboratory

Jeffrey Hittinger
University of Michigan
Aerospace Engineering & Scientific Computing
Fellowship Years: 1996-2000
Current Status: Lawrence Livermore National Laboratory

Gordon Hogenson
University of Washington
Physical Chemistry
Fellowship Years: 1993-1996
Current Status: Microsoft

William Humphrey
University of Illinois
Physics
Fellowship Years: 1992-1994
Current Status: TurboLabs Inc.

Jason Hunt
University of Michigan
Aerospace Engineering & Scientific Computing
Fellowship Years: 1999-2003
Current Status: Student, University of Michigan

E. McKay Hyde
California Institute of Technology
Applied & Computational Mathematics
Fellowship Years: 1999-2002
Current Status: University of Minnesota

I

Eugene Ingerman
University of California – Berkeley
Applied Mathematics
Fellowship Years: 1997-2001
Current Status: Staff, University of California – Davis

J

Nickolas Jovanovic
Yale University
Mechanical Engineering
Fellowship Years: 1992-1994
Current Status: Faculty, University of Arkansas – Little Rock

K

Jeremy Kepner
Princeton University
Computational Cosmology
Fellowship Years: 1993-1996
Current Status: Staff, Massachusetts Institute of Technology

Sven Khatri
California Institute of Technology
Electrical Engineering
Fellowship Years: 1993-1996
Current Status: VocalPoint Technology

Yury Krongauz
Northwestern University
Theoretical & Applied Mechanics
Fellowship Years: 1993-1996

L

Eric Lee
Rutgers University
Mechanical Engineering
Fellowship Years: 1999-2003
Current Status: Student, Rutgers University

Jack Lemmon
Georgia Institute of Technology
Mechanical Engineering
Fellowship Years: 1991-1994
Current Status: Medtronic, Inc.

Lars Liden
Boston University
Cognitive & Neural Systems
Fellowship Years: 1994-1998
Current Status: Staff, University of Washington

Tasha (Palmer) Lopez
University of California – Los Angeles
Chemical Engineering
Fellowship Years: 2000-2001
Current Status: IBM

Christie Lundy
University of Missouri – Rolla
Physics
Fellowship Years: 1991-1994
Current Status: State of Missouri Employee

M

William Marganski
Boston University
Biomedical Engineering
Fellowship Status: 1998-2002
Current Status: Boston Biomedical Research Institute

Daniel Martin
University of California – Berkeley
Mechanical Engineering
Fellowship Years: 1993-1996
Current Status: Lawrence Berkeley National Laboratory

Marcus Martin
University of Minnesota
Physical Chemistry
Fellowship Years: 1997-1999
Current Status: Sandia National Laboratories – New Mexico

Richard McLaughlin
Princeton University
Applied Mathematics
Fellowship Years: 1991-1994
Current Status: Faculty, University of North Carolina

Lisa Mesaros
University of Michigan
Aerospace Engineering & Scientific Computing
Fellowship Years: 1991-1995
Current Status: FLUENT, Inc.

Erik Monsen
Stanford University
Aerospace and Astronautical Engineering
Fellowship Years: 1991-1994
Current Status: Student, University of Colorado

Brian Moore
North Carolina State University
Nuclear Engineering
Fellowship Years: 1992-1995
Current Status: Global Nuclear Fuels

James (Dan) Morrow
Carnegie Mellon University
Robotics
Fellowship Years: 1992-1995
Current Status: Sandia National Laboratories – New Mexico

Michael Mysinger
Stanford University
Chemical Engineering
Fellowship Years: 1996-2000
Current Status: Arqule

N

Pauline Ng
University of Washington
Bioengineering
Fellowship Years: 2000-2002
Current Status: Illumina

Brian Nguyen Gunney
University of Michigan
Aerospace Engineering & Scientific Computing
Fellowship Years: 1993-1996
Current Status: Lawrence Livermore National Laboratory

Diem-Phuong Nguyen
University of Utah
Chemical Engineering
Fellowship Years: 1999-2003
Current Status: Student, University of Utah

Debra E. Nielsen
Colorado State University
Civil Engineering
Fellowship Years: 1992-1996

Joyce Noah
Stanford University
Theoretical Chemisty
Fellowship Years: 2001-2003
Current Status: Student, Stanford University

O

Christopher Oehmen
University of Memphis
Biomedical Engineering
Fellowship Years: 1999-2003
Current Status: Pacific Northwest National Laboratory

P

Steven Parker
University of Utah
Computer Science
Fellowship Years: 1994-1997
Current Status: Faculty, University of Utah

Joel Parriott
University of Michigan
Astronomy & Astrophysics
Fellowship Years: 1992-1996
Current Status: Office of Management and Budget

Virginia Pasour
North Carolina State University
Biomathematics
Fellowship Years: 1998-1999
Current Status: Student, Cornell University

Robert (Chris) Penland
Duke University
Biomedical Engineering
Fellowship Years: 1993-1997
Current Status: Physiome Science, Inc.

James Phillips
University of Illinois
Physics
Fellowship Years: 1995-1999
Current Status: Staff, University of Illinois

Todd Postma
University of California – Berkeley
Nuclear Engineering
Fellowship Years: 1994-1998
Current Status: Totality

Richard Propp
University of California – Berkeley
Mechanical Engineering
Fellowship Years: 1993-1996
Current Status: Real Time Solutions

Q

Alejandro Quezada
University of California – Berkeley
Geophysics
Fellowship Years: 1997-1998

R

Nathan Rau
University of Illinois
Civil Engineering
Fellowship Years: 2000-2001
Current Status: Hanson Professional Services

Clifton Richardson
Cornell University
Physics
Fellowship Years: 1991-1995

John Rittner
Northwestern University
Mechanical Engineering
Fellowship Years: 1991-1995
Current Status: Chicago Board Options Exchange

Courtney Roby
University of Colorado
Electrical Engineering
Fellowship Years: 2002-2003
Current Status: Student, University of Colorado

David Ropp
University of Arizona
Applied Mathematics
Fellowship Years: 1992-1995
Current Status: Sandia National Laboratories – New Mexico

Robin Rosenfeld
Scripps Research Institute
Biology
Fellowship Years: 1996-1997
Current Status: Staff, Scripps Research Institute

S

Robert Sedgewick
University of California – Santa Barbara
Physics
Fellowship Years: 2000-2003
Current Status: Staff, University of Pittsburgh

Susanne (Essig) Seefried
Massachusetts Institute of Technology
Aeronautics/Astronautics
Fellowship Years: 1997-2002
Current Status: Student, Massachusetts Institute of Technology

Marc Serre
University of North Carolina
Environmental Science & Engineering
Fellowship Years: 1996-1999
Current Status: Faculty, University of North Carolina

Elsie Simpson Pierce
University of Illinois
Nuclear Engineering
Fellowship Years: 1991-1993
Current Status: Lawrence Livermore National Laboratory

Melinda Sirman
University of Texas
Engineering Mechanics
Fellowship Years: 1994-1996

Steven Smith
North Carolina State University
Chemical Engineering
Fellowship Years: 1992-1994
Current Status: E.I. DuPont

Scott Stanley
University of California – San Diego
Mechanical Engineering
Fellowship Years: 1994-1998
Current Status: Persist Technologies, Inc.

James Strzelec
Stanford University
Computational Mathematics
Fellowship Years: 1992-1994

Rajeev Surati
Massachusetts Institute of Technology
Electrical Engineering & Computer Science
Fellowship Years: 1995-1997
Current Status: Nexaweb

Laura (Painton) Swiler
Carnegie Mellon University
Engineering & Public Policy
Fellowship Years: 1992-1995
Current Status: Sandia National Laboratories – New Mexico

T

Shilpa Talwar
Stanford University
Scientific Computing
Fellowship Years: 1992-1994
Current Status: Sandia National Laboratories – California

Mayya Tokman
California Institute of Technology
Applied Mathematics
Fellowship Years: 1996-2000
Current Status: Faculty, University of California – Berkeley

Mario Trujillo
University of Illinois
Mechanical Engineering
Fellowship Years: 1997-2000
Current Status: Staff, Pennsylvania State University

V

Anton Van Der Ven
Massachusetts Institute of Technology
Materials Science
Fellowship Years: 1996-2000
Current Status: Staff, Massachusetts Institute of Technology

Laura Vann Dominik
Florida Atlantic University
Electrical Engineering
Fellowship Years: 1993-1997
Current Status: Pratt & Whitney

Rajesh Venkataramani
Massachusetts Institute of Technology
Chemical Engineering
Fellowship Years: 1995-1999
Current Status: Goldman Sachs

Stephen Vinay
Carnegie Mellon University
Chemical Engineering
Fellowship Years: 1998-2000
Current Status: Bettis Laboratory

W

Phillip Weeber
University of North Carolina
Environmental Science & Engineering
Fellowship Years: 1994-1996

Adam Weller
Princeton University
Chemical Engineering
Fellowship Years: 2001-2002
Current Status: Student, Princeton University

Gregory Whiffen
Cornell University
Environmental Systems Engineering
Fellowship Years: 1991-1995
Current Status: NASA - Jet Propulsion Laboratory

Collin Wick
University of Minnesota
Computational Chemistry
Fellowship Years: 2000-2003
Current Status: National Technical University Athens, Greece

James Wiggs
University of Washington
Physical Chemistry
Fellowship Years: 1991-1994
Current Status: Novum

Jon Wilkening
University of California – Berkeley
Applied Mathematics
Fellowship Years: 1997-2001
Current Status: Staff, Courant Institute

Glenn Williams
University of North Carolina
Environmental Science & Engineering
Fellowship Years: 1993-1996
Current Status: Faculty, Old Dominion University

C. Eric Williford
Florida State University
Meteorology
Fellowship Years: 1993-1996

Lee Worden
Princeton University
Applied Mathematics
Fellowship Years: 1998-2002
Current Status: Staff, University of California – Davis

Peter Wyckoff
Massachusetts Institute of Technology
Chemical Engineering
Fellowship Years: 1992-1995
Current Status: Ohio Supercomputing Center

Z

Charles Zeeb
Colorado State University
Mechanical Engineering
Fellowship Years: 1993-1997
Current Status: Colorado State University

Scott Zoldi
Duke University
Theoretical & Computational Physics
Fellowship Years: 1996-1998
Current Status: Fair Issac Corporation

FELLOWS DIRECTORY

Matthew Anderson
University of Texas
Physics

Advisor:
Richard Matzner
Practicum:
Los Alamos National Laboratory
Contact:
astro@einstein.ph.utexas.edu
Notable:
Recipient of the David Bruton, Jr. Fellowship for 2001-2002
Research Synopsis:
Using the CVODE integrator from the SUNDIALS suite (Lawrence Livermore National Laboratory — LLNL) and the elliptic solving capabilities of PETSc (Argonne National Laboratory), I am exploring the long term numerical stability of isolated black hole simulations by performing "constrained evolutions", simulations which solve the constraint equations at each timestep of an evolution. In order to simulate black hole mergers on large computational domains, I have implemented the constrained evolution approach using SAMRAI (LLNL) for adaptive mesh refinement. Using these three packages together as a framework on which to solve the Einstein equations with hundreds of processors has provided a flexible and powerful way to research constrained evolution in numerical relativity.

Devin Balkcom
Carnegie Mellon University
Robotics

Advisor:
Matt Mason
Practicum:
Sandia National Laboratories – New Mexico
Contact:
devin@ri.cmu.edu
Research Synopsis:
I'm interested in simple models of the physical systems that arise in robotic locomotion and manipulation. I'm currently working on the problem of robotic origami folding. Origami folding is a good challenge problem for robotics because it requires complicated manipulation skills that are not well understood; for example, the manipulation of flexible objects, planning for closed chains, and planning for systems with a large number of degrees of freedom. My thesis focuses on a 'rigid body' model of origami. If there are crossing creases in the origami design, the 'origami mechanism' has the structure of a (typically compound) closed chain. The space of configurations for mechanisms of this type is quite complicated, and is usually not even a manifold.

Gavin Conant
University of New Mexico
Biology

Advisor:
Andreas Wagner
Practicum:
Sandia National Laboratories – New Mexico
Contact:
gconant@unm.edu
Research Synopsis:
"Why are duplicate genes so common in the genomes of eukaryotes?"; "How do protein interactions and transcription binding sites evolve?"; "Do duplicate genes often diverge in function?" The above questions indicate how limited our understanding of the process of converting information stored in our genes into the enzymes and structures necessary to maintain life actually is. Using sequence analysis tools and measurements of molecular divergence, our lab has studied whether duplicate genes often show asymmetric amino acid divergence, whether the circuits often seen in transcriptional regulatory networks are evolved features, and whether duplicate genes protect the nematode C. elegans from loss of function in its genes.

Ryan Elliott
University of Michigan
Aerospace Engineering

Advisor:
Nicolas Triantafyllidis
Practicum:
Los Alamos National Laboratory
Contact:
elliottr@umich.edu
Research Synopsis:
My research will advance current knowledge of the relationship between the nano-scale behavior and the micro-scale behavior of certain inter-metallic alloys, such as NiTi. Modeling the behavior and stability of single crystal martensitic transformations is the goal of my dissertation research. A temperature dependent atomic pair-potential model is used to represent the atomic interactions. The continuum energy density W(F,s;T) (where F is the deformation gradient, s is a set of internal atomic degrees of freedom, and T is the temperature) is derived from the consideration of a regular lattice of atoms constituting the atomic structure of the material. This formulation leads to a set of highly non-linear equations to determine the equilibrium configurations of the crystal as well as the stability of those configurations.

Catherine Grasso
Cornell University
Bioinformatics

Advisor:
Richard Durrett
Practicum:
Lawrence Berkeley National Laboratory
Contact:
cgrasso@mbi.ucla.edu
Research Synopsis:
We have developed a new data structure for multiple sequence alignment, which is a directed acyclic graph, called the partial order multiple sequence alignment (PO-MSA) representation, which accurately represents the true biological content of a multiple sequence alignment. We have also developed POA, an algorithm that directly aligns PO-MSAs to each other. This data structure and algorithm, which is an order of magnitude faster than other methods, have been successfully employed in EST sequence analysis in order to build databases of both splice sites and snps in the human genome. We are currently working on a version of POA designed for protein sequences. Our alignments, which will reflect the cut and paste events that constructed the multi-domain protein families, will provide us with data that will be the basis for re-formalizing the phylogenetics of protein families.

Boyce Griffith
New York University – Courant Institute
Applied Mathematics

Advisor:
Charles Peskin
Practicum:
Lawrence Livermore National Laboratory
Contact:
griffith@cims.nyu.edu
Research Synopsis:
I am interested in the development of numerical methods and computational tools required to simulate a three dimensional heart model that includes realistic electrophysiology, muscle mechanics, and blood-tissue interaction. The goal of my research is to develop the numerical methods and computational software required to provide a platform where excitation-contraction coupling can be examined in the context of the beating heart. One approach to modeling blood-tissue interaction is the immersed boundary (IB) method. I am developing a new parallel implementation of the IB method that aims to incorporate techniques such as adaptive mesh refinement and inexact Newton methods to improve the efficiency and accuracy of the computations. Additionally, I am involved in work aiming to extend the IB methodology to model the electrical activity of the heart. Ultimately, we intend that these two projects will be merged into a coupled electro-mechanical-fluidic model of the heart.

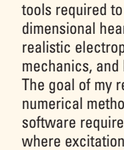

Ahmed Ismail
Massachusetts Institute of Technology
Chemical Engineering

Advisor:
Gregory Rutledge
Practicum:
Oak Ridge National Laboratory
Contact:
aei@mit.edu
Research Synopsis:
My research represents an initial step towards the modeling of physical systems using multiresolution. Wavelet transforms of the a lattice model, along with its Hamiltonian and its configuration space, allow for the development of a recursive, adaptive algorithm which can produce thermodynamic property information at any desired length scale. We have developed a new modeling paradigm which successively coarse-grains objects into "block" objects with properties that are given by the probability distributions of the mean of the objects that created them: spins yield block spins, while monomer "beads" of a polymer random walk yield coarse-grained "blobs." These block objects can be analyzed much more efficiently than the original systems, and show comparable scaling behavior, allowing us to simulate a range of behavior previously inaccessible to numerical simulation.

Daniel Horner
University of California – Berkeley
Chemistry

Advisor:
C. William McCurdy
Practicum:
Argonne National Laboratory
Contact:
dahorner@lbl.gov
Notable:
Invited to and attended the 2002 Noble Laureate Conference in Lindau, Germany
Research Synopsis:
My research focus is in electron scattering theory. Using both high performance machines as well as analytical methods, our goal is to be able to solve more, and more complex, scattering systems. Many approximate theories have been developed to ease the computational complexity of scattering calculations. However, many of these methods do not provide adequate accuracy. We are engaged in developing methods that make the exact formalism (solutions of the Schrödinger equation with appropriate scattering boundary conditions) more computationally feasible.

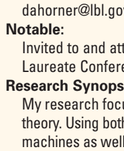

Benjamin Keen
University of Michigan
Mathematics

Advisor:
Smadar Karni
Practicum:
Lawrence Berkeley National Laboratory
Contact:
bkeen@umich.edu
Research Synopsis:
I'm interested in using level sets to achieve hard moving embedded boundaries in unsteady compressible flow calculations. Level sets are a technique for tracking material interfaces in flow calculations. The basic idea is to initialize a function on the domain whose zero level set denotes the boundary, and append an extra equation to the system of PDEs that describes the flow that advects this function with the flow. Then, the information provided by the level set can be fed back into the simulation, e.g. by using a different equation of state in the two regions.

Justin Koo
University of Michigan
Aerospace Engineering

Advisor:
Iain Boyd
Practicum:
Lawrence Livermore National Laboratory
Contact:
kooj@engin.umich.edu
Research Synopsis:
The present focus of my research is the simulation of the plasma dynamics inside operating Hall thrusters. Existing 1D and 2D computer models of the Hall thruster acceleration process based on a hybrid particle-fluid approach do not accurately describe the operating characteristics of real devices. From experimental measurements, it is known that the electron energy distribution function in Hall thrusters is not Maxwellian. This has a profound effect on both the rate of ionization and the electron mobility inside the acceleration channel. The research plan for my thesis is to develop a hierarchy of electron energy distribution models within the framework of a 2-D hybrid PIC-MCC Hall thruster code. These models will incorporate progressively more detailed physics into the energy transport equation and should provide significantly more accurate simulation results.

Michael Kowalok
University of Wisconsin
Medical Physics

Advisor:
Douglass Henderson
Practicum:
Oak Ridge National Laboratory
Contact:
mkowalok@students.wisc.edu
Research Synopsis:
My PhD work is centered on adapting radiation transport codes and computer optimization techniques for inverse radiation therapy treatment planning. The inverse process begins with a desired dose distribution and works backwards to determine an optimized configuration of radiation beams that will deliver that distribution. A main component of this work has been an investigation of adjoint Monte Carlo transport and the development of adjoint analytic methods to support the inverse treatment planning process.

Heather Netzloff
Iowa State University
Physical Chemistry

Advisor:
Mark Gordon
Practicum:
Lawrence Berkeley National Laboratory
Contact:
netzloff@iastate.edu
Research Synopsis:
My current research projects cover a wide variety of areas in quantum/computational chemistry. Since most biological and experimental processes occur in solution (versus gas), the study of solvation effects on the mechanisms of chemical reactions is very important. We approach solvation in terms of cluster studies and the Effective Fragment Potential (EFP) method for solvation. Currently, we are interfacing the EFP method with quantum chemistry codes, state averaged MCSCF and CI, as well as sampling techniques and molecular dynamics, which will allow solvent effects to be studied more easily. This will ultimately provide an opportunity to go from the study of relatively small clusters to intermediate clusters to bulk and supercritical properties of fluids.

Catherine Norman
Northwestern University
Applied Mathematics

Advisor:
Michael Miksis
Practicum:
Lawrence Berkeley National Laboratory
Contact:
c-norman@nwu.edu
Research Synopsis:
My research involves modeling flows which require the tracking of interfaces between fluids in various phases. Examples of such flows include predicting how an oil spill will spread along the ocean surface or following a gas bubble as it rises through a liquid. Current models of such phenomena have difficulties with the conservation of mass, tracking interfaces that break apart and reconnect, and determining the curvature of the interfaces. Many numerical methods also model an interface as a sharp line between two fluid phases, which is not physically accurate. There is actually a small region where the two phases coexist. I am particularly interested in developing efficient numerical algorithms for studying multiphase flows when three phases are involved (e.g. oil/water/air).

THIRD YEAR FELLOWS

Marcelo Alvarez
University of Texas
Computational Astrophysics
Advisor:
Paul Shapiro
Practicum:
Los Alamos National Laboratory
Contact:
marcelo@astro.as.utexas.edu

Kristopher Andersen
University of California – Davis
Physics
Advisor:
Warren Pickett
Practicum:
Oak Ridge National Laboratory
Contact:
keandersen@ucdavis.edu

Nathan Carstens
Massachusetts Institute of Technology
Nuclear Engineering
Advisor:
Ronald Ballinger
Practicum:
Los Alamos National Laboratory
Contact:
nate_carstens@yahoo.com

Annette Evangelisti
University of New Mexico
Computational Molecular Biology
Advisor:
Andreas Wagner
Practicum:
Los Alamos National Laboratory
Contact:
amevang@unm.edu

Sommer Gentry
Massachusetts Institute of Technology
Optimization/Control Theory
Advisor:
Eric Feron
Practicum:
Sandia National Laboratories – New Mexico
Contact:
sommerg@mit.edu
Notable:
Winner of IEEE Systems Man, Cybernetics Conference 2003 Best Student Paper award.

Ahna Girshick
University of California – Berkeley
Vision Science
Advisor:
Martin Banks
Practicum:
Lawrence Berkeley National Laboratory
Contact:
ahna@uclink.berkeley.edu

Kristen Grauman
Massachusetts Institute of Technology
Computer Science
Advisor:
Trevor Darrell
Practicum:
Lawrence Berkeley National Laboratory
Contact:
kgrauman@mit.edu
Notable:
Received the Boston University Computer Science department's annual award for the outstanding senior in the graduating class. Received the Albert McGuinn award from the Boston College, College of Arts and Sciences in 2001. Awarded Emerson Music Scholarship from the MIT Music department to support private piano study.

Heath Hanshaw
University of Michigan
Nuclear Engineering
Advisor:
Edward Larsen
Practicum:
Los Alamos National Laboratory
Contact:
heathhanshaw@usa.net
Notable:
Received a Navy Commendation Medal for accomplishments while serving as faculty of the US Naval Academy Physics Department. Awarded the American Nuclear Society Scholarship in 2001-2003.

Richard Katz
Columbia University
Geodynamics
Advisor:
Marc Spiegelman
Practicum:
Argonne National Laboratory
Contact:
katz@ldeo.columbia.edu

Benjamin Kirk
University of Texas
Aerospace Engineering
Advisor:
Graham Carey
Practicum:
Sandia National Laboratories – New Mexico
Contact:
benkirk@mail.utexas.edu

Seung Lee
Massachusetts Institute of Technology
Mechanical Engineering
Advisor:
Roger Kamm
Practicum:
Pacific Northwest National Laboratory
Contact:
selee@mit.edu
Notable:
Selected as a Gates Millennium Scholar scholarship recipient.

Mary Ann Leung
University of Washington
Theoretical Physical Chemistry
Advisor:
William Reinhardt
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mleung@u.washington.edu
Notable:
Awarded the Klaus A. and Mary Ann D. Saegebarth Graduate Student Fellowship by the University of Washington.

Randall McDermott
University of Utah
Chemical Engineering
Advisor:
Philip Smith
Practicum:
Sandia National Laboratories – California
Contact:
randy@crsim.utah.edu
Notable:
Invited to and attended the 2002 Noble Laureate Conference in Lindau, Germany. Named as a John Zink fellow for outstanding work in combustion.

Matthew McNenly
University of Michigan
Aerospace Engineering
Advisor:
Iain Boyd
Practicum:
Sandia National Laboratories – New Mexico
Contact:
mcnenly@engin.umich.edu

Richard Mills
College of William and Mary
Computer Science
Advisor:
Andreas Stathopoulos
Practicum:
Los Alamos National Laboratory
Contact:
rtm@cs.wm.edu

Julian Mintseris
Boston University
Bioinformatics
Advisor:
Zhiping Weng
Practicum:
Lawrence Berkeley National Laboratory
Contact:
julianm@bu.edu

Elijah Newren
University of Utah
Mathematics
Advisor:
Aaron Fogelson
Practicum:
Lawrence Livermore National Laboratory
Contact:
newren@math.utah.edu

Christopher Rinderspacher
University of Georgia
Chemistry
Advisor:
Henry Shaefer
Practicum:
Sandia National Laboratories – California
Contact:
crinders@chem.uga.edu

Samuel Schofield
University of Arizona
Applied Mathematics
Advisor:
Mary Poulton
Practicum:
Argonne National Laboratory
Contact:
sschofie@u.arizona.edu

Matthew Wolinsky
Duke University
Geomorphology
Advisor:
Lincoln Pratson
Practicum:
Oak Ridge National Laboratory
Contact:
maw@duke.edu

SECOND YEAR FELLOWS

Bree Aldridge
Massachusetts Institute of Technology
Computational Biology
Advisor:
Douglas Lauffenburger
Contact:
breea@mit.edu

Teresa Bailey
Texas A&M University
Engineering
Advisor:
Marvin Adams
Contact:
baileyte@tamu.edu

Michael Barad
University of California – Davis
Environmental Modeling
Advisor:
Geoffrey Schladow
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mfbarad@ucdavis.edu

Jaydeep Bardhan
Massachusetts Institute of Technology
Electrical Engineering
Advisor:
Jacob White
Practicum:
Argonne National Laboratory
Contact:
jbardhan@mit.edu

Mary Biddy
University of Wisconsin
Engineering
Advisor:
Juan de Pablo
Contact:
mbiddy@students.wisc.edu

Nawaf Bou-Rabee
California Institute of Technology
Applied and Computational Mathematics
Advisor:
Jerrold Marsden
Contact:
nawaf@acm.caltech.edu

Kevin Chu
Massachusetts Institute of Technology
Applied Mathematics
Advisor:
Martin Bazant
Contact:
kchu@math.mit.edu

Kristine Cochran
University of Illinois – Urbana – Champaign
Structures
Advisor:
Keith Hjelmstad
Practicum:
Sandia National Laboratories – New Mexico
Contact:
kbergero@uiuc.edu

Gregory Davidson
Oregon State University
Nuclear Engineering
Advisor:
Todd Palmer
Practicum:
Bettis Atomic Power Laboratory
Contact:
davidsgr@engr.orst.edu

Michael Driscoll
Boston University
Bioinformatics
Advisor:
James Collins
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mdriscol@bu.edu

Mary Dunlop
California Institute of Technology
Mechanical Engineering
Advisor:
Richard Murray
Contact:
mjdunlop@caltech.edu

Michael Greminger
University of Minnesota
Mechanical Engineering
Advisor:
Bradley Nelson
Practicum:
Sandia National Laboratories – New Mexico
Contact:
grem@me.umn.edu

Owen Hehmeyer
Princeton University
Chemical Engineering
Advisor:
Athanassios Z. Panagiotopoulos
Contact:
hehmeyer@princeton.edu

Yan Karklin
Carnegie Mellon University
Computational Neuroscience
Advisor:
Michael Lewicki
Practicum:
Lawrence Berkeley National Laboratory
Contact:
yan+fellowship@cs.cmu.edu

Benjamin Lewis
Massachusetts Institute of Technology
Computational Biology
Advisor:
Chris Burge
Contact:
benlewis@mit.edu

Alex Lindblad
University of Washington
Structural Engineering
Advisor:
George Turkiyyah
Contact:
alind@u.washington.edu

Nathaniel Morgan
Georgia Institute of Technology
Mechanical Engineering
Advisor:
Marc Smith
Practicum:
Los Alamos National Laboratory
Contact:
n_r_morgan@yahoo.com

Gregory Novak
University of California – Santa Cruz
Theoretical Astrophysics
Advisor:
Sandra Faber
Practicum:
Lawrence Berkeley National Laboratory
Contact:
novak@ucolick.org

David Schmidt
University of Illinois – Urbana – Champaign
Communications and Signal Processing
Advisor:
Bruce Hajek
Practicum:
Sandia National Laboratories – New Mexico
Contact:
dschmidt@uiuc.edu

Amoolya Singh
University of California – Berkeley
Computational Biology
Advisor:
Richard Karp
Contact:
agni@cs.berkeley.edu

Eric Sorin
Stanford University
Chemical Physics
Advisor:
Vijay Pande
Contact:
esorin@stanford.edu

Obioma Uche
Princeton University
Materials/Statistical Mechanics
Advisor:
Salvatore Torquato
Contact:
ouuche@yahoo.com

Joshua Waterfall
Cornell University
Biophysics
Advisor:
James Sethna
Practicum:
Lawrence Berkeley National Laboratory
Contact:
jjw36@cornell.edu

Michael Wu
University of California – Berkeley
Computational Neuroscience
Advisor:
Jack Gallant
Practicum:
Los Alamos National Laboratory
Contact:
wafting@uclink4.berkeley.edu

Paul Bauman
University of Texas
Computational and Applied Mathematics
Advisor:
J. Tinsley Oden
Contact:
pbauman@ices.utexas.edu

William Conley
Purdue University
Nanoscale Mechanics
Advisor:
Arvind Raman
Contact:
wconley@ecn.purdue.edu

Aron Cummings
Washington State University
Electrical Engineering
Advisor:
Mohamed Osman
Contact:
Aron@turbonet.com

Krzysztof Fidkowski
Massachusetts Institute of Technology
Computational Fluid Dynamics
Advisor:
David Darmofal
Contact:
kfid@mit.edu

Sarah Moussa
University of California – Berkeley
Applied Science and Technology
Advisor:
Michael Jordan
Contact:
smoussa@uclink.berkeley.edu

Tod Pascal
California Institute of Technology
Physical Chemistry
Advisor:
William Goddard III
Contact:
tpascal@wag.caltech.edu

Emma Rainey
California Institute of Technology
Theoretical Geophysics
Advisor:
David Stevenson
Contact:
emma@gps.caltech.edu

Mark Rudner
Massachusetts Institute of Technology
Physics
Advisor:
Wolfgang Ketterle
Contact:
rudner@mit.edu

Jason Sese
Stanford University
Chemical Engineering
Advisor:
Charles Musgrave
Contact:
jasese@stanford.edu

Christina Smith
Vanderbilt University
Chemical Engineering
Advisor:
Peter Cummings
Contact:
Christina.m.smith@vanderbilt.edu

Samuel Stechmann
New York University
Applied Mathematics
Advisor:
Marsha Berger
Contact:
stechman@cims.nyu.edu

Brian Taylor
University of Illinois – Urbana – Champaign
Detonation and Shock Dynamics
Advisor:
Scott Stewart
Contact:
bdtaylo1@uiuc.edu

William Triffo
Rice University
Bioengineering
Advisor:
Robert Raphael
Contact:
triffo@rice.edu

Michael Wolf
University of Illinois – Urbana – Champaign
Computer Science
Advisor:
Michael Heath
Contact:
mwolf@uiuc.edu

Brandon Wood
Massachusetts Institute of Technology
Computational Materials Science
Advisor:
Nicola Marzari
Contact:
brandonw@mit.edu